

Gatling: Rapid-Fire Consensus from Parallel Composition

Giulia Scaffino ✉ 

TU Wien, Austria
Common Prefix, Vienna, Austria

Max Resnick ✉ 

Anza, New York, NY, USA

Joachim Neu ✉ 

a16z Crypto Research, New York, NY, USA

Abstract

Consensus protocols form the core of blockchains and other replicated state machines, ensuring that all correct nodes process the same totally ordered log of input transactions. In fault-free executions, performance is driven by the good-case transaction latency – the time between a transaction becoming known to all nodes and its confirmation by the consensus protocol – which depends on both how frequently proposals are made and, once made, how quickly they are confirmed. While prior work has established tight lower bounds on confirmation latency that modern protocols already achieve, it remains open whether the inter-proposal time can be further reduced below the state-of-the-art of one network delay.

We introduce Gatling, an atomic broadcast protocol that achieves arbitrarily small inter-proposal times under rotating leader schedules; in particular, smaller than the network delay. Gatling runs multiple parallel instances of a black-box atomic broadcast protocol and staggers their proposal schedules to generate proposals in faster succession than state-of-the-art protocols. A deterministic interleaving rule merges the outputs of these instances into a single global log. We analyze the effects of head-of-line blocking caused by crashed leaders, and derive Gatling’s optimal number of parallel instances. We further study the impact of Gatling on predictable validity and present two variants that retain this property. Finally, our experiments confirm that Gatling can be used with off-the-shelf component protocols to achieve low latency without fine-tuning the component protocol for minimum latency.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases Distributed Systems, Byzantine Fault Tolerance Consensus, Parallel Composition, Inter-Proposal Time

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.4

Related Version *Extended Version*: <https://arxiv.org/abs/2606.18220> [40]

Supplementary Material *Software (Source Code)*: <https://github.com/scaffino/Gatling>
archived at `swh:1:dir:316ac6dbf49dd46e0a39e88b31573b63a9068dcb`

1 Introduction

After the launch of the Bitcoin network in 2009, block times fell exponentially for fifteen years. Bitcoin 2009: 10 minutes. Ethereum 2015: 15 seconds. Solana 2020: 400 ms [35, 11, 51]. However, more than five years later, blockchains still launch with almost the same block times Solana had in 2020. Why did fifteen years of exponential progress, which promised to drive block times to zero, suddenly stall in 2020?

The reason is structural. All of the blockchains listed above proceed one block at a time. Each block references its parent, and so before consensus can begin on the next block, the previous block must be known. Roughly speaking, this bounds block times to the time it takes



© Giulia Scaffino, Max Resnick, and Joachim Neu;
licensed under Creative Commons License CC-BY 4.0
8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyrpoulou; Article No. 4; pp. 4:1–4:24



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

for a single round of consensus – and that time has tight, well-known lower bounds. Prior work [2] proved a 3Δ confirmation bound under partial synchrony in the $3f + 1 \leq n < 5f - 1$ regime, and a smaller bound holds at $n \geq 5f + 1$, achieved by Minimmit [18] at the cost of resilience. The remaining knob is the inter-proposal time. Pipelining (e.g., in Moonshot [22] and Hydrangea++ [44]) pushed it to Δ . Hydrangea++ asserts, without proof, that Δ is the floor. Intuitively, they argue that the next leader cannot propose a block without seeing the previous block, and that takes at least Δ time. This argument is of course correct; however, in this paper we challenge the assumption that the next leader must see the previous block before proposing their own.

We present Gatling, an atomic broadcast protocol whose inter-proposal time can be made arbitrarily small, in particular smaller than the network delay Δ . Gatling runs K parallel instances of a closed-box atomic broadcast protocol with inter-proposal time T_{ipt} , staggered by T_{ipt}/K , and interleaves their outputs into a single log via a deterministic merge rule. Each component instance still respects the 3Δ confirmation bound; the speedup comes from composition. Setting K large drives the inter-proposal time toward zero. As a proof of concept, we instantiate Gatling on top of a Simplex [15] implementation and obtain 244 ms end-to-end transaction latency and 50 ms slot times on a global cluster with nodes in Asia, Europe, and North America.

Setting. Blockchains are distributed systems whose operation relies on atomic broadcast protocols to ensure consistent replication of state across participating nodes. In such protocols, nodes receive *transactions* as inputs from the environment – an abstraction representing all entities external to the protocol – and, after some time, they output to the environment a totally ordered sequence of *confirmed* transactions, commonly referred to as the *log*. At their core, atomic broadcast protocols typically consist of a proposal phase, in which a designated node proposes a payload of transactions to be appended to the log, followed by a number of voting phases (often two), during which nodes collectively decide whether to confirm the proposal. Phases repeat over time, allowing the protocol to extend the log incrementally as new transactions arrive.

In atomic broadcast executions in which all nodes follow the protocol (*good-case*), a key performance metric is the *average transaction latency*, defined as the average time required for a transaction to be confirmed after it has been received by all nodes. The good-case average transaction latency, hereafter for brevity referred to as *transaction latency*, naturally decomposes into two components: the *inter-proposal time*, which captures the time until a transaction is included in a proposal, and the *confirmation time*, which measures how long a proposal takes to become confirmed. Minimizing transaction latency therefore requires reducing both components simultaneously. Tab. 1 summarizes the inter-proposal time, confirmation time, and transaction latency of several state-of-the-art protocols. The remainder of this section unpacks where each component’s lower bound comes from and how prior protocols have approached it.

Confirmation Time. Prior work [2] fully characterizes the tight good-case bounds on confirmation time in the authenticated setting across synchronous, partially synchronous, and asynchronous network models. Under partial synchrony with $3f + 1 \leq n < 5f - 1$, the tight lower bound is three communication rounds. Traditional protocols such as PBFT [13], Tendermint [10], and Simplex [15] operate in conservative rounds of duration Δ and saturate this bound at 3Δ . The bound therefore offers no further room for improvement at the standard resilience threshold: any reduction in transaction latency must come from the inter-proposal time.

■ **Table 1** Inter-proposal time (T_{ipt}), confirmation time (T_{conf}), and average good-case transaction latency of different protocols. Gatling is parameterized by K such that $\epsilon = 2\Delta/K$ can be made arbitrarily small. Unlike the other protocols, Gatling is a closed-box construction: it uses the underlying protocol only through its interface, without regard to internals.

Protocol	Inter-Proposal Time	Confirmation Time	Transaction Latency	Closed Box
PBFT [13], Tendermint [10]	3Δ	3Δ	4.5Δ	No ✗
HotStuff [52]	2Δ	7Δ ^{a)}	8Δ	No ✗
HotStuff-1 [27], HotStuff-2 [32]	2Δ	5Δ ^{b)}	6Δ	No ✗
Simplex [15]	2Δ	3Δ	4Δ	No ✗
Moonshot [22], Hydrangea++ [44]	Δ	3Δ ^{c)}	3.5Δ	No ✗
Shoal++ [6] ^{d)}	Δ	4.5Δ	5Δ	No ✗
Gatling (this work)	ϵ	3Δ	$(3 + \epsilon/2)\Delta$	Yes ✓

^{a)} HotStuff achieves a minimum commit latency of 7Δ when the next leader aggregates the votes for the current leader’s proposal. ^{b)} HotStuff-1 has a 5Δ confirmation time on the regular path, and an optimistic (speculative) path that reduces latency to 3Δ by allowing nodes to send early execution (confirmation) responses to clients after a single QC. Clients need to handle responses with care, to avoid accepting provisional confirmations that may later be revoked or reordered. ^{c)} We ignore the 20%-resilience fast path of Hydrangea++ as this is an orthogonal technique. ^{d)} After completion of this work, it was brought to our attention that Shoal++ [6] employs a construction similar to Gatling in the context of DAG-based protocols to reduce the inter-proposal time.

Inter-Proposal Time: Pipelining. Traditional atomic broadcast protocols operate in discrete slots (or views), each associated with a leader pseudo-randomly elected from the participating nodes (rotating leaders). The leader of the slot is entitled to cast a proposal for its slot. Typically, proposals are strictly sequential: each proposal must complete all voting phases before the next can be issued. In PBFT and Tendermint, for instance, a leader proposes a block and waits for two full rounds of voting, yielding a “propose-vote-vote” structure and an inter-proposal time of 3Δ . The HotStuff family of protocols [52, 27, 32] and Jolteon [25] relax this strict sequentiality by cryptographically chaining proposals with quorum certificates (QCs), such that each proposal includes a QC that carries votes for the previous proposal: as a result, a leader can issue a new proposal while nodes are still voting on the previous one. This technique, called pipelining, produces an effective “propose-vote-propose-vote” pattern and reduces the protocols’ inter-proposal time to 2Δ .¹ Simplex [15] also achieves a 2Δ inter-proposal time: unlike the HotStuff family, which attains this bound via chained QCs in a single-vote pipeline, Simplex advances to the next slot immediately upon notarization of the current slot, while a distinct finalization vote proceeds in parallel with the subsequent proposal.

Optimistic designs push pipelining even further: under the assumption of consecutive honest leaders, Moonshot [22] and Hydrangea++ [44] allow leaders to issue new proposals even before the previous one has been quorum-certified. This aggressive pipelining results in a “propose-propose-propose” structure, enabling a new proposal at every round and achieving the state-of-the-art minimum inter-proposal time of Δ . Because transactions may arrive at any point between two consecutive proposals, the expected waiting time until inclusion in a proposal is half of the inter-proposal time; the lowest transaction latency achieved by existing protocols is the 3.5Δ of Moonshot and Hydrangea++.

¹ HotStuff [52] also advertises to advance in rounds matching the actual network delay δ , decoupling progress from the pessimistic bound Δ ; this property, known as *responsiveness*, is shared by several protocols [13, 27, 32, 25, 22]. Since the dominant strategy for profit-maximizing proposers is to exhaust their timeouts, blockchain protocols de-facto run non-responsively; we discuss this below and in Sec. 9.

Inter-Proposal Time: Stable Leaders. In classical atomic broadcast protocols, proposals are cryptographically chained: a leader can successfully propose a new batch of transactions only after receiving the previous leader’s proposal. As a result, when the leader rotates at every slot, the inter-proposal time must be at least Δ to give time to the previous proposal to propagate to the next leader. Protocols achieving a sub- Δ inter-proposal time typically allow a single leader to remain in control for multiple consecutive slots and issue several proposals in sequence (stable or slowly rotating leaders). In this regime, the leader has local access to its own prior proposals and can therefore bypass network propagation delays. This is the case, for instance, of the original PBFT [13]. Keeping the same leader in power for extended periods is, however, problematic in blockchains, as it grants the leader disproportionate control over transaction ordering, exacerbating MEV extraction, censorship, and eroding fairness among nodes. A middle ground is achieved by some PBFT-like protocols instantiated with slowly rotating leaders: they allow a single leader to produce multiple proposals within a fixed number of consecutive slots (called a window), separated only by local processing delays (a few tens of milliseconds). When the leader rotates across windows, though, they incur a larger inter-proposal time of at least Δ (in practice, a few hundreds of milliseconds). Another example of such protocols is Alpenglow [5].


On Responsiveness. A substantial body of work has been devoted to the design of *responsive* Byzantine fault-tolerant protocols. Responsiveness refers to the ability of a protocol to make progress as a function of the actual network conditions (rounds of length δ), rather than operating with conservative timeouts that reflect the worst-case network bound (Δ). The goal is to reduce latency under $\delta \ll \Delta$ network conditions.

Recent work [41, 17] shows that, in blockchains, the dominant strategy of profit-maximizing proposers is to delay their proposal until their timeout expires: this allows them to gain more information about incoming transactions and craft a proposal that maximizes the Maximal Extractable Value (MEV). This behavior is often referred to as *timing games*, and has been observed in production systems such as Ethereum [41] and Solana [17]. As a result of economic incentives, blockchain protocols de-facto run non-responsively at Δ speed, and not at the δ speed that responsiveness would allow. For this reason, in this work, we consider protocol rounds of length Δ .

1.1 Our Contribution

Gatling. Gatling (Sec. 3) runs K parallel and independent instances of a closed-box atomic broadcast component protocol – e.g., PBFT, Tendermint, Simplex, Moonshot, or a DAG-based protocol such as Bullshark [47], Shoal [46], or Mysticeti [8]. Each component instance is parameterized with a fixed, regular proposal schedule that generates a proposal every T_{ipt} time, so when all nodes are honest each instance extends its log every T_{ipt} . In Gatling, K is a tunable parameter, and the proposal schedules of consecutive instances are offset by T_{ipt}/K . As illustrated in Fig. 1 for $K = 4$, the component protocols together produce a T_{ipt}/K inter-proposal time. By appropriately choosing K , Gatling achieves an arbitrarily small inter-proposal time.

Since we want Gatling to output a single transaction log rather than K independent parallel logs, the K component logs are combined via a deterministic merge rule. All nodes assign the same index from 1 to K to each component instance, and each node locally constructs the Gatling output log by interleaving the K component logs in order of increasing slot number and instance number. For example, when $K = 2$, the Gatling log alternates slot



■ **Figure 1** The Gatling protocol depicted in this figure runs $K = 4$ instances of a component protocol Π^c . The Gatling output log (at the bottom) is formed by deterministically interleaving blocks from all instances, in increasing order of slot number and instance number. A block is confirmed once all preceding blocks in this global order are confirmed. Checkmarks indicate blocks that are confirmed by their instance, whereas clocks indicate blocks that are not yet confirmed by their instance. The last red block is not yet included in the Gatling log because, in the interleaved ordering, it comes after the third orange block, which has not yet been confirmed. This does not happen in the good-case but it could happen in case of, e.g., a network hiccup.

by slot: the slot-1 block of instance 1, then the slot-1 block of instance 2, then the slot-2 block of instance 1, and so on. Gatling confirms a block once all slots that come before it have been confirmed.

We prove that Gatling implements secure atomic broadcast with transaction latency $(T_{ipt}/K, T_{conf})$, where T_{ipt}/K is the inter-proposal time – which can be made arbitrarily small (Sec. 4) – and T_{conf} is the confirmation time. We perform a stochastic latency analysis (Sec. 5) to characterize the impact of occasional proposal confirmation failures and identify the optimal inter-proposal time that minimizes expected end-to-end latency. We implement a Gatling proof of concept on top of a production-grade Simplex deployment [15, 42, 19] with a deliberately unoptimized 500 ms per-instance inter-proposal time, and run it on a global cluster with nodes in Asia, Europe, and North America. End-to-end transaction latency falls as K grows, reaching 244 ms at $K = 9$ under a 1% proposal-drop rate while driving the slot time down to 50 ms (Sec. 8).

Discussion and Related Work. We further discuss how the design of vanilla Gatling inherently loses what is sometimes called *predictable validity* [36, 37, 29] (Sec. 6) – a property that enables proposers to reason at proposal time about transaction execution outcomes – and we propose two variants that retain it (Sec. 7) under a slowly rotating leader schedule. The first variant retains predictable validity by mimicking the behaviour of, e.g., PBFT-style protocols. The second retains predictable validity at the execution level by creating two distinct tiers of block space and interleaving their execution (Sec. 7). One tier, named *subprime block space*, is reserved for transactions that are not time-critical and cannot be front-run (e.g., layer-2 data blob transactions); the other, *prime block space*, is reserved for time-critical and highly contentious transactions (e.g., DeFi).

A long line of work has sought to reduce consensus latency through the design of responsive protocols complemented with pipelining [52, 27, 32, 25, 22], speculative execution [27], optimistic assumptions [22], or fast-path techniques that trade resilience for fewer voting rounds [18, 5, 43, 45]. We compare Gatling with these approaches in more detail

in Sec. 9. Gatling improves latency along an orthogonal dimension: it achieves arbitrarily small inter-proposal times through composition, while each component protocol operates under conservative worst-case network assumptions and traditional Byzantine fault-tolerance thresholds. Notably, *Gatling shows that even starting from a simple, unoptimized protocol with poor standalone performance* – e.g., Simplex with a 500 ms inter-proposal time – *parallel composition alone is sufficient to achieve latencies improving upon those of highly optimized protocols*. This highlights that *extreme protocol-level optimization is not a prerequisite for ultra-low latency*. Moreover, Gatling is designed to remain compatible with responsive, pipelined, speculative, and fast-path techniques implemented within its component protocols, enabling these mechanisms to be seamlessly layered on top of it.

Conceptually, Gatling can be viewed as extending the parallelism underlying DAG-based protocols [33, 28, 47, 21, 46, 8]. Whereas DAG protocols support multiple concurrent ordering streams over a shared graph of proposals that reference one another, Gatling removes these references entirely by composing independent consensus instances and imposing staggered proposal schedules. The Gatling log is recovered with a deterministic merge rule that does not need proposals to reference one another, thereby allowing the inter-proposal time to drop below the network delay.

After completion of this work, it was brought to our attention that Shoal++ [6] employs a construction similar to Gatling in the context of DAG-based protocols to reduce the inter-proposal time. In particular, it operates three parallel DAG instances, time-staggered by a fixed Δ offset, in a not-fully-generic “ajar-box” manner. By contrast, Gatling explores regimes with inter-proposal times below the network delay and provides variants that retain predictable validity even when the offsets become arbitrarily small. Gatling is closed-box, and our experiments allow us to isolate the latency reduction achieved purely from time-staggered parallel composition.

2 Model and Preliminaries

Our overarching goal is to design an atomic broadcast protocol. We consider a setting with n nodes, each of which equipped with a cryptographic identity (typically, a public/secret key pair for digital signatures) known to all other nodes (public key infrastructure, PKI). Over time, the nodes receive transactions from the *environment* and, over time, they return a log of transactions to the environment. The nodes send each other messages through a *network* to reach agreement on an ordering of their input transactions into their output logs. The nodes’ output logs are *confirmed*. Time proceeds in discrete *rounds*. Throughout the entire protocol execution, the nodes have *synchronized clocks*: all nodes always know which is the current round. In each round, each node receives messages from the network and possibly transactions from the environment. In every round, each node updates its internal state, produces messages to send to other nodes via the network, and outputs a log of confirmed transactions.

An *adversary* aims to disrupt the system by corrupting nodes and delaying network messages. We consider an adversary that corrupts a fixed set of at most f nodes at the start of the protocol execution, before any randomness is drawn (static corruption). The adversary learns the internal state of corrupted nodes and causes them to deviate arbitrarily from the protocol for the entire execution (permanent Byzantine faults). Corrupted nodes, also known as adversarial nodes, run any algorithm chosen by the adversary in time polynomial in the size of the inputs (computationally bounded adversary).

Nodes can exchange messages between each other via a fully-connected network of point-to-point links. The network is *partially synchronous* [23]: there exists an upper-bound message delay Δ that is known to all nodes and a global stabilization time $\text{GST} < \infty$ that is adaptively chosen by the adversary. Before GST , the adversary can delay messages arbitrarily (asynchronous network). After GST , the adversary can delay messages by at most Δ (synchronous network). During synchronous periods, when a node instructs the network to send a message to another node, the message is enqueued in the recipient's pending message queue together with a countdown initialized to Δ ; the countdown decreases by 1 with each round. The adversary can also decrease the countdown of every message at will. Once the countdown hits 0, the message is delivered to the recipient at the beginning of the next round. This means that every message sent by an honest node by time t is delivered to all honest nodes by time $\max(t, \text{GST}) + \Delta$.

Atomic Broadcast Interface. The nodes run a Byzantine fault-tolerant consensus protocol of the atomic broadcast flavour to reach agreement on their output logs, i.e., a common ordering of their input transactions.

► **Definition 1** (Atomic Broadcast Interface). *An atomic broadcast protocol is run by n nodes. In every round, nodes possibly receive transactions from the environment and they return to the environment a confirmed sequence of transactions called output log.*

We consider atomic broadcast executions in which, regardless of the strategy of the adversary, all correct nodes output a log λ that is safe and live. We write $\lambda \preceq \tilde{\lambda}$ to denote that λ is a prefix of $\tilde{\lambda}$. We denote by **broadcast** the function that a node invokes to broadcast a transaction to the network, and by **deliveredLog** the function that a node invokes to retrieve its output log.

► **Definition 2** (Atomic Broadcast Security). *An atomic broadcast protocol is secure with resilience τ in partially synchronous networks iff, in every partially-synchronous execution with up to $f \leq \tau$ adversarial nodes, except with probability $\text{negl}(\kappa)$ where κ is the security parameter, the following properties hold:*

- **Safety:** *For every two honest nodes p, q and for every two times t, t' , if λ is the log output by p at t and $\tilde{\lambda}$ is the log output by q at t' , then the two logs are consistent, i.e., $\lambda \preceq \tilde{\lambda}$ or vice versa.*
- **Liveness:** *For every t_0 and every transaction tx , if every honest node p calls **broadcast** with parameter tx by time t_0 , then there exists a $t_1 \geq t_0$ such that for every $t_2 \geq t_1$ and for every p that calls **deliveredLog** at t_2 , $\text{tx} \in \lambda$.*

Enriched Atomic Broadcast Interface. As components of Gatling, we need an atomic broadcast protocol that proceeds in slots. Each *slot* is associated with a leader node and a fixed proposal time T^{prop} , during which the leader is expected to propose a payload of transactions to be appended to the output log, and broadcast it to the network. Slots typically consists of three phases: a proposal phase, and two voting phases known as notarization and confirmation. Each payload that is included in the output log is annotated with the slot number in which it was proposed. The slot number is non-decreasing along the log entries.

We additionally require the atomic broadcast protocol to implement an enriched interface. This interface is obtained by parameterizing the protocol with a sequence of proposal times, denoted T^{prop} , and by enriching the classical atomic broadcast interface with two additional functions, **broadcast** and **deliveredAnnotatedLog**. Since the Gatling protocol runs several parallel instances of an atomic broadcast protocol, the two new functions make explicit when Gatling writes an input to an instance, or reads the output from an instance. For clarity, we now introduce the enriched interface of the atomic broadcast protocol.

► **Definition 3** (Enriched Atomic Broadcast Interface). *An enriched atomic broadcast protocol Π is run by n nodes, and parameterized by a sequence of proposal times T^{prop} . It provides a function `broadcast` that a node can call with a transaction tx as parameter and that returns nothing. It also provides a function `deliveredAnnotatedLog` which a node can call without any parameter and that returns an annotated output log, i.e., a sequence of transactions λ annotated with the slot number in which they were proposed.*

We emphasize that the enriched interface is merely a syntactic variant of the standard atomic broadcast interface, and does not alter its security semantics. We highlight that virtually all partially synchronous atomic broadcast protocols we are aware of can be straightforwardly adapted to provide the enriched interface. These protocols already operate over a notion of slot – called *view* in PBFT [13], Tendermint [10], HotStuff [52], and Moonshot [22], *heights* in Simplex [15], and *epochs* in Streamlet [16]. They can be easily modified to enforce a fixed, regular schedule of slot proposal times (e.g., every 2Δ time) and to annotate confirmed payloads with the slot number in which they were proposed, a value over which the protocols already establish consensus. Finally, let T_{conf} be the time it takes for a proposal to be confirmed after it has been proposed. We are interested in atomic broadcast protocols that fulfill a stronger notion of liveness, which can be articulated into two properties. We name the first property *slot-driven liveness*.

► **Definition 4** (Slot-Driven Liveness). *An atomic broadcast protocol has slot-driven liveness iff for any slot starting after GST whose associated leader is honest, a (possibly empty) proposal is produced and becomes confirmed by all honest nodes within T_{conf} time from the slot's start.*

The second property is *good-case $(T_{\text{ipt}}, T_{\text{conf}})$ -latency*. Before stating this property, we need to define what it means for a sequence of proposal times to be T_{ipt} -spaced.

► **Definition 5** (T_{ipt} -Spaced Sequence). *A sequence T is T_{ipt} -spaced iff, for every two distinct $t_1, t_2 \in T$, $|t_2 - t_1| \geq T_{\text{ipt}}$.*


Informally, a protocol has good-case $(T_{\text{ipt}}, T_{\text{conf}})$ -latency if, when all nodes are honest, transactions are confirmed $T_{\text{ipt}} + T_{\text{conf}}$ time after they are received by all nodes.

► **Definition 6** (Good-Case $(T_{\text{ipt}}, T_{\text{conf}})$ -Latency). *A protocol has good-case $(T_{\text{ipt}}, T_{\text{conf}})$ -latency iff, when the protocol is instantiated with a T_{ipt} -spaced proposal sequence T^{prop} , if $\text{GST} = 0$ and all nodes are honest, then for every $t \in T^{\text{prop}}$ and for every transaction tx , if tx is received by every honest node by t , then tx is confirmed by every honest node by $t + T_{\text{conf}}$.*

3 Protocol

The goal of this work is to design an atomic broadcast protocol with an arbitrarily small inter-proposal time – in particular, smaller than the bound on the network delay. To this end, we design Gatling with a compositional approach. We describe the protocol step by step and we provide the pseudocode executed by the nodes in Alg. 1.

The Gatling protocol is parameterized by a positive integer $K > 0$, a positive integer $T_{\text{ipt}} > 0$, and a black-box atomic broadcast protocol Π^c that exposes the enriched interface of Def. 3 and satisfies the slot-driven liveness and good-case $(T_{\text{ipt}}, T_{\text{conf}})$ -latency of Defs. 4 and 6. For simplicity, we refer to Π^c as the component protocol. Nodes running Gatling execute K concurrent and mutually independent instances of the component protocol. Nodes uniquely identify these instances by assigning them indices from 1 to K ; in particular, all nodes assign the same index to the same protocol instance. Each instance maintains its own



■ **Figure 2** The Gatling protocol Π^G depicted in this figure runs $K = 4$ instances of a component protocol Π^c . The union of the proposal sequences of Π_1^c (red), Π_2^c (blue), Π_3^c (magenta), Π_4^c (orange) is the $T_{\text{apt}}/4$ -spaced proposal sequence of Π^G . We denote by $\lambda_j[i]$ the payload confirmed in the i -th slot of the j -th instance. The Gatling log λ at time $t \geq \text{GST}$ is reconstructed at the top as the concatenation of the payloads of the component protocols, ordered by increasing slot and increasing instance number. Checkmarks indicate blocks that are confirmed, whereas clocks indicate blocks that are not yet confirmed. Note that $\lambda_4[3], \lambda_1[4], \lambda_2[4]$, etc., are not yet confirmed by Gatling (not yet part of λ) because, at time t , the corresponding component protocols have not yet confirmed the payloads for those slots.

leader schedule, produces its own annotated output log, and is parameterized by a proposal schedule that generates a proposal every T_{apt} time. The proposal schedule of each instance is offset by T_{apt}/K relative to the previous one (Alg. 1, ln. 4). Consequently, the union of the proposal sequences of all K component protocols forms a single proposal sequence with inter-proposal time T_{apt}/K .

When a node receives a transaction tx from the environment, it invokes the **broadcast** function (Alg. 1, ln. 18) that forwards tx to all K instances of the component protocol (Alg. 1, lns. 20 and 21). Then, tx is picked up by the leaders of the instances and included in a proposal. In particular, once the transaction has been received by all nodes, the instance proposing next has a proposal at most T_{apt}/K time away. Therefore, the transaction is guaranteed to appear in some proposal within at most T_{apt}/K time and confirmed after an additional T_{conf} time by the corresponding protocol instance.

It is not sufficient for tx to be confirmed by a component protocol: we need it to be confirmed by the Gatling protocol. Towards this, we describe how the output log of Gatling is built, and visually illustrate this in Fig. 2. The nodes construct the Gatling output log by concatenating the K output logs of the component protocols as follows: For each slot index i , the nodes define the i -th slot of the Gatling protocol as the tuple $s_i = (s_{i,1}, \dots, s_{i,K})$, where $s_{i,j}$ denotes the i -th slot of the j -th instance. Then, they construct the Gatling output log as the sequence of payloads confirmed in the Gatling slots. In other words, the Gatling output log concatenates the payloads of the component protocols output logs, ordered by increasing slot and instance number. In Gatling, proposals are not chained. With this in mind, we now define the Gatling confirmation rule, which consists of three conditions. The payload of slot $s_{i,j}$ is confirmed if: (1) instance j has confirmed the payload of slot s_i , (2) all instances $j = 1, \dots, K$ have decided for all slots $s_{i'}$ with $i' < i$, and (3) all instances j' with $j' < j$ have decided for slot s_i . As shown in Fig. 3, a component protocol that fails to confirm a proposal for a slot (e.g., due to an adversarial leader) still reaches consensus on the slot being empty.

■ **Algorithm 1** Algorithm run by a node of the Gatling protocol.

```

1 function new  $\Pi^G(\Pi^c, K, T_{\text{ipt}})$ 
2   // Start component protocols with proposal schedules s.t. the Gatling inter-proposal time is  $T_{\text{ipt}}/K$ 
3   for  $k = 1, \dots, K$ 
4      $\pi_k \leftarrow \text{new } \Pi^c(T_k^{\text{prop}} = \{(s + \frac{k-1}{K}) \cdot T_{\text{ipt}} \mid s \in \mathbb{N}_0\})$ 

5 function  $\Pi^G.\text{deliveredLog}()$ 
6   // Get component protocols' output logs annotated with slot number
7   for  $k = 1, \dots, K$ 
8      $\lambda_k \leftarrow \pi_k.\text{deliveredLogAnnotated}()$ 
9   // Find max slot number for which all component protocols have decided
10   $s^* \leftarrow \max s \in \mathbb{N}_0$  so that  $\forall s \leq s^* : \forall k \leq K : \lambda_k[s] \neq \perp$ 
11  // Find max instance number that has already decided for slot  $s^* + 1$ 
12   $k^* \leftarrow \max k \in \mathbb{N}_0$  so that  $\forall k \leq k^* : \lambda_k[s^* + 1] \neq \perp$ 
13  // Output Gatling protocol's current log
14  if  $k^* = 0$ 
15    return  $\lambda_1[1] \parallel \dots \parallel \lambda_k[1] \parallel \dots \parallel \lambda_1[s^*] \parallel \dots \parallel \lambda_K[s^*]$ 
16  else
17    return  $\lambda_1[1] \parallel \dots \parallel \lambda_k[1] \parallel \dots \parallel \lambda_1[s^* + 1] \parallel \dots \parallel \lambda_{k^*}[s^* + 1]$ 

18 function  $\Pi^G.\text{broadcast}(\text{tx})$ 
19   // Forward Gatling protocol's inputs to component protocols
20   for  $k = 1, \dots, K$ 
21      $\pi_k.\text{broadcast}(\text{tx})$ 

```

In the next section, we prove that the Gatling protocol in Alg. 1 is secure, and it has a good-case $(T_{\text{ipt}}/K, T_{\text{conf}})$ -latency.

4 Analysis

We now state the main results of this work, namely the security and the good-case $(\frac{T_{\text{ipt}}}{K}, T_{\text{conf}})$ -latency of Gatling. Proofs appear in the extended version of this work [40].


► **Theorem 7** (Gatling Security). *If Π^c is a secure atomic broadcast protocol as per Def. 2, then the Gatling protocol in Alg. 1 that uses Π^c as its component protocol is also secure as per Def. 2.*

► **Theorem 8** (Gatling Good-Case $(\frac{T_{\text{ipt}}}{K}, T_{\text{conf}})$ -Latency). *If Π^c has good-case $(T_{\text{ipt}}, T_{\text{conf}})$ -latency, then the composite protocol in Alg. 1 has good-case $(\frac{T_{\text{ipt}}}{K}, T_{\text{conf}})$ -latency.*

5 Optimal Inter-Proposal Time

In the previous sections we established that in the good-case, after GST, Gatling achieves $(T_{\text{ipt}}/K, T_{\text{conf}})$ -latency for any $K \geq 1$. Therefore, the Gatling inter-proposal time can be made arbitrarily small by adding component instances. In real world deployments, component instances may occasionally fail to confirm a proposal for a slot, both in the good-case due to network delays and in the bad-case due to the presence of adversarial nodes. The Gatling's confirmation rule turns these failures into head-of-line blocking events on the Gatling log: a delayed or failed confirmation in one component protocol slot can prevent subsequent proposals already confirmed at the component level from being appended to the log.

To see the effect, consider a single failure resolved at the affected instance's next slot (i.e., $T_{\text{ipt}} = K\epsilon$ later), where $\epsilon = T_{\text{ipt}}/K$ is the Gatling inter-proposal time. The proposal immediately following the failed one is ready ϵ after it, so it is held back by up to $T_{\text{ipt}} - \epsilon = (K - 1)\epsilon$, and its j -th successor by $(K - j)\epsilon$.



■ **Figure 3** Behavior of Gatling in the presence of an adversarial proposer. Left: Each component instance assigns a leader per slot; confirmed blocks (green checkmarks) are immediately appended to the Gatling log (bottom), if all leaders are honest. If a leader is adversarial (second leader of the red instance, semitransparent), no block is confirmed for that slot (red cross). Center: The other instances (blue, magenta, orange) continue confirming blocks independently, but Gatling delays including them in the log because it waits for the red instance to reach consensus on the adversarial slot. Right: Once the red instance progresses under a subsequent honest leader and confirms a proposal, Gatling resumes extending the log, appending an empty block for the adversarial slot together with the pending confirmed blocks from the other instances.

This section gives an intuition of how to identify the Gatling’s optimal inter-proposal time ϵ^* and, in turn, the optimal number of parallel instances K^* , that minimize the expected transaction latency in the presence of proposals that fail to confirm. The full analysis appears in the extended version of this work [40].

Stochastic Latency Model. The $(K - 1)\epsilon$ value above assumes a single failure resolved at the next slot; expected latency instead depends on how often failures occur and on how much slack separates consecutive positions, which is what we model here. Optimizing ϵ involves a tension: packing proposals closer together (ϵ small) shortens the time a transaction waits for the next available proposal, but it also shortens the head start that each proposal’s predecessors enjoy, thus holding up a longer queue of otherwise finalized proposals. When proposals are far apart, a slow component instance has plenty of time to decide before the proposals behind it are ready to be appended, and blocking is rare. Therefore, expected end-to-end latency is the sum of two terms that move in opposite directions as ϵ changes: the wait for inclusion, which grows with ϵ , and the delay caused by head-of-line blocking, which shrinks with ϵ .

To locate the sweet spot for ϵ , we adopt a stochastic model of component instances’ behaviour: each instance attempts to decide its slot within the good-case confirmation time T_{conf} , each attempt succeeds with probability s independently of all other attempts and of all other slots and instances, and a failed attempt is followed by a new attempt on the next slot. The resulting per-slot decision time is geometric: a slot is decided after one attempt, or two, or three, so its distribution is a staircase that jumps at each multiple of T_{conf} . Those jumps obstruct a closed form. We therefore approximate the staircase by a smooth exponential curve shifted by T_{conf} , whose rate is fixed by requiring that it passes through the staircase at every attempt boundary. The curve is shifted by T_{conf} , because no slot can be decided before one attempt has elapsed. The approximation is exact at the boundaries and interpolates in between. It also turns the expected-latency expression into a classical q -series, which admits a closed-form optimization.

The next theorem shows that this trade-off has a single sweet spot, and gives both the optimal inter-proposal time and the resulting latency in closed form. For the full analysis that led to the theorem, as well as its proof, we refer to the extended version of this work [40].

► **Theorem 9** (Optimal Gatling Inter-Proposal Time). *Consider the model above, with per-attempt failure probability $p := 1 - s$ and smooth-approximation rate $\lambda_p := -\log p/T_{\text{conf}}$, and let $L(\epsilon)$ denote the resulting expected end-to-end transaction latency, comprising the wait for the next available proposal and the time from that proposal until it is appended to the Gatling log. Then L has a unique minimizer*

$$\epsilon^*(s) = -\lambda_p^{-1} \log q_p^* = T_{\text{conf}} \cdot \frac{-\log q_p^*}{-\log p}, \quad (1)$$

where $q_p^* \in (0, 1)$ is the unique solution to

$$\int_0^1 (px; q_p^*)_\infty \sum_{j=1}^{\infty} \frac{p j (q_p^*)^j}{1 - px(q_p^*)^j} dx = \frac{1}{2} \quad (2)$$

and $(a; q)_\infty := \prod_{m \geq 0} (1 - aq^m)$ denotes the q -Pochhammer symbol. The corresponding minimum expected latency is


$$L^*(s) = T_{\text{conf}} + \lambda_p^{-1} \phi_p(q_p^*), \quad \text{with} \quad \phi_p(q) := -\frac{1}{2} \log q + \int_0^1 \frac{1 - (px; q)_\infty}{x} dx. \quad (3)$$

The theorem is stated in terms of q_p^* because the q -series form makes the optimization tractable; q is the decay of the smooth exponential curve over one inter-proposal time, and thus a reparametrization of ϵ . Equation (1) inverts it, giving an optimal ϵ^* that is a multiple of T_{conf} with a factor depending only on s . Equation (2) is the tension we described above, written as a balance. Its left-hand side is the rate at which increasing ϵ relieves head-of-line blocking, aggregated over all predecessors, and the $\frac{1}{2}$ on the right is the rate at which it grows the wait for inclusion, since a transaction waits on average $\epsilon/2$ for the next available proposal. The optimal ϵ^* is where the two rates meet. Finally, Equation (3) splits the achievable latency into the component protocol's good-case confirmation time T_{conf} and an overhead term that captures the cost of parallelism, which depends on s .

Numerical Results. We obtain q_p^* by numerically solving Equation (2) and then evaluate the closed-form expressions for $\epsilon^*(s)$ and $L^*(s)$ in Thm. 9. Fig. 4 showcases $\epsilon^*(s)$ and $L^*(s)$ in units of Δ as functions of the per-attempt success probability s . As $s \rightarrow 1$, the optimal inter-proposal time $\epsilon^*(s)$ shrinks to zero while the minimized expected latency $L^*(s)$ converges to T_{conf} (here 3Δ): Gatling recovers the component protocol's good-case confirmation latency with arbitrarily small inter-proposal time. For example, at $s = 0.99$ the optimal global inter-proposal time is $\epsilon^* \approx 0.09\Delta$ and the corresponding expected end-to-end latency is $L^* \approx 3.09\Delta$. At $s = 0.999$ these become $\epsilon^* \approx 0.02\Delta$ and $L^* \approx 3.02\Delta$. Translating ϵ^* into the design parameter K of Sec. 3, a Gatling protocol that runs parallel instances of a component protocol with inter-proposal time T_{ipt} should choose a number $K^*(s) = \lceil T_{\text{ipt}}/\epsilon^*(s) \rceil$ of parallel instances. Larger values of K increase head-of-line blocking without further reducing latency.

6 Transaction Deduplication

Since Gatling forwards every transaction to all K component protocol instances (Alg. 1, lns. 20 and 21), the same transaction may appear multiple times in the Gatling output log. As observed for multi-proposer protocols [39, 24, 47], duplicate transactions waste block space and reduce the effective throughput of the protocol. One might ask whether forwarding




■ **Figure 4** Optimal Gatling global inter-proposal spacing $\epsilon^*(s)$ and minimized expected end-to-end finality latency $L^*(s)$ under the model of Thm. 9.

each transaction only to the component instance next in line to propose, rather than to all K instances, would avoid duplicates and thereby improve what is sometimes called *goodput*. For practical purposes, this is indeed the case. Modern blockchains frequently rely on private mempools or relay networks, which disseminate transactions rapidly and preferentially to upcoming leaders, effectively ensuring that transactions are promptly available to the next leader to propose. In turn, for most consensus protocols, it suffices for liveness if an honest leader includes a transaction in its proposed block.

From a theoretical perspective, Def. 2 guarantees liveness of the component protocols only if a transaction is eventually input by all honest nodes. This requires that for Gatling liveness, honest nodes forward input transactions to all component protocols. On the other hand, strengthening the liveness guarantee so that input by one rather than all honest nodes suffices for liveness, would introduce complications in the definition of latency: under the current Def. 2, the latency clock starts when all honest nodes have broadcast the transaction, ensuring that the next leader has already received it. If the clock instead starts when only one honest node has broadcast the transaction, the transaction may still need to propagate to the next leader or wait for the node that broadcast it to become a leader, potentially adding up to the effective transaction latency. The theoretical latency metric then contains terms that in reality are not there due to how transaction dissemination and consensus are handled in modern systems.

7 Enhancing Gatling with Predictable Validity

An attentive reader may have observed that when Gatling is run with an inter-proposal time smaller than Δ and with successive blocks being proposed by different leader nodes (rotating leaders), a block does not have enough time to propagate from one leader to the next before the following proposal is issued. As a result, a leader cannot validate the transactions in its proposal against the most recent state of the log. Under these conditions, Gatling no longer satisfies a property commonly referred to as predictable validity [36, 37, 29]. Predictable validity guarantees that transactions deemed valid by a leader at proposal time



■ **Figure 5** Comparison between the empty-blocks Gatling (top) and a PBFT-style protocol (bottom), both instantiated with a slowly rotating leader schedule. In this picture, $K = 4$, $\Delta = 400\text{ms}$, and the Gatling inter-proposal time is 100ms . The PBFT-style protocol proposes multiple (non-empty) blocks during a leader’s window and separates consecutive windows by at least Δ to ensure block propagation and, therefore, predictable validity. The empty-blocks variant of Gatling reproduces this behavior by alternating proposal windows with explicit Δ -sized gaps with empty proposals.


remain semantically valid when they are eventually executed. Under predictable validity, proposers can therefore reason soundly about execution outcomes using only the locally available state. While most popular blockchain protocols such as Bitcoin, Ethereum, and Solana fulfill predictable validity, DAG-based logs such as Sui [9] and Aptos [46] as well as “lazy” logs such as Celestia [14] and Pod [3] have forgone it in favor of lower latency. For blockchains that want to preserve predictable validity while retaining the benefits of reduced inter-proposal time, we present two variants of Gatling— *empty-blocks* and *subprime-blocks* — with inter-proposal times smaller than Δ , operating under slowly rotating leaders.

7.1 Empty-Blocks Gatling Variant

Consider a Gatling protocol that achieves a sub- Δ inter-proposal time $\frac{T_{\text{ipt}}}{K}$ by running K parallel instances of a component protocol, each configured with an inter-proposal time of T_{ipt} . Time is partitioned into a sequence of contiguous slots, called windows, of length at least Δ . Each window is assigned to a randomly selected leader node which, within its assigned window, can issue a fixed number of consecutive proposals. For Gatling to achieve this slowly rotating leader schedule, the leader schedules of the component protocols are coordinated so that a single leader can issue K consecutive Gatling blocks by proposing in K consecutive component protocol instances. To enable this coordination, each atomic broadcast component protocol exposes an interface that accepts the proposer schedule as an explicit input, in addition to the sequence of proposal times.

Within a window, a sub- Δ inter-proposal time does not compromise predictable validity: because the same leader issues multiple consecutive proposals, it can locally maintain and update a consistent view of the log across proposals. Hence, when building a new proposal, the leader can validate the transactions against the most recent state of the log. The challenge arises at window boundaries. Since Gatling is instantiated with a sub- Δ inter-proposal time T_{ipt} , less than Δ time elapses (at least) between the last proposal of one window and the first proposal of the subsequent window. Consequently, the leader of the subsequent window has to issue a proposal(s) before having received one or more proposals of the previous window.

To preserve predictable validity across windows, this variant of Gatling introduces a “window-gap-window-gap” structure, with windows occurring in strict alternation with gaps of at least Δ (Fig. 5). In a window, the assigned leader includes a non-empty payload in all of its proposals. In contrast, in a gap, proposals carry empty payloads that do not modify the state of the log. These gaps provide sufficient time for proposal propagation between



■ **Figure 6** Subprime-blocks variant of Gatling that restores predictable validity at execution. At the top, the consensus confirmation ordering. At the bottom, we illustrate two possible orderings in which blocks are executed: in one (top), blocks of windows and gaps are executed contiguously, with a window being executed before the preceding gap. In the other (bottom), for each window and the preceding gap, every block from the window is executed before the corresponding block from the gap.

consecutive windows, allowing the leader responsible for the next window to observe the preceding blocks. This ensures predictable validity. The empty proposals could be issued similarly to the ones in a window, or they could be issued by a leader so that all proposals are received by the next leader before it begins issuing proposals for its window (e.g., unique slot proposal time T^{prop} at the beginning of the window for all proposals). Alternatively, the protocol could mandate nodes to implicitly insert empty proposals in their local logs, without requiring explicit consensus on those (no QCs).

Fig. 5 illustrates that the empty-blocks variant of Gatling (top of the figure) closely mirrors the proposal pattern of traditional PBFT-style protocols with slowly rotating leaders (bottom of the figure). In PBFT-like protocols, a leader proposes multiple blocks within a window, and successive windows are separated by at least Δ time to allow blocks to propagate, thereby preserving predictable validity. Similarly, this Gatling variant alternates between proposal windows and Δ -sized gaps: proposals are issued at sub- Δ intervals within a window, while the intervening gaps ensure sufficient time for propagation. Gatling provides a general black-box composition framework for atomic broadcast that encompasses stable and slowly rotating leader protocols, such as PBFT and Alpenglöw, as special cases. By adjusting its parameters, Gatling can reproduce their execution patterns, while also enabling alternative operating regimes in which proposals are issued at sub- Δ inter-proposal times and consistently carry transactions, trading predictable validity for higher proposal rates.

7.2 Subprime-Blocks Gatling Variant

While the empty-blocks variant of Gatling already captures many of the desired properties (sub- Δ inter-proposal times and predictable validity), the need to validate transactions against the most recent state introduces gaps that periodically stall progress and increase the inter-proposal time of non-empty blocks. To address this limitation, we introduce an advanced variant of Gatling named *subprime-blocks* Gatling – illustrated in Fig. 6 – that preserves predictable validity while eliminating the stalls.

The key observation is that empty blocks proposed in the gaps of the previous design can be repurposed as a distinct tier of block space, which we refer to as *subprime block space*. This tier is reserved for transactions that are not time-critical and are insensitive to ordering races, such as simple value transfers or layer-2 data blob transactions. In contrast, the *prime block space* tier is reserved for high-priority transactions – such as DeFi interactions – that

are sensitive to execution order and frequently participate in front-run races. In this variant, the designated leader of a window makes proposals for prime block space, by only including time-critical transactions that access highly contentious state. Instead, the designated leader of a window of a gap makes proposals for subprime block space, by only including non-time-critical transactions that access non-contentious state.

Additionally, we modify the interaction between the consensus and execution layers by allowing the log produced by consensus to differ from the log consumed by the execution layer. As per Fig. 6, consider window number 3, followed by gap number 3, followed by window number 4, followed by gap number 4, and so on. While the Gatling log output by consensus is the ordered sequence of blocks proposed in subsequent windows and gaps, the log consumed by the execution layer always executes a window before the immediately preceding gap. For example, the execution layer processes blocks from window 3, then blocks from gap 2, then blocks from window 4, followed by blocks from gap 3, and so on (see the execution order in Fig. 6). Because each gap has duration at least Δ , when the leader of a window issues its first proposal, all blocks that will be executed before it must have been proposed at least Δ time earlier and therefore already received and executed by that leader. Consequently, proposers in windows always construct their proposals on top of a known and up-to-date log state. In contrast, the leaders proposing during gaps have no knowledge of the state at the time their subprime blocks will be executed, since their blocks will be executed after the blocks proposed in the next window. Subprime-block transactions are not guaranteed to execute successfully. At proposal time, they only need to be well formed and pass all validity checks that do not depend on the eventual execution state (e.g., correct format, valid signatures, resource bounds, declared dependencies). State-dependent validity is evaluated later at execution time against the state determined by Gatling’s execution order; if the transaction is invalid against that state, it (soft-)fails without invalidating the block. We can further refine the execution sequence by ordering blocks individually rather than executing all blocks from windows and gaps contiguously, as illustrated in the optimized execution ordering of Fig. 6: execution alternates between blocks from a window and blocks from the preceding gap. For instance, the first block of window 3 is executed first, then the first block of gap 2; next, the second block of window 3, followed by the second block of gap 2, and so on until all blocks of window 3 and gap 2 have been executed. The same alternating pattern is applied to window 4 and gap 3, etc.

In the variant of Sec. 7.1, nodes execute blocks only in during windows, while during gaps they wait for new blocks. This subprime-blocks Gatling variant, besides higher proposal rates, also offers a more uniformly distributed execution load.

8 Evaluation

We implemented a proof-of-concept of Gatling to measure its reduced transaction latency and its practicality in a realistic environment. Our experiment confirms the theoretical findings in Sec. 5. Our implementation is available at: <https://github.com/scaffino/Gatling>.

Implementation Overview. Our implementation extends Commonware’s *alto* blockchain [19], a minimal but complete Rust implementation of a blockchain that relies on the Commonware *monorepo* [20], a collection of production-grade primitives including consensus, networking, and cryptographic libraries. We augment *alto* by enabling each validator to run $K \geq 1$ independent instances of the Simplex consensus. Each instance maintains its own log, follows an independent leader schedule, and communicates over dedicated network channels.

Importantly, each instance has an inter-proposal time of $T_{\text{ipt}} = 500\text{ms}$ and is configured to be non-responsive. We set the leader deadline to 225ms (the time by which the validators must have received the block; otherwise, they will cast a skip vote) and the notarize deadline to 375ms (the time by which a validator must have received a notarize certificate to cast a finalization certificate). We implement the Gatling log reconstruction and confirmation rules described in Sec. 3, allowing the outputs of these parallel instances to be deterministically merged into a single global log.

We do not model transaction dissemination via gossip. Instead, transactions arrive according to a Poisson process at a rate of approximately 100 transactions per second. This transaction rate is below the system's saturation point and therefore does not introduce queuing effects that could distort latency measurements. Transactions are delivered directly to the validator expected to propose the next block, effectively modeling instantaneous transaction submission and emulating the private-mempool behavior commonly observed in real-world blockchain deployments. We note that when the Gatling inter-proposal time ϵ is sufficiently small relative to the global network delay, bypassing the designated next leader may reduce latency when that leader is geographically distant.

To emulate crash failures, we randomly suppress proposals with probability p . Varying p allows us to evaluate the effect of proposal failures on head-of-line blocking and transaction latency.

Measuring Latency. We deployed Gatling on a global 10 validator cluster with 1 node in London, 3 in Tokyo, 3 in Singapore, 2 in Dallas, and 1 in Miami. The largest latencies in the cluster were around 110-112ms between validators in Miami and Singapore. We compute the transaction latency for varying values of K . We denote by skip rate the percentage of slots that resulted in a skip vote, i.e., the leader missed the 225ms deadline, aggregated over all instances during the measurement window. Fig. 7 presents the transaction latency as a function of K , measured across the three proposal-drop rate regimes: 0% (top left), 1% (top right), 5% (bottom left), and 10% (bottom right).² Each bar is a stacked bar chart decomposing total latency into its two components: the wait time (blue), and the confirmation time (orange).

In the good case, where no proposal failures occur, increasing K reduces the inter-proposal time and thus lowers transaction waiting time. As a consequence, we expect the transaction latency to decrease monotonically with larger values of K . This is confirmed by Fig. 7; for $K = \{15, 40\}$ the average transaction latency shows a slight increase due to anomalously high skip rates. The 0% proposal-drop rate in Fig. 7 also reveals that end-to-end latency stabilizes around 215ms, consistent with the theoretical 3Δ confirmation time floor for $\Delta \approx 70\text{ms}$.

In the presence of proposal failures, transaction latency is expected to exhibit a non-monotonic dependence on K : it initially decreases as the inter-proposal time shrinks, but eventually increases as head-of-line blocking becomes the dominant source of delay (Sec. 5). This is consistent with Fig. 7: the transaction latency is minimized at $K = 9$ for the 1% drop rate, at $K = 4$ for the 5% drop rate, and at $K = 3$ for the 10% drop rate. In the right panel of Fig. 7 the blue component shrinks steadily from $K = 1$ to $K = 9$, and total latency falls from 447ms to a minimum of 244ms at $K = 9$. Beyond $K = 9$, however, a 1% proposal-drop probability is sufficient to trigger head-of-line blocking frequently enough that the confirmation time wait (orange component) grows, pushing total latency back up to 306ms at $K = 10$. The center and left panels of Fig. 7 illustrate how reliability governs the optimal K . With a 5% drop rate (center), the optimal number of instances shifts to

² The corresponding numerical results appear in the extended version of this paper [40].




Figure 7 Transaction latency as a function of K across three proposal-drop rate regimes: 0% (top left), 1% (top right), 5% (bottom left), and 10% (bottom right). Each bar decomposes total latency into the average time transactions wait to be included in a proposal (proposal wait mean, blue), and the average time it takes for them to be confirmed after inclusion (finality wait mean, orange). The more reliable the validator set is, the more concurrent instances can be tolerated without introducing additional delay from head-of-line blocking.

$K = 4$, yielding a minimum latency of 296ms, while with a 10% drop rate the optimum is at $K = 3$ with 337ms. In both cases, the effect of head-of-line blocking increases with K . This confirms the core tradeoff identified in Sec. 5.

Our experiments also demonstrate that a blockchain can rotate leaders every $\epsilon = 50$ ms or less on a global cluster where communication time between nodes exceeds 50ms.

Consistency with the Theoretical Optimum. The model of Sec. 5 makes three predictions about end-to-end latency as a function of the Gatling inter-proposal time ϵ ; equivalently, as a function of the number of parallel instances $K = T_{\text{ipt}}/\epsilon$. First, the latency curve has a unique minimum at some $\epsilon^*(s)$. Second, the corresponding optimum $K^*(s)$ shifts toward smaller K as the per-attempt failure probability $p = 1 - s$ grows. Third, as $s \rightarrow 1$, the minimum expected latency $L^*(s)$ approaches the single-attempt confirmation floor T_{conf} . All three predictions are visible in Fig. 7: each panel exhibits a clear U-shape with a minimum; the optimum shifts from $K = 9$ at the 1% drop rate to $K = 4$ at the 5% drop rate and $K = 3$ at the 10% drop rate; and the minimum observed latency falls monotonically as the drop rate shrinks, from 337ms at 10% to 297ms at 5% and 244ms at 1%, approaching $T_{\text{conf}} \approx 197$ ms, the time of one round of the component protocol, measured at $K = 1$ where head-of-line blocking is absent.

9 Related Work

Responsive Protocols and Timing Games. An important line of work in distributed systems and blockchains has focused on a class of protocols, called *responsive*, which improve latency by reducing dependence on worst-case timing assumptions [1, 7]. Protocols such as

Tendermint [10] and Streamlet [16] advance according to fixed delays chosen to tolerate the maximum network delay (Δ), causing nodes to wait for entire Δ timeout periods even when the actual message delay (δ) is much smaller.

In contrast, responsive protocols such as PBFT [13], Simplex [15], HotStuff [52], HotStuff-1 [27], HotStuff-2 [32], Jolteon [25], and Moonshot [22] are designed so that progress is driven by message arrivals rather than by timeout expirations, allowing nodes to reach a decision as quickly as the network conditions permit. As pointed out in [1, 38, 26], responsiveness is execution-dependent and cannot be guaranteed in all runs: when a leader is adversarial, correct nodes may be forced to delay commitment, reintroducing latency that is independent of the actual network conditions. As a result, responsiveness holds only for a subset of optimistic executions: those with a correct leader and sufficiently many honest participants.

For a long time, the design of responsive protocols has been carried out without considering timing incentives of the nodes. Only recently, the work of [41, 17], analyzed the incentives of responsive protocols in the blockchain setting, showing that when elected as proposer the best strategy for profit-maximising nodes is to delay the proposal for as long as possible, while still ensuring the timely inclusion in the ledger. This strategic delay in block proposals, often called timing game, is exploited to acquire additional information about pending transactions and optimize the Maximal Extractable Value (MEV) capture. Timing games have been observed in major blockchains such as Ethereum and Solana. Designing responsive protocols that explicitly account for timing games and strategic proposer behavior was first explored by [4]; however, this composition remains largely underexplored and represents a promising direction for future work. Our work improves latency via a mechanism that is orthogonal and complementary to the one of responsive protocols. While responsive protocols achieve low latency under favorable network delays (δ), Gatling attains arbitrarily small inter-block times using only conservative worst-case assumptions (Δ) in its component protocols, providing latency improvements independent of optimistic network conditions.

Pipelined and Optimistic Protocols. In traditional BFT-style protocols, block proposals are strictly sequential: a leader must wait for the previous proposal to be confirmed through multiple voting phases before issuing the next proposal. PBFT [13], for example, proceeds through a leader proposal phase (pre-prepare) followed by two voting phases (prepare, commit), resulting in a “propose-vote-vote” structure per proposal. With Casper-FFG [12] and then the HotStuff family of protocols [52, 27, 32], a *pipelined, chain-based* formulation is introduced: proposals are cryptographically linked to each other in a chain, and each proposal carries a quorum certificate attesting that a quorum of nodes has agreed on the previous proposal. This structure allows consensus phases to be pipelined: while nodes vote to form a certificate for the current proposal, the leader can already issue a new proposal. Conceptually, this results in an alternating “propose-vote-propose-vote” pattern, where making a new proposal is overlapped with voting on the previous one.

A complementary line of work explores optimistic/speculative techniques that further reduce the number of network round trips to reach confirmation in BFT consensus protocols. For instance, Moonshot [22] and Hydrangea++ [44] adopt an *optimistic proposal* approach allowing leaders to propose new blocks before the previous block has been quorum-certified, overlapping proposal and voting phases. This reduces latency under the assumption of consecutive honest leaders. This design effectively results in a “propose-propose-propose” structure, enabling more aggressive proposal pipelining. We observe that for $K = 2$, Gatling recovers the behavior of Moonshot in the optimistic case: when the leader is honest and includes a transaction in its block, and the previous leader was also honest, then that transaction is viewed as confirmed by everyone 3Δ later, even if at most f nodes are adversarial.

A different approach is taken by HotStuff-1 [27]: next to the traditional HotStuff-style commit rule, in the optimistic case, it introduces a speculative technique that reduces commit latency by sending clients early execution responses after a single quorum certificate, without waiting for the usual multi-phase commit rule.

Carnot [30] obtains block times below the network delay by adopting a multi-view leader schedule and it shows that, under bounded bandwidth, the equilibrium proposal interval is limited by transaction dissemination and proposal metadata transmission. Unlike Gatling, which supports rapidly rotating leaders through parallel composition and derives its optimal instance count from head-of-line blocking, Carnot employs a stable leader schedule and uses erasure coding to reduce the bandwidth bottleneck of the leader. Carnot’s bandwidth analysis is complementary to ours and would impose an additional finite optimum on Gatling’s number of instances in a bandwidth-constrained network.

Fast-Path Protocols. A long line of work that started with [34] has explored *fast-path* techniques that reduce latency by decreasing the number of voting rounds required to confirm a block, trading resilience for lower round complexity. Classical BFT protocols require two rounds of voting to tolerate up to f Byzantine faults, ensuring safety even under adversarial conditions. Fast-path designs relax these guarantees under optimistic assumptions – such as honest leader and a reduced number of adversarial nodes – to decide on a proposal with a single round of voting. This approach reduces transaction latency when the assumptions hold, but retain a fallback slow-path with two rounds of voting that restores full fault-tolerance when the optimistic assumptions are violated. Recently, there has been renewed interest in fast-path designs in blockchains and distributed systems which has led to the designs of Minimmit [18], Alpenglow [5], Kudzu [43], and Hydrangea [45]. While fast-path protocols reduce latency by weakening fault-tolerance assumptions in optimistic executions, Gatling preserves full resilience while improving block production rates through composition. As such, fast-path techniques are complementary to Gatling rather than competing: they could be layered on top of the component protocols to further reduce Gatling’s confirmation latency.

Multiple Concurrent Proposals. Protocols that leverage multiple concurrent proposals (MCP) have recently gained renewed attention in the Ethereum and Solana communities, notably through the Ethereum’s EIP-7805 (FOCIL) [50] and this recently proposed protocol [24]. MCP approaches allow multiple nodes to simultaneously propose batches of transactions for inclusion in the next block; while each round still produces a single block, that block aggregates contributions from several proposers, thereby reducing the leader’s monopoly over transaction inclusion and ordering. While Gatling reduces the inter-proposal time, MCP protocols [24] require multiple rounds of communication among nodes when a proposal is constructed, which leads to longer inter-proposal times.

In DAG-based protocols [33, 28, 47, 21, 46, 8], in each round, every node can independently create and broadcast a proposal that back-references to typically $2f + 1$ proposals from the previous round, forming a directed acyclic graph (DAG). These protocols exploit parallelism by allowing multiple leaders to be active concurrently and they can be viewed as a form of parallel composition, in which multiple ordering streams with different leader schedules operate concurrently over a shared DAG. Unlike Gatling, however, these streams remain coupled through the DAG’s reference structure, which imposes dependencies and constrains progress to the network delay.

Comparing Gatling with Multi-BFT Protocols. Recent work on Multi-BFT protocols explores running multiple PBFT instances in parallel to overcome the throughput limitations of single-leader designs.

For example, Mir-BFT [48] increases throughput by running multiple PBFT proposers in parallel, each responsible for a subset of log positions. To distribute work across proposers, transactions are partitioned into buckets and assigned according to a deterministic mapping. Confirmed proposals are merged into a single total order according to their assigned slot numbers. Mir-BFT further employs an epoch-based reconfiguration mechanism in which a designated epoch leader determines the assignment of transactions, slots, and ordering responsibilities across proposers. As a result, Mir-BFT’s parallel streams share protocol state, including epoch configuration, slot ownership, and bucket ownership, and collectively realize a coordinated partitioning of a single global log. Gatling also exploits parallelism through multiple concurrent consensus instances, but does so through a black-box composition of atomic broadcast protocols. Each instance operates independently and need not share proposer assignments, transaction partitions, or reconfiguration state with other instances; the only coordination across instances is a deterministic merge rule applied to their outputs. Consequently, Gatling merges independent logs rather than by partitioning a single log. Furthermore, while Mir-BFT can achieve low effective inter-proposal times under sustained load as a consequence of having multiple active proposers, Gatling explicitly engineers a low and regular inter-proposal cadence by staggering proposal schedules across instances. This design enables reliable and arbitrarily small inter-proposal times, optimizing for low transaction latency.

ISS [49] generalizes Mir-BFT by introducing a protocol-independent framework for parallel consensus. Transactions are partitioned across multiple consensus instances, whose outputs are deterministically merged into a single replicated log. This abstraction removes the need for Mir’s dedicated epoch leader and orderer-assignment mechanism while retaining the scalability benefits of parallel execution. Like ISS, Gatling composes multiple consensus instances and combines their outputs into a single log. However, ISS and Gatling target different goals. ISS uses parallelism to partition workload across independent streams and thereby increase throughput, whereas Gatling uses staggered proposal schedules to reduce transaction latency.

Ladon [31] follows a related approach, running parallel BFT protocols and deriving a single global order through monotonic block ranks. Leaders determine ranks using information collected across instances, and nodes confirm blocks according to rank-based thresholds. This cross-instance coordination distinguishes Ladon from Gatling in which the parallel instances operate independently as closed boxes.

Comparing Gatling with Shoal++. After completing this work, we were made aware that Shoal++ [6] uses a construction similar to Gatling in the context of DAG-based protocols to reduce the inter-proposal time. In particular, it runs three DAG instances in parallel, time-staggered by a fixed Δ offset, in a not-fully-generic “ajar-box” composition. By contrast, Gatling explores regimes with inter-proposal intervals below the network delay and introduces variants that preserve predictable validity even when the offsets become arbitrarily small. Moreover, Gatling is closed-box and works for general atomic broadcast component protocols, and our experiments evaluate the latency reduction solely attributable to time-staggered parallel composition.

References

- 1 Ittai Abraham and Kartik Nayak. What is responsiveness? <https://decentralizedthoughts.github.io/2022-12-18-what-is-responsiveness/>, 2022.
- 2 Ittai Abraham, Kartik Nayak, Ling Ren, and Zhuolun Xiang. Good-case latency of Byzantine broadcast: a complete categorization. In *PODC*, pages 331–341. ACM, 2021. doi:10.1145/3465084.3467899.
- 3 Orestis Alpos, Bernardo David, Jakov Mitrovski, Odysseas Sofikitis, and Dionysis Zindros. pod: An optimal-latency, censorship-free, and accountable generalized consensus layer. In *DISC*, volume 356 of *LIPICs*, pages 4:1–4:24. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICs.DISC.2025.4.
- 4 Kaya Alpturer, Kushal Babel, and Aditya Saraf. Timing games in responsive consensus protocols, 2025. [arXiv:2510.25144v1](https://arxiv.org/abs/2510.25144).
- 5 Anza. Alpenglow: A new consensus for solana. <https://www.anza.xyz/blog/alpenglow-a-new-consensus-for-solana>, 2024.
- 6 Balaji Arun, Zekun Li, Florian Suri-Payer, Sourav Das, and Alexander Spiegelman. Shoal++: High throughput DAG BFT can be fast and robust! In *NSDI*, pages 813–826. USENIX Association, 2025. URL: <https://www.usenix.org/conference/nsdi25/presentation/arun>.
- 7 Hagit Attiya, Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Bounds on the time to reach agreement in the presence of timing uncertainty. *J. ACM*, 41(1):122–152, 1994. doi:10.1145/174644.174649.
- 8 Kushal Babel, Andrey Chursin, George Danezis, Anastasios Kichidis, Lefteris Kokoris-Kogias, Arun Koshy, Alberto Sonnino, and Mingwei Tian. Mysticeti: Reaching the latency limits with uncertified DAGs. In *NDSS*. The Internet Society, 2025. URL: <https://www.ndss-symposium.org/ndss-paper/mysticeti-reaching-the-latency-limits-with-uncertified-dags/>.
- 9 Sam Blackshear, Andrey Chursin, George Danezis, Anastasios Kichidis, Lefteris Kokoris-Kogias, Xun Li, Mark Logan, Ashok Menon, Todd Nowacki, Alberto Sonnino, Brandon Williams, and Lu Zhang. Sui Lutris: A blockchain combining broadcast and consensus. In *CCS*, pages 2606–2620. ACM, 2024. doi:10.1145/3658644.3670286.
- 10 Ethan Buchman. Tendermint: Byzantine fault tolerance in the age of blockchains. Master’s thesis, University of Guelph, 2016. URL: <https://atrium.lib.uoguelph.ca/items/5459099e-67aa-4a23-83ae-d3471d8d8336>.
- 11 Vitalik Buterin. Ethereum: A next-generation smart contract and decentralized application platform. <https://ethereum.org/whitepaper/>, 2014.
- 12 Vitalik Buterin and Virgil Griffith. Casper the friendly finality gadget, 2019. [arXiv:1710.09437v4](https://arxiv.org/abs/1710.09437).
- 13 Miguel Castro and Barbara Liskov. Practical Byzantine fault tolerance. In *OSDI*, pages 173–186. USENIX Association, 1999. URL: <https://dl.acm.org/citation.cfm?id=296824>.
- 14 Celestia. Celestia: Modular blockchain network. <https://celestia.org/>, 2023.
- 15 Benjamin Y. Chan and Rafael Pass. Simplex consensus: A simple and fast consensus protocol. In *TCC (4)*, volume 14372 of *Lecture Notes in Computer Science*, pages 452–479. Springer, 2023. doi:10.1007/978-3-031-48624-1_17.
- 16 Benjamin Y. Chan and Elaine Shi. Streamlet: Textbook streamlined blockchains. In *AFT*, pages 1–11. ACM, 2020. doi:10.1145/3419614.3423256.
- 17 ChorusOne. Post on X. <https://x.com/ChorusOne/status/1955615856580301252>, 2024.
- 18 Brendan Kobayashi Chou, Andrew Lewis-Pye, and Patrick O’Grady. Minimit: Fast finality with even faster blocks, 2026. [arXiv:2508.10862v1](https://arxiv.org/abs/2508.10862).
- 19 Commonware. Commonware alto. <https://github.com/commonwarexyz/alto>, 2024. GitHub repository.
- 20 Commonware. Commonware monorepo. <https://github.com/commonwarexyz/monorepo>, 2024. GitHub repository.

- 21 George Danezis, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. Narwhal and Tusk: a DAG-based mempool and efficient BFT consensus. In *EuroSys*, pages 34–50. ACM, 2022. doi:10.1145/3492321.3519594.
- 22 Isaac Doidge, Raghavendra Ramesh, Nibesh Shrestha, and Joshua Tobkin. Moonshot: Optimizing block period and commit latency in chain-based rotating leader BFT. In *DSN*, pages 470–482. IEEE, 2024. doi:10.1109/DSN58291.2024.00052.
- 23 Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, 1988. doi:10.1145/42282.42283.
- 24 Pranav Garimidi, Joachim Neu, and Max Resnick. Multiple concurrent proposers: Why and how. Cryptology ePrint Archive, Paper 2025/1772, 2025. URL: <https://eprint.iacr.org/2025/1772>.
- 25 Rati Gelashvili, Lefteris Kokoris-Kogias, Alberto Sonnino, Alexander Spiegelman, and Zhuolun Xiang. Jolteon and Ditto: Network-adaptive efficient consensus with asynchronous fallback. In *Financial Cryptography*, volume 13411 of *Lecture Notes in Computer Science*, pages 296–315. Springer, 2022. doi:10.1007/978-3-031-18283-9_14.
- 26 Amir Herzberg and Shay Kutten. Early detection of message forwarding faults. *SIAM J. Comput.*, 30(4):1169–1196, 2000. doi:10.1137/S0097539796312745.
- 27 Dakai Kang, Suyash Gupta, Dahlia Malkhi, and Mohammad Sadoghi. HotStuff-1: Linear consensus with one-phase speculation. *Proc. ACM Manag. Data*, 3(3):171:1–171:29, 2025. doi:10.1145/3725308.
- 28 Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. All you need is DAG. In *PODC*, pages 165–175. ACM, 2021. doi:10.1145/3465084.3467905.
- 29 Lucianna Kiffer, Joachim Neu, Srivatsan Sridhar, Aviv Zohar, and David Tse. Nakamoto consensus under bounded processing capacity. In *CCS*, pages 363–377. ACM, 2024. doi:10.1145/3658644.3670347.
- 30 Andrew Lewis-Pye and Patrick O’Grady. The carnot bound: Limits and possibilities for bandwidth-efficient consensus, 2026. arXiv:2603.11797v3.
- 31 Hanzheng Lyu, Shaokang Xie, Jianyu Niu, Chen Feng, Yinqian Zhang, and Ivan Beschastnikh. Ladon: High-performance multi-BFT consensus via dynamic global ordering. In *EuroSys*, pages 226–242. ACM, 2025. doi:10.1145/3689031.3696102.
- 32 Dahlia Malkhi and Kartik Nayak. Extended abstract: HotStuff-2: Optimal two-phase responsive BFT. Cryptology ePrint Archive, Paper 2023/397, 2023. URL: <https://eprint.iacr.org/2023/397>.
- 33 Dahlia Malkhi, Chrysoula Stathakopoulou, and Maofan Yin. BBKA-Chain: Low latency, high throughput BFT consensus on a DAG. In *FC (1)*, volume 14744 of *Lecture Notes in Computer Science*, pages 51–73. Springer, 2024. doi:10.1007/978-3-031-78676-1_4.
- 34 Jean-Philippe Martin and Lorenzo Alvisi. Fast Byzantine consensus. *IEEE Trans. Dependable Secur. Comput.*, 3(3):202–215, 2006. doi:10.1109/TDSC.2006.35.
- 35 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>, 2009.
- 36 Joachim Neu, Srivatsan Sridhar, Lei Yang, David Tse, and Mohammad Alizadeh. Longest chain consensus under bandwidth constraint. In *AFT*, pages 126–147. ACM, 2022. doi:10.1145/3558535.3559777.
- 37 Joachim Neu, Ertem Nusret Tas, and David Tse. The availability-accountability dilemma and its resolution via accountability gadgets. In *Financial Cryptography*, volume 13411 of *Lecture Notes in Computer Science*, pages 541–559. Springer, 2022. doi:10.1007/978-3-031-18283-9_27.
- 38 Rafael Pass and Elaine Shi. Hybrid consensus: Efficient consensus in the permissionless model. In *DISC*, volume 91 of *LIPIcs*, pages 39:1–39:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.DISC.2017.39.
- 39 Alejandro Ranchal-Pedrosa, Benjamin Marsh, Lefteris Kokoris-Kogias, and Alberto Sonnino. Sedna: Sharding transactions in multiple concurrent proposer blockchains, 2025. arXiv:2512.17045v1.

- 40 Giulia Scaffino, Max Resnick, and Joachim Neu. Gatling: Rapid-fire consensus from parallel composition, 2026. [arXiv:2606.18220v1](#).
- 41 Caspar Schwarz-Schilling, Fahad Saleh, Thomas Thiery, Jennifer Pan, Nihar Shah, and Barnabé Monnot. Time is money: Strategic timing games in proof-of-stake protocols. In *AFT*, volume 282 of *LIPIcs*, pages 30:1–30:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.AFT.2023.30.
- 42 Victor Shoup. Sing a song of Simplex. In *DISC*, volume 319 of *LIPIcs*, pages 37:1–37:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.DISC.2024.37.
- 43 Victor Shoup, Jakub Sliwinski, and Yann Vonlanthen. Kudzu: Fast and simple high-throughput BFT. In *DISC*, volume 356 of *LIPIcs*, pages 42:1–42:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.DISC.2025.42.
- 44 Nibesh Shrestha and Aniket Kate. Hydrangea++: Enhancing Hydrangea with optimistic proposals. <https://supra.com/documents/hydrangea-plus-plus.pdf>, 2025.
- 45 Nibesh Shrestha, Aniket Kate, and Kartik Nayak. Hydrangea: Optimistic two-round partial synchrony with improved fault resilience. Cryptology ePrint Archive, Paper 2025/1112, 2025. URL: <https://eprint.iacr.org/2025/1112>.
- 46 Alexander Spiegelman, Balaji Arun, Rati Gelashvili, and Zekun Li. Shoal: Improving DAG-BFT latency and robustness. In *FC (1)*, volume 14744 of *Lecture Notes in Computer Science*, pages 92–109. Springer, 2024. doi:10.1007/978-3-031-78676-1_6.
- 47 Alexander Spiegelman, Neil Giridharan, Alberto Sonnino, and Lefteris Kokoris-Kogias. Bullshark: DAG BFT protocols made practical. In *CCS*, pages 2705–2718. ACM, 2022. doi:10.1145/3548606.3559361.
- 48 Chrysoula Stathakopoulou, Tudor David, Matej Pavlovic, and Marko Vukolic. Mir-BFT: Scalable and robust BFT for decentralized networks. *J. Syst. Res.*, 2(1), 2022.
- 49 Chrysoula Stathakopoulou, Matej Pavlovic, and Marko Vukolic. State machine replication scalability made simple. In *EuroSys*, pages 17–33. ACM, 2022. doi:10.1145/3492321.3519579.
- 50 Thomas Thiery, Francesco D’Amato, Julian Ma, Barnabé Monnot, Terence Tsao, Jacob Kaufmann, and Jihoon Song. EIP-7805: Fork-choice enforced inclusion lists (FOCIL). Ethereum Improvement Proposals, 2024. URL: <https://eips.ethereum.org/EIPS/eip-7805>.
- 51 Anatoly Yakovenko. Solana: A new architecture for a high performance blockchain v0.8.13. <https://solana.com/solana-whitepaper.pdf>, 2018.
- 52 Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan-Gueta, and Ittai Abraham. HotStuff: BFT consensus with linearity and responsiveness. In *PODC*, pages 347–356. ACM, 2019. doi:10.1145/3293611.3331591.