Analysis of Deterministic Longest-Chain Protocols

Elaine Shi

November 6, 2018

Abstract

Most classical consensus protocols rely on a leader to coordinate nodes' voting efforts. One novel idea that stems from blockchain-style consensus is to rely, instead, on a "longest-chain" idea for such coordination. Such a longest-chain idea was initially considered in *randomized* protocols, where in each round, a node has some probability of being elected a leader who can propose the next block. Recently, well-known systems have started implementing the *deterministic* counterpart of such longest-chain protocols — the deterministic counterpart is especially attractive since it is even simpler to implement than their randomized cousins. A notable instantiation is the Aura protocol which is shipped with Parity's open-source Ethereum implementation.

Interestingly, mathematical analyses of deterministic, longest-chain protocols are lacking even though there exist several analyses of randomized versions. In this paper, we provide the first formal analysis of deterministic, longest-chain-style consensus. We show that a variant of the Aura protocol can defend against a Byzantine adversary that controls fewer than $\frac{1}{3}$ fraction of the nodes, and this resilience parameter is tight by some technical interpretation. Based on insights gained through our mathematical treatment, we point out that Aura's concrete instantiation actually fails to achieve the resiliene level they claim.

Finally, while our tight proof for the longest-chain protocol is rather involved and non-trivial; we show that a variant of the "longest-chain" idea which we call "largest-set" enables a textbook construction that admits a simple proof (albeit with slower confirmation).

1 Introduction

A central abstraction in distributed systems is called *consensus* or *state machine replication* [6, 13, 23]. In state machine replication, a distributed system of nodes reach agreement on an ever-growing, linearly ordered log. State machine replication protocols have been studied for three decades [6, 13, 15] and two major approaches exist: classical-style consensus (e.g., PBFT [6] and Paxos [13]) and blockchain-style consensus (e.g., Bitcoin and Ethereum's proof-of-work blockchains and proof-of-stake blockchain constructions [7, 8, 12, 21]).

In both classical and blockchain consensus, nodes vote to reach a decision. In classical consensus, the voting process is typically coordinated by a leader (possibly rotating over time); and whenever the leader is honest and proposes the same item (e.g., block or batch of transactions) to everyone, everyone will vote on the same item and a decision can be made very soon by collecting majority or super-majority votes (possibly over multiple rounds of voting). In blockchain-style protocols, the leader election and voting processes are coalesced into one: whenever one is elected leader it has the opportunity to cast a vote. To coordinate nodes voting effort, a new "longest-chain" idea is adopted — in some sense, people vote on the most-popular history observed so far. Such coordination is important for consistency: intuitively, if nodes fail to concentrate their voting efforts, it can lead to many forks and no chain will emerge as the dominant decision. Besides using the longest chain

idea as a coordination mechanism, the blockchain data structure is also elegant because it creates a hash-chain among the entire history: by cryptographically binding each block to the entire history, in some sense each decision reinforces the agreement on past blocks.

In the past few years, the community has made amazing progress in formally understanding the security properties of blockchain-style consensus [8, 10, 12, 19–21]. While most existing mathematical analyses focus on randomized blockchain protocols, some practical systems [2, 3] have started adopting *deterministic* blockchain protocols that also embody the "longest-chain" idea, possibly because the deterministic variants are even simpler to implement than their randomized cousins. Unfortunately, to the best of our knowledge, the security of such deterministic blockchain constructions is not yet well-understood.

A representative deterministic, longest-chain-style protocol is Aura [2] (short for Authority Round)[1]. Aura is well-known in the cryptocurrency community since it is one of the few blockchain consensus algorithms shipped with Parity's open-source Ethereum implementation [1]. Roughly speaking, Aura's blockchain protocol works as follows:

- There are n nodes each with a well-known public key. In round r, node ($r \mod n$) is the leader.
- In every round r , the leader chooses the longest chain seen so far, and extends the chain by signing the next block containing 1) the current round number r, henceforth also called the block's *timestamp*, 2) a batch of transactions, and 3) the hash of the parent block.
- In a valid blockchain, the blocks' timestamps must strictly increase, and moreover honest nodes reject blocks with timestamps in the future.
- At any time, a node takes its longest chain and chops off some number of blocks at the end, and the prefix is considered finalized.

Specifically, Aura makes the following choices and claims: for finalization, they choose to chop off $\lfloor \frac{n}{2} \rfloor + 1$ trailing blocks signed by distinct signers. Their documentation claims that the protocol defends against *minority* Byzantine corruptions.

1.1 Our Results and Contributions

Formal analyses of deterministic blockchains. To the best of our knowledge we are the first to provide mathematical proofs for deterministic, longest-chain-style protocols. More concretely, we make the following contributions:

- We present a variant of Aura's protocol that is provably secure. In this variant, in round r, all blocks in the longest chain whose timestamps are at most $r - n$ are considered finalized (*c.f.* Aura chops off only $\lfloor \frac{n}{2} \rfloor + 1$ blocks from distinct signers).
- We prove that this variant ensures consistency and liveness in the presence of any adversary that can control fewer than $\frac{1}{3}$ fraction of the nodes; further, transaction confirmation takes $\Theta(n)$ rounds since we need to chop off $\Theta(n)$ trailing blocks for finalization. Note that the

[1] Another notable example is EOS's consensus protocol called DPoS — EOS is among the top five cryptocurrencies by market cap at the time of the writing. Although in online blog posts and forums [3,4] there have been back-and-forth discussions about the DPoS algorithm, we have not found a formal specification or an open-source implementation from EOS. It also looks likely that EOS might implement a *partition-tolerant* variant of the longest-chain idea, and not the family of algorithms we analyze in this paper.

round complexity of this protocol is asymptotically optimal due to the well-known lower bound [5, 9, 11, 14, 16, 17] that any *deterministic* consensus protocol must incur at least $f + 1$ rounds where f is the number of corrupt nodes (and this lower bound holds even for crash fault).

- Unlike existing proofs for randomized blockchains that require stochastic analyses that reason about the good properties respected by an overwhelming fraction of execution traces (as opposed to the worst-case execution trace), our proof relies on elementary discrete math — nonetheless as we show, proving an asymptotically tight bound on transaction confirmation time is rather *non-trivial*. The most sophisticated part of the proof involves a combinatorics argument that a particular good event called a "pivot" appears every n rounds even against the worst-case attack.

- We show that the $\frac{1}{3}$ resilience parameter is tight for this protocol; specifically, we show that if the adversary can control how honest nodes break ties among multiple chains of the same length, then there is an explicit $\frac{1}{3}$ attack against this protocol (note that our security proof admits arbitrary tie-breaking).

- Last but not the least, inspired by the longest-chain idea, we construct a new consensus protocol that relies on a similar "largest-set" idea, i.e., nodes always vote on the item that has collected the most number of votes so far. This protocol is of particular interest for pedagogy (e.g., to illustrate why and how a longest-chain-type idea works), since it admits a simple and elementary proof; moreover the proof is in some sense a deterministic counterpart of the recent analyses of randomized blockchains [10, 12, 19, 21].

We point out that it is an interesting observation that the most natural embodiment of the "deterministic, longest-chain" idea defends only against less than $\frac{1}{3}$ corruptions while in comparison, the randomized counterparts can secure against up to minority corruption [10, 12, 19, 21]. Thus for future work, a technically intriguing question is whether there exist other natural embodiments of the "deterministic, longest-chain" idea that achieves stronger resilience.

Analysis of Aura's instantiation. Based on our mathematical analysis, we reflect on Aura's specific choices and claims for their deterministic blockchain instantiation.

As mentioned, our deterministic, longest-chain protocol described above defends against any $< \frac{1}{3}$ attack, and this is tight when tie-breaking can be arbitrary (for scoring two chains of the same length). Aura's instantiation employs a specific tie-breaking rule: when two chains are of the same length, they prefer the "earlier" one, i.e., whose last block has an earlier timestamp. We show that even under their tie-breaking rule, and even if we chop off all blocks with the most recent n rounds for finalization (i.e., more than what they actually chop off), there is a $\frac{3}{7}$ attack against consistency. Further, if we actually chop off only $\lfloor \frac{n}{2} \rfloor + 1$ trailing blocks with distinct signers for finalization (and using their tie-breaking rule), there is an explicit attack that breaks consistency when the adversary controls slightly more than $\frac{3}{8}$ fraction of nodes. We conclude that Aura's claim of defending against minority corruptions is incorrect under the classical notion of a Byzantine adversary. Our analysis also implies that their concrete instantiation actually achieves a resilience parameter that is between $\frac{1}{6}$ and $\frac{3}{8}$.

Disclosure and recommendations to Parity. We have emailed Parity regarding our findings, and have suggested that they chop off more blocks from the end for higher resilience, and weaken their public claim that they defend against upto minority Byzantine corruptions.

2 Preliminaries

2.1 Model

We consider a standard, *synchronous* model of execution, where an adversary $\mathcal{A}$ is allowed to adaptively corrupt nodes during the protocol execution. In this section, we describe the model in more detail and introduce relevant notations.

Protocol execution. We use the standard Interactive Turing Machine (ITM) approach to model a protocol's execution. Each *node* in the protocol (also referred to as a party or a protocol participant) is modeled as an ITM. Besides the protocol participants, there are also two special ITMs, the adversary denoted $\mathcal{A}$, and the environment denoted $\mathcal{Z}$ respectively.

The environment $\mathcal{Z}$ spawns a set of n nodes at the start of protocol execution; further $\mathcal{Z}$ is responsible for providing inputs to honest nodes and receiving output from honest nodes. At any time during protocol execution, a node is either *honest* or has been *corrupt*. All corrupt nodes are in control of the adversary $\mathcal{A}$, i.e., all messages received by corrupt nodes are forwarded to $\mathcal{A}$, and $\mathcal{A}$ controls all corrupt nodes' actions. The adversary $\mathcal{A}$ and the environment $\mathcal{Z}$ are allowed to communicate freely at any time during the protocol.

Communication model. We assume a *fully synchronous* communication model, where honest nodes' messages are delivered to honest recipients at the beginning of the next round.

Our protocol adopts a *multicast* communication pattern — whenever an honest node multicasts a message, it is destined for all nodes including itself. Corrupt nodes, on the other hand, may send messages to only a subset of the nodes.

We assume a rushing adversary, i.e., the adversary $\mathcal{A}$ can observe messages sent by honest nodes in a round before deciding what messages corrupt nodes will send in the same round.

Corruption model. We assume that $\mathcal{A}$ may adaptively corrupt nodes during protocol execution. Given a protocol that has a fixed termination time T_{end} (which may be a function dependent on total number of nodes), if a node becomes corrupt in or before round T_{end}, we say that the node is *eventually-corrupt*; else we say that the node is *honest-forever*. The protocols described in this paper will have provable security when $n \geq 3f + 1$ where f is the number of eventually-corrupt nodes, and n denotes the total number of nodes.

Henceforth we say that $(\mathcal{A}, \mathcal{Z})$ is (n, f)-respecting w.r.t. some protocol Π iff with probability 1 during an execution of Π, $(\mathcal{A}, \mathcal{Z})$ spawns n nodes and adaptively corrupts up to f number of nodes by the end of protocol execution.

2.2 Byzantine Agreement: Definitions

For simplicity, we begin our exposition not with state machine replication which aims to confirm a linearly ordered log through repeatedly reaching consensus, but with the 1-bit single-shot version, i.e., byzantine agreement. We formally define byzantine agreement below. Later in Section 4, we describe how to naturally extend our protocol to support state machine replication.

In a byzantine agreement protocol, there is a designated sender that is part of the common knowledge. Henceforth, without loss of generality, we may assume that node 0 is the designated sender.

Syntax. Prior to protocol start, the sender receives an input $b \in \{0, 1\}$ from the environment $\mathcal{Z}$. At the end of the protocol, every node i (including the sender) outputs a bit b_i to the environment $\mathcal{Z}$.

Security definition. A byzantine agreement protocol must satisfy consistency, validity, and T_{end}-termination. Specifically, let $T_{\text{end}} := \mathsf{poly}(\kappa, n)$ be a polynomial in κ and n we say that a protocol Π satisfies consistency, validity, or T_{end}-termination w.r.t. $(\mathcal{A}, \mathcal{Z})$ iff there exists a negligible function $\mathsf{negl}(\cdot)$ such that for every $\kappa \in \mathbb{N}$, except with $\mathsf{negl}(\kappa)$ probability over the choice of $\mathsf{view} \leftarrow_{\$} \mathsf{EXEC}^{\Pi}(\mathcal{A}, \mathcal{Z}, \kappa)$, the following properties hold:

- *Consistency.* If an honest node outputs b_i and another honest node outputs b_j in view, then it must hold that $b_i = b_j$.
- *Validity.* If the sender remains honest till the end of the execution and the sender's input is b in view, then all honest nodes must output b in view.
- T_{end}*-termination.* By the end of round $T_{\text{end}}(\kappa, n)$, all honest nodes output a bit.

We say that a protocol Π satisfies consistency, validity, or T_{end}-termination in (n, f)-environments iff iff for every p.p.t. $(\mathcal{A}, \mathcal{Z})$ that is (n, f)-respecting w.r.t. Π, Π satisfies consistency, validity, or T_{end}-termination w.r.t. $(\mathcal{A}, \mathcal{Z})$.

2.3 Notations and Assumptions

We assume a permissioned, classical setting with a public-key infrastructure (PKI). In other words, assume that there are n nodes numbered $0, 1, 2, \ldots, n-1$ whose public keys are common knowledge. Henceforth we use the notation pk_i to denote node i's public key for $i \in [n]$.

Without loss of generality, we shall assume that node 0 is the designated sender.

3 A Determinisitic, Longest-Chain Protocol

As mentioned, for clarity, we begin our exposition with byzantine agreement (i.e., 1-bit single-shot consensus). Later in Section 4 we show how our protocol naturally extends to state machine replication (i.e., agreeing on a linearly ordered log through repeated consensus).

3.1 Useful Definitions and Notations

Leader. In every round i, node $(i \mod n)$ is the leader. Thus each round has a unique leader.

Valid blockchain. A blockchain of length L is an ordered list of blocks where each block is of the form $\mathsf{chain} := \{(t_\ell, \mathsf{data}_\ell, \sigma_\ell)\}_{\ell \in [L]}$ where $t_\ell \in \mathbb{N}$ denotes the purported round in which this block was mined (also referred to the block's *timestamp*), data_ℓ denotes an arbitrary payload string, σ_ℓ denotes a signature of the chain up to length ℓ.

Henceforth given a chain $\mathsf{chain} := \{(t_\ell, \mathsf{data}_\ell, \sigma_\ell)\}_{\ell \in [L]}$, we define $\mathsf{extract}(\mathsf{chain})$ to be the operator removes the signature from every block outputting $\{(t_i, \mathsf{data}_i)\}_{i \in [L]}$. We use $\mathsf{chain}[i]$ to denote the i-th block in chain and we use $\mathsf{chain}[:i]$ to denote the prefix of chain upto the i-th block (inclusive).

A blockchain $\mathsf{chain} := \{(t_\ell, \mathsf{data}_\ell, \sigma_\ell)\}_{\ell \in [L]}$ is said to be valid iff

- For any $1 \le i < j \le L$, it holds that $t_i < t_j$;
- For any $\ell \in [L]$, σ_ℓ is a valid signature of the prefix $\mathsf{extract}(\mathsf{chain}[:\ell - 1])$ under pk_i where $i = t_\ell \mod n$ is the leader of round t_ℓ.

Remark 1. *Note that a practical optimization is to build a hash chain: let* $h_0 = 0$ *and let* $h_i := H(h_{i-1}, t_i, \mathsf{data}_i)$*. In this case, the signature for the* i*-th block would be computed over* h_i*; and further we include* h_i *in the block* $\mathsf{chain}[i]$*.*

3.2 Protocol Π_{chain}

- Round 0: sender (i.e., node 0) signs its input and multicasts the input along with the signature.
- For every round $r = 1, 2, \ldots$, every node (including the sender) performs the following:
  1. Receive chains from the network. Discard any invalid chains and any chains containing timestamps greater than or equal to r. Chains that are not discarded are considered part of the nodes' view so far.
  2. Let chain be the longest valid blockchain that has been observed in the node's view so far[2]. If there are multiple such longest chains, break ties arbitrarily. If no such chain has been observed, let chain be the empty chain.
  3. If the current node is the leader of round r, then let data be the set

     $$\{(b, \sigma_b) : b \in \{0,1\} \text{ and at least one valid signature } \sigma_b \text{ has been observed for } b\}$$

     Sign $\mathsf{extract}(\mathsf{chain})||(r, \mathsf{data})$ thus obtaining the signature σ; and multicast the new chain

     $$\mathsf{chain}||(r, \mathsf{data}, \sigma)$$

- At the beginning of round T_{end} (the choice of T_{end} will be stated later), do the following. Let chain be the current longest valid blockchain in view. In all blocks in chain whose timestamp is at most n, look for the first bit b with a valid signature from the sender, and output the bit b. If such a bit is not found, output 0.

Parameters. We will prove that the above protocol satisfies the definition of byzantine agreement for $T_{\text{end}} \geq 2n + 1$, as long as the adversary controls fewer than $\frac{n}{3}$ nodes.

Note that our consistency proof actually shows that in round $r > n$, all blocks in an honest node's chain with round numbers $r - n$ or smaller must have stabilized. However, here the termination parameter T_{end} is set to $2n+1$ to account for the fact that the earliest block containing the sender's signature may not be the first block in a blockchain.

3.3 Proofs

For clarity, in the main body of the paper, we first present a simpler proof but for worse termination parameters. Specifically, let n denote the total number of nodes; let f denote the number of corrupt nodes; and let $\epsilon := \frac{1}{3} - \frac{f}{n}$. In this looser but simpler proof, we will set $T_{\text{end}} := (1 + \lceil \frac{1}{\epsilon} \rceil)n$. In particular, if $n = 3f + 1$, then the protocol will run for $\Theta(n^2)$ number of rounds.

Later in Appendix B, we will present a tighter proof for the termination parameter $T_{\text{end}} = 2n+1$. However, this tighter proof requires several additional non-trivial techniques.

[2]Henceforth this is referred to as the node's longest chain in the current round r. Note that here we refer to the longest chain at the *beginning* of the round.

3.3.1 Additional Terminology

We define the following helpful terminology.

- *Blocks.* For convenience, when the context is clear, we also use the term "the i-th block in chain" as an alias for the prefix of chain upto the i-th block, i.e., $\mathsf{chain}[:i]$. Given two blockchains chain and chain', the two length-ℓ blocks $\mathsf{chain}[:\ell]$ and $\mathsf{chain}'[:\ell]$ are said to be distinct iff $\mathsf{extract}(\mathsf{chain}[:\ell])$ and $\mathsf{extract}(\mathsf{chain}'[:\ell])$ are distinct (i.e., we ignore the signature when defining distinctness for blocks).

- *Honest-forever and eventually-corrupt blocks.* A block whose timestamp corresponds to an honest-forever leader round is said to be an *honest-forever block*; otherwise it is said to be an *eventually-corrupt block*.

- *Honest chain.* If in some execution trace view, some honest node's longest chain in round r is chain, then we say that chain is an *honest chain* in view (or an honest chain in round r in view).

- *Honest-forever and eventually corrupt leader round.* A round t is said to be an honest-forever leader round if node $(t \mod n)$ remains honest forever in view; else t is said to be an eventually-corrupt leader round.

In our proofs, we often ignore the negligible fraction of views where signature forgery occur, and prove properties for the remaining good views. This is formalized in the fact below.

Fact 1. *Suppose that the signature scheme is secure. Then, except with negligible probability over the choice of* view*, if the leader for round* t*, i.e., node* $i := (t \mod n)$ *is honest-forever in* view*, then no honest node will ever have in its view a signature* σ *on a message that is valid under* pk_i *— but node* i *never signed the message in* view*.*

Proof. Straightforward by signature security. □

3.3.2 Validity

Special boundary blocks. In the remainder of this section, we shall assume that an honest node's longest chain in some round r consists of *i)* an imaginary "end-of-chain block" whose timestamp is r, and whose length is the original chain length plus 1; and *ii)* an imaginary genesis block whose timestamp is 0 and whose chain length is 0.

The special genesis and end-of-chain blocks are assumed to be honest-forever blocks.

Lemma 1 (Chain growth)**.** *With probability* 1*, the following holds. Let* chain *denote some node's longest chain in some round. Suppose that* $\mathsf{chain}[\ell] := (t, \_, \_)$ *and* $\mathsf{chain}[\ell'] := (t', \_, \_)$ *are both honest-forever blocks (including genesis or end-of-chain blocks) in* chain *and moreover* $t' - t - 1 \geq kn$ *for some positive integer* k*. It holds that* $\ell' - \ell - 1 \geq k(n - f)$*.*

Proof. Consider any execution trace view: suppose that in some round r in view an honest node i is the leader and suppose that node i's longest chain is of length ℓ_r at the beginning of round r. By definition of the protocol, every honest-forever node's longest chain will be of length at least $\ell_r + 1$ at the beginning of round $r + 1$.

The remainder of the proof follows by observing that with probability 1, in every n consecutive rounds, at least $(n - f)$ rounds have leaders that are honest-forever, and for honest node $(r \mod n)$, its longest chain is of length $\ell' - 1$ in round r (and thus it appended a new block at length ℓ'). □

Lemma 2 (Chain quality)**.** *With probability* 1*, the following holds. Let* chain *denote an honest node's longest chain in round* $r+1$*. Then, for any* n *consecutive rounds* $[t+1, t+n] \subseteq [r]$*, there must be an honest-forever block in* chain *with a timestamp in the range* $[t+1, t+n]$*.*

Proof. For the sake of contradition, suppose that in some execution trace view, in some honest chain chain, there exist honest-forever blocks $\mathsf{chain}[\ell] := (t, \_, \_)$ and $\mathsf{chain}[\ell'] := (t', \_, \_)$ such that 1) there are no honest-forever blocks in between; and 2) $t' - t - 1 \geq n$. Let $k \geq 1$ be the largest integer such that $t' - t - 1 \geq kn$. By Lemma 1, $\ell' - \ell - 1 \geq k(n-f)$. Now, all blocks $\mathsf{chain}[\ell+1..\ell'-1]$ are eventually-corrupt blocks with distinct timestamps in the range (t, t'). Since there are at most f eventually-corrupt nodes, we have that at most $(k+1)f$ of the timestamps in the range (t, t') correspond to eventually-corrupt leaders.

To reach a contradiction, it suffices to observe that the following cannot be true given that $n > 3f+1$ and $\ell' - \ell - 1 \geq k(n-f)$:

$$\ell' - \ell - 1 \leq (k+1)f$$

□

Theorem 1 (Validity and liveness)**.** *Suppose that the signature scheme is secure and that* $n \geq 3f+1$*. Then, except with negligible probability, every honest-forever node will output a bit (i.e., liveness holds). Moreover, if the sender is honest-forever, then every honest-forever node must output the sender's input bit (i.e., validity holds).*

Proof. Liveness trivially holds by protocol definition. For validity, if the sender is honest-forever, it will sign its input bit and multicast it in round 0. Due to Lemma 2, at the beginning of round T_{end}, any honest node's longest chain must have an honest-forever block whose timestamp is between $[1, n]$. By protocol definition, this honest-forever block must contain the sender's signature on its input bit. The remainder of the validity proof follows in a straightforward fashion by Fact 1. □

3.3.3 Consistency

Fact 2. *Assume that the signature scheme is secure. Then, except with negligible probability, the following holds. Suppose that node* i *is an honest-forever leader in some round* t *and signed a chain of length* ℓ *in round* t*; further, suppose that node* j *is an honest-forever leader in some round* $t' \neq t$ *and it signed a chain of length* ℓ' *in round* t'*: we have that* $\ell \neq \ell'$*.*

Proof. Straightforward by protocol definition. When node i signs a chain of length ℓ in round t, in every round $r > t$, every honest node's longest chain must be of length at least ℓ and thus no honest node will ever sign a chain of length ℓ in any round $r > t$. □

Fact 3. *Suppose that the signature scheme is secure. Then except with negligible probability over the choice of* view*, the following holds. Let* chain *and* chain' *denote two honest chains in* view*. Suppose that* $\mathsf{chain}[\ell]$ *and* $\mathsf{chain}'[\ell]$ *are both honest-forever blocks. Then, it holds that* $\mathsf{chain}[:\ell] = \mathsf{chain}'[:\ell]$*.*

Proof. Straightforward if we ignore the negligible fraction of views where bad events related to Fact 1 take place, and from Fact 2. □

Given two blockchains chain and chain', if $\mathsf{extract}(\mathsf{chain}[:\ell]) = \mathsf{extract}(\mathsf{chain}'[:\ell])$, we say that chain and chain' share the block $\mathsf{chain}[:\ell]$.

Lemma 3. *Except with negligible probability over the choice of* view*, the following holds. Let* chain *and* chain′ *be the longest chains of two honest nodes at the beginning of round* T_{end}*. Suppose that the last honest-forever block shared by* extract(chain) *and* extract(chain′) *has the timestamp* t^**. It must hold that* $T_{\text{end}} - t^* - 1 < n \cdot \lceil \frac{1}{\epsilon} \rceil$*.*

Proof. Suppose not, i.e., suppose that $T_{\text{end}} - t^* - 1 \geq n \cdot \lceil \frac{1}{\epsilon} \rceil$. Let $k \geq \lceil \frac{1}{\epsilon} \rceil$ be the largest integer such that $T_{\text{end}} - t^* - 1 \geq kn$. Recall that "block" is used as an alias of a prefix of a blockchain and that distinctness of two blocks is defined disregarding the signatures (see Section 3.3.1).

Due to chain growth (Lemma 1), there are at least $k(n-f)$ blocks after time t^* in either chain or chain′ (not counting the end-of-chain special boundary block). Due to chain quality (Lemma 2) and Fact 3, chain and chain′ must diverge in round $t^* + n$ — otherwise the block at t^* cannot be the latest shared honest-forever block of chain and chain′. This means that the number of distinct blocks in chain and chain′ after time t^* must be at least $2k(n-f) - n$.

On the other hand, during every honest-forever leader round, only a unique block will be signed by the honest-forever leader. For every eventually-corrupt round t, chain and chain′ can each contain a block with the timestamp t. Thus, we have that the total number of distinct blocks after time t^* in chain and chain′ must be upper bounded by $(k+1)f + (k+1)n$.

Our goal is to show the following inequality which would suffice for reaching a contradiction:

$$2k(n-f) - n > (k+1)f + (k+1)n$$

Plugging in $f = (\frac{1}{3} - \epsilon)n$, it suffices to show that $2k(\frac{2}{3} + \epsilon)n - n \geq (k+1)(\frac{1}{3} - \epsilon)n + (k+1)n$. Simplifying the above, it suffices to show that $2k(\frac{2}{3} + \epsilon) - 1 \geq (k+1)(\frac{1}{3} - \epsilon) + (k+1)$, i.e., $\frac{4}{3}k + 2k\epsilon - 1 \geq \frac{1}{3}k + \frac{1}{3} - k\epsilon - \epsilon + (k+1)$, i.e., $\epsilon + 3k\epsilon > 2 + \frac{1}{3}$. The last inequality holds as long as $k \geq \lceil \frac{1}{\epsilon} \rceil$. □

Theorem 2 (Consistency for Π_{chain})**.** *Assume that the signature scheme is secure and let* $T_{\text{end}} := (1 + \lceil \frac{1}{\epsilon} \rceil)n$*. Then, the longest-chain protocol* Π_{chain} *satisfies consistency in* (n, f)*-environments as long as* $n \geq 3f + 1$*.*

Proof. Straightforward from Lemma 3 and the definition of the protocol. □

Tighter proof. In Appendix B, we present a tighter but more sophisticated proof for $T_{\text{end}} := 2n+1$. The most technical part of the proof is a combinatorics argument that a certain good event called a "pivot" happens every n rounds.

3.4 An Explicit Attack with $\frac{1}{3}$ Corruption

We now show that our analysis is tight in resilience by demonstrating a $\frac{1}{3}$ attack against the deterministic blockchain protocol Π_{chain}. We consider 8 nodes numbered $0, 1, \ldots, 8$, and every node i that is a multiple of 3 is corrupt; every remaining node is honest. Note that exactly $\frac{1}{3}$ of the nodes are corrupt. The goal of the adversary is to maintain two equal-length chains forever; and he does so by diverging the honest nodes' mining attempts. The attack is graphically illustrated below where red blocks denote corrupt blocks and green blocks denote honest blocks — we will explain the attack in detail below.

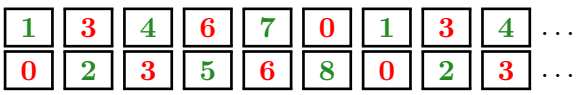

We now explain how the attack works — we assume that when there are two chains of equal length, the adversary can control how honest nodes break ties, e.g., by delivering one chain slightly

earlier than the other within the same round. Note that our proofs earlier admit arbitrary tie-breaking. For convenience, we number the rounds starting from 0 too. Henceforth, without risk of ambiguity, we use a block's timestamp to denote the block.

- Round 0: corrupt node 0 signs a block but he does not release this block yet.
- Round 1: the honest node 1 mines a block. At this moment, the adversary releases the 0-block.
- Round 2: the honest node 2 sees two longest chains, the "0"-chain and the "1"-chain respectively. He chooses to extend the "0"-chain and mines the next block 2.
- Round 3: corrupt node 3 extends the "1"-chain and releases this block. Additionally, he also extends the "0-2"-chain, but he withholds this block.
- Round 4: the honest node 4 extends the "1-3"-chain. In the same round, the adversary releases the "0-2-3"-chain.
- Round 5: the honest node 5 sees two longest chains, the "1-3-4"-chain and the "0-2-3"-chain. He chooses to extend the "0-2-3"-chain.
- Round 6: corrupt node 6 extends the "1-3-4"-chain and releases this block. Additionally, he also extends the "0-2-3-5"-chain but he withholds this block.
- Round 7: the honest node 7 extends the "1-3-4-6"-chain. In the same round, the adversary releases the "0-2-3-5-6"-chain.
- Round 8: honest node 8 sees two longest chains ending at the blocks 6 and 7 respectively. He chooses to extend the chain ending at 6.
- Round 9: corrupt node 0 extends the "1-3-4-6-7" chain with the next block, and releases this block. Additionally, he extends the "0-2-3-5-6-8" chain with the next block but withholds this block.
- This goes on.

4 State Machine Replication

Byzantine agreement allows us to reach agreement only once. In practical applications, we often need to reach agreement many times establishing a *linearly ordered log* (e.g., of transactions). A protocol that allows a set of nodes to agree on an ever-growing, linearly ordered log is called a *state machine replication* protocol.

4.1 Definitions: State Machine Replication

A state machine replication protocol can be defined as below.

Syntax. In a state machine replication protocol, in every round, an honest node receives as input a set of transactions txs from $\mathcal{Z}$ at the beginning of the round, and outputs a LOG collected thus far to $\mathcal{Z}$ at the end of the round.

Security. Let T_{confirm} be a polynomial function in the security parameter κ and possibly other parameters of the execution such as the number of nodes participating, the corrupt fraction, the network delay, etc.

Definition 1 (Security of a state machine replication protocol)**.** We say that a state machine replication protocol Π satisfies consistency (or T_{confirm}-liveness resp.) w.r.t. some $(\mathcal{A}, \mathcal{Z})$, iff there exists a negligible function $\mathsf{negl}(\cdot)$, such that for any $\kappa \in \mathbb{N}$, except with $\mathsf{negl}(\kappa)$ probability over the choice of $\mathsf{view} \leftarrow \mathsf{EXEC}^{\Pi}(\mathcal{A}, \mathcal{Z}, \kappa)$, consistency (or T_{confirm}-liveness resp.) is satisfied:

- *Consistency*: A view satisfies consistency iff the following holds:
  - *Common prefix.* Suppose that in view, an honest node i outputs LOG to $\mathcal{Z}$ at time t, and an honest node j outputs LOG' to $\mathcal{Z}$ at time t' (i and j may be the same or different), it holds that either $\mathsf{LOG} \prec \mathsf{LOG}'$ or $\mathsf{LOG}' \prec \mathsf{LOG}$. Here the relation $\prec$ means "is a prefix of". By convention we assume that $\emptyset \prec x$ and $x \prec x$ for any x.
  - *Self-consistency.* Suppose that in view, a node i is honest during $[t, t']$, and outputs LOG and LOG' at times t and t' respectively, it holds that $\mathsf{LOG} \prec \mathsf{LOG}'$.
- *Liveness*: A view satisfies T_{confirm}-liveness iff the following holds: if in some round $t \leq |\mathsf{view}| - T_{\text{confirm}}$, some node honest in round t either received from $\mathcal{Z}$ an input set txs that contains some transaction tx or has tx in its output log to $\mathcal{Z}$ in round t, then, for any node i honest at any time $t' \geq t + T_{\text{confirm}}$, let LOG be the output of node i at time t', it holds that $\mathsf{tx} \in \mathsf{LOG}$.

  Intuitively, liveness says that transactions input to an honest node get included in honest nodes' LOGs within T_{confirm} time; and further, if a transaction appears in some honest node's LOG, it will appear in every honest node's LOG within T_{confirm} time.

We say that a state machine replication protocol Π satisfies consistency and T_{confirm}-liveness in (n, f)-environments iff for every p.p.t. $(\mathcal{A}, \mathcal{Z})$ that is (n, f)-respecting w.r.t. Π, Π satisfies consistency and T_{confirm}-liveness w.r.t. $(\mathcal{A}, \mathcal{Z})$.

4.2 A Longest-Chain-Based SMR Protocol Π_{smr}

Although our earlier Byzantine agreement protocols are described for the binary case (i.e., when the sender's input is a bit), it is easy to modify the protocols to support *multi-valued Byzantine agreement* (i.e., when the sender's input is an arbitrary string rather than a single bit). In the synchronous model, it is relatively easy to construct state machine replication (SMR) from *multi-valued Byzantine agreement* (MV-BA), e.g., in each round, fork n instances of MV-BA where node i is the sender in the i-th instance; and the final log would be the concatenation of the output of all these MV-BA instances [22]. This simple transformation from MV-BA to SMR, however, incurs a rather large overhead.

One advantage of longest-chain-style protocols is that these protocols themselves naturally give rise to state machine replication protocols. For example, our longest-chain protocol described in Section 3 requires only a couple simple modifications to become a state machine replication protocol (henceforth denoted Π_{smr}):

- *Transaction propagation.* Upon receiving a set of transactions txs from the environment $\mathcal{Z}$, multicast txs to everyone.
- *Block content.* Whenever creating a block, instead of placing in the data field any bit that has a signature from the sender, place in the data field any transaction the node has observed but has not appeared in the current prefix of the blockchain.

- *Output log.* At the end of every round r, let chain be the longest chain at the beginning of this round — output the prefix of chain with timestamps at most $r - n$.

Theorem 3. *Assume that the signature scheme is secure. Then, the state machine replication protocol Π_{smr} satisfies consistency and $(2n+1)$-liveness in (n, f)-environments as long as $n \geq 3f+1$.*

The proofs are a somewhat straightforward extension of our proofs for byzantine agreement. We present the proofs in the next sub-section.

4.3 Proofs for State Machine Replication

We now prove Theorem 3. Note that to obtain confirmation time $T_{\text{confirm}} := n + 1$, the proofs here build on the conclusions from our tighter analysis in Appendix B.

The consistency proof is implied by Theorem 4. We now prove the liveness of Π_{smr} for $T_{\text{confirm}} := 2n + 1$. It suffices to prove the following lemma.

Lemma 4. *Suppose that the signature scheme is secure. Then, except with negligible probability over the choice of view, the following holds:*

a) Suppose that $\mathcal{Z}$ inputs a set containing the transaction tx to some honest node in round r, then in round $r + 2n + 1$, every honest node's output to $\mathcal{Z}$ must contain tx; and further,

b) Suppose that in some round r some honest node's output to $\mathcal{Z}$ includes the transaction tx, then in any round $r' \geq r + n + 1$, any honest node's output to $\mathcal{Z}$ must contain tx as well

Proof. Claim (a) holds due to the following. If some honest node i receives tx from $\mathcal{Z}$ in some round r, it will mulicast tx to everyone in round r. Thus all honest nodes will have received tx at the beginning of round $r + 1$. In round $r' = r + 1 + 2n$, let chain denote the longest chain of some honest node i, i will will output to $\mathcal{Z}$ the prefix of chain whose timestamps are at most $r + n + 1$ by protocol definition. By chain quality[3] (Lemma 2), in chain, there must be an honest block whose timestamp is in the range $[r + 1, r + n + 1]$ — by honest protocol definition, this honest block (or the prefix chain till this block) must contain tx.

Claim (b) holds due to the following. Suppose that some honest node's longest chain in round r is chain, and some block in chain with the timestamp $t \leq r - n$ contains tx (and thus the node outputs tx to $\mathcal{Z}$ in round r). Now, in round $r' = r + n + 1$, every honest node will output a prefix of their longest chain then containing all blocks whose timestamps are no greater than $r + 1$. Due to chain quality, there is at least one honest-forever block whose timestamp is between $r - n$ and $r + 1$. The remainder of the proofs follows due to consistency.

□

5 Analysis of Aura's Instantiation

Based on the insights gained from our earlier mathematical analysis, we turn our attention to Aura's specific instantiation. As mentioned, Aura's specific choice and claim are

1. they chop off only $\lfloor \frac{n}{2} \rfloor + 1$ blocks for finalization;

[3] Note that Lemma 2 holds for Π_{smr} too in the same way it holds for Π_{chain}.

2. they employ the following tie-breaking rule: for two equal-length chains, they prefer the chain whose last block has an earlier timestamp and

3. they claim to defend against any minority, Byzantine adversary [2].

We now reflect on Aura's choices and claims.

A $\frac{3}{7}$-attack even with tie breaking. Recall that earlier we argued that $\frac{1}{3}$ is the best we can hope for when tie breaking is arbitrary and can be adversarially manipulated. We now demonstrate a $\frac{3}{7}$-attack against consistency even with their tie-breaking. The attack is depicted below — based on this illustration, the detailed description of the attack is similar to Section 3.4. Note that this attack works even if we chop of n rounds at the end for finalization (i.e., more than what Aura actually chops off).

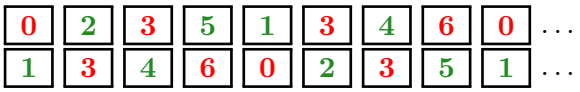

Reduced resilience when chopping off fewer blocks. We now consider what resilience Aura can achieve by chopping off only $\lfloor \frac{n}{2} \rfloor + 1$ blocks with distinct signers (and using their specific tie-breaking rule). We show that a consistency attack is possible if the adversary controls just slightly more than $\frac{3}{8}$ fraction of the nodes. The attack depicted above has the following generalization: $3f$ corrupt nodes can sustain two equal-length forks of length $4f$ (each fork having distinct signers). Thus with $3f$ corrupt nodes we have to chop off at least $4f$ blocks (if not more) for consistency. Now suppose $n = 8f - 3$, then Aura would actually chop off only $4f - 1$ blocks and this would lead to a consistency attack when $3f$ out of $8f - 3$ nodes are corrupt.

Finally, it is not difficult to generalize our proof to see that if one actually chops off only $\lfloor \frac{n}{2} \rfloor + 1$ trailing blocks with distinct signers, we can still guarantee consistency and liveness as long as $f < \frac{n}{6}$. This means that Aura's concrete instantiation actually gives a resilience parameter that is somewhere between $\frac{1}{6}$ and $\frac{3}{8}$, and they cannot defend against corrupt minority as they claim.

Acknowledgments

We thank Danlu Huang, Xia Yu, Yizhou Yu, Jiaheng Zhang, for helpful discussions. Through a course project in "Distributed Consensus and Blockchains" class, Danlu, Yizhou, Jiaheng looked into the Aura protocol and performed a security analysis. Xia Yu kindly helped the author understand the Aura implementation.

References

[1] `https://github.com/paritytech/parity-ethereum/tree/master/ethcore/src/engines/authority_round`.

[2] Aura - authority round. `https://wiki.parity.io/Aura`.

[3] DPOS consensus algorithm - the missing white paper. `https://steemit.com/dpos/@dantheman/dpos-consensus-algorithm-this-missing-white-paper`.

[4] Fix dpos loss of consensus due to conflicting last irreversible block. `https://github.com/EOSIO/eos/issues/2718`.

[5] M. K. Aguilera and S. Toueg. A simple bivalency proof that t-resilient consensus requires $t + 1$ rounds. *Inf. Process. Lett.*, 71(3-4):155–158, 1999.

[6] M. Castro and B. Liskov. Practical byzantine fault tolerance. In *OSDI*, 1999.

[7] P. Daian, R. Pass, and E. Shi. Snow white: Provably secure proofs of stake. Cryptology ePrint Archive, Report 2016/919, 2016.

[8] B. David, P. Gaži, A. Kiayias, and A. Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake protocol. Cryptology ePrint Archive, Report 2017/573, 2017. `http://eprint.iacr.org/2017/573`.

[9] C. Dwork and Y. Moses. Knowledge and common knowledge in a byzantine environment i: Crash failures. In *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning About Knowledge*, TARK '86, pages 149–169, San Francisco, CA, USA, 1986. Morgan Kaufmann Publishers Inc.

[10] J. A. Garay, A. Kiayias, and N. Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Eurocrypt*, 2015.

[11] V. Hadzilacos. A lower bound for byzantine agreement with fail-stop processors, 1983.

[12] A. Kiayias, A. Russell, B. David, and R. Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Crypto*, 2017.

[13] L. Lamport. Fast paxos. *Distributed Computing*, 19(2):79–103, 2006.

[14] L. Lamport and M. Fischer. Byzantine generals and transaction commit protocols. Technical report, 1982.

[15] L. Lamport, R. Shostak, and M. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, July 1982.

[16] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.

[17] Y. Moses and S. Rajsbaum. The unified structure of consensus: A *Layered Analysis* approach. In *PODC*, pages 123–132. ACM, 1998.

[18] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008.

[19] R. Pass, L. Seeman, and A. Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Eurocrypt*, 2017.

[20] R. Pass and E. Shi. Rethinking large-scale consensus. In *CSF*, 2017.

[21] R. Pass and E. Shi. The sleepy model of consensus. In *Asiacrypt*, 2017.

[22] R. Pass and E. Shi. Thunderella: Blockchains with optimistic instant confirmation. Manuscript, 2017.

[23] F. B. Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Comput. Surv.*, 22(4):299–319, Dec. 1990.

A A Largest-Set Protocol: Textbook Construction and Proof

In this section, we describe a variant of a deterministic, longest-chain protocol. In this variant, when a node is elected as a leader in a round, it votes on the item with the most number of votes seen so far. We call this protocol a "largest-set" protocol (as opposed to "longest-chain"). As we shall see, this variant enables a very simple proof and we recommend it for pedagogical purposes. The proofs for this protocol can be viewed as a simplified, deterministic variant of the analysis of Nakamoto's blockchain [18–20].

A.1 Useful Definitions

Leader. In every round i, node ($i \mod n$) is the leader. Thus each round has a unique leader.

Valid votes. A valid vote for a normal bit $b \in \{0,1\}$ is a tuple (b, t, i, σ) where $i \in [0..n-1]$ is a node identifier, $t \geq 1$ is a round number (also called the vote's *timestamp*), and σ is a valid signature under pk_i for the tuple (b, t). A valid vote for the dummy bit $\perp$ is a tuple $(\perp, t, i, \sigma)$ where $i \in [0..n-1]$ is a node identifier, $t \geq T_\perp$ is a round number no earlier than $T_\perp$, and σ is a valid signature under pk_i for the tuple $(\perp, t)$.

Two votes (b, t, i, σ) and (b', t', i', σ') are said to be *distinct* iff $(b, t, i) \neq (b', t', i')$.

Vouch for. Given an execution trace view, we define "vouch for" as below:

- If $b \in \{0,1\}$ is a normal bit and node i does not have the sender's signature on $(\texttt{propose}, b)$ in its view at the beginning of round r, then we say that there are no votes vouching for b w.r.t. node i and round r.
- Else, the number of votes vouching for $\widetilde{b} \in \{0, 1, \perp\}$ w.r.t. node i and round r is defined to be the number of distinct valid votes for $\widetilde{b}$ in node i's view at the beginning of round r.

If $\widetilde{b} \in \{0, 1, \perp\}$ has X votes vouching for it w.r.t. node i and round r in some execution trace view, we sometimes also say: at the beginning of round r, node i has X votes vouching for $\widetilde{b}$ in view. If $\widetilde{b}$ has the most number of votes vouching for it (in comparison with any other bit) w.r.t. node i and round r, we say that node i perceives $\widetilde{b}$ as the *most popular* bit in round r.

A.2 A Largest-Set Protocol Π_{set}

We now describe a simple largest-set protocol that realizes byzantine agreement as defined in Section 2.2. In a longest-chain protocol such as Nakamoto's blockchain [18] (and see also Section 3), nodes vote by extending the longest chain; in this largest-set protocol, there is no such concept of a chain, and nodes vote on the bit that so far has the largest number of votes vouching for it (i.e., the most popular bit).

Imprecisely speaking, in every round, every node looks for a bit signed by the sender and carrying the most number of votes, and then signs that bit. At the end of the protocol, whichever bit has the most number of signatures will be the bit output (i.e., decided). One issue is that if the sender is corrupt, it may not sign any bit. To handle this case, we introduce a dummy bit denoted $\perp$ sometime $T_\perp$ into the protocol — nodes are allowed to sign $\perp$ $T_\perp$ time into the protocol start. $T_\perp$ is chosen to be large enough such that if the sender did sign a bit upfront, the dummy bit $\perp$ cannot outnumber a normal bit in terms of number of votes and thus honest nodes will agree on a normal bit. On the other hand, $T_\perp$ must be chosen to be small enough (w.r.t. to the protocol duration T_{end}) such that the protocol retains consistency.

More concretely, the protocol works as follows.

- *Round 0:* let b be the sender's input bit. The sender signs $(\texttt{propose}, b)$ and multicasts the bit and the resulting signature.
- *For round* $r = 1, 2, \ldots, T_{\text{end}}$*:*
  - Receive all messages from the network and discard every vote whose timestamp is at least the r (discarded votes do not become part of the node's view in the protocol).
  - If the current node is the leader of this round, perform the following steps (otherwise skip). Let $\widetilde{b} \in \{0, 1, \perp\}$ be the most popular bit (w.r.t. the current node and round r), sign $(\widetilde{b}, r)$, and multicast 1) the resulting vote; 2) all valid votes the node has for $\widetilde{b}$ so far in its view; and 3) a sender's signature on $(\texttt{propose}, \widetilde{b})$ if applicable. If there are multiple most popular bits, break ties arbitrarily.
- At the beginning of round $T_{\text{end}} + 1$: Find the most popular bit $\widetilde{b} \in \{0, 1, \perp\}$ breaking ties arbitrarily. If $\widetilde{b} \in \{0, 1\}$, output $\widetilde{b}$. Otherwise output 0.

Parameters. Let $\epsilon = \frac{1}{3} - \frac{f}{n}$, we will choose $T_\perp = n + 1$, and $T_{\text{end}} = 2n \cdot \lceil \frac{1}{\epsilon} \rceil$. For example, if $n = 3f + 1$, i.e., $\epsilon = \frac{1}{3n}$, then the protocol will run for $\Theta(n^2)$ number of rounds.

A.3 Proofs

If a round's leader is honest-forever, we say that this round is an honest-forever leader round. Otherwise, we say that it is an eventually-corrupt leader round.

Fact 4. *Suppose that the signature scheme is secure, then except with negligible probability, no honest node's view in the protocol will ever contain a signature valid under some honest-forever node i but the signature was not signed by this node i during the protocol.*

Proof. Straightforward by signature security. □

Lemma 5 (Vote growth)**.** *Suppose that $T_{\text{end}} - T_\perp + 1 \geq kn$. With probability 1, the following holds: at the beginning of round $T_{\text{end}} + 1$, any honest node i's most popular bit must have at least $k(n - f)$ votes vouching for it.*

Proof. In every honest-forever leader round $r \geq T_\perp$, suppose that the leader $(r \mod n)$ of round r's most popular bit at the beginning of r has X votes vouching for it, then, at the beginning of round $r + 1$, every honest node's most popular bit has at least $X + 1$ votes vouching for it — since in round r node $(r \mod n)$ will add another vote to this most popular bit and multicast it to everyone else. The remainder of the proof follows in a straightforward manner by observing that there are at most f eventually-corrupt leader rounds among every n rounds. □

Lemma 6 (Consistency)**.** *Assume that the signature scheme is secure and that $f = (\frac{1}{3} - \epsilon)n$. Then, except with negligible probability over the choice of the execution trace* view*, at the beginning of round $T_{\text{end}} + 1$, all honest nodes must output the same bit $\widetilde{b} \in \{0, 1, \perp\}$.*

Proof. Due to Lemma 5, at the beginning of round $T_{\text{end}} + 1$, every honest node's most popular bit must have at least $k(n - f)$ votes where $k = \lceil \frac{1}{\epsilon} \rceil - 1$ is the largest integer such that $T_{\text{end}} - T_\perp + 1 \geq kn$. Suppose, for the sake of contradiction, that for a non-negligible fraction of the views, at the beginning of round $T_{\text{end}} + 1$, two honest node's most popular bit are different bits henceforth

denoted b and b'. Then it holds that there are at least $2k(n-f) = 2k(\frac{2}{3}+\epsilon)n$ distinct votes whose timestamps are upper bounded by T_{end} in view.

On the other hand, note that in every honest-forever leader round, no more than one vote may be cast for either b or b' by the end of the protocol; whereas in every eventually-corrupt leader round, one vote can be cast for each of b and b' by the end of the protocol. Due to Fact 4, except for a negligible fraction of the execution traces views, the total number of votes cast for either b or b' by the end of the protocol is at most $(k+1)(n+f) = (k+1)(\frac{4}{3}-\epsilon)n$. We conclude that

$$(k+1)\left(\frac{4}{3}-\epsilon\right)n \geq 2k\left(\frac{2}{3}+\epsilon\right)n$$

Simplifying the above equation, we have that

$$\frac{4}{3}-\epsilon \geq 3\epsilon\left(\lceil\frac{1}{\epsilon}\rceil - 1\right) \geq 3-3\epsilon \geq 2$$

This is impossible and thus we reach a contradiction.

□

Fact 5. *If the sender is honest in round* 0*, then at the beginning of round* $t \geq kn+1$ *in the protocol, any honest node's most popular bit must have at least* $k(n-f)$ *votes vouching for it.*

Proof. Due to the same argument as Lemma 5 and the fact that the sender is honest and will sign its input bit in round 0. □

Lemma 7 (Validity)**.** *If the sender is honest-forever, then, all honest nodes must output the sender's input bit.*

Proof. Henceforth we ignore the negligible fraction of views where Fact 4 fails. If the sender is honest-forever, due to Fact 5, we have that at the beginning of round $T_\perp$, any honest node's most popular bit must have at least $n-f$ votes vouching for it; moreover, by Fact 4, its most popular bit must be the sender's input bit b^*.

Suppose for the sake of contradiction that some honest node did not output the sender's input bit b^* at the end of the protocol. By Fact 4, this differing output bit must be $\perp$. Let $t^* > T_\perp$ be the *first* round in which some honest node's most popular bit is $\perp$. By definition of the protocol, before round t^*, no honest node will vote for $\perp$. Thus all valid votes for $\perp$ at the beginning of round t^* must come from eventually-corrupt nodes and they must have timestamps between $[T_\perp, t^*-1]$. The total number of eventually-corrupt votes for $\perp$ with timestamps between $[T_\perp, t^*-1]$ must be upper bounded by the number of eventually-corrupt leader rounds between $[T_\perp, t^*-1]$. Let k be the largest integer such that $t^* - T_\perp \geq kn$. Then, the total number of eventually-corrupt votes for $\perp$ with timestamps between $[T_\perp, t^*-1]$ must be upper bounded by $(k+1)f$.

On the other hand, due to Fact 5, at the beginning of round t^*, the most popular bit in any honest node's view must have at least $(k+1)(n-f)$ votes vouching for it.

We have that

$$(k+1)f \geq (k+1)(n-f)$$

However, this is impossible if $n \geq 3f+1$. □

B A Tighter Proof for the Longest-Chain Protocol

In this section, we present a tight proof for our deterministic longest-chain protocol. In the main body, we presented a much simpler proof that requires chopping off up to $O(\frac{n}{\epsilon})$ round numbers for finalization where $\epsilon := \frac{1}{3} - \frac{f}{n}$. In particular, if $n = 3f + 1$, then we need to chop off quadratic number of blocks. In this section, we prove that in fact, chopping off the most recent n round numbers from the end of the chain is sufficient. This tighter proof is much more technical than the earlier proof in Section 3.3.

We begin by defining a good event called a *pivot*. We then argue why a pivot contributes to consistency. Finally, we present a combinatorics argument why a pivot happens every n rounds.

Pivot. Given an execution trace view, a pivot is a point of time t such that for any window $[t_0, t_1]$ that contains t, there are strictly more honest-forever leader rounds than eventually-corrupt leader rounds. Our definition of a pivot is inspired by the earlier work by Pass and Shi [21].

Note that by definition, if t is a pivot in view, then node $t \mod n$ (i.e., the leader for round t) must be an honest-forever node. Further, it is easy to see that the following fact holds:

Fact 6. *Assume that $n \geq 2f + 1$. Then the following holds for any execution trace view: t is a pivot in view iff for any window $[t_0, t_1]$ that contains t and $t_1 - t_0 + 1 \leq n$, there are strictly more honest-forever leader rounds than eventually-corrupt leader rounds in view.*

Proof. Straightforward due to the following observation: for any consecutive n rounds, there must be more honest-forever leader rounds than eventually-corrupt leader rounds. □

Lemma 8 (Pivot leads to convergence)**.** *Assume that the signature scheme is secure, then the following holds except for a negligible fraction of views. Suppose that t is a pivot in view; further, suppose that* chain *is the longest chain belonging to an honest node in round $r > t$ in view and* chain′ *is the longest chain belonging to an honest node in round $s > t$ in view. It holds that* extract(chain) *and* extract(chain′) *must agree for the prefix whose timestamps are at most t.*

Proof. Ignore the negligible fraction of views where bad events related to Fact 1 take place.

For the remaining good views, suppose that the honest-forever node ($t \mod n$) mined a block at length ℓ in view. It suffices to prove that chain$[\ell]$ and chain′$[\ell]$ must both be this honest-forever block. We prove it for chain since the argument for chain′ is identical.

If chain$[\ell]$ is an honest-forever block, the the proof follows directly by Fact 3. We thus focus on the case when chain$[\ell]$ is an eventually-corrupt block. In this case, let chain$[\ell_L] := (t_L, _, _)$ be the latest honest-forever block to the left of chain$[\ell]$; if no such block exists, simply let $\ell_L := 0$ and $t_L = 0$. Similarly, let chain$[\ell_R] := (t_R, _, _)$ be the earliest honest-forever block to the right of chain$[\ell]$; if no such block exists, simply let $\ell_R := |\text{chain}| + 1$ and let $t_R := r$. Henceforth, we imagine that chain has two imaginary blocks at length 0 and $|\text{chain}|+1$ denoting beginning and end of chain respectively. For any $\widetilde{t} \in (t_L, t_R)$, if some honest-forever leader mined a block at length ℓ in round $\widetilde{t}$, then chain$[\ell]$ must be an eventually-corrupt block. Since all blocks in chain must have distinct timestamps, we have that between (t_L, t_R), there are at least as many eventually-corrupt leader-rounds as there are honest-forever leader-rounds. This contradicts the fact that t is a pivot. □

Lemma 9 (Existence of a pivot every n rounds)**.** *Suppose that $n \geq 3f + 1$. With probability 1, for every n consecutive rounds $(t, t + n]$, there must exist a pivot $t^* \in (t, t + n]$.*

Proof. **Placing numbers on a circle.** Imagine we place m numbers, either $+1$ or -1 on a circle such that fewer than $\frac{m}{3}$ of them are -1 — such a placement is henceforth said to be a *valid configuration.*

Pivot. For any $i, j \in [0..m-1]$ such that $i \neq j$, we use the notation $s^{\rightarrow}[i..j]$ to denote the cumulative sum going from the i-th number clockwise until we reach the j-th number (inclusive of both the i-th and the j-th number). We define $s^{\rightarrow}[i..i]$ to be the i-th number. $s^{\leftarrow}[i..j]$ is similarly defined but going counter-clockwise.

The i-th number where $i \in [0..m-1]$ is said to be a *pivot* iff for any $j \in [0..m-1]$, we have that $s^{\rightarrow}[i..j] > 0$ and $s^{\leftarrow}[i..j] > 0$. It is not difficult to see that this definition is equivalent to the following: the i-th number is said to be a pivot iff for any clockwise segment spanning positions $a..b$ that includes i $s^{\rightarrow}[a..b] > 0$.

Configuration induced by an execution. Any execution trace view will induce a configuration as explained below. We place n numbers on a circle based on view: for $i \in [0..n-1]$, if node i is honest-forever in view, the i-th number is $+1$; otherwise it is -1. Since we assume that $n \geq 3f+1$, it must hold that the configuration induced by view is a valid configuration.

To prove the above lemma, it suffices to prove that the configuration induced by any view has a pivot. We prove a generalization of the statement:

Claim 1. *Suppose that we place m numbers, either $+1$ or -1, on a circle as mentioned above. In any valid configuration (i.e., fewer than $\frac{m}{3}$ of them are -1), there must exist a pivot.*

Proof. We now prove the above claim. A *negative group* is a sequence of consecutive -1s on the circle sandwiched by $+1$s on both sides. A *positive group* is sequence of consecutive $+1$s on the circle sandwiched by -1s on both sides. If all numbers on the circle are $+1$ (or -1 resp.), all the numbers form a single positive (or negative resp.) group. If a positive (or negative resp.) group does not contain all numbers on the circle, it is said to be a *proper* positive (or negative resp.) group.

Any *proper* positive (or negative resp.) has a *left-most* number and a *right-most* number: essentially the group spans the left-most number going clockwise to the right-most number.

We now prove by induction.

Base case. All numbers form a single positive group. In this case the proof is trivial.

Induction step. If a valid configuration is not the base case, then it must contain a sequence of alternating proper positive and proper negative groups — let T denote the total number of such alternating positive and negative groups (the base case can be regarded as $T = 1$). We need to prove the following: if a pivot exists for every valid configuration where $T \leq \tau$ (i.e., induction hypothesis), then a pivot exists for every valid configuration where $T \leq \tau + 1$.

We now perform the following elimination procedure:

- *Start:* Say we start from a valid configuration C that is not all positive.
- *Find negative group:* Find an arbitrary proper negative group, let its right-most number be position i_R and the left-most number be position i_L.
- *Counter-clockwise scan:* Scan counter-clockwise from i_R until we reach the first position j such that $s^{\rightarrow}[j..i_R] = 0$.
- *Clockwise scan:* Start from the left-most number of the negative group, say the i_L-th number, and scan clockwise until we reach the first position k such that $s^{\rightarrow}[i_L..k] = 0$.

- *Elimination:* Eliminate all numbers between positions $[j..k]$, and the ending configuration is called C'.

Henceforth all numbers eliminated due to the counter-clockwise scan together comprise the "counter-clockwise elimination"; and all numbers eliminated due to the clockwise scan together comprise the "clockwise elimination". The clockwise and counter-clockwise eliminations overlap in the negative group chosen.

We now prove that the following facts hold for the above elimination procedure.

Fact 7. *A number that is pivot in C cannot be eliminated by the above procedure.*

Proof. It suffices to argue that a position p is a pivot in C cannot be eliminated during a counter-clockwise elimination (the argument is symmetric for clockwise). For p to be eliminated, it must hold that $s^{\rightarrow}[p..i] \leq 0$ where i is the right-most number of some proper negative group. This cannot hold by the definition of a pivot. □

Fact 8. *A number that is pivot in C' must be a pivot in C.*

Proof. For simplicity, we assume that the elimination simply rewrites eliminated numbers as 0s — this way, we can use the same position labels for C and C'.

Henceforth, we denote the cumulative sum $s^{\rightarrow}[a..b]$ in C as $S[a..b]$, and the cumulative sum $s^{\rightarrow}[a..b]$ in C' as $S'[a..b]$. Due to Fact 7, it suffices to prove that for any position p that is not eliminated during the elimination procedure, for any p^*, we have that $S[p..p^*] \geq S'[p..p^*]$ (and the other direction, i.e., counter-clockwise, is symmetric). We consider the following cases:

- Case 1: p^* is not among the eliminated positions. In this case, the proof is obvious since the eliminated positions sum up to 0.
- Case 2: p^* is among the eliminated positions. Recall that j is the left-most among the eliminated positions; and that i_R is the right-most number of the negative group involved in the elimination procedure. It suffices to prove that $S[j..p^*] \geq 0$. This holds since otherwise j cannot the first number scanning counter-clockwise from i_R such that $S[j..i_R] = 0$ — there must exist such a number in between positions $[p^* + 1..i_R]$.

□

Fact 9. *If the elimination procedure eliminates x number of -1s, it can eliminate no more than $2x$ number of $+1$s.*

Proof. Straightforward by observing that in both the clockwise and the counter-clockwise eliminations, there are as many -1s elimanted as there are $+1$s. Further, the -1s eliminated in both directions can overlap whereas the $+1$s cannot overlap in both directions (since it is not difficult to see that there must be numbers left after the elimination). □

Due to Facts 8 to prove the induction step, it suffices to prove that in C' there is a pivot — the latter holds due to Fact 9 and the induction hypothesis. □

□

Theorem 4 (Consistency)**.** *Assume that the signature scheme is secure. Then, except with negligible probability, the following holds: let* chain *be some honest node's longest chain in round r; and let* chain' *be some honest node's longest chain in round $t \geq r$, it holds that* chain *and* chain' *agree in the part of their prefixes with timestamps $r - n$ or smaller.*

Proof. Straightforward due to Lemma 8 and Lemma 9. □

B.1 Proofs for State Machine Replication

We now prove Theorem 3. Note that to obtain confirmation time $T_{\text{confirm}} := n + 1$, the proofs here build on the conclusions from our tighter analysis in Appendix B.

The consistency proof is implied by Theorem 4. We now prove the liveness of Π_{smr} for $T_{\text{confirm}} := 2n + 1$. It suffices to prove the following lemma.

Lemma 10. *Suppose that the signature scheme is secure. Then, except with negligible probability over the choice of* view*, the following holds:*

a) Suppose that $\mathcal{Z}$ inputs a set containing the transaction tx *to some honest node in round r, then in round $r + 2n + 1$, every honest node's output to $\mathcal{Z}$ must contain* tx*; and further,*

b) Suppose that in some round r some honest node's output to $\mathcal{Z}$ includes the transaction tx*, then in any round $r' \geq r + n + 1$, any honest node's output to $\mathcal{Z}$ must contain* tx *as well*

Proof. Claim (a) holds due to the following. If some honest node i receives tx from $\mathcal{Z}$ in some round r, it will mulicast tx to everyone in round r. Thus all honest nodes will have received tx at the beginning of round $r + 1$. In round $r' = r + 1 + 2n$, let chain denote the longest chain of some honest node i, i will will output to $\mathcal{Z}$ the prefix of chain whose timestamps are at most $r + n + 1$ by protocol definition. By chain quality[4] (Lemma 2), in chain, there must be an honest block whose timestamp is in the range $[r + 1, r + n + 1]$ — by honest protocol definition, this honest block (or the prefix chain till this block) must contain tx.

Claim (b) holds due to the following. Suppose that some honest node's longest chain in round r is chain, and some block in chain with the timestamp $t \leq r - n$ contains tx (and thus the node outputs tx to $\mathcal{Z}$ in round r). Now, in round $r' = r + n + 1$, every honest node will output a prefix of their longest chain then containing all blocks whose timestamps are no greater than $r + 1$. Due to chain quality, there is at least one honest-forever block whose timestamp is between $r - n$ and $r + 1$. The remainder of the proofs follows due to consistency.

□

[4]Note that Lemma 2 holds for Π_{smr} too in the same way it holds for Π_{chain}.