Ouroboros: A Provably Secure Proof-of-Stake Blockchain Protocol

Aggelos Kiayias*   Alexander Russell†   Bernardo David‡   Roman Oliynykov§

April 3, 2017

Abstract

We present "Ouroboros", the first blockchain protocol based on *proof of stake* with rigorous security guarantees. We establish security properties for the protocol comparable to those achieved by the bitcoin blockchain protocol. As the protocol provides a "proof of stake" blockchain discipline, it offers qualitative efficiency advantages over blockchains based on proof of physical resources (e.g., proof of work). We also present a novel reward mechanism for incentivizing Proof of Stake protocols and we prove that, given this mechanism, honest behavior is an approximate Nash equilibrium, thus neutralizing attacks such as selfish mining. We also present initial evidence of the practicality of our protocol in real world settings by providing experimental results on transaction processing time.

1 Introduction

A primary consideration regarding the operation of blockchain protocols based on proof of work (PoW)—such as bitcoin [24]—is the energy required for their execution. At the time of this writing, generating a single block on the bitcoin blockchain requires a number of hashing operations exceeding 2^{60}, which results in striking energy demands. Indeed, early calculations indicated that the energy requirements of the protocol were comparable to that of a small country [26].

This state of affairs has motivated the investigation of alternative blockchain protocols that would obviate the need for proof of work by substituting it with another, more energy efficient, mechanism that can provide similar guarantees. It is important to point out that the proof of work mechanism of bitcoin facilitates a type of randomized "leader election" process that elects one of the miners to issue the next block. Furthermore, provided that all miners follow the protocol, this selection is performed in a randomized fashion proportionally to the computational power of each miner. (Deviations from the protocol may distort this proportionality as exemplified by "selfish mining" strategies [16, 30].)

A natural alternative mechanism relies on the notion of "proof of stake" (PoS). Rather than miners investing computational resources in order to participate in the leader election process, they instead run a process that randomly selects one of them proportionally to the *stake* that each possesses according to the current blockchain ledger.

In effect, this yields a self-referential blockchain discipline: maintaining the blockchain relies on the stakeholders themselves and assigns work to them (as well as rewards) based on the amount

---

of stake that each possesses as reported in the ledger. Aside from this, the discipline should make no further "artificial" computational demands on the stakeholders. In some sense, this sounds ideal; however, realizing such a proof-of-stake protocol appears to involve a number of definitional, technical, and analytic challenges.

Previous work. The concept of PoS has been discussed extensively in the bitcoin forum.[1] Proof-of-stake based blockchain design has been more formally studied by Bentov et al., both in conjunction with PoW [5] as well as the sole mechanism for a blockchain protocol [4]. Although Bentov et al. showed that their protocols are secure against some classes of attacks, they do not provide a formal model for analysing PoS based protocols or security proofs relying on precise definitions. Heuristic proof-of-stake based blockchain protocols have been proposed (and implemented) for a number of cryptocurrencies.[2] Being based on heuristic security arguments, these cryptocurrencies have been frequently found to be deficient from the point of view of security. See [4] for a discussion of various attacks.

It is also interesting to contrast a PoS-based blockchain protocol with a classical consensus blockchain that relies on a *fixed* set of authorities (see, e.g., [13]). What distinguishes a PoS-based blockchain from those which assume static authorities is that stake changes over time and hence the trust assumption evolves with the system.

Another alternative to PoW is the concept of *proof of space* [2, 15], which has been specifically investigated in the context of blockchain protocols [27]. In a proof of space setting, a "prover" wishes to demonstrate the utilization of space (storage / memory); as in the case of a PoW, this utilizes a physical resource but can be less energy demanding over time. A related concept is *proof of space-time* (PoST) [23]. In all these cases, however, an expensive physical resource (either storage or computational power) is necessary.

The PoS Design challenge. A fundamental problem for PoS-based blockchain protocols is to simulate the leader election process. In order to achieve a randomized election among stakeholders, entropy must be introduced into the system, and mechanisms to introduce entropy may be manipulated by the adversary. For instance, independently of the solution, an adversary controlling a set of stakeholders may choose to simulate the protocol execution trying different sequences of stakeholder participants so that it finds a favorable chain continuation that biases the leader election. To prevent this manipulation, honest stakeholders must be able to add sufficient entropy and counter any lookahead performed by the adversary.

Our Results. We present "Ouroboros", a provably secure proof of stake system. To the best of our knowledge this is the first blockchain protocol of its kind with a rigorous security analysis. In more detail, our results are as follows.

First, we provide a model that formalizes the problem of realizing a PoS-based blockchain protocol. The model we introduce is in the spirit of [19], focusing on *persistence* and *liveness*, two formal properties of a robust transaction ledger. *Persistence* states that once a node of the system proclaims a certain transaction as "stable", the remaining nodes, if queried and responding honestly, will also report it as stable. Here, stability is to be understood as a predicate that will be parameterized by some security parameter k that will affect the certainty with which the property holds. (E.g., "more than k blocks deep".) *Liveness* ensures that once an honestly generated transaction

---

[1]See "Proof of stake instead of proof of work", Bitcoin forum thread. Posts by user "QuantumMechanic" and others. (`https://bitcointalk.org/index.php?topic=27787.0.`).

[2]A non-exhaustive list includes NXT, Neucoin, Blackcoin, Tendermint, Bitshares.

has been made available for a sufficient amount of time to the network nodes, say u time steps, it will become stable. The conjunction of liveness and persistence provides a robust transaction ledger in the sense that honestly generated transactions are adopted and become immutable. Our model is suitably amended to facilitate PoS-based dynamics.

Second, we describe a novel blockchain protocol based on PoS. Our protocol assumes that parties can freely create accounts and receive and make payments, and that stake shifts over time. We utilize a (very simple) secure multiparty implementation of a coin-flipping protocol to produce the randomness for the leader election process. This distinguishes our approach (and prevents so called "grinding attacks") from other previous solutions that either defined such values deterministically based on the current state of the blockchain or used collective coin flipping as a way to introduce entropy [4]. Also, unique to our approach is the fact that the system ignores round-to-round stake modifications. Instead, a snapshot of the current set of stakeholders is taken in regular intervals called *epochs*; in each such interval a secure multiparty computation takes place utilizing the blockchain itself as the broadcast channel. Specifically, in each epoch a set of randomly selected stakeholders form a committee which is then responsible for executing the coin-flipping protocol. The outcome of the protocol determines the set of next stakeholders to execute the protocol in the next epoch as well as the outcomes of all leader elections for the epoch.

Third, we provide a set of formal arguments establishing that no adversary can break persistence and liveness. Our protocol is secure under a number of plausible assumptions: (1) the network is highly synchronous, (2) the majority of the selected stakeholders is available as needed to participate in each epoch, (3) the stakeholders do not remain offline for long periods of time, (4) the adaptivity of corruptions is subject to a small delay that is measured in rounds linear in the security parameter (or alternatively, the players have access to a sender-anonymous broadcast channel). At the core of our security arguments is a probabilistic argument regarding a combinatorial notion of "forkable strings" which we formulate, prove and verify experimentally. In our analysis we also distinguish *covert attacks*, a special class of general forking attacks. "Covertness" here is interpreted in the spirit of covert adversaries against secure multiparty computation protocols, cf. [3], where the adversary wishes to break the protocol but prefers not to be caught doing so. We show that covertly forkable strings are a subclass of the forkable strings with much smaller density; this permits us to provide two distinct security arguments that achieve different trade-offs in terms of efficiency and security guarantees.

Fourth, we turn our attention to the incentive structure of the protocol. We present a novel reward mechanism for incentivizing the participants to the system which we prove to be an (approximate) Nash equilibrium. In this way, attacks like *block withholding* and *selfish-mining* [16, 30] are mitigated by our design. The core idea behind the reward mechanism is to provide positive payoff for those protocol actions that cannot be stifled by a coalition of parties that diverges from the protocol. In this way, it is possible to show that, under plausible assumptions, namely that certain protocol execution costs are small, following the protocol faithfully is an equilibrium when all players are rational.

Fifth, we introduce a stake delegation mechanism that can be seamlessly added to our blockchain protocol. Delegation is particularly useful in our context as we would like to allow our protocol to scale even in a setting where the set of stakeholders is highly fragmented. In such cases, the delegation mechanism can enable stakeholders to delegate their "voting rights", i.e., the right of participating in the committees running the leader selection protocol in each epoch. As in *liquid democracy*, (a.k.a. delegative democracy [18]), stakeholders have the ability to revoke their delegative appointment when they wish independently of each other.

Sixth, given our model and protocol description we also explore how various attacks considered in practice can be addressed within our framework. Specifically, we discuss double spending attacks,

transaction denial attacks, 51% attacks, nothing-at-stake, desynchronization attacks and others.

Finally, we present evidence regarding the efficiency of our design. First we consider double spending attacks. For illustrative purposes, we perform a comparison with Nakamoto's analysis for bitcoin regarding transaction confirmation time with assurance 99.9%. Against covert adversaries, the transaction confirmation time is from 10 to 16 times faster than that of bitcoin, depending on the adversarial hashing power; for general adversaries confirmation time is from 5 to 10 times faster. Moreover, our concrete analysis of double-spending attacks relies on our combinatorial analysis of forkable and covertly forkable strings and applies to a much broader class of adversarial behavior than Nakamoto's more simplified analysis.[3] We then survey our prototype implementation and report on benchmark experiments run in the Amazon cloud that showcase the power of our proof of stake blockchain protocol in terms of performance.

Paper overview. We lay out the basic model in Sec. 2. To simplify the analysis of our protocol, we present it in four stages that are outlined in Sec. 3. In short, in Sec. 4 we describe and analyze the protocol in the static setting; we then transition to the dynamic setting in Sec. 5. Our incentive mechanism and the equilibrium argument are presented in Sec. 7. We then present the protocol enhancement with anonymous channels in Sec. **??** and with a delegation mechanism in Sec. 8. Following this, in Sec. 9 we discuss the resilience of the protocol under various particular attacks of interest. Finally, in Sec. 10 we discuss transaction confirmation times as well as general performance results obtained from a prototype implementation running in the Amazon cloud.

2 Model

Time, slots, and synchrony. We consider a setting where time is divided into discrete units called *slots*. A ledger, described in more detail below, associates with each time slot (at most) one ledger *block*. Players are equipped with (roughly synchronized) clocks that indicate the current slot. This will permit them to carry out a distributed protocol intending to collectively assign a block to this current slot. In general, each slot sl_r is indexed by an integer $r \in \{1, 2, \ldots\}$, and we assume that the real time window that corresponds to each slot has the following properties.

- The current slot is determined by a publicly-known and monotonically increasing function of current time.

- Each player has access to the current time. Any discrepancies between parties' local time are insignificant in comparison with the length of time represented by a slot.

- The length of the time window that corresponds to a slot is sufficient to guarantee that any message transmitted by an honest party at the beginning of the time window will be received by any other honest party by the end of that time window (even accounting for small inconsistencies in parties' local clocks). In particular, while network delays may occur, they never exceed the slot time window.

Transaction Ledger Properties. A protocol Π implements a robust transaction ledger provided that the ledger that Π maintains is divided into "blocks" (assigned to time slots) that determine the order with which transactions are incorporated in the ledger. It should also satisfy the following two properties.

---

[3]Nakamoto's simplifications are pointed out in [19]: the analysis considers only the setting where a block withholding attacker acts without interaction as opposed to a more general attacker that tries strategically to split the honest parties in more than one chains during the course of the double spending attack.

- **Persistence.** Once a node of the system proclaims a certain transaction tx as *stable*, the remaining nodes, if queried, will either report tx in the same position in the ledger or will not report as stable any transaction in conflict to tx. Here the notion of stability is a predicate that is parameterized by a security parameter k; specifically, a transaction is declared *stable* if and only if it is in a block that is more than k blocks deep in the ledger.

- **Liveness.** If all honest nodes in the system attempt to include a certain transaction, then after the passing of time corresponding to u slots (called the transaction confirmation time), all nodes, if queried and responding honestly, will report the transaction as stable.

In [22, 29] it was shown that persistence and liveness can be derived from the following three elementary properties provided that protocol Π derives the ledger from a data structure in the form of a blockchain.

- **Common Prefix (CP); with parameters** $k \in \mathbb{N}$**.** The chains $\mathcal{C}_1, \mathcal{C}_2$ possessed by two honest parties at the onset of the slots $sl_1 < sl_2$ are such that $\mathcal{C}_1^{\lceil k} \preceq \mathcal{C}_2$, where $\mathcal{C}_1^{\lceil k}$ denotes the chain obtained by removing the last k blocks from $\mathcal{C}_1$, and $\preceq$ denotes the prefix relation.

- **Chain Quality (CQ); with parameters** $\mu \in (0,1] \to (0,1]$ **and** $k \in \mathbb{N}$**.** For any subset S of (possibly malicious) stakeholders with relative stake α and any portion of length k in a chain possessed by an honest party at the onset of a certain slot, the ratio of blocks originating from members of S can be at most $\mu(\alpha)$.

- **Chain Growth (CG); with parameters** $\tau \in (0,1], s \in \mathbb{N}$**.** Consider the chains $\mathcal{C}_1, \mathcal{C}_2$ possessed by two honest parties at the onset of two slots sl_1, sl_2 with sl_2 at least s slots ahead of sl_1. Then it holds that $\text{len}(\mathcal{C}_2) - \text{len}(\mathcal{C}_1) \geq \tau \cdot s$. We call τ the speed coefficient.

Some remarks are in place. Regarding common prefix, we capture a strong notion of common prefix, cf. [22]. Regarding chain quality, the function μ satisfies $\mu(\alpha) \geq \alpha$ for protocols of interest. In an ideal setting, μ would be the identity function: in this case, the percentage of malicious blocks in any sufficiently long chain segment is proportional to the cumulative stake of a set of (malicious) stakeholders.

It is worth noting that for bitcoin we have $\mu(\alpha) = \alpha/(1 - \alpha)$, and this bound is in fact tight—see [19], which argues this guarantee on chain quality. The same will hold true for our protocol construction. As we will show, this will still be sufficient for our incentive mechanism to work properly.

Finally chain growth concerns the rate at which the chain grows (for honest parties). As in the case of bitcoin, the *longest* chain plays a preferred role in our protocol; this provides an easy guarantee of chain growth.

Security Model. We adopt the model introduced by [19] for analysing security of blockchain protocols enhanced with an ideal functionality $\mathcal{F}$. We denote by $\mathsf{VIEW}_{\Pi,\mathcal{A},\mathcal{Z}}^{P,\mathcal{F}}(\kappa)$ the view of party P after the execution of protocol Π with adversary $\mathcal{A}$, environment $\mathcal{Z}$, security parameter κ and access to ideal functionality $\mathcal{F}$. We note that multiple different "functionalities" can be encompassed by $\mathcal{F}$.

We stress that contrary to [19], our analysis is in the "standard model", and without a random oracle functionality. Nevertheless we do employ a "diffuse" and "Key and Transaction" functionality with the following interfaces described below.

- Diffuse functionality. It maintains a incoming string for each party U_i that participates. A party, if activated, is allowed at any moment to fetch the contents of its incoming string hence one may think of this as a mailbox. Furthermore, parties can give the instruction to the functionality to diffuse a message. The functionality keeps rounds and all parties are allowed to diffuse once in a round. Rounds do not advance unless all parties have diffused a message. The adversary, when activated, can also interact with the functionality and is allowed to read all inboxes and all diffuse requests and deliver messages to the inboxes in any order it prefers. At the end of the round, the functionality will ensure that all inboxes contain all messages that have been diffused (but not necessarily in the same order they have been requested to be diffused).

- Key and Transaction functionality. The key registration functionality is initialized with n users, $U_1, \ldots, U_n$ and their respective stake $s_1, \ldots, s_n$; given such initialization, the functionality will consult with the adversary and will accept a (possibly empty) sequence of $(\mathsf{Corrupt}, U)$ messages and mark the corresponding users U as corrupt. For the corrupt users without a public-key registered the functionality will allow the adversary to set their public-keys while for honest users the functionality will sample public/secret-key pairs and record them. Public-keys of corrupt users will be marked as such. Subsequently, any sequence of the following actions may take place: (i) A user may request to retrieve its public and secret-key, whereupon, the functionality will return it to the user. (ii) The whole directory of public-keys may be required in whereupon, the functionality will return it to the requesting user. (iii) A new user may be requested to be created by a message $(\mathsf{Create}, U, \mathcal{C})$ from the environment, in which case the functionality will follow the same procedure as before: it will consult the adversary regarding the corruption status of U and will set its public and possibly secret-key depending on the corruption status; moreover it will store $\mathcal{C}$ as the suggested initial state. The functionality will return the public-key back to the environment upon successful completion of this interaction. (v) A transaction may be requested on behalf of a certain user by the environment, by providing a template for the transaction (which should contain a unique nonce) and a recipient. The functionality will adjust the stake of each stakeholder accordingly. (iv) An existing user may be requested to be corrupted by the adversary via a message $(\mathsf{Corrupt}, U)$. A user can only be corrupted after a delay of D slots; specifically, after a corruption request is registered the secret-key will be released after D slots have passed according to the clock maintained in the Diffuse interface.

Given the above we will assume that the execution of the protocol is with respect to a functionality $\mathcal{F}$ that is incorporating the above two functionalities as well as possibly additional functionalities to be explained below. Note that a corrupted stakeholder U will relinquish its entire state to $\mathcal{A}$; from this point on, the adversary will be activated in place of the stakeholder U. Beyond any restrictions imposed by $\mathcal{F}$, the adversary can only corrupt a stakeholder if it is given permission by the environment $\mathcal{Z}$ running the protocol execution. The permission is in the form of a message $(\mathsf{Corrupt}, U)$ which is provided to the adversary by the environment. In summary, regarding activations we have the following.

- At each slot sl_j, the environment $\mathcal{Z}$ is allowed to activate any subset of stakeholders it wishes. Each one of them will possibly produce messages that are to be transmitted to other stakeholders.

- The adversary is activated at least as the last entity in each sl_j, (as well as during all adversarial party activations).

- If a stakeholder does not fetch in a certain slot the messages written to its incoming string in the diffuse functionality they are flushed.

It is easy to see that the model above confers such sweeping power on the adversary that one cannot establish any significant guarantees on protocols of interest. It is thus important to restrict the environment suitably (taking into account the details of the protocol) so that we may be able to argue security. With foresight, the restrictions we will impose on the environment are as follows.

Restrictions imposed on the environment. The environment, which is responsible for activating the honest parties in each round, will be subject to the following constraints regarding the activation of the honest parties running the protocol.

- In each slot there will be at least one honest activated party (independently of whether it is a slot leader).

- There will be a parameter $k \in \mathbb{Z}$ that will signify the maximum number of slots that an honest shareholder can be offline. In case an honest stakeholder is spawned after the beginning of the protocol via $(\mathsf{Create}, U, \mathcal{C})$ its initialization chain $\mathcal{C}$ provided by the environment should match an honest parties' chain which was active in the previous slot.

- In each slot sl_r, and for each active stakeholder U_j there will be a set $\mathbb{S}_j(r)$ of public-keys and stake pairs of the form $(\mathsf{vk}_i, s_i) \in \{0,1\}^* \times \mathbb{N}$, for $j = 1, \ldots, n_r$ where n_r is the number of users introduced up to that slot. Public-keys will be marked as "corrupted" if the corresponding stakeholder has been corrupted. We will say the adversary is restricted to less than 50% relative stake if it holds that the total stake of the corrupted keys divided by the total stake $\sum_i s_i$ is less than 50% in all possible $\mathbb{S}_j(r)$. In case the above is violated an event $\mathsf{Bad}^{1/2}$ becomes true for the given execution.

We note that the offline restriction stated above is very conservative and our protocol can tolerate much longer offline times depending on the way the course of the execution proceeds; nevertheless, for the sake of simplicity, we use the above restriction. Finally, we note that in all our proofs, whenever we say that a property Q holds with high probability over all executions, we will in fact argue that $Q \vee \mathsf{Bad}^{1/2}$ holds with high probability over all executions. This captures the fact that we exclude environments and adversaries that trigger $\mathsf{Bad}^{1/2}$ with non-negligible probability.

3   Our Protocol: Overview

We first provide a general overview of our protocol design approach. The protocol's specifics depend on a number of parameters as follows: (i) k is the number of blocks a certain message should have "on top of it" in order to become part of the immutable history of the ledger, (ii) ϵ is the advantage in terms of stake of the honest stakeholders against the adversarial ones; (iii) D is the corruption delay that is imposed on the adversary, i.e., an honest stakeholder will be corrupted after D slots when a corrupt message is delivered by the adversary during an execution; (iv) L is the lifetime of the system, measured in slots; (v) R is the length of an epoch, measured in slots.

We present our protocol description in four *stages* successively improving the adversarial model it can withstand. In all stages an "ideal functionality" $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ is available to the participants. The functionality captures the resources that are available to the parties as preconditions for the secure operation of the protocol (e.g., the genesis block will be specified by $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$).

Stage 1. (case for static stake; $D = L$). In the first stage, the trust assumption is static and remains with the initial set of stakeholders. There is an initial stake distribution which is hardcoded into the genesis block that includes the public-keys of the stakeholders, $\{(\mathsf{vk}_i, s_i)\}_{i=1}^n$. Based on our restrictions to the environment, honest majority with advantage ϵ is assumed among those initial stakeholders. Specifically, the environment initially will allow the corruption of a number of stakeholders whose relative stake represents $\frac{1-\epsilon}{2}$ for some $\epsilon > 0$. The environment allows party corruption by providing tokens of the form $(\mathsf{Corrupt}, U)$ to the adversary; note that due to the corruption delay imposed in this first stage any further corruptions will be against parties that have no stake initially and hence the corruption model is akin to "static corruption." $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ will subsequently sample ρ which will seed a "weighted by stake" stakeholder sampling and in this way lead to the election of a subset of m keys $\mathsf{vk}_{i_1}, \ldots, \mathsf{vk}_{i_m}$ to form the committee that will possess honest majority with overwhelming probability in m, (this uses the fact that the relative stake possessed by malicious parties is $\frac{1-\epsilon}{2}$; a linear dependency of m to ϵ^{-2} will be imposed at this stage). In more detail, the committee will be selected implicitly by appointing a stakeholder with probability proportional to its stake to each one of the L slots. Subsequently, stakeholders will issue blocks following the schedule that is determined by the slot assignment. The longest chain rule will be applied and it will be possible for the adversary to fork the blockchain views of the honest parties. Nevertheless, we will prove with a Markov chain argument that the probability that a fork can be maintained over a sequence of n slots drops exponentially with at least $\sqrt{n}$, cf. Theorem 4.18 against general adversaries. An even more favorable analysis can be made against covert adversaries, i.e., adversaries that prefer to remain "under the radar" cf. Theorem 4.28.

Stage 2. (dynamic state with a beacon, epoch period of R slots long, $D = R \ll L$). The central idea for the extension of the lifetime of the above protocol is to consider the sequential composition of several invocations of it. We detail a way to do that, under the assumption that a trusted beacon emits a uniformly random string in regular intervals. More specifically, the beacon, during slots $\{j \cdot R + 1, \ldots, (j+1)R\}$, reveals the j-th random string that seeds the leader election function. The critical difference compared to the static state protocol is that the stake distribution is allowed to change and is drawn from the blockchain itself. This means that at a certain slot sl that belongs to the j-th epoch (with $j \geq 2$), the stake distribution that is used is the one reported in the most recent block with time stamp less than $j \cdot R - 2k$.

Regarding the evolving stake distribution, transactions will be continuously generated and transferred between stakeholders via the environment and players will incorporate posted transactions in the blockchain based ledgers that they maintain. In order to accomodate the new accounts that are being created, the $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ functionality enables a new $(\mathsf{vk}, \mathsf{sk})$ to be created on demand and assigned to a new party U_i. Specifically, the environment can create new parties who will interact with $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ for their public/secret-key in this way treating it as a trusted component that maintains the secret of their wallet. Note that the adversary can interfere with the creation of a new party, corrupt it, and supply its own (adversarially created) public-key instead. As before, the environment, may request transactions between accounts from stakeholders and it can also generate transactions in collaboration with the adversary on behalf of the corrupted accounts. Recall that our assumption is that at any slot, in the view of any honest player, the stakeholder distribution satisfies honest majority with advantage ϵ (note that different honest players might perceive a different stakeholder distribution in a certain slot). Furthermore, the stake can shift by at most σ statistical distance over a certain number of slots. The statistical distance here will be measured considering the underlying distribution to be the weighted-by-stake sampler and how it changes over the specified time interval. The security proof can be seen as an induction in the number of epochs L/R with

the base case supplied by the proof of the static stake protocol. In the end we will argue that in this setting, a $\frac{1-\epsilon}{2} - \sigma$ bound in adversarial stake is sufficient for security of a single draw (and observe that the size of committee, m, now should be selected to overcome also an additive term of size $\ln(L/R)$ given that the lifetime of the systems includes such a number of successive epochs). The corruption delay remains at $D = R$ which can be selected arbitrarily smaller than L, thus enabling the adversary to perform adaptive corruptions as long as this is not instantaneous.

Stage 3. (dynamic state without a beacon, epoch period of R slots long, $R = \Theta(k)$ and delay $D \in (R, 2R) \ll L$). In the third stage, we remove the dependency to the beacon, by introducing a secure multiparty protocol with "guaranteed output delivery" that simulates it. In this way, we can obtain the long-livedness of the protocol as described in the stage 2 design but only under the assumption of the stage 1 design, i.e., the mere availability of an initial random string and an initial stakeholder distribution with honest majority. The core idea is the following: given we guarantee that an honest majority among elected stakeholders will hold with very high probability, we can further use this elected set as participants to an instance of a secure multiparty computation (MPC) protocol. This will require the choice of the length of the epoch to be sufficient so that it can accommodate a run of the MPC protocol. From a security point of view, the main difference with the previous case, is that the output of the beacon will become known to the adversary before it may become known to the honest parties. Nevertheless, we will prove that the honest parties will also inevitably learn it after a short number of slots. To account for the fact that the adversary gets this headstart (which it may exploit by performing adaptive corruptions) we increase the wait time for corruption from R to a suitable value in $(R, 2R)$ that negates this advantage and depends on the secure MPC design. A feature of this stage from a cryptographic design perspective is the use of the ledger itself for the simulation of a reliable broadcast that supports the MPC protocol.

Stage 4. (input endorsers, stakeholder delegates, anonymous communication). In the final stage of our design, we augment the protocol with two new roles for the entities that are running the protocol and consider the benefits of anonymous communication. Input-endorsers create a second layer of transaction endorsing prior to block inclusion. This mechanism enables the protocol to withstand deviations such as selfish mining and enables us to show that honest behaviour is an approximate Nash equilibrium under reasonable assumptions regarding the costs of running the protocol. Note that input-endorsers are assigned to slots in the same way that slot leaders are, and inputs included in blocks are only acceptable if they are endorsed by an eligible input-endorser. Second, the *delegation* feature allows stakeholders to transfer committee participation to selected delegates that assume the responsibility of the stakeholders in running the protocol (including participation to the MPC and issuance of blocks). Delegation naturally gives rise to "stake pools" that can act in the same way as mining pools in bitcoin. Finally, we observe that by including an anonymous communication layer we can remove the corruption delay requirement that is imposed in our analysis. This is done at the expense of increasing the online time requirements for the honest parties.

4 Our Protocol: Static State

4.1 Basic Concepts and Protocol Description

We begin by describing the blockchain protocol π_{SPoS} in the "static stake" setting, where leaders are assigned to blockchain slots with probability proportional to their (fixed) initial stake which

will be the effective stake distribution throughout the execution. To simplify our presentation, we abstract this leader selection process, treating it simply as an "ideal functionality" that faithfully carries out the process of randomly assigning stakeholders to slots. In the following section, we explain how to instantiate this functionality with a specific secure computation.

We remark that—even with an ideal leader assignment process—analyzing the standard "longest chain" preference rule in our PoS setting appears to require significant new ideas. The challenge arises because large collections of slots (epochs, as described above) are assigned to stakeholders at once; while this has favorable properties from an efficiency (and incentive) perspective, it furnishes the adversary a novel means of attack. Specifically, an adversary in control of a certain population of stakeholders can, at the beginning of an epoch, choose when standard "chain update" broadcast messages are delivered to honest parties with full knowledge of future assignments of slots to stakeholders. In contrast, adversaries in typical PoW settings are constrained to make such decisions in an online fashion. We remark that this can have a dramatic effect on the ability of an adversary to produce alternate chains; see the discussion on "forkable strings" below for detailed discussion.

In the static stake case, we assume that a fixed collection of n stakeholders $U_1, \ldots, U_n$ interact throughout the protocol. Stakeholder U_i possesses s_i stake before the protocol starts. For each stakeholder U_i a verification and signing key pair $(\mathsf{vk}_i, \mathsf{sk}_i)$ for a prescribed signature scheme is generated; we assume without loss of generality that the verification keys $\mathsf{vk}_1, \ldots$ are known by all stakeholders. Before describing the protocol, we establish basic definitions following the notation of [19].

Definition 4.1 (Genesis Block)**.** *The* genesis block B_0 *contains the list of stakeholders identified by their public-keys, their respective stakes* $(\mathsf{vk}_1, s_1), \ldots, (\mathsf{vk}_n, s_n)$ *and auxiliary information* ρ.

With foresight we note that the auxiliary information ρ will be used to seed the slot leader election process.

Definition 4.2 (State)**.** *A* state *is a string* $st \in \{0,1\}^\lambda$.

Definition 4.3 (Block)**.** *A block* B *generated at a slot* $sl_i \in \{sl_1, \ldots, sl_R\}$ *contains the current state* $st \in \{0,1\}^\lambda$, *data* $d \in \{0,1\}^*$, *the slot number* sl_i *and a signature* $\sigma = \mathsf{Sign}_{\mathsf{sk}_i}(st, d, sl)$ *computed under* sk_i *corresponding to the stakeholder* U_i *generating the block.*

Definition 4.4 (Blockchain)**.** *A* blockchain *(or simply* chain*) relative to the genesis block* B_0 *is a sequence of blocks* $B_1, \ldots, B_n$ *associated with a strictly increasing sequence of slots for which the state* st_i *of* B_i *is equal to* $H(B_{i-1})$, *where* H *is a prescribed collision-resistant hash function. The* length *of a chain* $\mathrm{len}(\mathcal{C}) = n$ *is its number of blocks. The block* B_n *is the* head *of the chain, denoted* $\mathrm{head}(\mathcal{C})$. *We treat the empty string* ε *as a legal chain and by convention set* $\mathrm{head}(\varepsilon) = \varepsilon$.

Let $\mathcal{C}$ be a chain of length n and k be any non-negative integer. We denote by $\mathcal{C}^{\lceil k}$ the chain resulting from removal of the k rightmost blocks of $\mathcal{C}$. If $k \geq \mathrm{len}(\mathcal{C})$ we define $\mathcal{C}^{\lceil k} = \varepsilon$. We let $\mathcal{C}_1 \preceq \mathcal{C}_2$ indicate that the chain $\mathcal{C}_1$ is a prefix of the chain $\mathcal{C}_2$.

Definition 4.5 (Epoch)**.** *An* epoch *is a set of* R *adjacent slots* $S = \{sl_1, \ldots, sl_R\}$.

(The value R is a parameter of the protocol we analyze in this section.)

Definition 4.6 (Adversarial Stake Ratio)**.** *Let* $U_\mathcal{A}$ *be the set of stakeholders controlled by an adversary* $\mathcal{A}$. *Then the* adversarial stake ratio *is defined as*

$$\alpha = \frac{\sum_{j \in U_\mathcal{A}} s_j}{\sum_{i=1}^n s_i},$$

where n *is the total number of stakeholders and* s_i *is stakeholder* U_i*'s stake.*

Slot Leader Selection.   In the protocol described in this section, for each $0 < j \leq R$, a *slot leader* E_j is determined who has the (sole) right to generate a block at sl_j. Specifically, for each slot a stakeholder U_i is selected as the slot leader with probability p_i proportional to its stake registered in the genesis block B_0; these assignments are independent between slots. In this static stake case, the genesis block as well as the procedure for selecting slot leaders are determined by an ideal functionality $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$, defined in Figure 1. This functionality is parameterized by the list $\{(\mathsf{vk}_1, s_1), \ldots, (\mathsf{vk}_n, s_n)\}$ assigning to each stakeholder its respective stake, a distribution $\mathcal{D}$ that provides auxiliary information ρ and a *leader selection function* F defined below.

Definition 4.7 (Leader Selection Process). *A leader selection process *with respect to stakeholder distribution* $\mathbb{S} = \{(\mathsf{vk}_1, s_1), \ldots, (\mathsf{vk}_n, s_n)\}$, $(\mathcal{D}, \mathsf{F})$ *is a pair consisting of a distribution and a deterministic function such that, when $\rho \leftarrow \mathcal{D}$ it holds that for all $sl_j \in \{sl_1, \ldots, sl_R\}$, $\mathsf{F}(\mathbb{S}, \rho, sl_j)$ outputs $U_i \in \{U_1, \ldots, U_n\}$ with probability*

$$p_i = \frac{s_i}{\sum_{k=1}^{n} s_k}$$

where s_i is the stake held by stakeholder U_i (we call this "weighing by stake"); furthermore the family of random variables $U_i = \mathsf{F}(\mathbb{S}, \rho, sl_j)$ are independent.

We note that weighing by stake sampling can be implemented in a straightforward manner. For instance, a simple process operates as follows. Let $\tilde{p}_i = s_i / \sum_{j=i}^{n} s_i$. First, the process flips a $\tilde{p}_1$-biased coin to see whether the first stakeholder is selected; subsequently, for all $j \geq 2$, it flips a $(1 - \tilde{p}_1) \ldots (1 - \tilde{p}_{j-1}) \tilde{p}_j$ biased coin to see whether the j-th stakeholder is selected. When we implement it as a function $F(\cdot)$ sufficient randomness must be allocated to simulate the biased coin flips. If we implement the above with λ precision for each individual coin flip, then selecting a stakeholder will require $n \lceil \log \lambda \rceil$ random bits in total. Note that using a pseudorandom number generator (PRG) one may use a shorter "seed" string and then stretch it using the PRG to the appropriate length.
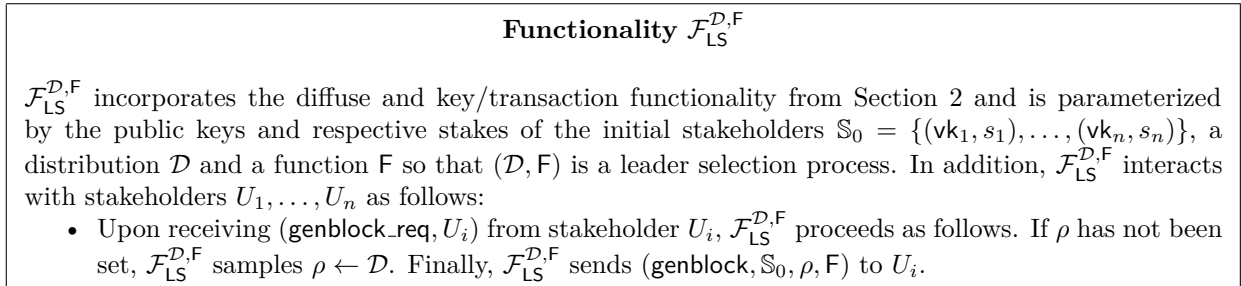

Functionality $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$

$\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ incorporates the diffuse and key/transaction functionality from Section 2 and is parameterized by the public keys and respective stakes of the initial stakeholders $\mathbb{S}_0 = \{(\mathsf{vk}_1, s_1), \ldots, (\mathsf{vk}_n, s_n)\}$, a distribution $\mathcal{D}$ and a function F so that $(\mathcal{D}, \mathsf{F})$ is a leader selection process. In addition, $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ interacts with stakeholders $U_1, \ldots, U_n$ as follows:

- Upon receiving $(\mathsf{genblock\_req}, U_i)$ from stakeholder U_i, $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ proceeds as follows. If ρ has not been set, $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ samples $\rho \leftarrow \mathcal{D}$. Finally, $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ sends $(\mathsf{genblock}, \mathbb{S}_0, \rho, \mathsf{F})$ to U_i.

Figure 1: Functionality $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$.

A Protocol in the $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$-hybrid model.   We start by describing a simple PoS based blockchain protocol considering static stake in the $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$-hybrid model, i.e., where the genesis block B_0 (and consequently the slot leaders) are determined by the ideal functionality $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$. The stakeholders $U_1, \ldots, U_n$ interact among themselves and with $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ through Protocol π_{SPoS} described in Figure 2.

The protocol relies on a $\mathsf{maxvalid}_S(\mathcal{C}, \mathbb{C})$ function that chooses a chain given the current chain $\mathcal{C}$ and a set of valid chains $\mathbb{C}$ that are available in the network. In the static case we analyze the simple "longest chain" rule. (In the dynamic case the rule is parameterized by a common chain length; see Section 5.)

Function maxvalid($\mathcal{C}, \mathbb{C}$): Returns the longest chain from $\mathbb{C} \cup \{\mathcal{C}\}$. Ties are broken in favor of $\mathcal{C}$, if it has maximum length, or arbitrarily otherwise.
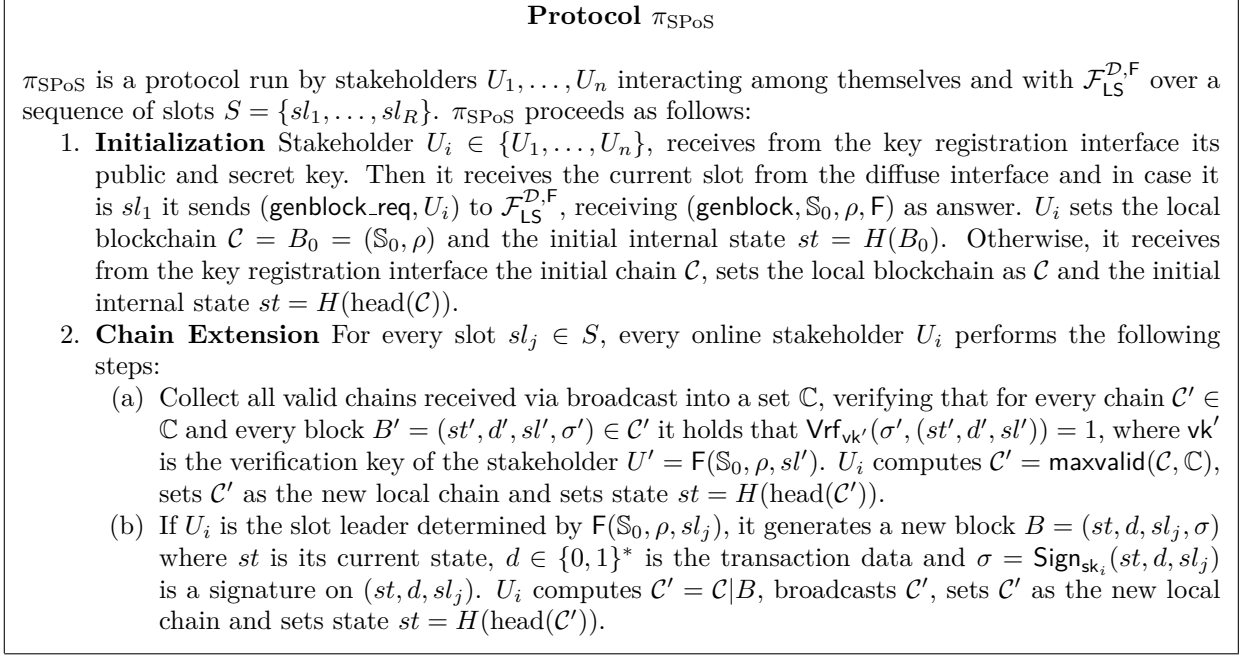

Protocol π_{SPoS}

π_{SPoS} is a protocol run by stakeholders $U_1, \ldots, U_n$ interacting among themselves and with $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D}, \mathsf{F}}$ over a sequence of slots $S = \{sl_1, \ldots, sl_R\}$. π_{SPoS} proceeds as follows:

1. **Initialization** Stakeholder $U_i \in \{U_1, \ldots, U_n\}$, receives from the key registration interface its public and secret key. Then it receives the current slot from the diffuse interface and in case it is sl_1 it sends (genblock_req, U_i) to $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D}, \mathsf{F}}$, receiving (genblock, $\mathbb{S}_0, \rho, \mathsf{F}$) as answer. U_i sets the local blockchain $\mathcal{C} = B_0 = (\mathbb{S}_0, \rho)$ and the initial internal state $st = H(B_0)$. Otherwise, it receives from the key registration interface the initial chain $\mathcal{C}$, sets the local blockchain as $\mathcal{C}$ and the initial internal state $st = H(\mathrm{head}(\mathcal{C}))$.

2. **Chain Extension** For every slot $sl_j \in S$, every online stakeholder U_i performs the following steps:

    (a) Collect all valid chains received via broadcast into a set $\mathbb{C}$, verifying that for every chain $\mathcal{C}' \in \mathbb{C}$ and every block $B' = (st', d', sl', \sigma') \in \mathcal{C}'$ it holds that $\mathsf{Vrf}_{\mathsf{vk}'}(\sigma', (st', d', sl')) = 1$, where vk' is the verification key of the stakeholder $U' = \mathsf{F}(\mathbb{S}_0, \rho, sl')$. U_i computes $\mathcal{C}' = \mathsf{maxvalid}(\mathcal{C}, \mathbb{C})$, sets $\mathcal{C}'$ as the new local chain and sets state $st = H(\mathrm{head}(\mathcal{C}'))$.

    (b) If U_i is the slot leader determined by $\mathsf{F}(\mathbb{S}_0, \rho, sl_j)$, it generates a new block $B = (st, d, sl_j, \sigma)$ where st is its current state, $d \in \{0, 1\}^*$ is the transaction data and $\sigma = \mathsf{Sign}_{\mathsf{sk}_i}(st, d, sl_j)$ is a signature on (st, d, sl_j). U_i computes $\mathcal{C}' = \mathcal{C}|B$, broadcasts $\mathcal{C}'$, sets $\mathcal{C}'$ as the new local chain and sets state $st = H(\mathrm{head}(\mathcal{C}'))$.

Figure 2: Protocol π_{SPoS}.

4.2 Forkable Strings

In our security arguments we routinely use elements of $\{0, 1\}^n$ to indicate which slots—among a particular window of slots of length n—have been assigned to adversarial stakeholders. When strings have this interpretation we refer to them as *characteristic strings*.

Definition 4.8 (Characteristic String). *Fix an execution with genesis block B_0, adversary $\mathcal{A}$, and environment $\mathcal{Z}$. Let $S = \{sl_{i+1}, \ldots, sl_{i+n}\}$ denote a sequence of slots of length $|S| = n$. The characteristic string $w \in \{0, 1\}^n$ of S is defined so that $w_k = 1$ if and only if the adversary controls the slot leader of slot sl_{i+k}. For such a characteristic string $w \in \{0, 1\}^*$ we say that the index i is adversarial if $w_i = 1$ and honest otherwise.*

We start with some intuition on our approach to analyze the protocol. Let $w \in \{0, 1\}^n$ be a characteristic string for a sequence of slots S. Consider two observers that (i.) go offline immediately prior to the commencement of S, (ii.) have the same view $\mathcal{C}_0$ of the current chain prior to the commencement of S, and (iii.) come back online at the last slot of S and request an update of their chain. A fundamental concern in our analysis is the possibility that such observers can be presented with a "diverging" view over the sequence S: specifically, the possibility that the adversary can force the two observers to adopt two different chains $\mathcal{C}_1, \mathcal{C}_2$ whose common prefix is $\mathcal{C}_0$.

We observe that not all characteristic strings permit this. For instance the (entirely honest) string 0^n ensures that the two observers will adopt the same chain $\mathcal{C}$ which will consist of n new blocks on top of the common prefix $\mathcal{C}_0$. On the other hand, other strings do not guarantee such common extension of $\mathcal{C}_0$; in the case of 1^n, it is possible for the adversary to produce two completely

different histories during the sequence of slots S and thus furnish to the two observers two distinct chains $\mathcal{C}_1, \mathcal{C}_2$ that only share the common prefix $\mathcal{C}_0$. In the remainder of this section, we establish that strings that permit such "forkings" are quite rare—indeed, we show that they have density $2^{-\Omega(\sqrt{n})}$ so long as the fraction of adversarial slots is $1/2 - \epsilon$.

To reason about such "forkings" of a characteristic string $w \in \{0, 1\}^n$, we define below a formal notion of "fork" that captures the relationship between the chains broadcast by *honest* slot leaders during an execution of the protocol π_{SPoS}. In preparation for the definition, we recall that honest players always choose to extend a maximum length chain among those available to the player on the network. Furthermore, if such a maximal chain $\mathcal{C}$ includes a block B previously broadcast by an honest player, the prefix of $\mathcal{C}$ prior to B must entirely agree with the chain (terminating at B) broadcast by this previous honest player. This "confluence" property follows immediately from the fact that the state of any honest block effectively commits to a unique chain beginning at the genesis block. To conclude, any chain $\mathcal{C}$ broadcast by an honest player must begin with a chain produced by a previously honest player (or, alternatively, the genesis block), continue with a possibly empty sequence of adversarial blocks and, finally, terminate with an honest block. It follows that the chains broadcast by honest players form a natural directed tree. The fact that honest players reliably broadcast their chains and always build on the longest available chain introduces a second important property of this tree: the "depths" of the various honest blocks added by honest players during the protocol must all be distinct.

Of course, the actual chains induced by an execution of π_{SPoS} are comprised of blocks containing a variety of data that are immaterial for reasoning about forking. For this reason the formal notion of fork below merely reflects the directed tree formed by the relevant chains and the *identities of the players*—expressed as indices in the string w—responsible for generating the blocks in these chains.

Forks and forkable strings. We define, below, the basic combinatorial structures we use to reason about the possible views observed by honest players during a protocol execution with this characteristic string.


Figure 3: A fork F for the string $w = 010100110$; vertices appear with their labels and honest vertices are highlighted with double borders. Note that the depths of the (honest) vertices associated with the honest indices of w are strictly increasing. Two tines are distinguished in the figure: one, labeled $\hat{t}$, terminates at the vertex labeled 9 and is the longest tine in the fork; a second tine t terminates at the vertex labeled 3. The quantity $\mathrm{gap}(t)$ indicates the difference in length between t and $\hat{t}$; in this case $\mathrm{gap}(t) = 4$. The quantity $\mathrm{reserve}(t) = |\{i \mid \ell(v) < i \leq |w| \text{ and } w_i = 1\}|$ indicates the number of adversarial indices appearing after the label of the last honest vertex v of the tine; in this case $\mathrm{reserve}(t) = 3$. As each leaf of F is honest, F is closed.

Definition 4.9 (Fork)**.** *Let $w \in \{0,1\}^n$ and let $H = \{i \mid w_i = 0\}$ denote the set of honest indices. A* fork *for the string w is a directed, rooted tree $F = (V, E)$ with a labeling $\ell : V \to \{0, 1, \ldots, n\}$ so that*

- *each edge of F is directed away from the root;*

- *the root $r \in V$ is given the label $\ell(r) = 0$;*

- *the labels along any directed path in the tree are strictly increasing;*

- *each honest index $i \in H$ is the label of exactly one vertex of F;*

- *the function $\mathbf{d} : H \to \{1, \ldots, n\}$, defined so that $\mathbf{d}(i)$ is the depth in F of the unique vertex v for which $\ell(v) = i$, is strictly increasing. (Specifically, if $i, j \in H$ and $i < j$, then $\mathbf{d}(i) < \mathbf{d}(j)$.)*

As a matter of notation, we write $F \vdash w$ to indicate that F is a fork for the string w. We say that a fork is trivial *if it contains a single vertex, the root.*

Definition 4.10 (Tines and height)**.** *A path in a fork F originating at the root is called a* tine. *For a tine t we let $\mathrm{length}(t)$ denote its* length, *equal to the number of edges on the path. The* height *of a fork (as usual for a tree) is defined to be the length of the longest tine. For two tines t_1 and t_2 of a fork F, we write $t_1 \sim t_2$ if they share an edge. Note that $\sim$ is an equivalence relation on the set of nontrivial tines; on the other hand, if t_ϵ denotes the "empty" tine consisting solely of the root vertex then $t_\epsilon \not\sim t$ for any tine t.*

If a vertex v of a fork is labeled with an adversarial index (i.e., $w_{\ell(v)} = 1$) we say that the vertex is *adversarial*; otherwise, we say that the vertex is *honest*. For convenience, we declare the root vertex to be honest. We extend this terminology to tines: a tine is *honest* if it terminates with an honest vertex and *adversarial* otherwise. By this convention the empty tine t_ϵ is honest.

See Figure 3 for an example, which also demonstrates some of the quantities defined above and in the remainder of this section. The fork shown in the figure reflects an execution in which *(i.)* the honest player associated with the first slot builds directly on the genesis block (as it must), *(ii.)* the honest player associated with the third slot is shown a chain of length 1 produced by the adversarial player of slot 2 (in addition to the honestly generated chain of step *(i.)*), which it elects to extend, *(iii.)* the honest player associated with slot 5 is shown a chain of length 2 building on the chain of step *(i.)* augmented with a further adversarial block produced by the player of slot 4, etc.

Definition 4.11. *We say that a fork is* flat *if it has two tines $t_1 \not\sim t_2$ of length equal to the height of the fork. A string $w \in \{0,1\}^*$ is said to be* forkable *if there is a flat fork $F \vdash w$.*

Note that in order for an execution of π_{SPoS} to yield two entirely disjoint chains of maximum length, the characteristic string associated with the execution must be forkable. Our goal is to establish the following upper bound on the number of forkable strings.

Structural features of forks: closed forks, prefixes, reach, and margin. We begin by defining a natural notion of inclusion for two forks:

Definition 4.12 (Fork prefixes)**.** *If w is a prefix of the string $w' \in \{0,1\}^*$, $F \vdash w$, and $F' \vdash w'$, we say that F is a* prefix *of F', written $F \sqsubseteq F'$, if F is a consistently-labeled subgraph of F'. Specifically, every vertex and edge of F appears in F' and, furthermore, the labels given to any vertex appearing in both F and F' are identical.*

If $F \sqsubseteq F'$, each tine of F appears as the prefix of a tine in F'. In particular, the labels appearing on any tine terminating at a common vertex are identical and, moreover, the depth of any honest vertex appearing in both F and F' is identical.

In many cases, it is convenient to work with forks that do not "commit" anything beyond final honest indices.

Definition 4.13 (Closed forks). *A fork is* closed *if each leaf is honest. By convention the trivial fork, consisting solely of a root vertex, is closed.*

Note that a closed fork has a unique longest tine (as all maximal tines terminate with an honest vertex, and these must have distinct depths). Note, additionally, that if w is a prefix of w' and $F' \vdash w'$, then there is a unique closed fork $F \vdash w$ for which $F \sqsubseteq F'$.

Definition 4.14 (Gap, reserve and reach). *Let $F \vdash w$ be a closed fork and let $\hat{t}$ denote the (unique) tine of maximum length in F. We define the* gap *of a tine t, denoted $\mathrm{gap}(t)$, to be the difference in length between $\hat{t}$ and t; thus*

$$\mathrm{gap}(t) = \mathrm{length}(\hat{t}) - \mathrm{length}(t) \,.$$

We define the reserve *of a tine t to be the number of adversarial indices appearing in w after the last index in t; specifically, if t is given by the path $(r, v_1, \ldots, v_k)$, where r is the root of F, we define*

$$\mathrm{reserve}(t) = |\{i \mid w_i = 1 \text{ and } i > \ell(v_k)\}| \,.$$

We remark that this quantity depends both on F and the specific string w associated with F. Finally, for a tine t we define

$$\mathrm{reach}(t) = \mathrm{reserve}(t) - \mathrm{gap}(t) \,.$$

Definition 4.15 (Margin). *For closed fork $F \vdash w$ we define $\lambda(F)$ to be the maximum reach taken over all tines in F:*

$$\lambda(F) = \max_t \mathrm{reach}(t) \,.$$

Likewise, we define the margin *of F, denoted $\mu(F)$, to be the "penultimate" reach taken over edge-disjoint tines of F: specifically,*

$$\mathrm{margin}(F) = \mu(F) = \max_{t_1 \nsim t_2} \Big(\min\{\mathrm{reach}(t_1), \mathrm{reach}(t_2)\} \Big) \,. \tag{1}$$

We remark that the maxima above can always obtained by honest tines. Specifically, if t is an adversarial tine of a fork $F \vdash w$, $\mathrm{reach}(t) \leq \mathrm{reach}(\bar{t})$, where $\bar{t}$ is the longest honest prefix of t.

As $\sim$ is an equivalence relation on the nonempty tines. It follows that there is always a pair of (edge-disjoint) tines t_1 and t_2 achieving the maximum in the defining equation (11) which satisfy $\mathrm{reach}(t_1) = \lambda(F) \geq \mathrm{reach}(t_2) = \mu(F)$.

The relevance of margin to the notion of forkability is reflected in the following proposition.

Proposition 4.16. *A string w is forkable if and only if there is a closed fork $F \vdash w$ for which $\mathrm{margin}(F) \geq 0$.*

Proof. If w has no honest indices, then the trivial fork consisting of a single root node is flat, closed, and has non-negative margin; thus the two conditions are equivalent. Consider a forkable string w with at least one honest index and let $\hat{i}$ denote the largest honest index of w. Let F be a flat fork for w. As mentioned above, there is a unique closed fork $\overline{F} \vdash w$ obtained from F by removing any adversarial vertices from the ends of the tines of F. Note that the tine $\hat{t}$ containing $\hat{i}$ is the longest

tine in $\overline{F}$, as this is the largest honest index of w. On the other hand, F is flat, in which case there are two edge-disjoint tines t_1 and t_2 with length at least that of $\hat{t}$. The prefixes of these two tines in $\overline{F}$ must clearly have reserve no less than gap (and hence non-negative reach); thus $\mathrm{margin}(\overline{F}) \geq 0$ as desired.

On the other hand, suppose w has a closed fork with $\mathrm{margin}(F) \geq 0$, in which case there are two edge-disjoint tines of F, t_1 and t_2, for which $\mathrm{reach}(t_i) \geq 0$. Then we can produce a flat fork by simply adding to each t_i a path of $\mathrm{gap}(t_i)$ vertices labeled with the subsequent adversarial indices promised by the definition of reserve(). $\qquad\square$

In light of this proposition, for a string w we focus our attention on the quantities

$$\lambda(w) = \max_{\substack{F \vdash w, \\ F \text{ closed}}} \lambda(F)\,, \qquad \mu(w) = \max_{\substack{F \vdash w, \\ F \text{ closed}}} \mu(F)\,,$$

and, for convenience,

$$\mathbf{m}(w) = (\lambda(w), \mu(w))\,.$$

Note that this overloads the notation $\lambda(\cdot)$ and $\mu(\cdot)$ so that they apply to both forks and strings, but the setting will be clear from context. We remark that the definitions do not guarantee a priori that $\lambda(w)$ and $\mu(w)$ can be achieved by the same fork, though this will be established in the lemma below. In any case, it is clear that $\lambda(w) \geq 0$ and $\lambda(w) \geq \mu(w)$ for all strings w; furthermore, by Proposition 4.23 a string w is forkable if and only if $\mu(w) \geq 0$. We refer to $\mu(w)$ as the *margin* of the string w.

In preparation for the proof of Theorem 4.18, we establish a recursive description for these quantities.

Lemma 4.17. $\mathbf{m}(\epsilon) = (0, 0)$ *and, for all nonempty strings* $w \in \{0, 1\}^*$,

$$\mathbf{m}(w1) = (\lambda(w) + 1, \mu(w) + 1)\,, \text{ and}$$

$$\mathbf{m}(w0) = \begin{cases} (\lambda(w) - 1, 0) & \text{if } \lambda(w) > \mu(w) = 0, \\ (0, \mu(w) - 1) & \text{if } \lambda(w) = 0, \\ (\lambda(w) - 1, \mu(w) - 1) & \text{otherwise.} \end{cases}$$

Furthermore, for every string w, there is a closed fork $F_w \vdash w$ for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$.

Proof. The proof proceeds by induction. If $w = \epsilon$, define F_ϵ to be the trivial fork; $F_\epsilon \vdash w$ is the unique closed fork for this string and $\mathbf{m}(\epsilon) = (0, 0) = (\lambda(F_\epsilon), \mu(F_\epsilon))$, as desired.

In general, we consider $\mathbf{m}(w')$ for a string $w' = wx$—where $w \in \{0, 1\}^*$ and $x \in \{0, 1\}$; the argument recursively expands $\mathbf{m}(w')$ in terms of $\mathbf{m}(w)$ and the value of the last symbol x. In each case, we consider the relationship between two closed forks $F \sqsubseteq F'$ where $F \vdash w$ and $F' \vdash w' = wx$.

In the case where $x = 1$, we must have $F = F'$ as graphs, because the forks are assumed to be closed; it is easy to see that the reach of any tine t of $F \vdash w$ has increased by exactly one when viewed as a tine of $F' \vdash w'$. We write $\mathrm{reach}_{F'}(t) = \mathrm{reach}_F(t) + 1$, where we introduce the notation $\mathrm{reach}_\square()$ to denote the reach in a particular fork. It follows that $\lambda(F') = \lambda(F) + 1$ and $\mu(F') = \mu(F) + 1$. If $F^* \vdash w'$ is a closed fork for which $\lambda(F^*) = \lambda(w')$, note that F^* may be treated as a fork for w and, applying the argument above, we find that $\lambda(w') \leq \lambda(w) + 1$. A similar argument implies that $\mu(w') \leq \mu(w) + 1$. On the other hand, by induction there is a fork F_w for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$ and hence $\mathbf{m}(w') \geq (\lambda(w) + 1, \mu(w) + 1)$. We conclude that

$$\mathbf{m}(w') = (\lambda(w) + 1, \mu(w) + 1)\,. \tag{2}$$

16

Moreover, $\mathbf{m}(w') = (\lambda(F_w), \mu(F_w))$, where F_w is treated as a fork for $w' = w1$.

The case when $x = 0$ is more delicate. As above, we consider the relationship between two closed forks $F \vdash w$ and $F' \vdash w' = w0$ for which $F \sqsubseteq F'$. Here F' is necessarily obtained from F by appending a path labeled with a string of the form 1^a0 to the end of a tine t of F. (In fact, it is easy to see that we may always assume that this is appended to an honest tine.) In order for this to be possible, $\mathrm{gap}(t) \leq \mathrm{reserve}(t)$ (which is to say that $\mathrm{reach}(t) \geq 0$) and, in particular, $\mathrm{gap}(t) \leq a \leq \mathrm{reserve}(t)$: for the first inequality, note that the depth of the new honest vertex must exceed that of the deepest (honest) vertex in F and hence $a \geq \mathrm{gap}(t)$; as for the second inequality, there are only $\mathrm{reserve}(t)$ possible adversarial indices that may be added to t and hence $a \leq \mathrm{reserve}(t)$. We define the quantity $\tilde{a} \geq 0$ by the equation $a = \mathrm{gap}(t) + \tilde{a}$ and let t' denote the tine (of F') resulting by extending t in this way. We say that $\tilde{a}$ is the *parameter* for this pair of forks $F \sqsubseteq F'$.

Of course, every honest tine t of F is an honest tine of F' and it is clear that $\mathrm{reach}_{F'}(t) = \mathrm{reach}_F(t) - (\tilde{a} + 1)$, as the length of the longest tine t' in F' exceeds the length of the longest tine of F by exactly $\tilde{a} + 1$. Note that the reach of the new honest tine t' (in F') is always 0, as both $\mathrm{gap}(t')$ and $\mathrm{reserve}(t')$ are zero. It remains to describe how $\mu(w)$ and $\lambda(w)$ are determined by this process.

The case $\lambda(w) > \mu(w) = 0$. By induction, there is a fork F_w for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$. Let t_1 and t_2 be edge-disjoint tines of F_w for which $\lambda(F_w) = \mathrm{reach}(t_1)$ and $\mu(F_w) = \mathrm{reach}(t_2)$. Define $F' \vdash w'$ to be the fork obtained by extending the tine t_2 of F_w with parameter $\tilde{a} = 0$ to yield a new tine t'_2 in F'. Then $\mathrm{reach}_{F'}(t_1) = \lambda(w) - 1$ and $\mathrm{reach}_{F'}(t'_2) = 0$. It follows that $\lambda(w0) \geq \lambda(w) - 1$ and $\mu(w0) \geq 0$. We will show that $\lambda(w0) \leq \lambda(w) - 1$ and that $\mu(w0) \leq 0$, in which case we can conclude that

$$\lambda(w0) = \lambda(w) - 1 \qquad \text{and} \qquad \mu(w0) = 0\,.$$

Moreover, the fork $F_{w'} = F'$ achieves these statistics, as desired.

We return to establish that $\lambda(w0) \leq \lambda(w) - 1$ and that $\mu(w0) \leq 0$. Let $F^* \vdash w0$ be a closed fork for which $\lambda(w0) = \lambda(F^*)$ and let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$; as above, let $\tilde{a}$ denote the parameter for this extension. Let t^* be an honest tine of F^* so that $\mathrm{reach}_{F^*}(t^*) = \lambda(w0)$. If t^* is a tine of F, $\mathrm{reach}_{F^*}(t^*) = \mathrm{reach}_F(t^*) - (\tilde{a} + 1) \leq \lambda(w) - 1$. Otherwise t^* was obtained by extension and $\mathrm{reach}_{F^*}(t^*) = 0 \leq \lambda(w) - 1$ by assumption. In either case $\lambda(w0) \leq \lambda(w) - 1$, as desired. It remains to show that $\mu(w0) \leq 0$. Now consider $F^* \vdash w0$ to be a closed fork for which $\mu(F^*) = \mu(w0)$. Let t_1^* and t_2^* be two edge-disjoint honest tines of F^* so that $\mathrm{reach}_{F^*}(t_1^*) = \lambda(F^*)$ and $\mathrm{reach}_{F^*}(t_2^*) = \mu(F^*) = \mu(w0)$. Let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$ and let $\tilde{a}$ be the parameter for this extension. If both t_1^* and t_2^* are tines of F, $\mathrm{reach}_{F^*}(t_i^*) = \mathrm{reach}_F(t_i^*) - (\tilde{a} + 1)$ and, in particular, $\mathrm{reach}_F(t_1^*) \geq \mathrm{reach}_F(t_2^*)$. It follows that $\mathrm{reach}_F(t_2^*) \leq \mu(F) \leq \mu(w) = 0$ and hence that $\mu(w0) < 0$. Otherwise, one of the two tines was the result of extension and has zero $\mathrm{reach}()$ in F^*. As $\mathrm{reach}_{F^*}(t_1^*) \geq \mathrm{reach}_{F^*}(t_2^*)$, in either case it follows that $\mu(F^*) = \mathrm{reach}_{F^*}(t_2^*) \leq 0$, as desired.

The case $\lambda(w) = 0$. By induction, there is a fork F_w for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$. Let t_1 and t_2 be edge-disjoint tines of F_w for which $\lambda(F_w) = \mathrm{reach}(t_1)$ and $\mu(F_w) = \mathrm{reach}(t_2)$. Define $F' \vdash w'$ to be the fork obtained by extending the tine t_1 of F_w with parameter $\tilde{a} = 0$ to yield a new tine t'_1 in F'. Then $\mathrm{reach}_{F'}(t'_1) = 0$ and $\mathrm{reach}_{F'}(t_2) = \mathrm{reach}_F(t_2) - 1$. It follows that $\lambda(w0) \geq 0$ and $\mu(w0) \geq \mu(w) - 1$. We will show that $\lambda(w0) \leq 0$ and that $\mu(w0) \leq \mu(w) - 1$,

17

in which case we can conclude that

$$\lambda(w0) = 0 \qquad \text{and} \qquad \mu(w0) = \mu(w) - 1\,.$$

Moreover, the fork $F_{w'} = F'$ achieves these statistics, as desired.

We return to establish that $\lambda(w0) \leq 0$ and that $\mu(w0) \leq \mu(w) - 1$. Let $F^* \vdash w0$ be a closed fork for which $\lambda(w0) = \lambda(F^*)$ and let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$; as above, let $\tilde{a}$ denote the parameter for this extension. Let t^* be an honest tine of F^* so that $\mathrm{reach}_{F^*}(t^*) = \lambda(w0)$. Note that t^* cannot be a tine of F; if it were then $\mathrm{reach}_{F^*}(t^*) = \mathrm{reach}_F(t^*) - (\tilde{a} + 1) \leq \lambda(w) - 1 < 0$ which contradicts $\lambda(w0) \geq 0$. Thus t^* was obtained by extension and $\mathrm{reach}_{F^*}(t^*) = 0$. It remains to show that $\mu(w0) \leq 0$. Now let $F^* \vdash w0$ be a closed fork for which $\mu(F^*) = \mu(w0)$. Let t_1^* and t_2^* be two edge-disjoint honest tines of F^* so that $\mathrm{reach}_{F^*}(t_1^*) = \lambda(F^*)$ and $\mathrm{reach}_{F^*}(t_2^*) = \mu(F^*) = \mu(w0)$. Let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$ and let $\tilde{a}$ be the parameter for this extension. Similarly, t_1^* cannot be a tine of F; if it were, $\lambda(F^*) = \mathrm{reach}_{F^*}(t_1^*) = \mathrm{reach}_F(t_1^*) - (\tilde{a} + 1) \leq \lambda(F) - 1 \leq \lambda(w) - 1 < 0$ which contradicts $\lambda(F) \geq 0$. It follows that t_1^* must extend a tine t_1 of F for which $\mathrm{reach}_F(t_1) = 0$, because extension can only occur for tines of non-negative reach and $\lambda(F) = 0 = \lambda(w)$. Thus t_2^* is a tine of F and $t_1 \not\sim t_2^*$ so that $\mathrm{reach}_F(t_2^*) \leq \mu(F) \leq \mu(w)$ and we conclude that $\mu(w0) = \mathrm{reach}_{F^*}(t_2^*) \leq \mathrm{reach}_F(t_2^*) - 1 \leq \mu(w) - 1$, as desired.

The case $\lambda(w) > 0, \mu(w) \neq 0$. By induction, there is a fork F_w for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$. Let t_1 and t_2 be edge-disjoint tines of F_w for which $\lambda(F_w) = \mathrm{reach}(t_1)$ and $\mu(F_w) = \mathrm{reach}(t_2)$. In fact, any extension of F_w will suffice for the construction; for concreteness, define $F' \vdash w'$ to be the fork obtained by extending the tine t_1 of F_w with parameter $\tilde{a} = 0$. Then $\mathrm{reach}_{F'}(t_i) = \mathrm{reach}_{F_w}(t_i) - 1$. It follows that $\lambda(w0) \geq \lambda(w) - 1$ and $\mu(w0) \geq \mu(w) - 1$. We will show that $\lambda(w0) \leq \lambda(w) - 1$ and that $\mu(w0) \leq \mu(w) - 1$, in which case we can conclude that

$$\lambda(w0) = \lambda(w) - 1 \qquad \text{and} \qquad \mu(w0) = \mu(w) - 1\,.$$

Moreover, the fork $F_{w'} = F'$ achieves these statistics, as desired.

We return to establish that $\lambda(w0) \leq \lambda(w) - 1$ and that $\mu(w0) \leq \mu(w) - 1$. Let $F^* \vdash w0$ be a closed fork for which $\lambda(w0) = \lambda(F^*)$ and let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$; as above, let $\tilde{a}$ denote the parameter for this extension. Let t^* be an honest tine of F^* so that $\mathrm{reach}_{F^*}(t^*) = \lambda(w0)$. Note that if t^* is a tine of F then $\mathrm{reach}_{F^*}(t^*) = \mathrm{reach}_F(t^*) - (\tilde{a} + 1) \leq \lambda(w) - 1$; otherwise t^* is obtained by extension and $\mathrm{reach}_{F^*}(t^*) = 0 \leq \lambda(w) - 1$, as desired. (Recall that $\lambda(w) > 0$.) It remains to show that $\mu(w0) \leq \mu(w) - 1$. Now let $F^* \vdash w0$ be a closed fork for which $\mu(F^*) = \mu(w0)$. Let t_1^* and t_2^* be two edge-disjoint honest tines of F^* so that $\mathrm{reach}_{F^*}(t_1^*) = \lambda(F^*)$ and $\mathrm{reach}_{F^*}(t_2^*) = \mu(F^*) = \mu(w0)$. Let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$ and let $\tilde{a}$ be the parameter for this extension. If both t_1^* and t_2^* are tines of F then $\mathrm{reach}_{F^*}(t_i^*) = \mathrm{reach}_F(t_i^*) - (\tilde{a} + 1)$ and, in particular, $\mathrm{reach}_F(t_1^*) \geq \mathrm{reach}_F(t_2^*)$ so that $\mathrm{reach}_F(t_2^*) \leq \mu(w)$ and $\mathrm{reach}_{F^*}(t_2^*) \leq \mu(w) - 1$, as desired.

To complete the argument, we consider the case that one of the tines t_i^* arises by extension. Note that in this case $\mathrm{reach}_{F^*}(t_2^*) \leq 0$, as either t_2^* is obtained by extension so that it has zero reach, or t_1^* is obtained by extension so that $\mathrm{reach}_{F^*}(t_2^*) \leq \mathrm{reach}_{F^*}(t_1^*) = 0$. Here we further separate the analysis into two cases depending on the sign of $\mu(w)$:

- If $\mu(w) > 0$, then $\mathrm{reach}_{F^*}(t_2^*) \leq 0 \leq \mu(w) - 1$, as desired.

18

- If $\mu(w) < 0$ then t_2^* cannot be the extension of a tine in F. To see this, suppose to the contrary that t_2^* extends a tine t_2 of F; then $\mathrm{reach}_F(t_2) \geq 0$. Additionally, t_1^* must be a tine of F, edge-disjoint from t_2, and $\mathrm{reach}_F(t_1^*) = \mathrm{reach}_{F^*}(t_1^*) + (\tilde{a}+1) > 0$. It follows that $\mu(w) \geq \mu(F) \geq 0$, a contradiction.

  The other possibility is that t_1^* is an extension of a tine t_1 of F in which case $\mathrm{reach}_F(t_1) \geq 0$. Note that t_2^* is a tine of F and edge-disjoint from t_1; thus $\min(\mathrm{reach}_F(t_2^*), \mathrm{reach}_F(t_1)) \leq \mu(F) < 0$ and $\mathrm{reach}_F(t_2^*) \leq \mu(F)$. We conclude that $\mathrm{reach}_{F^*}(t_2^*) = \mathrm{reach}_F(t_2^*) - (\tilde{a}+1) \leq \mu(w) - 1$, as desired. $\qquad\square$

With this recursive description in place, we return to the proof of Theorem 4.18, which we restate below for convenience.

Theorem 4.18, restated. *Let $\epsilon \in (0,1)$ and let w be a string drawn from $\{0,1\}^n$ by independently assigning each $w_i = 1$ with probability $(1-\epsilon)/2$. Then*

$$\Pr[w \text{ is forkable}] = 2^{-\Omega(\sqrt{n})}.$$

Proof of Theorem 4.18. The theorem concerns the probability distribution on $\{0,1\}^n$ given by independently selecting each $w_i \in \{0,1\}$ so that

$$\Pr[w_i = 0] = \frac{1+\epsilon}{2} = 1 - \Pr[w_i = 1].$$

when w is drawn with this distribution. For the string $w_1 \ldots w_n$ chosen with the probability distribution above, define the random variables

$$\Lambda_t = \lambda(w_1 \ldots w_t) \qquad \text{and} \qquad M_t = \mu(w_1 \ldots w_t).$$

Our goal is to establish that

$$\Pr[w \text{ forkable}] = \Pr[M_n \geq 0] = 2^{-\Omega(\sqrt{n})}.$$

We extract from the statement of Lemma 4.24 some facts about these random variables.

$$\Lambda_t > 0 \implies \begin{cases} \Lambda_{t+1} = \Lambda_t + 1 & \text{if } w_{t+1} = 1, \\ \Lambda_{t+1} = \Lambda_t - 1 & \text{if } w_{t+1} = 0; \end{cases} \tag{3}$$

$$M_t < 0 \implies \begin{cases} M_{t+1} = M_t + 1 & \text{if } w_{t+1} = 1, \\ M_{t+1} = M_t - 1 & \text{if } w_{t+1} = 0; \end{cases} \tag{4}$$

$$\Lambda_t = 0 \implies \begin{cases} \Lambda_{t+1} = 1 & \text{if } w_{t+1} = 1, \\ \Lambda_{t+1} = 0 & \text{if } w_{t+1} = 0, \\ M_{t+1} < 0 & \text{if } w_t = 0. \end{cases} \tag{5}$$

In light of the properties (13) above, the random variables Λ_t are quite well-behaved when positive—in particular, considering the distribution placed on each w_i, they simply follow the familiar biased random walk of Figure 6. Likewise, considering the properties (14), the random variables M_t follow a biased random walk when negative. The remainder of the proof combines these probability laws with (15) and the fact that $M_t() \leq \Lambda_t()$ to establish that $M_n < 0$ with high probability.

We recall two basic facts about the standard biased walk associated with the Markov chain of Figure 6. Let $Z_i \in \{\pm 1\}$ (for $i = 1, 2, \ldots$) denote a family of independent random variables for which $\Pr[Z_i = 1] = (1-\epsilon)/2$. Then the biased walk given by the variables $Y_t = \sum_i^t Z_i$ has the following properties.

Figure 4: The simple biased walk where $p = (1 + \epsilon)/2$ and $q = 1 - p$.

Constant escape probability; gambler's ruin. With constant probability, depending only on ϵ, $Y_t \neq 1$ for all $t > 0$. In general, for each $k > 0$,

$$\Pr[\exists t, Y_t = k] = \alpha^k \,, \tag{6}$$

for a constant $\alpha < 1$ depending only on ϵ. (In fact, the constant α is $(1 - \epsilon)/(1 + \epsilon)$; see, e.g., [20, Chapter 12] for a complete development.)

Concentration (the Chernoff bound). Consider T steps of the biased walk beginning at state 0; then the resulting value is tightly concentrated around $-\epsilon T$. Specifically, $\mathbb{E}[Y_T] = -\epsilon T$ and

$$\Pr\left[Y_T > -\frac{\epsilon T}{2}\right] = 2^{-\Omega(T)} \,. \tag{7}$$

(The constant hidden in the $\Omega()$ notation depends only on ϵ. See, e.g., [1, Cor. A.1.14].)

Partitioning the string w, we write $w = w^{(1)} \cdots w^{(\sqrt{n})}$ where $w^{(t)} = w_{1+a_{t-1}} \ldots w_{a_t}$ and $a_t = \lceil t\sqrt{n} \rceil$. Let $\Lambda_{(0)} = 0$ and $\Lambda_{(t)} = \Lambda_{a_t}$; similarly define $M_{(0)} = 0$ and $M_{(t)} = M_{a_t}$. Fix $\delta \ll \epsilon$ to be a small constant.

We define three events based on the random variables $\Lambda_{(t)}$ and $M_{(t)}$:

Hot We let Hot_t denote the event that $\Lambda_{(t)} \geq \delta\sqrt{n}$ and $M_{(t)} \geq -\delta\sqrt{n}$.

Volatile We let Vol_t denote the event that $-\delta\sqrt{n} \leq M_{(t)} \leq \Lambda_{(t)} < \delta\sqrt{n}$.

Cold We let Cold_t denote the event that $M_{(t)} < -\delta\sqrt{n}$.

Note that for each t, exactly one of these events occurs—they partition the probability space. Then we will establish that

$$\Pr[\mathsf{Cold}_{t+1} \mid \mathsf{Cold}_t] \geq 1 - 2^{-\Omega(\sqrt{n})} \,, \tag{8}$$

$$\Pr[\mathsf{Cold}_{t+1} \mid \mathsf{Vol}_t] \geq \Omega(\epsilon) \,, \tag{9}$$

$$\Pr[\mathsf{Hot}_{t+1} \mid \mathsf{Vol}_t] \leq 2^{-\Omega(\sqrt{n})} \,. \tag{10}$$

Note that the event Vol_0 occurs by definition. Assuming these inequalities, we observe that the system is very likely to eventually become cold, and stay that way. In this case, $\mathsf{Cold}_{\sqrt{n}}$ occurs, $M_n^\circ < \delta\sqrt{n} < 0$, and w is not forkable. Specifically, note that the probability that the system ever transitions from volatile to hot is no more than $2^{-\Omega(\sqrt{n})}$ (as transition from Vol to Hot is bounded above by $2^{-\Omega(\sqrt{n})}$, and there are no more than $\sqrt{n}$ possible transition opportunities). Note, also, that while the system is volatile, it transitions to cold with constant probability during each period. In particular, the probability that the system is volatile for the entire process is no more that $2^{-\Omega(\sqrt{n})}$. Finally, note that the probability that the system ever transitions out of the cold state is no more than $2^{-\Omega(\sqrt{n})}$ (again, there are at most $\sqrt{n}$ possible times when this could happen, and any individual transition occurs with probability $2^{-\Omega(\sqrt{n})}$). It follows that the system is cold at the end of the process with probability $1 - 2^{-\Omega(\sqrt{n})}$.

It remains to establish the three inequalities (18), (19), and (20).

20

Figure 5: An illustration of the transitions between Cold, Vol, and Hot.

Inequality (18)**:** This follows directly from (13) and (16). Specifically, in light of (14) the random variables M_i follow the probability law of the simple biased walk when they are negative. Conditioned on $M_{(t)} = M_{a_t} < -\delta\sqrt{n}$, the probability that any future M_s ever climbs to the value -1 is no more than $\alpha^{-\delta\sqrt{n}} = 2^{-\Omega(n)}$, as desired. (Here $\alpha < 1$ is a fixed constant that depends only on ϵ.)

Inequality (19)**:** This follows from (13), (15), (16), and (17). Specifically, conditioned on Vol_t, $\Lambda_{(t)} \leq \delta\sqrt{n}$. Recall from (13) that the random variables Λ_i follow the probability law of the simple biased walk when they are positive. Let D be the event that $\Lambda_i > 0$ for all $a_t \leq i < a_t + 2\delta\sqrt{n}$. According to (17), then, where we take $T = 2\delta\sqrt{n}$, $\Pr[D] \leq 2^{-\Omega(\sqrt{n})}$. With near certainty, then, the random variables Λ_i visit the value 0 during this period. Observe that if $\Lambda_i = 0$ then, by (15), $M_{i+1} \leq -1$ with constant probability and (conditioned on this), by (16), with constant probability the subsequent random variables M_j do not return to the value 0. Additionally, in light of (17), the probability that there is a sequence $w_i \ldots w_j$ of length at least $2(\delta/\epsilon)\sqrt{n}$ for which

$$\left(\sum_{k=i}^{j} \begin{cases} 1 & \text{if } w_k = 1, \\ -1 & \text{if } w_k = 0. \end{cases} \right) \geq -\delta\sqrt{n}$$

is no more than $(\sqrt{n})^n 2^{-\Omega(\sqrt{n})}$. It follows that with constant probability, the walk (of Λ_i) hits 0, as described above, and then M_i terminates at a value less than $-\delta\sqrt{n}$.

Inequality (20)**:** This follows from (13), (15), (16), and (17). Specifically, conditioned on Vol_t, $\Lambda_{(t)} \leq \delta\sqrt{n}$. Recall from (13) that the random variables Λ_i follow the probability law of the simple biased walk when they are positive. Let D be the event that $\Lambda_i > 0$ for all $a_t \leq i < a_t + 2\delta\sqrt{n}$. According to (17), then, where we take $T = 2\delta\sqrt{n}$, $\Pr[D] \leq 2^{-\Omega(\sqrt{n})}$. With near certainty, then, the random variables Λ_i visit the value 0 during this period. Conditioned on $\overline{D}$, in order for $\Lambda_{a_{t+1}} \geq \delta\sqrt{n}$ there must be a sequence of these random variables $0 = \Lambda_i, \Lambda_{i+1}, \ldots, \Lambda_j = \lfloor \delta\sqrt{n} \rfloor$ so that none of these take the value 0 except the first. (Such a sequence arises by taking i to be the last time the variables $\Lambda_{a_t}, \ldots$ visit 0 and j the first subsequent time that the sequence is larger than $\delta\sqrt{n}$.) In light of (16), the probability of such a subsequence appearing at a particular value for i is no more than $\alpha^{-\delta\sqrt{n}}$. It follows that the probability that $\Lambda_{a_{t+1}} \geq \delta\sqrt{n}$ is less than $\sqrt{n}\alpha^{-\delta\sqrt{n}} = 2^{-\Omega(\sqrt{n})}$, as desired. $\qquad\square$

Theorem 4.18. *Let $\epsilon \in (0, 1)$ and let w be a string drawn from $\{0, 1\}^n$ by independently assigning each $w_i = 1$ with probability $(1 - \epsilon)/2$. Then $\Pr[w \text{ is forkable}] = 2^{-\Omega(\sqrt{n})}$.*

Structural features of forks: closed forks, prefixes, reach, and margin. We begin by defining a natural notion of inclusion for two forks:

Definition 4.19 (Fork prefixes). *If w is a prefix of the string $w' \in \{0, 1\}^*$, $F \vdash w$, and $F' \vdash w'$, we say that F is a* prefix *of F', written $F \sqsubseteq F'$, if F is a consistently-labeled subgraph of F'. Specifically, every vertex and edge of F appears in F' and, furthermore, the labels given to any vertex appearing in both F and F' are identical.*

If $F \sqsubseteq F'$, each tine of F appears as the prefix of a tine in F'. In particular, the labels appearing on any tine terminating at a common vertex are identical and, moreover, the depth of any honest vertex appearing in both F and F' is identical.

In many cases, it is convenient to work with forks that do not "commit" anything beyond final honest indices.

Definition 4.20 (Closed forks). *A fork is* closed *if each leaf is honest. By convention the trivial fork, consisting solely of a root vertex, is closed.*

Note that a closed fork has a unique longest tine (as all maximal tines terminate with an honest vertex, and these must have distinct depths). Note, additionally, that if w is a prefix of w' and $F' \vdash w'$, then there is a unique closed fork $F \vdash w$ for which $F \sqsubseteq F'$.

Definition 4.21 (Gap, reserve and reach). *Let $F \vdash w$ be a closed fork and let $\hat{t}$ denote the (unique) tine of maximum length in F. We define the* gap *of a tine t, denoted $\mathrm{gap}(t)$, to be the difference in length between $\hat{t}$ and t; thus*

$$\mathrm{gap}(t) = \mathrm{length}(\hat{t}) - \mathrm{length}(t) \,.$$

We define the reserve *of a tine t to be the number of adversarial indices appearing in w after the last index in t; specifically, if t is given by the path $(r, v_1, \ldots, v_k)$, where r is the root of F, we define*

$$\mathrm{reserve}(t) = |\{i \mid w_i = 1 \text{ and } i > \ell(v_k)\}| \,.$$

We remark that this quantity depends both on F and the specific string w associated with F. Finally, for a tine t we define

$$\mathrm{reach}(t) = \mathrm{reserve}(t) - \mathrm{gap}(t) \,.$$

Definition 4.22 (Margin). *For closed fork $F \vdash w$ we define $\lambda(F)$ to be the maximum reach taken over all tines in F:*

$$\lambda(F) = \max_t \mathrm{reach}(t) \,.$$

Likewise, we define the margin *of F, denoted $\mu(F)$, to be the "penultimate" reach taken over edge-disjoint tines of F: specifically,*

$$\mathrm{margin}(F) = \mu(F) = \max_{t_1 \nsim t_2} \Big(\min\{\mathrm{reach}(t_1), \mathrm{reach}(t_2)\} \Big) \,. \tag{11}$$

We remark that the maxima above can always obtained by honest tines. Specifically, if t is an adversarial tine of a fork $F \vdash w$, $\mathrm{reach}(t) \leq \mathrm{reach}(\bar{t})$, where $\bar{t}$ is the longest honest prefix of t.

As $\sim$ is an equivalence relation on the nonempty tines. It follows that there is always a pair of (edge-disjoint) tines t_1 and t_2 achieving the maximum in the defining equation (11) which satisfy $\mathrm{reach}(t_1) = \lambda(F) \geq \mathrm{reach}(t_2) = \mu(F)$.

The relevance of margin to the notion of forkability is reflected in the following proposition.

Proposition 4.23. *A string w is forkable if and only if there is a closed fork $F \vdash w$ for which* $\mathrm{margin}(F) \geq 0$.

Proof. If w has no honest indices, then the trivial fork consisting of a single root node is flat, closed, and has non-negative margin; thus the two conditions are equivalent. Consider a forkable string w with at least one honest index and let $\hat{i}$ denote the largest honest index of w. Let F be a flat fork for w. As mentioned above, there is a unique closed fork $\overline{F} \vdash w$ obtained from F by removing any adversarial vertices from the ends of the tines of F. Note that the tine $\hat{t}$ containing $\hat{i}$ is the longest tine in $\overline{F}$, as this is the largest honest index of w. On the other hand, F is flat, in which case there are two edge-disjoint tines t_1 and t_2 with length at least that of $\hat{t}$. The prefixes of these two tines in $\overline{F}$ must clearly have reserve no less than gap (and hence non-negative reach); thus $\mathrm{margin}(\overline{F}) \geq 0$ as desired.

On the other hand, suppose w has a closed fork with $\mathrm{margin}(F) \geq 0$, in which case there are two edge-disjoint tines of F, t_1 and t_2, for which $\mathrm{reach}(t_i) \geq 0$. Then we can produce a flat fork by simply adding to each t_i a path of $\mathrm{gap}(t_i)$ vertices labeled with the subsequent adversarial indices promised by the definition of reserve(). $\qquad\square$

In light of this proposition, for a string w we focus our attention on the quantities

$$\lambda(w) = \max_{\substack{F \vdash w, \\ F \text{ closed}}} \lambda(F), \qquad \mu(w) = \max_{\substack{F \vdash w, \\ F \text{ closed}}} \mu(F),$$

and, for convenience,

$$\mathbf{m}(w) = (\lambda(w), \mu(w)).$$

Note that this overloads the notation $\lambda(\cdot)$ and $\mu(\cdot)$ so that they apply to both forks and strings, but the setting will be clear from context. We remark that the definitions do not guarantee a priori that $\lambda(w)$ and $\mu(w)$ can be achieved by the same fork, though this will be established in the lemma below. In any case, it is clear that $\lambda(w) \geq 0$ and $\lambda(w) \geq \mu(w)$ for all strings w; furthermore, by Proposition 4.23 a string w is forkable if and only if $\mu(w) \geq 0$. We refer to $\mu(w)$ as the *margin* of the string w.

In preparation for the proof of Theorem 4.18, we establish a recursive description for these quantities.

Lemma 4.24. $\mathbf{m}(\epsilon) = (0, 0)$ *and, for all nonempty strings $w \in \{0, 1\}^*$,*

$$\mathbf{m}(w1) = (\lambda(w) + 1, \mu(w) + 1), \text{ and}$$

$$\mathbf{m}(w0) = \begin{cases} (\lambda(w) - 1, 0) & \text{if } \lambda(w) > \mu(w) = 0, \\ (0, \mu(w) - 1) & \text{if } \lambda(w) = 0, \\ (\lambda(w) - 1, \mu(w) - 1) & \text{otherwise.} \end{cases}$$

Furthermore, for every string w, there is a closed fork $F_w \vdash w$ for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$.

Proof. The proof proceeds by induction. If $w = \epsilon$, define F_ϵ to be the trivial fork; $F_\epsilon \vdash w$ is the unique closed fork for this string and $\mathbf{m}(\epsilon) = (0, 0) = (\lambda(F_\epsilon), \mu(F_\epsilon))$, as desired.

In general, we consider $\mathbf{m}(w')$ for a string $w' = wx$—where $w \in \{0, 1\}^*$ and $x \in \{0, 1\}$; the argument recursively expands $\mathbf{m}(w')$ in terms of $\mathbf{m}(w)$ and the value of the last symbol x. In each case, we consider the relationship between two closed forks $F \sqsubseteq F'$ where $F \vdash w$ and $F' \vdash w' = wx$.

In the case where $x = 1$, we must have $F = F'$ as graphs, because the forks are assumed to be closed; it is easy to see that the reach of any tine t of $F \vdash w$ has increased by exactly one

when viewed as a tine of $F' \vdash w'$. We write $\mathrm{reach}_{F'}(t) = \mathrm{reach}_F(t) + 1$, where we introduce the notation $\mathrm{reach}_\square()$ to denote the reach in a particular fork. It follows that $\lambda(F') = \lambda(F) + 1$ and $\mu(F') = \mu(F) + 1$. If $F^* \vdash w'$ is a closed fork for which $\lambda(F^*) = \lambda(w')$, note that F^* may be treated as a fork for w and, applying the argument above, we find that $\lambda(w') \leq \lambda(w) + 1$. A similar argument implies that $\mu(w') \leq \mu(w) + 1$. On the other hand, by induction there is a fork F_w for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$ and hence $\mathbf{m}(w') \geq (\lambda(w) + 1, \mu(w) + 1)$. We conclude that

$$\mathbf{m}(w') = (\lambda(w) + 1, \mu(w) + 1). \tag{12}$$

Moreover, $\mathbf{m}(w') = (\lambda(F_w), \mu(F_w))$, where F_w is treated as a fork for $w' = w1$.

The case when $x = 0$ is more delicate. As above, we consider the relationship between two closed forks $F \vdash w$ and $F' \vdash w' = w0$ for which $F \sqsubseteq F'$. Here F' is necessarily obtained from F by appending a path labeled with a string of the form 1^a0 to the end of a tine t of F. (In fact, it is easy to see that we may always assume that this is appended to an honest tine.) In order for this to be possible, $\mathrm{gap}(t) \leq \mathrm{reserve}(t)$ (which is to say that $\mathrm{reach}(t) \geq 0$) and, in particular, $\mathrm{gap}(t) \leq a \leq \mathrm{reserve}(t)$: for the first inequality, note that the depth of the new honest vertex must exceed that of the deepest (honest) vertex in F and hence $a \geq \mathrm{gap}(t)$; as for the second inequality, there are only $\mathrm{reserve}(t)$ possible adversarial indices that may be added to t and hence $a \leq \mathrm{reserve}(t)$. We define the quantity $\tilde{a} \geq 0$ by the equation $a = \mathrm{gap}(t) + \tilde{a}$ and let t' denote the tine (of F') resulting by extending t in this way. We say that $\tilde{a}$ is the *parameter* for this pair of forks $F \sqsubseteq F'$.

Of course, every honest tine t of F is an honest tine of F' and it is clear that $\mathrm{reach}_{F'}(t) = \mathrm{reach}_F(t) - (\tilde{a} + 1)$, as the length of the longest tine t' in F' exceeds the length of the longest tine of F by exactly $\tilde{a} + 1$. Note that the reach of the new honest tine t' (in F') is always 0, as both $\mathrm{gap}(t')$ and $\mathrm{reserve}(t')$ are zero. It remains to describe how $\mu(w)$ and $\lambda(w)$ are determined by this process.

The case $\lambda(w) > \mu(w) = 0$. By induction, there is a fork F_w for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$. Let t_1 and t_2 be edge-disjoint tines of F_w for which $\lambda(F_w) = \mathrm{reach}(t_1)$ and $\mu(F_w) = \mathrm{reach}(t_2)$. Define $F' \vdash w'$ to be the fork obtained by extending the tine t_2 of F_w with parameter $\tilde{a} = 0$ to yield a new tine t_2' in F'. Then $\mathrm{reach}_{F'}(t_1) = \lambda(w) - 1$ and $\mathrm{reach}_{F'}(t_2') = 0$. It follows that $\lambda(w0) \geq \lambda(w) - 1$ and $\mu(w0) \geq 0$. We will show that $\lambda(w0) \leq \lambda(w) - 1$ and that $\mu(w0) \leq 0$, in which case we can conclude that

$$\lambda(w0) = \lambda(w) - 1 \qquad \text{and} \qquad \mu(w0) = 0.$$

Moreover, the fork $F_{w'} = F'$ achieves these statistics, as desired.

We return to establish that $\lambda(w0) \leq \lambda(w) - 1$ and that $\mu(w0) \leq 0$. Let $F^* \vdash w0$ be a closed fork for which $\lambda(w0) = \lambda(F^*)$ and let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$; as above, let $\tilde{a}$ denote the parameter for this extension. Let t^* be an honest tine of F^* so that $\mathrm{reach}_{F^*}(t^*) = \lambda(w0)$. If t^* is a tine of F, $\mathrm{reach}_{F^*}(t^*) = \mathrm{reach}_F(t^*) - (\tilde{a} + 1) \leq \lambda(w) - 1$. Otherwise t^* was obtained by extension and $\mathrm{reach}_{F^*}(t^*) = 0 \leq \lambda(w) - 1$ by assumption. In either case $\lambda(w0) \leq \lambda(w) - 1$, as desired. It remains to show that $\mu(w0) \leq 0$. Now consider $F^* \vdash w0$ to be a closed fork for which $\mu(F^*) = \mu(w0)$. Let t_1^* and t_2^* be two edge-disjoint honest tines of F^* so that $\mathrm{reach}_{F^*}(t_1^*) = \lambda(F^*)$ and $\mathrm{reach}_{F^*}(t_2^*) = \mu(F^*) = \mu(w0)$. Let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$ and let $\tilde{a}$ be the parameter for this extension. If both t_1^* and t_2^* are tines of F, $\mathrm{reach}_{F^*}(t_i^*) = \mathrm{reach}_F(t_i^*) - (\tilde{a} + 1)$ and, in particular, $\mathrm{reach}_F(t_1^*) \geq \mathrm{reach}_F(t_2^*)$. It follows that $\mathrm{reach}_F(t_2^*) \leq \mu(F) \leq \mu(w) = 0$ and hence that $\mu(w0) < 0$. Otherwise, one of the two tines was the result of extension and has zero $\mathrm{reach}()$

24

in F^*. As $\text{reach}_{F^*}(t_1^*) \geq \text{reach}_{F^*}(t_2^*)$, in either case it follows that $\mu(F^*) = \text{reach}_{F^*}(t_2^*) \leq 0$, as desired.

The case $\lambda(w) = 0$. By induction, there is a fork F_w for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$. Let t_1 and t_2 be edge-disjoint tines of F_w for which $\lambda(F_w) = \text{reach}(t_1)$ and $\mu(F_w) = \text{reach}(t_2)$. Define $F' \vdash w'$ to be the fork obtained by extending the tine t_1 of F_w with parameter $\tilde{a} = 0$ to yield a new tine t_1' in F'. Then $\text{reach}_{F'}(t_1') = 0$ and $\text{reach}_{F'}(t_2) = \text{reach}_F(t_2) - 1$. It follows that $\lambda(w0) \geq 0$ and $\mu(w0) \geq \mu(w) - 1$. We will show that $\lambda(w0) \leq 0$ and that $\mu(w0) \leq \mu(w) - 1$, in which case we can conclude that

$$\lambda(w0) = 0 \qquad \text{and} \qquad \mu(w0) = \mu(w) - 1 \,.$$

Moreover, the fork $F_{w'} = F'$ achieves these statistics, as desired.

We return to establish that $\lambda(w0) \leq 0$ and that $\mu(w0) \leq \mu(w) - 1$. Let $F^* \vdash w0$ be a closed fork for which $\lambda(w0) = \lambda(F^*)$ and let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$; as above, let $\tilde{a}$ denote the parameter for this extension. Let t^* be an honest tine of F^* so that $\text{reach}_{F^*}(t^*) = \lambda(w0)$. Note that t^* cannot be a tine of F; if it were then $\text{reach}_{F^*}(t^*) = \text{reach}_F(t^*) - (\tilde{a} + 1) \leq \lambda(w) - 1 < 0$ which contradicts $\lambda(w0) \geq 0$. Thus t^* was obtained by extension and $\text{reach}_{F^*}(t^*) = 0$. It remains to show that $\mu(w0) \leq 0$. Now let $F^* \vdash w0$ be a closed fork for which $\mu(F^*) = \mu(w0)$. Let t_1^* and t_2^* be two edge-disjoint honest tines of F^* so that $\text{reach}_{F^*}(t_1^*) = \lambda(F^*)$ and $\text{reach}_{F^*}(t_2^*) = \mu(F^*) = \mu(w0)$. Let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$ and let $\tilde{a}$ be the parameter for this extension. Similarly, t_1^* cannot be a tine of F; if it were, $\lambda(F^*) = \text{reach}_{F^*}(t_1^*) = \text{reach}_F(t_1^*) - (\tilde{a} + 1) \leq \lambda(F) - 1 \leq \lambda(w) - 1 < 0$ which contradicts $\lambda(F) \geq 0$. It follows that t_1^* must extend a tine t_1 of F for which $\text{reach}_F(t_1) = 0$, because extension can only occur for tines of non-negative reach and $\lambda(F) = 0 = \lambda(w)$. Thus t_2^* is a tine of F and $t_1 \nsim t_2^*$ so that $\text{reach}_F(t_2^*) \leq \mu(F) \leq \mu(w)$ and we conclude that $\mu(w0) = \text{reach}_{F^*}(t_2^*) \leq \text{reach}_F(t_2^*) - 1 \leq \mu(w) - 1$, as desired.

The case $\lambda(w) > 0, \mu(w) \neq 0$. By induction, there is a fork F_w for which $\mathbf{m}(w) = (\lambda(F_w), \mu(F_w))$. Let t_1 and t_2 be edge-disjoint tines of F_w for which $\lambda(F_w) = \text{reach}(t_1)$ and $\mu(F_w) = \text{reach}(t_2)$. In fact, any extension of F_w will suffice for the construction; for concreteness, define $F' \vdash w'$ to be the fork obtained by extending the tine t_1 of F_w with parameter $\tilde{a} = 0$. Then $\text{reach}_{F'}(t_i) = \text{reach}_{F_w}(t_i) - 1$. It follows that $\lambda(w0) \geq \lambda(w) - 1$ and $\mu(w0) \geq \mu(w) - 1$. We will show that $\lambda(w0) \leq \lambda(w) - 1$ and that $\mu(w0) \leq \mu(w) - 1$, in which case we can conclude that

$$\lambda(w0) = \lambda(w) - 1 \qquad \text{and} \qquad \mu(w0) = \mu(w) - 1 \,.$$

Moreover, the fork $F_{w'} = F'$ achieves these statistics, as desired.

We return to establish that $\lambda(w0) \leq \lambda(w) - 1$ and that $\mu(w0) \leq \mu(w) - 1$. Let $F^* \vdash w0$ be a closed fork for which $\lambda(w0) = \lambda(F^*)$ and let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$; as above, let $\tilde{a}$ denote the parameter for this extension. Let t^* be an honest tine of F^* so that $\text{reach}_{F^*}(t^*) = \lambda(w0)$. Note that if t^* is a tine of F then $\text{reach}_{F^*}(t^*) = \text{reach}_F(t^*) - (\tilde{a} + 1) \leq \lambda(w) - 1$; otherwise t^* is obtained by extension and $\text{reach}_{F^*}(t^*) = 0 \leq \lambda(w) - 1$, as desired. (Recall that $\lambda(w) > 0$.) It remains to show that $\mu(w0) \leq \mu(w) - 1$. Now let $F^* \vdash w0$ be a closed fork for which $\mu(F^*) = \mu(w0)$. Let t_1^* and t_2^* be two edge-disjoint honest tines of F^* so that $\text{reach}_{F^*}(t_1^*) = \lambda(F^*)$ and $\text{reach}_{F^*}(t_2^*) = \mu(F^*) = \mu(w0)$. Let $F \vdash w$ be the unique closed fork for which $F \sqsubseteq F^*$ and let $\tilde{a}$ be the parameter for this extension. If both t_1^* and t_2^* are tines of F then $\text{reach}_{F^*}(t_i^*) = \text{reach}_F(t_i^*) - (\tilde{a} + 1)$ and, in particular, $\text{reach}_F(t_1^*) \geq \text{reach}_F(t_2^*)$ so that $\text{reach}_F(t_2^*) \leq \mu(w)$ and $\text{reach}_{F^*}(t_2^*) \leq \mu(w) - 1$, as desired.

25

To complete the argument, we consider the case that one of the tines t_i^* arises by extension. Note that in this case $\text{reach}_{F^*}(t_2^*) \leq 0$, as either t_2^* is obtained by extension so that it has zero reach, or t_1^* is obtained by extension so that $\text{reach}_{F^*}(t_2^*) \leq \text{reach}_{F^*}(t_1^*) = 0$. Here we further separate the analysis into two cases depending on the sign of $\mu(w)$:

- If $\mu(w) > 0$, then $\text{reach}_{F^*}(t_2^*) \leq 0 \leq \mu(w) - 1$, as desired.
- If $\mu(w) < 0$ then t_2^* cannot be the extension of a tine in F. To see this, suppose to the contrary that t_2^* extends a tine t_2 of F; then $\text{reach}_F(t_2) \geq 0$. Additionally, t_1^* must be a tine of F, edge-disjoint from t_2, and $\text{reach}_F(t_1^*) = \text{reach}_{F^*}(t_1^*) + (\tilde{a} + 1) > 0$. It follows that $\mu(w) \geq \mu(F) \geq 0$, a contradiction.

  The other possibility is that t_1^* is an extension of a tine t_1 of F in which case $\text{reach}_F(t_1) \geq 0$. Note that t_2^* is a tine of F and edge-disjoint from t_1; thus $\min(\text{reach}_F(t_2^*), \text{reach}_F(t_1)) \leq \mu(F) < 0$ and $\text{reach}_F(t_2^*) \leq \mu(F)$. We conclude that $\text{reach}_{F^*}(t_2^*) = \text{reach}_F(t_2^*) - (\tilde{a}+1) \leq \mu(w) - 1$, as desired. $\qquad \square$

With this recursive description in place, we return to the proof of Theorem 4.18, which we restate below for convenience.

Theorem 4.18, restated. *Let $\epsilon \in (0,1)$ and let w be a string drawn from $\{0,1\}^n$ by independently assigning each $w_i = 1$ with probability $(1-\epsilon)/2$. Then*

$$\Pr[w \text{ is forkable}] = 2^{-\Omega(\sqrt{n})}.$$

Proof of Theorem 4.18. The theorem concerns the probability distribution on $\{0,1\}^n$ given by independently selecting each $w_i \in \{0,1\}$ so that

$$\Pr[w_i = 0] = \frac{1 + \epsilon}{2} = 1 - \Pr[w_i = 1].$$

when w is drawn with this distribution. For the string $w_1 \ldots w_n$ chosen with the probability distribution above, define the random variables

$$\Lambda_t = \lambda(w_1 \ldots w_t) \qquad \text{and} \qquad M_t = \mu(w_1 \ldots w_t).$$

Our goal is to establish that

$$\Pr[w \text{ forkable}] = \Pr[M_n \geq 0] = 2^{-\Omega(\sqrt{n})}.$$

We extract from the statement of Lemma 4.24 some facts about these random variables.

$$\Lambda_t > 0 \implies \begin{cases} \Lambda_{t+1} = \Lambda_t + 1 & \text{if } w_{t+1} = 1, \\ \Lambda_{t+1} = \Lambda_t - 1 & \text{if } w_{t+1} = 0; \end{cases} \tag{13}$$

$$M_t < 0 \implies \begin{cases} M_{t+1} = M_t + 1 & \text{if } w_{t+1} = 1, \\ M_{t+1} = M_t - 1 & \text{if } w_{t+1} = 0; \end{cases} \tag{14}$$

$$\Lambda_t = 0 \implies \begin{cases} \Lambda_{t+1} = 1 & \text{if } w_{t+1} = 1, \\ \Lambda_{t+1} = 0 & \text{if } w_{t+1} = 0, \\ M_{t+1} < 0 & \text{if } w_t = 0. \end{cases} \tag{15}$$

In light of the properties (13) above, the random variables Λ_t are quite well-behaved when positive—in particular, considering the distribution placed on each w_i, they simply follow the

26

Figure 6: The simple biased walk where $p = (1 + \epsilon)/2$ and $q = 1 - p$.

familiar biased random walk of Figure 6. Likewise, considering the properties (14), the random variables M_t follow a biased random walk when negative. The remainder of the proof combines these probability laws with (15) and the fact that $M_t() \leq \Lambda_t()$ to establish that $M_n < 0$ with high probability.

We recall two basic facts about the standard biased walk associated with the Markov chain of Figure 6. Let $Z_i \in \{\pm 1\}$ (for $i = 1, 2, \ldots$) denote a family of independent random variables for which $\Pr[Z_i = 1] = (1 - \epsilon)/2$. Then the biased walk given by the variables $Y_t = \sum_i^t Z_i$ has the following properties.

Constant escape probability; gambler's ruin. With constant probability, depending only on ϵ, $Y_t \neq 1$ for all $t > 0$. In general, for each $k > 0$,

$$\Pr[\exists t, Y_t = k] = \alpha^k \,, \tag{16}$$

for a constant $\alpha < 1$ depending only on ϵ. (In fact, the constant α is $(1 - \epsilon)/(1 + \epsilon)$; see, e.g., [20, Chapter 12] for a complete development.)

Concentration (the Chernoff bound). Consider T steps of the biased walk beginning at state 0; then the resulting value is tightly concentrated around $-\epsilon T$. Specifically, $\mathbb{E}[Y_T] = -\epsilon T$ and

$$\Pr\left[Y_T > -\frac{\epsilon T}{2}\right] = 2^{-\Omega(T)} \,. \tag{17}$$

(The constant hidden in the $\Omega()$ notation depends only on ϵ. See, e.g., [1, Cor. A.1.14].)

Partitioning the string w, we write $w = w^{(1)} \cdots w^{(\sqrt{n})}$ where $w^{(t)} = w_{1+a_{t-1}} \ldots w_{a_t}$ and $a_t = \lceil t\sqrt{n} \rceil$. Let $\Lambda_{(0)} = 0$ and $\Lambda_{(t)} = \Lambda_{a_t}$; similarly define $M_{(0)} = 0$ and $M_{(t)} = M_{a_t}$. Fix $\delta \ll \epsilon$ to be a small constant.

We define three events based on the random variables $\Lambda_{(t)}$ and $M_{(t)}$:

Hot We let Hot_t denote the event that $\Lambda_{(t)} \geq \delta\sqrt{n}$ and $M_{(t)} \geq -\delta\sqrt{n}$.

Volatile We let Vol_t denote the event that $-\delta\sqrt{n} \leq M_{(t)} \leq \Lambda_{(t)} < \delta\sqrt{n}$.

Cold We let Cold_t denote the event that $M_{(t)} < -\delta\sqrt{n}$.

Note that for each t, exactly one of these events occurs—they partition the probability space. Then we will establish that

$$\Pr[\mathsf{Cold}_{t+1} \mid \mathsf{Cold}_t] \geq 1 - 2^{-\Omega(\sqrt{n})} \,, \tag{18}$$

$$\Pr[\mathsf{Cold}_{t+1} \mid \mathsf{Vol}_t] \geq \Omega(\epsilon) \,, \tag{19}$$

$$\Pr[\mathsf{Hot}_{t+1} \mid \mathsf{Vol}_t] \leq 2^{-\Omega(\sqrt{n})} \,. \tag{20}$$

Note that the event Vol_0 occurs by definition. Assuming these inequalities, we observe that the system is very likely to eventually become cold, and stay that way. In this case, $\mathsf{Cold}_{\sqrt{n}}$ occurs,

Figure 7: An illustration of the transitions between Cold, Vol, and Hot.

$M_n^\circ < \delta\sqrt{n} < 0$, and w is not forkable. Specifically, note that the probability that the system ever transitions from volatile to hot is no more than $2^{-\Omega(\sqrt{n})}$ (as transition from Vol to Hot is bounded above by $2^{-\Omega(\sqrt{n})}$, and there are no more than $\sqrt{n}$ possible transition opportunities). Note, also, that while the system is volatile, it transitions to cold with constant probability during each period. In particular, the probability that the system is volatile for the entire process is no more that $2^{-\Omega(\sqrt{n})}$. Finally, note that the probability that the system ever transitions out of the cold state is no more than $2^{-\Omega(\sqrt{n})}$ (again, there are at most $\sqrt{n}$ possible times when this could happen, and any individual transition occurs with probability $2^{-\Omega(\sqrt{n})}$). It follows that the system is cold at the end of the process with probability $1 - 2^{-\Omega(\sqrt{n})}$.

It remains to establish the three inequalities (18), (19), and (20).

Inequality (18)**:** This follows directly from (13) and (16). Specifically, in light of (14) the random variables M_i follow the probability law of the simple biased walk when they are negative. Conditioned on $M_{(t)} = M_{a_t} < -\delta\sqrt{n}$, the probability that any future M_s ever climbs to the value -1 is no more than $\alpha^{-\delta\sqrt{n}} = 2^{-\Omega(n)}$, as desired. (Here $\alpha < 1$ is a fixed constant that depends only on ϵ.)

Inequality (19)**:** This follows from (13), (15), (16), and (17). Specifically, conditioned on Vol_t, $\Lambda_{(t)} \le \delta\sqrt{n}$. Recall from (13) that the random variables Λ_i follow the probability law of the simple biased walk when they are positive. Let D be the event that $\Lambda_i > 0$ for all $a_t \le i < a_t + 2\delta\sqrt{n}$. According to (17), then, where we take $T = 2\delta\sqrt{n}$, $\Pr[D] \le 2^{-\Omega(\sqrt{n})}$. With near certainty, then, the random variables Λ_i visit the value 0 during this period. Observe that if $\Lambda_i = 0$ then, by (15), $M_{i+1} \le -1$ with constant probability and (conditioned on this), by (16), with constant probability the subsequent random variables M_j do not return to the value 0. Additionally, in light of (17), the probability that there is a sequence $w_i \dots w_j$ of length at least $2(\delta/\epsilon)\sqrt{n}$ for which

$$\left(\sum_{k=i}^{j} \begin{cases} 1 & \text{if } w_k = 1, \\ -1 & \text{if } w_k = 0. \end{cases} \right) \ge -\delta\sqrt{n}$$

is no more than $(\sqrt{n})^n 2^{-\Omega(\sqrt{n})}$. It follows that with constant probability, the walk (of Λ_i) hits 0, as described above, and then M_i terminates at a value less than $-\delta\sqrt{n}$.

28

Inequality (20)**:** This follows from (13), (15), (16), and (17). Specifically, conditioned on Vol_t, $\Lambda_{(t)} \leq \delta\sqrt{n}$. Recall from (13) that the random variables Λ_i follow the probability law of the simple biased walk when they are positive. Let D be the event that $\Lambda_i > 0$ for all $a_t \leq i < a_t + 2\delta\sqrt{n}$. According to (17), then, where we take $T = 2\delta\sqrt{n}$, $\Pr[D] \leq 2^{-\Omega(\sqrt{n})}$. With near certainty, then, the random variables Λ_i visit the value 0 during this period. Conditioned on $\overline{D}$, in order for $\Lambda_{a_{t+1}} \geq \delta\sqrt{n}$ there must be a sequence of these random variables $0 = \Lambda_i, \Lambda_{i+1}, \dots, \Lambda_j = \lfloor\delta\sqrt{n}\rfloor$ so that none of these take the value 0 except the first. (Such a sequence arises by taking i to be the last time the variables $\Lambda_{a_t}, \dots$ visit 0 and j the first subsequent time that the sequence is larger than $\delta\sqrt{n}$.) In light of (16), the probability of such a subsequence appearing at a particular value for i is no more than $\alpha^{-\delta\sqrt{n}}$. It follows that the probability that $\Lambda_{a_{t+1}} \geq \delta\sqrt{n}$ is less than $\sqrt{n}\alpha^{-\delta\sqrt{n}} = 2^{-\Omega(\sqrt{n})}$, as desired. $\qquad\square$

4.2.1 Covert adversaries, covert forks, and covertly forkable strings

The general notion of fork defined in Definition 4.9 above reflects the possibility that adversarial slot leaders may broadcast multiple blocks for a single slot; such adversaries may simultaneously extend many different chains. While this provides the adversary significant opportunities to interfere with the protocol, it leaves a suspicious "audit trail"—multiple signed blocks for the same slot—which conspicuously deviates from the protocol.

This motivates our consideration of a restricted class of *covert* adversaries, who broadcast no more than one block per slot (though not necessarily in the expected slot). Such an adversary may still deviate from the protocol by extending short chains, but do not produce such suspicious evidence and hence its strategy is more "deniable": it can blame network delays for its actions.[4]

Such an adversary yields a restricted notion of fork, defined below:

Definition 4.25. *Let $F \vdash w$ be a fork for a string $w \in \{0,1\}^*$. We say that F is* covert *if the labeling $\ell : V \to \{0, 1, \dots, \}$ is injective. In particular, no adversarial index is labeled by more than one node.*

As in the general case, we define a notion of forkable string for such adversaries.

Definition 4.26. *We say that a string w is* covertly forkable *if there is a flat covert fork $F \vdash w$.*

Covert adversaries and forks have much simpler structure than general adversaries. In particular, a string is covertly forkable if and only if a majority of its indices are adversarial. This provides an analogue of Proposition 4.23 for covertly forkable strings.

Proposition 4.27. *A string $w \in \{0,1\}^n$ is covertly forkable if and only if $\mathrm{wt}(w) \geq n/2$.*

Proof. Let w be a covertly forkable string and $F \vdash w$ a flat covert fork. As F is flat, there are two edge disjoint tines, t_1 and t_2, with length equal to $\mathrm{height}(F)$ and it follows that the number of vertices in F is at least $2 \cdot \mathrm{height}(F) + 1$. In this covert case the labeling function is injective, and it follows that $n \geq 2 \cdot \mathrm{height}(F)$. (Recall that the root vertex is labeled by 0, which is not an index into w.) On the other hand, the height of F is at least the number of honest indices of w. We conclude that the length of w is at least twice the number of honest indices, as desired.

If $\mathrm{wt}(w) \geq n/2$, we can produce a flat covert fork $F \vdash w$ by placing all honest indices on a common tine t_1 and selecting $\mathrm{length}(t_1)$ adversarial indices to form an edge-disjoint second tine t_2. $\qquad\square$

---

[4]Contrast this with a more general adversary that attempts to fork by signing two different blocks for the same slot; such an adversary cannot merely blame the network for such a deviation.

As the structure of covertly forkable strings is so simple, an analogue of Theorem 4.18 for the density of covertly forkable strings follows directly from standard large deviation bounds.

Theorem 4.28. *Let $\epsilon \in (0,1)$ and let w be a string drawn from $\{0,1\}^n$ by independently assigning each $w_i = 1$ with probability $(1 - \epsilon)/2$. Then*

$$\Pr[w \text{ is covertly forkable}] = 2^{-\Theta(n)}.$$

Proof. This follows from standard estimates for the cumulative density function of the binomial distribution. □

Experiments In order to gain further insight regarding the density of forkable strings, we exactly computed the probability that a string w drawn from the binomial distribution with parameter $p \in \{.40, .41, \ldots, .50\}$ is forkable for several different lengths. These results are presented in Figure 9. For comparison, we likewise computed the probability that a string drawn from the binomial distribution is covertly forkable. These results are presented in Figure 8. (Note that these second probabilities are simply appropriate evaluations of the cumulative density function of the binomial distribution.)


Figure 8: Graphs of the probability that a string drawn from the binomial distribution is covertly forkable. Graphs for string lengths $n = 500, 1000, 1500, 2000$ are shown with parameters $.40, .41, \ldots, .49, .50$.

4.3 Common Prefix

Recall that the chains constructed by honest players during an execution of π_{SPoS} correspond to tines of a fork, as defined and studied in the previous sections. The random assignment of slots to stakeholders given by $\mathcal{F}_{\mathrm{LS}}^{\mathcal{D},\mathsf{F}}$ guarantees that the coordinates of the associated characteristic string w follow the binomial distribution with probability equal to the adversarial stake. Thus Theorem 4.18 establishes that no execution of the protocol π_{SPoS} can induce two tines (chains) of maximal length with no common prefix.

Figure 9: Graphs of the probability that a string drawn from the binomial distribution is forkable. Graphs for string lengths $n = 500, 1000, 1500, 2000$ are shown with parameters $.40, .41, \ldots, .49, .50$.

In the context of π_{SPoS}, however, we wish to establish a much stronger common prefix property: The chains reported by any two honest players must have a "recent" common prefix, in the sense that removing a small number of blocks from the shorter chain results in a prefix of the longer chain.

To formally articulate and prove this property, we introduce some further definitions regarding tines and forks. We borrow the "truncation operator," described earlier in the paper for chains: for a tine t we let $t^{\lceil k}$ denote the tine obtained by removing the last k edges; if $\mathrm{length}(t) \leq k$, we define $t^{\lceil k}$ to consist solely of the root. Finally, let F be a fork for a string $w \in \{0,1\}^*$. For two honest tines t_1 and t_2 of F, define their *divergence* to be the quantity

$$\mathrm{div}(t_1, t_2) = \min_i (\mathrm{length}(t_i) - \mathrm{length}(t_1 \cap t_2)),$$

where $t_1 \cap t_2$ denotes the common prefix of t_1 and t_2. (Observe that if $\mathrm{div}(t_1, t_2) \leq k$ and, say $\mathrm{length}(t_1) \leq \mathrm{length}(t_2)$, the tine $t_1^{\lceil k}$ is a prefix of t_2.) Then define the divergence of w to be the maximum such divergence over all pairs of honest tines over all possible forks for w:

$$\mathrm{div}(w) = \max_{F \vdash w} \max_{\substack{t_1, t_2 \\ \text{honest} \\ \text{tines of } F}} \mathrm{div}(t_1, t_2).$$

We first establish that a string with large divergence must have a large forkable substring. We then apply this in Theorem 4.30 below to conclude that characteristic strings arising from π_{SPoS} are unlikely to have large divergence and, hence, possess the common prefix property.

Theorem 4.29. *Let $w \in \{0,1\}^*$ with $\mathrm{div}(w) \geq k$. Then there is a forkable substring $\breve{w}$ of w with $|\breve{w}| \geq k$.*

Proof. Consider a fork $F \vdash w$ and a pair of honest tines (t_1, t_2) so that $\mathrm{div}(t_1, t_2) = \mathrm{div}(w)$; we assume the tines are identified so that $\mathrm{length}(t_1) < \mathrm{length}(t_2)$. Let z_i denote the last vertex on the tine t_i and let y denote the last vertex on the tine $t_1 \cap t_2$ as in the diagram below.

31

Without loss of generality we may assume that $\ell(z_2)$, the honest index labeling z_2, is in fact the first honest index in w appearing after $\ell(z_1)$. To justify this, let β denote this first honest index of w after $\ell(z_1)$ and let x denote the unique vertex of F for which $\ell(x) = \beta$. Note that $\mathbf{d}(x) > \mathbf{d}(z_1)$. If the tine t ending at x shares an edge with t_1 after y, then t is disjoint from t_2 after y and it follows that $\mathrm{div}(t, t_2) > \mathrm{div}(t_1, t_2)$, a contradiction. Thus t shares no edges with t_1 after y, and $\mathrm{length}(t) > \mathrm{length}(t_1)$; it follows that $\mathrm{div}(t_1, t) = \mathrm{div}(t_1, t_2)$ and we may assume $t_2 = t$ in the remainder of the argument.

Let $\alpha = \ell(y)$ denote the label of y and, as indicated above, let $\beta = \ell(z_2)$. We wish to show that the string $\check{w} = w_{\alpha+1} \ldots w_{\beta-1}$ is forkable. Our strategy will be to construct a flat fork for $\check{w}$ by showing that—after applying some minor restructuring to F—treating the vertex y as the root for the portion of F labeled with indices of $\check{w}$ yields a flat fork.

Observe, that the vertex y cannot be adversarial: otherwise it is easy to construct a fork $\tilde{F} \vdash w$ and a pair of tines that achieve larger divergence. Specifically, construct $\tilde{F}$ from F by adding a new (adversarial) vertex $\tilde{y}$ to F for which $\ell(\tilde{y}) = \ell(y)$, adding an edge to $\tilde{y}$ from the vertex preceding y, and replacing the edge of t_1 following y with one from $\tilde{y}$; then the other relevant properties of the fork are maintained, but the divergence of the resulting tines has increased by one. (See the diagram below.)

A similar argument implies that the fork $F_0 \vdash w_1 \ldots w_\alpha$, obtained by including only the vertices with labels less than or equal to $\alpha = \ell(y)$, does not contain two distinct vertices of depth $\mathbf{d}(y)$. In the presence of another vertex $\tilde{y}$ (of F_0) with depth $\mathbf{d}(y)$, "redirecting" t_1 through $\tilde{y}$ (as in the argument above) would likewise result in a fork with larger divergence. As α is the last index of the string, this additionally implies that F_0 has no vertices of depth exceeding $\mathbf{d}(y)$.

In light of the remarks above, we observe that the fork F may be "pinched" at y so that any tine t of length exceeding $\mathbf{d}(y)$ is transformed to a new tine $t^{\triangleright y \triangleleft}$ that passes through the vertex y. Specifically, for a tine t of F we define $t^{\triangleright y \triangleleft}$ as follows:

- If $\mathrm{length}(t) \leq \mathbf{d}(y)$, define $t^{\triangleright y \triangleleft} = t$.

- Otherwise, $\mathrm{length}(t) > \mathbf{d}(y)$ and t has a vertex z of depth $\mathbf{d}(y) + 1$; note that $\ell(z) > \ell(y)$ because F_0 contains no vertices of depth exceeding $\mathbf{d}(y)$. Then $t^{\triangleright y \triangleleft}$ is defined to be the path given by the tine terminating at y, a (new) edge from y to z, and the suffix of t beginning at z. (As $\ell(z) > \ell(y)$ this has the increasing label property.)

Then define $F^{\triangleright y \triangleleft} \vdash w$ to be the fork given by the union of $t^{\triangleright y \triangleleft}$ over all tines t of F. To be precise: $F^{\triangleright y \triangleleft}$ is defined with the same set of vertices and labels as F; the edges of $F^{\triangleright y \triangleleft}$ are the union of those appearing in $t^{\triangleright y \triangleleft}$ for every tine t of F. (Alternatively, $F^{\triangleright y \triangleleft}$ can be described as the graph obtained from F by changing every edge of F directed towards a vertex of depth $\mathbf{d}(y) + 1$ so that

it originates from y.) Note that the graph $F^{\triangleright y \triangleleft}$ defined in this way is a legal fork and that the depths of vertices in F and $F^{\triangleright y \triangleleft}$ are identical.

By excising the tree rooted at y from this pinched fork $F^{\triangleright y \triangleleft}$ we may extract a fork for the string $w_{\alpha+1} \ldots w_n$. Specifically, consider the induced subgraph $F^{y \triangleleft}$ of $F^{\triangleright y \triangleleft}$ given by the vertices $\{y\} \cup \{z \mid \mathbf{d}(z) > \mathbf{d}(y)\}$. By treating y as a root vertex and suitably (re-)defining the labels $\ell^{y \triangleleft}$ of $F^{y \triangleleft}$ so that $\ell^{y \triangleleft}(z) = \ell(z) - \ell(y)$, this subgraph has the defining properties of a fork for $w_{\alpha+1} \ldots w_n$. Note, also, that the suffixes of the two tines t_1 and t_2 (beginning at y) appear in $F^{y \triangleleft}$ and share no edges in this fork.

Finally, let $\check{F}$ denote the fork obtained from $F^{y \triangleleft}$ by retaining only those vertices associated with the prefix $\check{w} = w_{\alpha+1} \ldots w_{\beta-1}$ which furthermore have depth no more than $\mathbf{d}(z_1)$. Note that z_1 is the deepest honest vertex in this fork (as $\ell(z_1)$ is the last honest index of $\check{w}$) so that every honest index indeed appears in $\check{F}$: hence $\check{F} \vdash \check{w}$. Observe, also, that the suffix of t_2 appearing in $F^{y \triangleleft}$ is entirely labeled by indices of $\check{w}$ and hence is trimmed by this process to have resulting length exactly $\mathbf{d}(z_1)$. Thus $\check{F}$ is a flat fork, as desired.

Note, finally, that $|\check{w}|$ is at least the length of any particular tine in $\check{F}$; thus $|\check{w}| \geq \text{length}(t_1) - \mathbf{d}(y) \geq k$. $\qquad\square$

Theorem 4.30. *Let $k, R \in \mathbb{N}$ and $\epsilon \in (0, 1)$. The probability that the π_{SPoS} protocol, when executed with a $(1 - \epsilon)/2$ fraction of adversarial stake, violates the common prefix property with parameter k throughout a period of R is no more than $\exp(-\Omega(\sqrt{k}) + \ln R)$; the constant hidden by the $\Omega()$ notation depends only on ϵ.*

Proof. The characteristic string $w \in \{0, 1\}^R$ for such an execution of π_{SPoS} is determined by assigning each $w_i = 1$ independently with probability $(1 - \epsilon)/2$.

Observe that an execution of π_{SPoS} violates the common prefix property with parameters k, R precisely when the fork F induced by this execution has $\text{div}(F) \geq k$. Thus we wish to show that the probability that $\text{div}(w) \geq k$ is no more than $\exp(-\Omega(\sqrt{k}) + \log R)$. Let Bad denote the event that $\text{div}(w) \geq k$.

It follows from Theorem 4.29 that if $\text{div}(w) \geq k$, there is a forkable substring $\check{w}$ of length at least k. Thus

$$\Pr[\mathsf{Bad}] \leq \Pr\left[\exists \alpha, \beta \in \{1, \ldots, R\} \text{ so that } \alpha + k - 1 \leq \beta \text{ and } w_\alpha \ldots w_\beta \text{ is forkable}\right]$$

$$\leq \sum_{1 \leq \alpha \leq R} \underbrace{\sum_{\alpha + k - 1 \leq \beta \leq R} \Pr[w_\alpha \ldots w_\beta \text{ is forkable}]}_{(*)} .$$

According to Theorem 4.18 the probability that a string of length t drawn from this distribution is forkable is no more than $\exp(-c\sqrt{t})$ for a positive constant c. Note that for any $\alpha \geq 1$,

$$\sum_{t=\alpha+k-1}^{R} e^{-c\sqrt{t}} \leq \int_{k-1}^{\infty} e^{-c\sqrt{t}} \, dt = (2/c^2)(1 + c\sqrt{k-1})e^{-c\sqrt{k-1}} = e^{-\Omega(\sqrt{k})}$$

and it follows that the sum $(*)$ above is $\exp(-\Omega(\sqrt{t}))$. Thus

$$\Pr[\mathsf{Bad}] \leq R \cdot \exp(-\Omega(\sqrt{k})) \leq \exp(\ln R - \Omega(\sqrt{k})),$$

as desired. $\qquad\square$

4.3.1 Common prefix with covert adversaries

We revisit the notion of common prefix in the setting of covert adversaries. We define the *covert divergence* of w to be the maximum divergence over all pairs of honest tines over all possible covert forks for w:

$$\text{cdiv}(w) = \max_{\substack{F \vdash w \\ F \text{ covert}}} \max_{\substack{t_1, t_2 \\ \text{honest} \\ \text{tines of } F}} \text{div}(t_1, t_2).$$

As in the setting with general adversaries, we wish to establish that a string with large covert divergence must have a large covertly forkable substring. A direct analogue of Theorem 4.30 then implies that characteristic strings arising from π_{SPoS} are unlikely to have large covert divergence and, hence, possess the common prefix property against covert adversaries.

We record an analogue of Theorem 4.29 for covert adversaries.

Theorem 4.31. *Let $w \in \{0,1\}^*$ with $\text{cdiv}(w) \geq k$. Then there is a covertly forkable substring $\widecheck{w}$ of w with $|\widecheck{w}| \geq k$.*

Proof. Consider a covert fork $F \vdash w$ and a pair of honest tines (t_1, t_2) so that $\text{div}(t_1, t_2) = \text{cdiv}(w)$; we assume the tines are identified so that $\text{length}(t_1) < \text{length}(t_2)$. Let z_i denote the last vertex on the tine t_i and let y denote the last vertex on the tine $t_1 \cap t_2$. As in the proof for the general case, we may assume without loss of generality that $\ell(z_2)$ is the first honest index in w appearing after $\ell(z_1)$. Define $\alpha = \ell(y)$ and let $\beta = \ell(z_2)$.

As in the proof for the general case, we will conclude that the string $\widecheck{w} = w_{\alpha+1} \ldots w_{\beta-1}$ is covertly forkable. Note that $\widecheck{w}$ must have length at least $\text{cdiv}(w) = \text{length}(t_1) - \text{length}(t_1 \cap t_2)$, as every vertex of t_1 after y labels an index of this string. Note further that $\widecheck{w}$ has no more than $\text{cdiv}(w)$ honest vertices as $\ell(t_1)$ is the index of the last honest vertex of $\widecheck{w}$. As in the proof of Proposition 4.27, the two tines t_1 and t_2 account for at least $2 \cdot \text{cdiv}(w)$ vertices labeled with indices of $\widecheck{w}$; as ℓ is invective, it follows that at least half the indices of $\widecheck{w}$ are adversarial, as desired. □

Finally, we remark that the proof of Theorem 4.30 applies with minor adaptations to the covert case.

Theorem 4.32. *Let $k, R \in \mathbb{N}$ and $\epsilon \in (0, 1)$. The probability that the π_{SPoS} protocol, when executed with a $(1 - \epsilon)/2$ fraction of adversarial stake and a covert adversary, violates the common prefix property with parameter k throughout a period of R is no more than $\exp(-\Omega(k) + \ln R)$; the constant hidden by the $\Omega()$ notation depends only on ϵ.*

Proof. The proof of Theorem 4.30 applies directly; in this case the asymptotics rely on Theorem 4.28 and the following bound applied in a way that the constant c depends only on ϵ.

$$\sum_{t=k}^{\infty} e^{-ct} \leq \int_{k-1}^{\infty} e^{-ct} \, dt = e^{-\Theta(k)}. \quad \square$$

□

4.4 Chain Growth and Chain Quality

We will start with the chain growth property.

Theorem 4.33. *The π_{SPoS} protocol satisfies the chain growth property with parameters $\tau = 1 - \alpha, s \in \mathbb{N}$ throughout an epoch of R slots with probability at least $1 - \exp(-\epsilon^2 s + \ln R)$ against an adversary holding an $\alpha - \epsilon$ portion of the total stake.*

Proof. Define $\mathsf{Ham}_a(\alpha)$ to be the event that the Hamming weight ratio of the characteristic string that corresponds to the slots $[a, a + s - 1]$ is up to α. Given that the adversarial stake is $\alpha - \epsilon$, each of the k slots has probability $\alpha - \epsilon$ being assigned to the adversary and thus the probability that the Hamming weight is more than αs drops exponentially in s. Specifically, using the additive version of the Chernoff bound, we have that $\mathbf{Pr}[\neg\mathsf{Ham}_a(\alpha)] \leq \exp(-2\epsilon^2 s)$. It follows that,

$$\mathbf{Pr}[\mathsf{Ham}_\alpha] \geq 1 - \exp(-2\epsilon^2 s).$$

Given the above we know that when Ham_α happens there will be at least $(1 - \alpha)s$ honest slots in the period of s rounds. Given that each honest slot enables an honest party to produce a block, all honest parties will advance by at least that many blocks. Using a union bound, it follows that the speed coefficient can be set to $\tau = (1 - \alpha)$ and it is satisfied with probability at least $1 - \exp(-2\epsilon^2 s + \ln(R))$. $\qquad\square$

Having established the chain growth property we now turn our attention to the chain quality property. Recall that the chain-quality property parameterized with k and it states that every k blocks in a chain observed at a certain slot the blocks corresponding to a set of stakeholders that hold cumulative stake ratio β are $\tau\beta$. In the next theorem we establish bounds for the parameter τ.

Theorem 4.34. *The π_{SPoS} protocol satisfies the chain quality property with parameters $\mu = \alpha/(1 - \alpha)$, $k \in \mathbb{N}$ throughout an epoch of R slots with probability at least*

$$1 - \exp\big(-\epsilon^2(1 - \alpha)^{-1}k + \ln R\big)$$

where $\alpha - \epsilon$ is the ratio of the cumulative stake of the set of malicious stakeholders.

Proof. First, using a similar argumentation as in the chain growth Theorem 4.33, we know that in a segment of s rounds the honest parties would advance by at least $(1 - \alpha)s$ blocks. Furthermore the adversary can produce at most αs blocks in the same period. It follows that in the chain of any honest party one would find at most $\alpha/(1 - \alpha)$ ratio of blocks originating from the adversary with probability $1 - \exp(-\epsilon^2 s + \ln R)$ among the blocks produced in the period that corresponds to that segment. It suffices to choose $s \geq (1 - \alpha)^{-1}k$. In this case we know that there will be at least k blocks produced in any segment of s rounds. $\qquad\square$

5 Our Protocol: Dynamic Stake

5.1 Using a Trusted Beacon

In the static version of the protocol in the previous section, we assumed that stake was static during the whole execution (i.e., one epoch), meaning that stake changing hands inside a given epoch does not affect leader election. Now we put forth a modification of protocol π_{SPoS} that can be executed over multiple epochs in such a way that each epoch's leader election process is parameterized by the stake distribution at a certain designated point of the previous epoch, allowing for change in the stake distribution across epochs to affect the leader election process. As before, we construct the protocol in a hybrid model, enhancing the $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ ideal functionality to now provide randomness and auxiliary information for the leader election process throughout the epochs (the enhanced functionality will be called $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$). We then discuss how to implement $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ using only $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ and in this way reduce the assumption back to the simple common random string selected at setup.

Before describing the protocol for the case of dynamic stake, we need to explain the modification of $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ so that multiple epochs are considered. The resulting functionality, $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$, allows stakeholders to query it for the leader selection data specific to each epoch. $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ is parameterized by the initial stake of each stakeholder before the first epoch e_1 starts; in subsequent epochs, parties will take into consideration the stake distribution after the previous epoch's first $R - k$ slots, where k is the number of slots needed to achieve common prefix. Notice that it is necessary to consider the stake distribution of previous epochs only in the slots where it is guaranteed that common prefix is achieved. In any case, the functionality $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ provides only a random string and leaves the interpretation according to the stakeholder distribution to the party that is calling it. The effective stakeholder distribution is the sequence $\mathbb{S}_1, \mathbb{S}_2, \ldots$ defined as follows: $\mathbb{S}_1$ is the initial stakeholder distribution; for slots $\{(j-1)R + 1, \ldots, jR\}$ for $j \geq 2$ the effective stakeholder $\mathbb{S}_j$ is determined by the stake allocation that is found in the chain of all honest parties at slot $(j-1)R - k$, or is undefined if the honest parties disagree on it. The functionality $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ is defined in Figure 10.
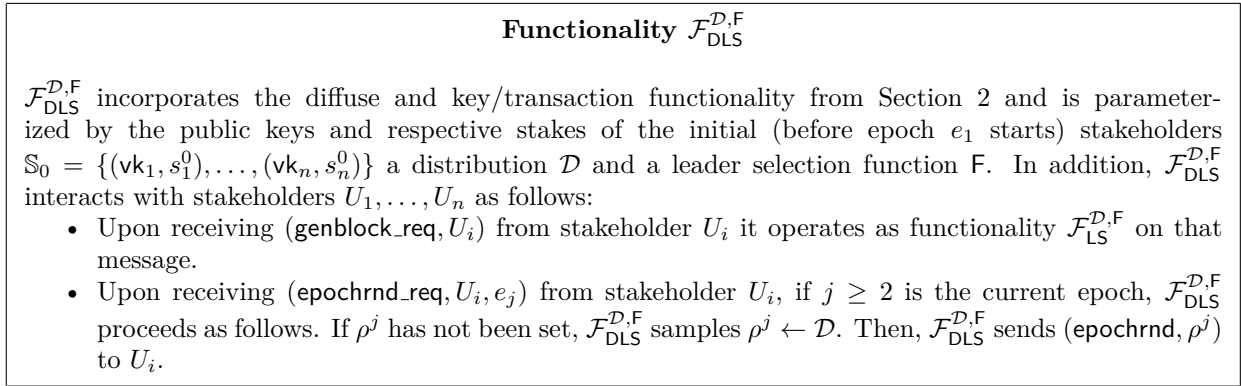

Functionality $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$

$\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ incorporates the diffuse and key/transaction functionality from Section 2 and is parameterized by the public keys and respective stakes of the initial (before epoch e_1 starts) stakeholders $\mathbb{S}_0 = \{(\mathsf{vk}_1, s_1^0), \ldots, (\mathsf{vk}_n, s_n^0)\}$ a distribution $\mathcal{D}$ and a leader selection function F. In addition, $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ interacts with stakeholders $U_1, \ldots, U_n$ as follows:

- Upon receiving $(\mathsf{genblock\_req}, U_i)$ from stakeholder U_i it operates as functionality $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ on that message.
- Upon receiving $(\mathsf{epochrnd\_req}, U_i, e_j)$ from stakeholder U_i, if $j \geq 2$ is the current epoch, $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ proceeds as follows. If ρ^j has not been set, $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ samples $\rho^j \leftarrow \mathcal{D}$. Then, $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ sends $(\mathsf{epochrnd}, \rho^j)$ to U_i.

Figure 10: Functionality $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$.

We now describe protocol π_{DPoS}, which is a modified version of π_{SPoS} that updates its genesis block B_0 (and thus the leader selection process) for every new epoch. The protocol also adopts an adaptation of the static $\mathsf{maxvalid}_S$ function, defined so that it narrows selection to those chains which share common prefix. Specifically, it adopts the following rule, parameterized by a prefix length k:

Function $\mathsf{maxvalid}(\mathcal{C}, \mathbb{C})$. Returns the longest chain from $\mathbb{C} \cup \{\mathcal{C}\}$ that does not fork from $\mathcal{C}$ more than k blocks. If multiple exist it returns $\mathcal{C}$, if this is one of them, or it returns the one that is listed first in $\mathbb{C}$.

Protocol π_{DPoS} is described in Figure 11 and functions in the $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$-hybrid model.

5.2 Simulating a Trusted Beacon

While protocol π_{DPoS} handles multiple epochs and takes into consideration changes in the stake distribution, it still relies on $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ to perform the leader selection process. In this section, we show how to implement $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ through Protocol π_{DLS}, which allows the stakeholders to compute the randomness and auxiliary information necessary in the leader election.

Recall, that the only essential difference between $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ and $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ is the continuous generation of random strings $\rho^2, \rho^3, \ldots$ for epochs $e_2, e_3, \ldots$. The idea is simple, protocol π_{DLS} will use a coin tossing protocol to generate unbiased randomness that can be used to define the values $\rho^j, j \geq 2$

Protocol π_{DPoS}

π_{DPoS} is a protocol run by a set of stakeholders, initially equal to $U_1, \ldots, U_n$, interacting among themselves and with $\mathcal{F}_{\text{DLS}}^{\mathcal{D},\mathsf{F}}$ over a sequence of L slots $S = \{sl_1, \ldots, sl_L\}$. π_{DPoS} proceeds as follows:

1. **Initialization** Stakeholder $U_i \in \{U_1, \ldots, U_n\}$, receives from the key registration interface its public and secret key. Then it receives the current slot from the diffuse interface and in case it is sl_1 it sends $(\mathsf{genblock\_req}, U_i)$ to $\mathcal{F}_{\text{LS}}^{\mathcal{D},\mathsf{F}}$, receiving $(\mathsf{genblock}, \mathbb{S}_0, \rho, \mathsf{F})$ as answer. U_i sets the local blockchain $\mathcal{C} = B_0 = (\mathbb{S}_0, \rho)$ and the initial internal state $st = H(B_0)$. Otherwise, it receives from the key registration interface the initial chain $\mathcal{C}$, sets the local blockchain as $\mathcal{C}$ and the initial internal state $st = H(\text{head}(\mathcal{C}))$.

2. **Chain Extension** For every slot $sl \in S$, every online stakeholder U_i performs the following steps:
   (a) If a new epoch e_j, with $j \geq 2$, has started, U_i defines $\mathbb{S}_j$ to be the stakeholder distribution drawn from the most recent block with time stamp less than $jR - 2k$ as reflected in $\mathcal{C}$ and sends $(\mathsf{epochrnd\_req}, U_i, e_j)$ to $\mathcal{F}_{\text{LS}}^{\mathcal{D},\mathsf{F}}$, receiving $(\mathsf{epochrnd}, \rho^j)$ as answer.
   (b) Collect all valid chains received via broadcast into a set $\mathbb{C}$, verifying that for every chain $\mathcal{C}' \in \mathbb{C}$ and every block $B' = (st', d', sl', \sigma') \in \mathcal{C}'$ it holds that $\mathsf{Vrf}_{\mathsf{vk}'}(\sigma', (st', d', sl')) = 1$, where vk' is the verification key of the stakeholder $U' = \mathsf{F}(\mathbb{S}_{j'}, \rho^{j'}, sl')$ with $e_{j'}$ being the epoch in which the slot B' belongs (as determined by sl'). U_i computes $\mathcal{C}' = \mathsf{maxvalid}(\mathcal{C}, \mathbb{C})$, sets $\mathcal{C}'$ as the new local chain and sets state $st = H(\text{head}(\mathcal{C}'))$.
   (c) If U_i is the slot leader determined by $\mathsf{F}(\mathbb{S}_j, \rho^j, sl)$ in the current epoch e_j, it generates a new block $B = (st, d, sl, \sigma)$ where st is its current state, $d \in \{0, 1\}^*$ is the data and $\sigma = \mathsf{Sign}_{\mathsf{sk}_i}(st, d, sl)$ is a signature on (st, d, sl). U_i computes $\mathcal{C}' = \mathcal{C}|B$, broadcasts $\mathcal{C}'$, sets $\mathcal{C}'$ as the new local chain and sets state $st = H(\text{head}(\mathcal{C}'))$.

Figure 11: Protocol π_{DPoS}

bootstrapping on the initial random string and initial honest stakeholder distribution. However, notice that the adversary could cause a simple coin tossing protocol to fail by aborting. Thus, we build a coin tossing scheme with "guaranteed output delivery."

Commitments and Coin Tossing. A coin tossing protocol allows two or more parties to obtain a uniformly random string. A classic approach to construct such a protocol is by using commitment schemes. In a commitment scheme, a *committer* carries out a *commitment phase*, which sends evidence of a given value to a *receiver* without revealing it; later on, in an *opening phase*, the committer can send that value to the receiver and convince it that the value is identical to the value committed to in the commitment phase. Such a scheme is called *binding* if it is hard for the committer to convince the receiver that he was committed to any value other than the one for which he sent evidence in the commitment phase, and it is called *hiding* if it is hard for the receiver to learn anything about the value before the opening phase. We denote the commitment phase with randomness r and message m by $\mathsf{Com}(r, m)$ and the opening as $\mathsf{Open}(r, m)$.

In a standard two-party coin tossing protocol [6], one party starts by sampling a uniformly random string u_1 and sending $\mathsf{Com}(r, u_1)$. Next, the other party sends another uniformly random string u_2 in the clear. Finally, the first party opens u_1 by sending $Open(r, u_1)$ and both parties compute output $u = u_1 \oplus u_2$. Note, however, that in this classical protocol the committer may selectively choose to "abort" the protocol (by not opening the commitment) once he observes the value u_2. While this is an intrinsic problem of the two-party setting, we can avoid this problem in the multi-party setting by relying on a verifiable secret sharing scheme and an honest majority amongst the protocol participants.

Figure 12: The two stages of the protocol π_{DPoS} that use the blockchain as a broadcast channel.

Verifiable Secret Sharing (VSS). A secret sharing scheme allows a *dealer* P_D to split a secret σ into n *shares* distributed to parties $P_1, \ldots, P_n$, such that no adversary corrupting up to t parties can recover σ. In a Verifiable Secret Sharing (VSS) scheme [17], there is the additional guarantee that the honest parties can recover σ even if the adversary corrupts the shares held by the parties that it controls and even if the dealer itself is malicious. We define a VSS scheme as a pair of efficient dealing and reconstruction algorithms $(\mathsf{Deal}, \mathsf{Rec})$. The dealing algorithm $\mathsf{Deal}(n, \sigma)$ takes as input the number of shares to be generated n along with the secret σ and outputs shares $\sigma_1, \ldots, \sigma_n$. The reconstruction algorithm Rec takes as input shares $\sigma_1, \ldots, \sigma_n$ and outputs the secret σ as long as no more than t shares are corrupted (unavailable shares are set to $\perp$ and considered corrupted). Schoenmakers [31] developed a simple VSS scheme based on discrete logarithms suitable for our purposes.

Constructing Protocol π_{DLS}. The main problem to be solved when realizing $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ with a protocol run by the stakeholders is that of generating uniform randomness for the leader selection process while tolerating adversaries that may try to interfere by aborting or feeding incorrect information to parties. In order to generate uniform randomness ρ^j for epoch e_j, $j \geq 2$, the elected stakeholders for epoch e_{j-1} will employ a coin tossing scheme for which all honest parties are guaranteed to receive output as long as there is an honest majority. The protocol has two stages, commit and reveal which are split into phases. The stages of the protocol are presented in Figure 12. The *Commitment Phase* covers the whole commitment stage, and proceeds as follows: for $1 \leq i \leq n$, stakeholder U_i samples a uniformly random string $u_i \in \{0, 1\}^{R \log \tau}$ and randomness r_i for the underlying commitment scheme, generates shares $\sigma_1^i, \ldots, \sigma_n^i$, and posts $\mathsf{Com}(r_i, u_i)$ to the blockchain together with the encryptions of the all the shares under the public-key of each respective shareholder. After $4k$ slots, players remove the k most recent blocks of their chain, and if commitments from a majority of stakeholders are posted on the blockchain and shares from a majority of stakeholders have been received, the reveal stage starts (in the other case the protocol halts). In the reveal stage there are two phase: the *Reveal Phase* and the *Recovery Phase*. In the reveal phase, for $1 \leq i \leq n$, stakeholder U_i posts $\mathsf{Open}(r_i, u_i)$ to the blockchain. After $4k$ slots players remove the most recent k blocks and identify all stakeholders that have issued openings of the form $\mathsf{Open}(r_i, u_i)$. In the final *Recovery Phase*, lasting $2k$ slots, if a stakeholder U^a that initially submitted a commitment is identified as not posting an opening to its commitment, the honest parties can post all shares $\sigma_1^a, \ldots, \sigma_n^a$ in order to use $\mathsf{Rec}(\sigma_1^a, \ldots, \sigma_n^a)$ to reconstruct u^a. Finally, each stakeholder uses the values u_i obtained in the second round to compute $\rho^j = \sum_i u_i$. Protocol π_{DLS} is described in Figure 13. We remark that it is possible to run the reveal and recovery phases in parallel, however for improved efficiency we choose to run them sequentially.

Security Proof Sketch. (Reducing $\mathcal{F}_{\mathsf{DLS}}^{\mathcal{D},\mathsf{F}}$ to $\mathcal{F}_{\mathsf{LS}}^{\mathcal{D},\mathsf{F}}$ via protocol π_{DLS}). As before, we consider an adversary who holds a strict minority portion of the total stake. We set the epoch to be

<div style="border: 1px solid black; padding: 10px;">

Protocol π_{DLS}

π_{DLS} is a protocol run by stakeholders $U_1, \ldots, U_n$ interacting among themselves over a sequence of L slots $S = \{sl_1, \ldots, sl_L\}$ and proceeds as follows for every epoch e_j that lasts $R = 10k$ slots:

1. **Commitment Phase** ($4k$ slots) When epoch e_j starts, for $1 \leq i \leq n$, stakeholder U_i samples a uniformly random string u_i and randomness r_i for the underlying commitment scheme, generates shares $\sigma_1^i, \ldots, \sigma_n^i \leftarrow \text{Deal}(n, u_i)$ and encrypts each share σ_k^i under stakeholder U_k's public-key. Finally, U_i posts the encrypted shares and commitments $\text{Com}(r_i, u_i)$ to the blockchain.

2. **Reveal Phase** ($4k$ slots) After slot $4k$, for $1 \leq i \leq n$, stakeholder U_i opens its commitment by posting $\text{Open}(r_i, u_i)$ to the blockchain, provided that the chain $\mathcal{C}^{\lceil k}$ contains commitments from the majority of U_i.

3. **Recovery Phase** ($2k$ slots) After slot $8k$, for any stakeholder U^a that has not participated in the reveal phase, i.e., it has not posted in $\mathcal{C}^{\lceil k}$ an $\text{Open}(r_a, u_a)$ message, for $1 \leq i \leq n$, U_i submits its share σ_i^a for insertion to the blockchain. When all shares $\sigma_1^a, \ldots, \sigma_n^a$ are available, each stakeholder U_i can compute $\text{Rec}(\sigma_1^a, \ldots, \sigma_n^a)$ to reconstruct u_a (independently of whether U^a opens the commitment or not).

   - Given input $(\text{genblock\_req}, U_i, e_j, \mathbb{S}_j)$, U_i uses the values above to compute $\rho^j = \sum_{l \in \mathbb{L}} u_l$ where $\mathbb{L}$ is the subset of stakeholders that participated in epoch e_j. It returns $(\text{genblock}, B_0, \mathbb{S}_j)$ with $B_0 = (\mathbb{S}_j, \rho^j)$.

</div>

Figure 13: Protocol π_{DLS}.

$R = 10k$ and observe that in the first $4k$ slots, i.e., the Commitment Phase, the Chain Growth property proven in Theorem 4.33 guarantees that at least $2k$ blocks will be added to the chain of all honest parties, the common prefix property, proven in Theorem 4.30, ensures that the remaining blockchain after chopping off k blocks will have at least k blocks and the chain quality property proven in Theorem 4.34 guarantees that a constant fraction of these blocks were generated by honest stakeholders with overwhelming probability. Thus, we know that at least one honest block containing all the honest parties' commitments was generated and included in the joint view of all honest parties. By a similar argument, by the end of the Reveal Phase, we are guaranteed to have all openings included in the blockchain and agreed upon by the honest stakeholders. In case a stakeholder does not post a valid opening to its commitment, yet by the same argument (but without needing common prefix), we have a guarantee that by the end of the Recovery at least $n/2 + 1$ shares will be posted (due to the honest majority) and agreed upon by the honest stakeholders, allowing them to recover the original input to the unopened commitments. Notice that if a majority of the stakeholders are honest, they either obtain enough values u_l to compute a uniformly random string ρ^j by the end of the reveal phase or manage to recover such value if the respective stakeholders do not open their commitments.

5.3 Robust Transaction Ledger

Recall that in the dynamic stake case, we would have to conceive a way to prevent deep forks. To see this, consider a player who is offline and joins the system after a number of epochs have passed. Even if in the system execution the current set of stakeholders satisfies honest majority, it could be the case that honest majority is violated in one of the previous epochs by this time and hence the adversary may produce an alternative history consistent with the view of an honest party. In order to capture the interaction between security and the modification of stake we introduce the following property.

Definition 5.1. *Consider two slots sl_1, sl_2 and an execution $\mathcal{E}$. The stake shift between sl_1, sl_2 is the maximum possible statistical distance of the two weighted-by-stake distributions that are defined using the stake reflected in the chain $\mathcal{C}_1$ of some honest stakeholder active at sl_1 and the chain $\mathcal{C}_2$ of some honest stakeholder active at sl_2 respectively.*

Given the definition above we can now express the following theorem about the common prefix property.

Theorem 5.2. *Fix parameters $k, R, L \in \mathbb{N}, \epsilon, \sigma \in (0,1)$. Let R be the epoch length and L the total lifetime of the system. Assume the adversary is restricted to $\frac{1-\epsilon}{2} - \sigma$ relative stake and that the π_{SPOS} protocol satisfies the common prefix property with parameters R, k and probability of error ϵ_{CP}, the chain quality property with parameters $\mu \geq 1/k, k$ and probability of error ϵ_{CQ} and the chain growth property with parameters $\tau \geq 1/2, k$ and probability of error ϵ_{CG}.*

Then, the π_{DPOS} protocol satisfies persistence with parameters k and liveness with parameters $u = 2k$ throughout a period of L slots (or $\mathsf{Bad}^{1/2}$ happens) with probability $1 - R(\epsilon_{\text{CQ}} + \epsilon_{\text{CP}} + \epsilon_{\text{CG}})$, assuming that σ is the maximum stake shift over $10k$ slots, corruption delay $D \geq 2R - 4k$ and no honest player is offline for more than k slots.

Proof. (sketch) Observe that with probability of error $\epsilon_{\text{CQ}} + \epsilon_{\text{CP}} + \epsilon_{\text{CG}}$ the π_{SPOS} protocol executed in the first epoch, given the assumptions imposed to the environment, will enable the parties to use the blockchain as a broadcast channel to simulate the trusted beacon and produce the randomness required to seed the leader election in the next epoch (this combines Theorems 4.30, 4.34, 4.33 and the results of the previous section). The corruption delay of $D \geq 2R - 4k$ will ensure that the adversary does not receive an adaptive corruption advantage for learning the next beacon value ahead of time. It follows that with probability $1 - (L/R)(\epsilon_{\text{CQ}} + \epsilon_{\text{CP}} + \epsilon_{\text{CG}})$ all epochs in the lifetime of the system will be seeded correctly and the π_{SPOS} protocol can be bootstrapped and will continue to operate properly in each next epoch. $\qquad\square$

5.4 Input Endorsers

We next present an extension of our basic protocol that assigns two different roles to stakeholders and introduces our incentive structure. As before in each epoch there is a set of elected stakeholders that runs the secure multiparty coin flipping protocol and are the slot leaders of the epoch. Together with those there is a (not necessarily disjoint) set of stakeholders called the endorsers. Now each slot has two types of stakeholders associated with it; the slot leader who will issue the block as before and the slot *endorser* who will endorse the input to be included in the block. Moreover, contrary to slot leaders, we can elect multiple slot endorsers, say m, for each slot. While this seems like an insignificant modification it gives us a room for improvement because of the following reason: endorsers' contributions will be acceptable even if they are d slots late, where $d \in \mathbb{N}$ is a parameter.

Note that in case no valid endorser input is available when the slot leader is about to issue the block, the leader will go ahead and issue an empty block, i.e., a block without any actual inputs (e.g., transactions in the case of a transaction ledger). Note that slot endorsers just like slot leaders are selected by weighing by stake and thus they are a representative sample of the stakeholder population. In the case of a transaction ledger the same transaction might be included by many input endorsers simultaneously. In case that a transaction is multiply present in the blockchain its first occurrence only will be its "canonical" position in the legder. The enhanced protocol, π_{DPOSwE}, can be easily seen to have the same persistence and liveness behaviour as π_{DPOS}: the modification with endorsers does not provide any possibility for the adversary to prevent the chain

from growing, accepting inputs, or being consistent. However, if we measure chain quality in terms of number of *endorsed inputs* included this produces a more favorable result: it is easy to see that the number of endorsed inputs originating from a set of stakeholders S in any k-long portion of the chain is proportional to the relative stake of S with high probability. This stems from the fact that it is sufficient that a single honest block is created for all the endorsed inputs of the last d slots to be included in it. Assuming $d \geq 2k$, any set of stakeholders S will be an endorser in a subset of the d slots with probability proportional to its cumulative stake, and thus the result follows.

6   Anonymous Communication and Stronger Adversaries

The protocols constructed in the previous section are proven secure against *delayed* adaptive corruptions, meaning that, after requesting to corrupt a given party U_i, the adversary has to wait for D slots before the corruption actually happens. However it is desirable to make D as small as possible, or even eliminate it altogether to achieve security against a standard adaptive adversary.

The delay is required because the adversary must not be able to corrupt parties once it knows that they are the slot leaders for a given slot. However, notice that the slot leaders are selected by weighting public keys by stake, while the adversary can only choose to corrupt a user U_i without knowing its public key. Thus, the adversary must be able to observe communication between U_i and the Diffuse functionality in order to determine which public key is associated with user U_i and detect when U_i is selected as a slot leader. We will show that we can eliminate the delay by extending our model with a sender anonymous broadcast channel (provided by the Diffuse functionality) and having the environment activate all parties in every round. We introduce the following modifications in the ideal functionalities:

- Diffuse Functionality: The functionality will work as described in Section 2 except that it will remove all information about the sender U_s of every message before delivering it to the receiver U_r's inbox (input tape), thus ensuring that the sender remain anonymous.[5]

- Key and Transaction Functionality: The functionality will work as described in Section 2 except that it will allow immediate corruption of a user U upon receiving a message $(\mathsf{Corrupt}, U)$ from the adversary.

Apart from these modifications in the ideal functionalities, we also change the environment behavior by requiring that it activates *all* users at every slot sl_j. Having all parties being activated at every slot results in an anonymity set of size equal to the number of honest parties, making it difficult for the adversary to associate a given public key with a user (*i.e.* any of the honest parties could be associated with a given public key that is not associated with a corrupted party). In this extended model we can reprove Theorem 5.2 without a delay D by strengthening the restrictions that are imposed on the environment in the following way.

- We will say the adversary is *restricted to less than* 50% *relative stake for windows of length D* if for all sets of consecutive slots of length D, the sum over all corrupted keys of the maximum stake held by each key during this period of D slots (in any possible $\mathbb{S}_j(r)$ where U_j is an honest party) is no more than 50% of the minimum total stake during this period. In case the above is violated an event $\mathsf{Bad}_D^{1/2}$ becomes true for the given execution.

Using the above strengthened condition, we can remove the corruption delay requirement D in Theorem 5.2 by assuming that $\mathsf{Bad}^{1/2}$ is substituted with $\mathsf{Bad}_D^{1/2}$.

---

[5] In practice, a sender anonymous broadcast channel with properties akin to those of the Diffuse functionality can be implemented by Mix-networks [11] or DC-networks [12] that can be executed by the nodes running the protocol.

7 Incentive Structure

While the fundamental blockchain analysis we perform is in the cryptographic setting of [19], we include in this section a discussion regarding the incentive structure of our system. For these purposes we adopt a standard game-theoretic framework, analogous to those successfully applied for other blockchains (see [21] for a recent analysis of bitcoin).

As in bitcoin, stakeholders that issue blocks are incentivized to participate in the protocol by collecting transaction fees. Contrary to bitcoin, of course, one does not need to incentivize stakeholder to invest computational resources to issue blocks. Rather, *availability* and *transaction verification* should be incentivized. Any stakeholder, at minimum, must be online and operational in the following circumstances.

- In the slot prior to a slot she is the elected shareholder so that she queries the network and obtains the currently longest blockchain as well as any endorsed inputs to include in the block.

- In the slot during which she is the elected shareholder so that she issues the block containing the endorsed inputs.

- In a slot during the commit stage of an epoch where she is supposed to issue the VSS commitment of her random string.

- In a slot during the reveal stage of an epoch where she is supposed to issue the required opening shares as well as the opening to her commitment.

- In general, in sufficient frequency, to check whether she is an elected shareholder for the next or current epoch.

- In a slot during which she is the elected input endorser so that she issues the endorsed input (e.g., the set of transactions) that requires processing all available transactions and verifying them.

In order to incentivize the above actions in the setting of a transaction ledger, fees can be collected from those that issue transactions to be included in the ledger which can then be transfered to the block issuers. In bitcoin, for instance, fees can be collected by the miner that produces a block of transactions as a reward. In our setting, similarly, a reward can be given to the parties that are issuing blocks and endorsing inputs. The reward mechanism does not have to be immediate as advocated in [28]. In our setting, it is possible to collect all fees of transactions included in a sequence of k blocks in a pool and then distribute that pool to all shareholders that participated during these k slots. For example, all input endorsers that were active may receive reward proportional to the number of inputs they endorsed during the period of k rounds (independently of the actual number of transactions they endorsed).

Other ways to distribute transaction fees are also feasible (including the one that is used by bitcoin itself—even though the bitcoin method is known to be vulnerable to attacks, e.g., the selfing-mining attack).

We proceed to make the reward mechanism more explicit in the case that the endorsed inputs to the ledger are sequences of transactions. First we set the reward interval to match the epoch duration, i.e., $d = 10k$. Let $\mathcal{C}$ be a chain consisting of blocks $B_0, B_1, \ldots$. Consider the sequence of blocks of the μ-th epoch denoted by $B_1, \ldots, B_s$ with timestamps in $\{\mu d + 1, \ldots, (\mu + 1)d\}$ that contain an $r \geq 0$ sequence of endorsed inputs. Suppose that the total reward pool P_R, as defined by transaction fees, is equal to the sum of the transaction fees that are included in the endorsed inputs.

If a transaction occurs multiple times (as part of different endorsed inputs) or even in conflicting versions, only the first occurrence of the transaction is taken into account (and is considered to be part of the ledger) in the calculation of P_R, where the total order used is induced by the order the endorsed inputs are included in $\mathcal{C}$. In the sequence of these blocks, we identify by $L_1, \ldots, L_d$ the slot leaders corresponding to the slots of the epoch and by $E_1, \ldots, E_r$ the input endorsers that contributed the sequence of r endorsed inputs. Subsequently, the i-th stakeholder U_i can claim[6] a reward up to the amount $(|\{j \mid U_i = E_j\}|/(2r) + |\{j \mid U_i = L_j\}|/(2d))P_R$.

Observe that the above reward mechanism has the following features: (i) it rewards elected committee members for just being committee members, independently of whether they issued a block or not, (ii) it rewards the input endorsers with the inputs that they have contributed.

We proceed to show that our system is a δ-Nash (approximate) equilibrium, cf. [25, Section 2.6.6]. Specifically, the theorem states that any party deviating from the protocol can get at most $(1 + \delta)$ of the rewards compared to the participants following the protocol.

Proposition 7.1. *The honest strategy in the protocol is a δ-Nash equilibrium provided that all players command a stake less than $(1 - \epsilon)/2 - \sigma$ for some constants $\epsilon, \sigma \in (0, 1)$ as in Theorem 5.2, for $\delta > 0$.*

Proof. (sketch) Consider a rational player U restricted as in the statement of the theorem, that engages in a protocol execution together with a number of other players that follow the protocol faithfully for a total number of L epochs. We will show that any deviation from the protocol will not result in substantially higher rewards for U. Observe that based on Theorem 5.2, no matter the strategy of U, the protocol will enable all honest users to get at least $(1 - \delta)$ fraction of the rewards they are entitled to as slot leaders and input endorsers with overwhelming probability in k. The latter stems from persistence and liveness: at least a number of honest blocks will be included in each epoch and as a result all input endorsers will have an opportunity to get their rewards. Given the above, player U will be rewarded proportionally to how many times she is selected as a slot leader (independent of her strategy), and moreover proportionally to how many inputs she contributes as an endorser. It thus follows that any other strategy that U may follow will lead to the same, or smaller amount of rewards. $\square$

Remark 1. *In the above theorem, for simplicity, we assumed that protocol costs are negligible. Nevertheless, it is straightforward to extend the proof to cover a setting where a negative term is introduced in the payoff function that is proportional to the number of times inputs are endorsed and the number of messages transmitted for the MPC protocol. The proof is resilient to these costs because endorsed inputs and MPC protocol messages cannot be stifled by the adversary and hence the reward function can be designed with suitable weights for such actions that offsets their cost.*

8 Stake Delegation

As discussed in the previous section, stakeholders must be online in order to generate blocks when they are selected as slot leaders. However, this might be unattractive to stakeholders with a small stake in the system. Moreover, requiring that a majority of elected stakeholders participate in the coin tossing protocol for refreshing randomness introduces a strain on the on the stakeholders and the network, since it might require broadcasting and storing a large number of commitments and shares.

---

[6]Claiming a reward is performed by issuing a "coinbase" type of transaction at any point after k blocks in a subsequent epoch to the one that a reward is being claimed from.

We mitigate these issues by providing a method for reducing the size of the group of stakeholders that engage in the coin tossing protocol. Instead of the elected stakeholders directly forming the committee that will run coin tossing, a group of delegates will act on their behalf. In more detail, we put forth a *delegation scheme*, whereby stakeholders will authorize other entities, called delegates, who may be stakeholders themselves, to represent them in the coin tossing protocol. A delegate may participate in the protocol only if it represents a certain number of stakeholders whose aggregate stake exceeds a given threshold. Such a participation threshold ensures that a "fragmentation" attack, that aims to increase the delegate population in order to hurt the performance of the protocol, cannot incur a large penalty.

8.1 Minimum Committee Size

To appreciate the benefits of delegation, recall that in the basic protocol (π_{DPoS}) a committee member selected by weighing by stake is honest with probability $1/2 + \epsilon$ (this being the fraction of the stake held by honest players). Thus, the number of honest players selected by k invocations of weighing by stake is a binomial distribution. We are interested in the probability of a malicious majority, which can be directly controlled by a Chernoff bound. Specifically, if we let Y be the number of times that a malicious committee member is elected then

$$\Pr[Y \geq k/2] = \Pr[Y \geq (1 + \delta)(1/2 - \epsilon)k]$$
$$\leq \exp(-\min\{\delta^2, \delta\}(1/2 - \epsilon)k/4)$$
$$< \exp(-\delta^2(1/2 - \epsilon)k/4)$$

for $\delta = 2\epsilon/(1 - 2\epsilon)$. Assuming $\epsilon < 1/4$, it follows that $\delta < 1$.

Consider the case that $\epsilon = 0.05$; then we have the bound $\exp(-0.00138 \cdot k)$ which provides an error of $1/1000$ as long as $k \geq 5000$. Similarly, in the case $\epsilon = 0.1$, we have the bound $\exp(-0.00625k)$ which provides the same error for $k \geq 1100$.

We observe that in order to withstand a significant number of epochs, say 2^{15} (which, if we equate a period with one day, will be 88 years), and require error probability 2^{-40}, we need that $k \geq 32648$.

In cases where the wealth in the system is not concentrated among a small set of stakeholders the above choice is bound to create a very large committee. (Of course, the maximum size of the committee is k.)

8.2 Delegation Scheme.

To facilitate a smaller committee size we introduce a simple delegation idea that can enable stakeholders to delegate the committee participation rights to other entities that they trust. The concept of delegation is simple: any stakeholder can allow a *delegate* to generate blocks on her behalf. In the context of our protocol, where a slot leader signs the block it generates for a certain slot, such a scheme can be implemented in a straightforward way based on *proxy signatures* [7].

A stakeholder can transfer the right to generate blocks by creating a *proxy signing key* that allows the delegate to sign messages of the form (st, d, sl_j) (i.e., the format of messages signed in Protocol π_{DPoS} to authenticate a block). In order to limit the delegate's block generation power to a certain range of epochs/slots, the stakeholder can limit the proxy signing key's valid message space to strings ending with a slot number sl_j within a specific range of values. The delegate can use a proxy signing key from a given stakeholder to simply run Protocol π_{DPoS} on her behalf, signing the blocks this stakeholder was elected to generate with the proxy signing key. This simple scheme is secure due to the *Verifiability* and *Prevention of Misuse* properties of proxy signature

schemes, which ensure that any stakeholder can verify that a proxy signing key was actually issued by a specific stakeholder to a specific delegate and that the delegate can only use these keys to sign messages inside the key's valid message space, respectively.

We remark that while proxy signatures can be described as a high level generic primitive, it is easy to construct such schemes from standard digital signature schemes through delegation-by-proxy as shown in [7]. In this construction, a stakeholder signs a certificate specifying the delegates identity (e.g., its public key) and the valid message space. Later on, the delegate can sign messages within the valid message space by providing signatures for these messages under its own public key along with the signed certificate. As an added advantage, proxy signature schemes can also be built from aggregate signatures in such a way that signatures generated under a proxy signing key have essentially the same size as regular signatures [7].

An important consideration in the above setting is the fact that a stakeholder may want to withdraw her support to a stakeholder prior to its proxy signing key expiration. Observe that proxy signing keys can be uniquely identified and thus they may be revoked by keeping a certificate revocation list within the blockchain.

8.2.1 Eligibility threshold

Delegation as described above can ameliorate fragmentation that may occur in the stake distribution. Nevertheless, this does not prevent a malicious stakeholder from dividing its stake to multiple accounts and, by refraining from delegation, induce a very large committee size. To address this, as mentioned above, a threshold T, say 1%, may be applied. This means that any delegate representing less a fraction less than T of the total stake is automatically barred from being a committee member. This can be facilitated by redistributing the voting rights of delegates representing less than T to other delegates in a deterministic fashion (e.g., starting from those with the highest stake and breaking ties according to lexicographic order).

8.2.2 Size of the committee of delegates

Suppose that a committee has been formed, $C_1, \ldots, C_m$, from a total of k draws of weighing by stake. Each committee member will hold k_i such votes where $\sum_{i=1}^{m} k_i = k$. Based on the eligibility threshold above it follows that $m \leq T^{-1}$ (the maximum possible value is the case when all stake is distributed in T^{-1} delegates each holding T of the stake).

9   Attacks Discussion

We next discuss a number of practical attacks and indicate how they are reflected by our modeling and mitigated.

Double spending attacks   In a double spending attack, the adversary wishes to revert a transaction that is confirmed by the network. The objective of the attack is to issue a transaction, e.g., a payment from an adversarial account holder to a victim recipient, have the transaction confirmed and then revert the transaction by, e.g., including in the ledger a second conflicting transaction. Such an attack is not feasible under the conditions of Theorem 5.2. Indeed, persistence ensures that once the transaction is confirmed by an honest player, all other honest players from that point on will never disagree regarding this transaction. Thus it will be impossible to bring the system to a state where the confirmed transaction is invalidated (assuming all preconditions of the theorem hold). See the next section for an experimental discussion about double spending.

Grinding attacks In stake grinding attacks, the adversary tries to influence the slot leader selection process to improve its chances of being selected to generate blocks (which can be used to perform other attacks such as double spending). Basically, when generating a block that is taken as input by the slot leader selection process, the adversary first tests several possible block headers and content in order to find the one that gives it the best chance of being selected as a slot leader again in the future. While this attack affects PoS based cryptocurrencies that collect randomness for the slot leader selection process from raw data in the blockchain itself (*i.e.* from block headers and content), our protocol uses a standard coin tossing protocol that is proven to generate unbiased uniform randomness as discussed in Section 5.2. We show that an adversary cannot influence the randomness generated in Figure 13, which is guaranteed to be uniformly random, thus guaranteeing that slot leaders are selected with probability proportional to their stake.

Transaction denial attacks In a transaction denial attack, the adversary wishes to prevent a certain transaction from becoming confirmed. For instance, the adversary may want to target a specific account and prevent the account holder from issuing an outgoing transaction. Such an attack is not feasible under the conditions of Theorem 5.2. Indeed, liveness ensures that, provided the transaction is attempted to be inserted for a sufficient number of slots by the network, it will be eventually confirmed.

Desynchronization attacks In a desynchronization attack, a shareholder behaves honestly but is nevertheless incapable of synchronizing correctly with the rest of the network. This leads to ill-timed issuing of blocks and being offline during periods when the shareholder is supposed to participate. Such an attack can be mounted by preventing the party's access to a time server or any other mechanism that allows synchronization between parties. Moreover, a desynchronization may also occur due to exceedingly long delays in message delivery. Our model allows parties to become desynchronized by incorporating them into the adversary. No guarantees of liveness and persistence are provided for desynchronized parties and thus we can get security as long as parties with less than 50% of stake get desynchronized. If more than parties get desynchronized our protocol can fail. More general models like partial synchrony [14, 29] are interesting to consider in the PoS design setting.

Eclipse attacks In an eclipse attack, message delivery to a shareholder is violated due to a subversion in the peer-to-peer message delivery mechanism. As in the case of desynchronization attacks, our model allows parties to be eclipse attacked by incorporating them into the adversary. No guarantees of liveness or persistence are provided for such parties.

51% attacks A 51% attack occurs whenever the adversary controls more than the majority of the stake in the system. It is easy to see that any sequence of slots in such a case is with very high probability forkable and thus once the system finds itself in such setting the honest stakeholders may be placed in different forks for long periods of time. Both persistence and liveness can be violated.

Bribery Attacks In bribery attacks [8], an adversary deliberately pays miners (through cryptocurrencies or fiat money) to work on specific blocks and forks, aiming at generating an arbitrary fork that benefits the adversary (*e.g.* by supporting a double spending attack). Miners of PoW based cryptocurrencies do not have to own any stake in order to mine blocks, which makes this attack strategy feasible. In this setting, if the adversary offers a bribe higher than the reward

for correctly generating a block, any rational miner has a clear incentive to accept the bribe and participate in the attack since it increases the miner's financial outcome. However, in our PoS based protocol, malicious slot leaders who agree to deliberately attack the system not only risk to forego any potential profit they would earn from behaving honestly but may also risk to lose equity. Notice that slot leaders must have money invested in the system in order to be able to generate blocks and if an attack against the system is observed it might bring currency value down. Even if the bribe is higher than the reward for correct behavior, the loss from currency devaluation can easily offset any additional profits made by participating in this attack. Hence, bribery attacks may be be less effective against a PoS based consensus protocol than a PoW based one. Currently our rationality model does not formally encompass this attack strategy and investigating its efficacy against PoS based consensus protocols is left as a future work.

Long-range attacks An attacker who wishes to double spend at a later point in time can mount a long-range attack [9] by computing a longer valid chain that starts right after the genesis block where it is the single stakeholder actively participating in the protocol. Even if this attacker owns a small fraction of the total stake, it can locally compute this chain generating only the blocks for slots where it is elected the slot leader and keep generating blocks ahead of current time until its alternative chain has more blocks than the main chain. Now, the attacker can post a transaction to the main chain, wait for it to be confirmed (and for goods to be delivered in exchange for the transaction) and present the longer alternative chain to invalidate its previously confirmed transaction. This attack is ineffective against Ouroboros for two reasons: Protocol $\pi_D LS$ will only output valid leader selection data allowing for the protocol to continue if a majority of the stakeholders participate (or have delegates participate in their behalf) and stakeholders will reject blocks generated for slots that are far ahead of time. Since the alternative chain is generated artificially with blocks and protocol messages generated solely by an attacker who controls a small fraction of the stake the leader selection data needed to start new epochs will be considered invalid by other nodes. Even if the attacker could find a strategy to generate an alternative chain with valid leader selection data, presenting this chain and its blocks generated at slots that are far ahead of time would not result in a successful attack since those blocks far ahead of time would be rejected by the honest stakeholders and the final alternative chain would be shorter than the main chain.

Nothing at stake attacks The "nothing at stake" problem refers in general to attacks against PoS blockchain systems that are facilitated by shareholders continuing simultaneously multiple blockchains exploiting the fact that little computational effort is needed to build a PoS blockchain. Provided that stakeholders are frequently online, nothing at stake is taken care of by our analysis of forkable strings (even if the adversary brute-forces all possible strategies to fork the evolving blockchain in the near future, there is none that is viable), and our chain selection rule that instructs players to ignore very deep forks that deviate from the block they received the last time they were online. It is also worth noting that, contrary to PoW-based blockchains, in our protocol it is infeasible to have a fork generated in earnest by two shareholders. This is because slots are uniquely assigned and thus at any given moment there is a single uniquely identified shareholder that is elected to advance the blockchain. Players following the longest chain rule will adopt the newly minted block (unless the adversary presents at that moment an alternative blockchain using older blocks). It is remarked in [10] that the "tragedy of commons" might lead stakeholders in some PoS based schemes to adhere to attacks because they do not have the power to deter attacks by themselves and would incur financial losses even if they did not join the attack. This would lead rational stakeholder to accept small bribes in alternative currencies that might at least obtain

some financial gain. However, in the incentive structure of Ouroboros, slot leaders and endorsers who could potentially join an attack would receive rewards in both the main and the adversarial chain, resulting in those stakeholders not achieving higher profits by joining the attack.

Past majority attacks As stake moves our assumption is that only the *current* majority of stakeholders is honest. This means that past account keys (which potentially do not hold any stake at present) may be compromised. This leads to a potential vulnerability for any PoS system since a set of malicious shareholders from the past can build an alternative blockchain exploiting such old accounts and the fact that it is effortless to build such a blockchain. In light of Theorem 5.2 such attack can only occur against shareholders who are not frequently online to observe the evolution of the system or in case the stake shifts are higher than what is anticipated by the preconditions of the theorem. This can be seen a special instance of the nothing at stake problem, where the attacker no longer owns any stake in the system and is thus free from any financial losses when conducting the attack.

Selfish-mining In this type of attack, an attacker withholds blocks and releases them strategically attempting to drop honestly generated blocks from the main chain. In this way the attacker reduces chain growth and increases the relative ratio of adversarially generated blocks. In conventional reward schemes, as that of bitcoin, this has serious implications as it enables the attacker to obtain a higher rate of rewards compared to the rewards it would be receiving in case it was following the honest strategy. Using our reward mechanism however, selfish mining attacks are neutralized. The intuition behind this, is that input endorsers, who are the entities that receive rewards proportionally to their contributions, cannot be stifled because of block withholding: any input endorser can have its contribution accepted for a sufficiently long period of time after its endorsement took place, thus ensuring it will be incorporated into the blockchain (due to sufficient chain quality and chain growth). Given that input endorsers' contributions are (approximately) proportional to their stake this ensures that reward distribution cannot be affected substantially by block withholding.

10 Experimental Results

We have implemented a prototype instantiation of Ouroboros in Haskell in order to evaluate its concrete performance. More specifically, we have implemented Protocol π_{DPoS} using Protocol π_{DLS} to generate leader selection parameters (basically generating fresh randomness for the weighed stake sampling procedure). For this instantiation, we use the PVSS scheme of [31] implemented over the elliptic curve secp256r1. This PVSS scheme's share verification information includes a commitment to the secret, which is also used as the commitment specified in protocol π_{DLS}; this eliminates the need for a separate commitment to be generated and stored in the blockchain. In order to obtain better efficiency, the final output ρ of Protocol π_{DLS} is a uniformly random binary string of 32 bytes. This string is then used as a seed for a PRG (ChaCha in our implementation) and stretched into R random labels of $\log \tau$ bits corresponding to each slot in an epoch. The weighing by stake leader selection process is then implemented by using the random binary string associated to each epoch to perform the sequence of coin-flips for selecting a stakeholder. The signature scheme used for signing blocks is ECDSA, also implemented over curve secp256r1.

Experimental Setup. In our experiments we consider a total number of stakeholders that varies from 1 (for the sake of a centralized baseline) to 100 communicating through a peer-to-peer network.

Although these numbers of stakeholders might seem low, notice that delegation can be used to limit the size of the committee of stakeholders actively participating in the protocol while still supporting a much larger number of stakeholders. We consider both situations where the stake is equally distributed among all stakeholders and situations with skewed stake distributions where a small number of stakeholders have most of the stake (resembling the distribution of computational power across mining pools in Bitcoin). The prototype implementation for each stakeholder is run on t2.large Amazon EC2 nodes equally spread among Asia, Europe, North America and South America. The slot duration is set to 20 seconds since at this level it was observed that almost always all nodes received messages transmitted by other nodes.

10.1 Transaction Confirmation Time Under Optimal Network Conditions

We first examine the time required for confirming a transaction in a setting where the network is not under substantial load and transactions are processed as they appear.

Adversary	BTC	OB Covert	OB General
0.10	50	3	5
0.15	80	5	8
0.20	110	7	12
0.25	150	11	18
0.30	240	18	31
0.35	410	34	60
0.40	890	78	148
0.45	3400	317	663

Figure 14: Transaction confirmation times in minutes that achieve assurance 99.9% against a hypothetical double spending attack with different levels of adversarial power for Bitcoin and Ouroboros (both covert and general adversaries).

In Fig. 14 we lay out a comparison in terms of transaction confirmation time between Bitcoin and Ouroboros showing how much a verifier has to wait to be sure that the best possible[7] double-spending attack succeeds with probability less than 0.1%. In the case of Bitcoin, we consider a double-spending attacker that commands a certain percentage of total hashing power and wishes to revert a transaction. The attacker attempts to double-spend via a block-witholding attack as described in the same paper (the attacker mines a private fork and releases it when it is long enough). In the case of Ouroboros we consider a double spending attacker that attempts to brute force the space of all possible forks for the current slot leader distribution in a certain segment of the protocol and commands a certain percentage of the total stake. We consider both the covert and the general adversarial setting for Ouroboros.

In all of the scenarios, we measure the number of minutes that one has to wait in order to achieve probability of double spending less than 0.1%. In Fig. 15 we present a graph that illustrates the speedup graphically.

We note that the above measurements compare our Ouroboros implementation with Bitcoin in the way the two systems are parameterized (with 10 minute blocks for Bitcoin). Exploring alternative parameterizations for Bitcoin (such as making the proof-of-work easier) can speed up the transaction processing, nevertheless this cannot be done without carefully measuring the impact on overall security.

---

[7]The "best possible" is only in the the case of Ouroboros, for Bitcoin we use the best known attack.

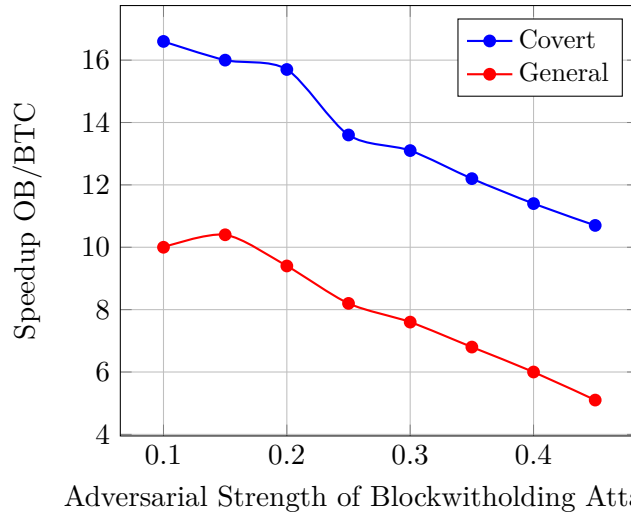Confirmation time speed up of Ouroboros over BTC

Figure 15: Ouroboros vs. Bitcoin speedup of transaction confirmation time against a hypothetical double spending attacker for assurance level 99.9%. Ouroboros is at least 10 to 5 times faster for regular adversaries and 16 to 10 times faster for covert adversaries.

10.2 Stress-testing Experiments
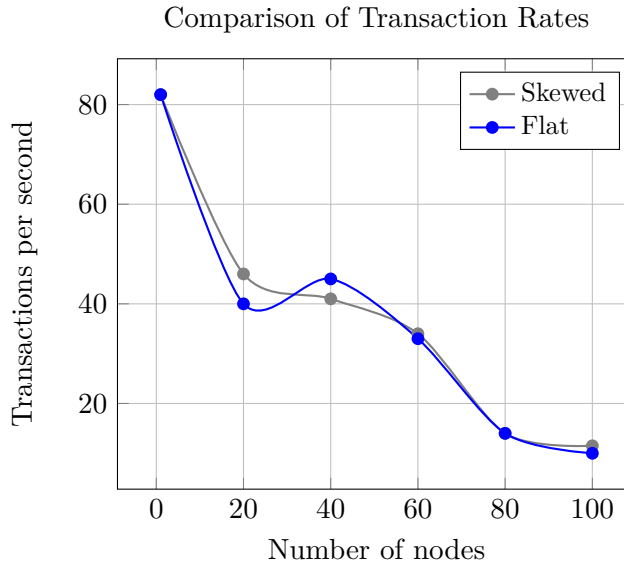
Comparison of Transaction Rates

Figure 16: Graph of transaction rates for an Ouroboros prototype run over a wide area network with flat and skewed stake distributions.

In the next experiment, we measure the maximum Transaction Per Second (TPS) rates that our protocol sustains when subjected to a stress-test when the nodes running the protocol are "bombarded" with transactions. In this experiment, apart from the nodes that implement stakeholders running the protocol, we employ an extra node that acts as a Transaction Generator (henceforth referred to as TxGen). The stakeholders participating in the experiment are registered in the

genesis block and TxGen starts by generating a number of initial transactions from each of the stakeholders distributing their initial stake in number of $maxTPS$ outputs to themselves, which form an output pool for each stakeholder. In the first phase, called *warm-up*, TxGen uses the outputs of each transaction pool as inputs to generate a number of $currentTPS$ transactions from each stakeholder to itself in regular intervals. These transactions (as well as future transactions in this experiments) are sent into the peer-to-peer network connecting the stakeholders. Once a transaction is confirmed (meaning that it appears in a block of depth k in the current blockchain), TxGen uses the output of this transaction as input to generate a new transaction with an output to the same stakeholder that owns the input. After the warm-up phase, the system is populated with a number of transactions that are both generated from the output pool of each stakeholder and from the outputs of confirmed transactions. In the next phase, called *interim*, the same transaction creation procedure continues but now TxGen increments a variable *counter* (initially set to 0) for each transaction (from/to any stakeholder) that is confirmed in the blockchain. In a final phase, called *cool-down*, no new transactions are generated at all but *counter* continues to be incremented for each confirmed transaction. By the end of the cool-down phase, a variable $realTPS$ is calculated by dividing *counter* by the total time elapsed from the beginning of the interim phase to the end of the cool-down phase. If $realTPS$ is bigger than $currentTPS$, the experiment is rerun with an incremented value of $currentTPS$ in order to measure the maximum TPS rate the system can sustain. Otherwise, the past $currentTPS$ value is considered as the highest TPS rate the system can achieve. Our results are summarized in Figures 16 and 17.

Known Prototype Caveats. Our implementation is a first rendering of a working prototype and we anticipate substantial improvements can be still attained. When analysing our current experimental results, one should keep in mind that our prototype implementation currently incurs several overheads that are not intrinsic to our protocol. One of the main issues is the naive peer-to-peer layer, which is not optimized and causes a heavy communication overhead as the number of nodes grows. These overheads explain why the TPS rates fall rapidly as the number of nodes increases. In an ideal network, the TPS of our system would approximate the average local TPS across all participating nodes, where local TPS corresponds to the performance of a node when is ran as a central transaction processor. Notice that the PVSS scheme from [31] requires $O(n^2)$ exponentiations for share generation and validation (where n is the size of the committee), which results in a concrete execution time that increases quadratically with committee size, limiting the size of the committees that our prototype implementation supports given the chosen slot duration. This limitation can be easily mitigated by using our delegation scheme with a threshold, cf. Section 8.2.1, or devising a PVSS scheme with better computational complexity, which is an interesting open question.

Nodes	TPS; flat	TPS; skewed
1	82	82
20	40	46
40	45	41
60	33	34
80	14	14
100	10	11.5

Figure 17: Transaction rates for an Ouroboros prototype run over a wide area network with flat and skewed stake distributions.

11 Acknowledgements

We thank Ioannis Konstantinou who contributed in a preliminary version of our protocol. We thank Lars Brünjes, Duncan Coutts, Peter Gaži, Kawin Worrasangasilpa for comments on previous drafts of the article. We thank George Agapov for the prototype implementation of our protocol.

References

[1] Noga Alon and Joel Spencer. *The Probabilistic Method.* Wiley, 3rd edition, 2008.

[2] Giuseppe Ateniese, Ilario Bonacina, Antonio Faonio, and Nicola Galesi. Proofs of space: When space is of the essence. In Michel Abdalla and Roberto De Prisco, editors, *Security and Cryptography for Networks - 9th International Conference, SCN 2014, Amalfi, Italy, September 3-5, 2014. Proceedings*, volume 8642 of *Lecture Notes in Computer Science*, pages 538–557. Springer, 2014.

[3] Yonatan Aumann and Yehuda Lindell. Security against covert adversaries: Efficient protocols for realistic adversaries. *J. Cryptology*, 23(2):281–343, 2010.

[4] Iddo Bentov, Ariel Gabizon, and Alex Mizrahi. Cryptocurrencies without proof of work. *CoRR*, abs/1406.5694, 2014.

[5] Iddo Bentov, Charles Lee, Alex Mizrahi, and Meni Rosenfeld. Proof of activity: Extending bitcoin's proof of work via proof of stake [extended abstract]y. *SIGMETRICS Performance Evaluation Review*, 42(3):34–37, 2014.

[6] Manuel Blum. Coin flipping by telephone. In Allen Gersho, editor, *Advances in Cryptology: A Report on CRYPTO 81, CRYPTO 81, IEEE Workshop on Communications Security, Santa Barbara, California, USA, August 24-26, 1981.*, pages 11–15. U. C. Santa Barbara, Dept. of Elec. and Computer Eng., ECE Report No 82-04, 1981.

[7] Alexandra Boldyreva, Adriana Palacio, and Bogdan Warinschi. Secure proxy signature schemes for delegation of signing rights. *J. Cryptology*, 25(1):57–115, 2012.

[8] Joseph Bonneau. Why buy when you can rent? - bribery attacks on bitcoin-style consensus. In Jeremy Clark, Sarah Meiklejohn, Peter Y. A. Ryan, Dan S. Wallach, Michael Brenner, and Kurt Rohloff, editors, *Financial Cryptography and Data Security - FC 2016 International Workshops, BITCOIN, VOTING, and WAHC, Christ Church, Barbados, February 26, 2016, Revised Selected Papers*, volume 9604 of *Lecture Notes in Computer Science*, pages 19–26. Springer, 2016.

[9] Vitalik Buterin. Long-range attacks: The serious problem with adaptive proof of work. https://blog.ethereum.org/2014/05/15/long-range-attacks-the-serious-problem-with-adaptive-proof-of-work/, 2014.

[10] Vitalik Buterin. Proof of stake faq. https://github.com/ethereum/wiki/wiki/Proof-of-Stake-FAQ, 2016.

[11] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–88, 1981.

[12] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *J. Cryptology*, 1(1):65–75, 1988.

[13] George Danezis and Sarah Meiklejohn. Centrally banked cryptocurrencies. In *23nd Annual Network and Distributed System Security Symposium, NDSS 2016, San Diego, California, USA, February 21-24, 2016*. The Internet Society, 2016.

[14] Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, 1988.

[15] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*, volume 9216 of *Lecture Notes in Computer Science*, pages 585–605. Springer, 2015.

[16] Ittay Eyal and Emin Gun Sirer. Majority is not enough: Bitcoin mining is vulnerable. In Angelos D. Keromytis, editor, *Financial Cryptography*, volume 7397 of *Lecture Notes in Computer Science*. Springer, 2014.

[17] Paul Feldman. A practical scheme for non-interactive verifiable secret sharing. In *28th Annual Symposium on Foundations of Computer Science, Los Angeles, California, USA, 27-29 October 1987*, pages 427–437. IEEE Computer Society, 1987.

[18] Bryan Ford. Delegative democracy. http://www.brynosaurus.com/deleg/deleg.pdf, 2002.

[19] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, volume 9057 of *Lecture Notes in Computer Science*, pages 281–310. Springer, 2015.

[20] Charles M Grinstead and J Laurie Snell. *Introduction to Probability*. American Mathematical Society, 2nd edition, 1997.

[21] Aggelos Kiayias, Elias Koutsoupias, Maria Kyropoulou, and Yiannis Tselekounis. Blockchain mining games. In Vincent Conitzer, Dirk Bergemann, and Yiling Chen, editors, *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 365–382. ACM, 2016.

[22] Aggelos Kiayias and Giorgos Panagiotakos. Speed-security tradeoffs in blockchain protocols. Cryptology ePrint Archive, Report 2015/1019, 2015. http://eprint.iacr.org/2015/1019.

[23] Tal Moran and Ilan Orlov. Proofs of space-time and rational proofs of storage. Cryptology ePrint Archive, Report 2016/035, 2016. http://eprint.iacr.org/2016/035.

[24] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. http://bitcoin.org/bitcoin.pdf, 2008.

[25] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[26] Karl J. O'Dwyer and David Malone. Bitcoin mining and its energy footprint. *ISSC 2014 / CIICT 2014, Limerick, June 26–27*, 2014.

[27] Sunoo Park, Krzysztof Pietrzak, Albert Kwon, Joël Alwen, Georg Fuchsbauer, and Peter Gazi. Spacemint: A cryptocurrency based on proofs of space. *IACR Cryptology ePrint Archive*, 2015:528, 2015.

[28] Rafael Pass. Cryptography and game theory. Securty and Cryptography for Networks, 2016, invited talk., 2016.

[29] Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. *IACR Cryptology ePrint Archive*, 2016:454, 2016.

[30] Ayelet Sapirshtein, Yonatan Sompolinsky, and Aviv Zohar. Optimal selfish mining strategies in bitcoin. *CoRR*, abs/1507.06183, 2015.

[31] Berry Schoenmakers. A simple publicly verifiable secret sharing scheme and its application to electronic voting. In Michael J. Wiener, editor, *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, volume 1666 of *Lecture Notes in Computer Science*, pages 148–164. Springer, 1999.