Practical Settlement Bounds for Longest-Chain Consensus

Peter Gaži[1], Ling Ren[2], and Alexander Russell[3(✉)]

[1] IOG, Bratislava, Slovakia
peter.gazi@iohk.io
[2] University of Illinois at Urbana-Champaign, Urbana, IL, USA
renling@illinois.edu
[3] University of Connecticut and IOG, Storrs, CT, USA
acr@uconn.edu

Abstract. Nakamoto's longest-chain consensus paradigm now powers the bulk of the world's cryptocurrencies and distributed finance infrastructure. An emblematic property of longest-chain consensus is that it provides probabilistic settlement guarantees that strengthen over time. This makes the exact relationship between settlement error and settlement latency a critical aspect of the protocol that both users and system designers must understand to make informed decisions. A recent line of work has finally provided a satisfactory rigorous accounting of this relationship for proof-of-work longest-chain protocols, but those techniques do not appear to carry over to the proof-of-stake setting.

This article develops a new analytic approach for establishing such settlement guarantees that yields explicit, rigorous settlement bounds for proof-of-stake longest-chain protocols, placing them on equal footing with their proof-of-work counterparts. Our techniques apply with some adaptations to the proof-of-work setting where they provide improvements to the state-of-the-art settlement bounds for proof-of-work protocols.

1 Introduction

Satoshi Nakamoto introduced the longest-chain consensus paradigm in the 2008 Bitcoin whitepaper [21]. Since its original proposal, the framework has been extended and generalized, and variants of longest-chain protocols now support the bulk of the world's cryptocurrencies and decentralized finance infrastructure.

The fundamental dynamics of the algorithm—in particular, the rate at which participants converge to achieve consensus—depend primarily on three critical parameters: r_h, the rate at which honest players are elected to advance the system; r_a, the rate at which adversarial players are elected to advance the system; and Δ, the maximum network delay. Despite the visible prominence of the algorithm and over a decade of concerted effort by the research community, the relationship between these critical parameters and the resulting consensus guarantee is still not well understood.

The last few years have witnessed rapid progress on this question. In 2020, two independent articles [8,12] precisely determined the region of triples

H. Handschuh and A. Lysyanskaya (Eds.): CRYPTO 2023, LNCS 14081, pp. 107–138, 2023.
https://doi.org/10.1007/978-3-031-38557-5_4

(r_h, r_a, Δ) for which longest-chain consensus *eventually* provides consistency, which is to say that participants in the system eventually converge on a finite prefix of the ledger. These results apply to both proof-of-work and proof-of-stake longest-chain protocols and, somewhat surprisingly, prove that their fundamental "regime of security" is the same.

Practice, however, demands *explicit* settlement guarantees, as blockchain users in the real world must be able to determine when transactions in the ledger have in fact settled with known risk. Likewise, deployed systems must explicitly calibrate block production rate against (estimated) network delays to yield reasonable settlement latency. Such explicit settlement guarantees in the proof-of-work setting have been the subject of an active thread of research [13,19,20]. These works have succeeded in providing satisfactory results for proof-of-work systems with conservative parameters similar to those used in Bitcoin. But significant gaps still remain for more aggressive parameters such as those used in Ethereum (before its switch to proof of stake). Furthermore, very little is known about the proof-of-stake setting, where the only explicit result makes the unrealistic assumption that the network has zero delay ($\Delta = 0$) [16]. This is particularly concerning as it seems that in recent years we have been witnessing the sentiment of preference for PoS over PoW due to the environmental impact of PoW, and longest-chain PoS represents a fair share of PoS deployments.

The main purpose of this article is to develop a new analytic approach for rigorous settlement guarantees for longest-chain rule protocols in the presence of network delays. While the new approach is somewhat simpler than previous techniques, the chief advantage is that it provides estimates that are both tight enough to directly inform practice and can be explicitly calculated in time polynomial in the relevant parameters. Our new techniques provide improvements over the state-of-the-art settlement bounds for proof-of-work longest-chain protocols [13]; more importantly, they also yield the first concrete settlement bounds for proof-of-stake longest-chain, placing them on equal footing with their proof-of-work counterparts. Finally, our analysis in both cases is the first to apply to the entire security regime: in particular, if longest-chain consensus possesses eventual security for a triple of parameters (r_h, r_a, Δ), our approach provides explicit bounds of security that converge exponentially quickly.

Our techniques and results apply to a wide family of longest-chain protocols, including all proof-of-work protocols following Nakamoto's Bitcoin white paper [21] and all proof-of-stake protocols axiomatized in [16] (such as variants of Ouroboros [2,4,7,17] and Snow White [6]). Deployed systems based on these protocols include Bitcoin [21], Ethereum,[1] Dogecoin, Cardano,[2] Polkadot,[3] and Mina.[4]

[1] https://ethereum.org/, prior to its shift to PoS in September 2022. The analysis also applies to currently deployed Ethereum Classic (ETC) and PoW Ethereum (ETHW) blockchains. In the rest of the paper, we refer to all these three instances together as "PoW-based Ethereum," or simply Ethereum if no confusion can arise.

[2] https://cardano.org/.

[3] https://polkadot.network/.

[4] https://minaprotocol.com/.

Our Techniques. Our analysis provides a family of recurrence relations that determine, for a fixed transcript of the leader-election lotteries during the protocol's execution, a sufficient condition for transaction settlement in any execution with this sequence of lottery outcomes. Coupling this with the stochastic process that governs leader election yields an efficient procedure for computing explicit upper bounds on settlement failure probabilities. An analogous procedure can provide lower bounds on these probabilities which we use to demonstrate the tightness of our upper bounds.

It is most convenient to discuss our approach in the context of recent related works, viz. [8,12,13]. The main difficulty in the analysis arises from accounting for network delays, as honest players may fail to see each other's latest messages (blocks) and end up undermining each other's contributions. Tackling this requires analysis of the complex sequencing of honest and adversarial blocks when a sequence of elected leaders repeatedly fall within Δ time of a previous leader. The combinatorics and resulting stochastic process are particularly difficult during "close races," i.e., when the adversary possesses a private chain that is about as long as the public chain. In such circumstances, honest leaders may be manipulated to contribute to the adversary's (now revealed) chain. Roughly speaking, the articles that settled the security regime [8,12] did so by focusing on the more tractable case where the protocol is *not* in a close race, which is sufficient to characterize the asymptotic behavior of the protocol.

In more detail, [8] focuses on a special type of blocks they call "Nakamoto blocks." The definition of Nakamoto blocks depends on the indefinite future, making them a powerful tool to analyze asymptotic security. But the distribution of Nakamoto block instances is highly complicated (and self-correlated) and thus difficult to tightly estimate, making these appear unsuitable for our goal of exact analysis. Similarly, the analysis in [12], roughly speaking, accounts for the close-race situation by considering a sequence of about Δ^2 back-to-back sequences of Δ-long silence followed by a unique honest lottery success, an event with a constant yet extremely small probability. This is again sufficient for an asymptotic analysis but spoils any chance of obtaining concrete tight and practical bounds. Unfortunately, this looseness in the close-race analysis appears to be a necessary consequence of their approach where the execution is seen as a sequence of steps with potentially significant inter-step interactions.

The recent article [13] that achieved practical settlement estimates for the proof-of-work setting made progress on exactly this issue with a new method of "deferrals." Intuitively, time is divided into periods of Δ and message delivery is restricted to occur at the end of each Δ period. The adversary is allowed to either deliver a message at the end of a period or "defer" its delivery to the end of the next period. This significantly simplifies the analysis as it reduces the large space of adversarial strategies to a single choice per block: whether or not to defer it. Unfortunately, this method of deferrals is not applicable in the proof-of-stake setting because a proof-of-stake adversary can produce as many blocks as it wishes from a single leadership election success. In particular, an optimal deferral strategy may choose to defer *a part of* same-success adversarial

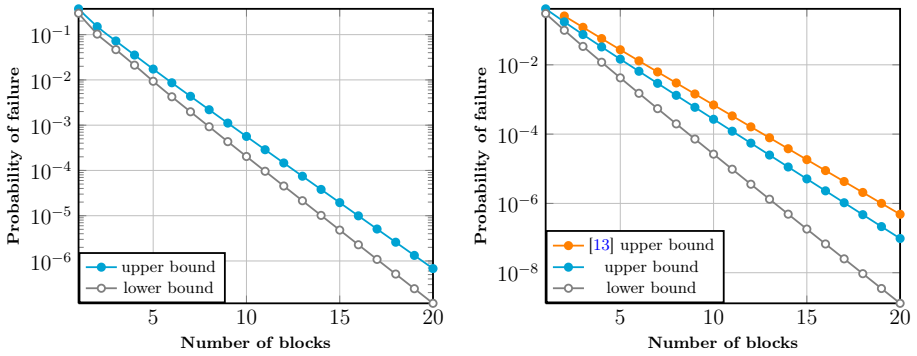

Fig. 1. Left: Cardano (PoS) block-based settlement failure for a 10% adversary and 2 seconds delay. **Right:** Ethereum (PoW) block-based settlement failure for a 10% adversary and 2 s delay, and results from [13] for comparison.

blocks, leading again to a large, complex strategy space that appears analytically intractable.

Our approach in this paper is, in some sense, the opposite of the deferral framework. Conceptually, our key idea is to divide time into judiciously defined periods—called *phases* in the rest of the paper—that are separated by Δ-long intervals of honest silence. Observe that when honest players carry out protocol execution in a given phase, they are aware of all honest messages (blocks) from all previous phases. From the perspective of the previous analysis of [12], this can be viewed as an aggressive expansion of the notion of a "step" so that they are long enough that troublesome inter-step interactions disappear. Indeed, it follows fairly easily that analysis of the protocol can be cleanly decomposed into a phase-by-phase analysis, without complicated interactions between phases. The natural concern with such an approach is that phases themselves are much more complex, both from a combinatorial perspective and in terms of the resulting stochastic process. Our principal contribution is to show that phases can be analyzed with high precision by examining certain *global properties* of the phase that were not available to analyses that operated on a "symbol-by-symbol" (or "step-by-step") level. For example, we show that a particular combinatorial quantity, "minimal honest depth," can capture most of the information necessary to characterize the relevant behavior of honest parties during the phase; this corresponds to the minimal possible maximal-chain growth in the phase over all blocktrees consistent with the phase; furthermore, this quantity gives rise to an analytically manageable stochastic process. These techniques significantly simplify the combinatorial treatment of longest chain rules and, aside from providing the first practical settlement bounds for the proof-of-stake setting, they also deliver improved guarantees in the proof-of-work case.

Example Results. Figure 1 shows some example results for both PoW-based Ethereum and Cardano (the largest longest-chain PoS blockchain at the time of writing). We assume an adversary that controls 10% of the total mining power

or stake, and a network with maximum 2 s delay. The expected block interval in Ethereum and Cardano is around 13 s and 20 s, respectively. Our results for Cardano are within two confirmation blocks of optimality. Specifically, with a confirmation depth of 15 blocks, we can bound the settlement error probability at within 4.811×10^{-6} and 1.943×10^{-5}. Furthermore, the settlement error probability with 13 blocks is at least 2.143×10^{-5} (the lower bound), which is larger than the settlement error upper bound at 15 blocks. Similarly, our results for PoW-based Ethereum are within three or four confirmation blocks of optimality.

Additionally, our results allow us to compare for the first time the settlement speed of longest-chain protocols based on PoW versus PoS. We provide such comparison in Fig. 5 for Cardano's parametrization, quantifying the tradeoff between these two approaches.

2 Preliminaries and Model

Basic Notation. We use $\mathbb{N}$ to denote the set of natural numbers with zero, i.e., $\mathbb{N} = \{0, 1, 2, \ldots\}$. Throughout this paper, we use the symbol $\Delta \in \mathbb{N}$ to denote the maximum delay of a message, expressed in slots. Most of the notions defined below depend on Δ, but we keep this dependence implicit for the sake of lighter notation. When we want to refer to maximum delay expressed in *seconds* rather than slots, in line with previous work we use the symbol Δ_{r}, where "r" stands for "real".

2.1 Modeling Blockchains with Network Delay

Our modeling of the protocol and its execution environment adopts and extends the model from [12,13] and applies to both PoW and PoS. We summarize the model here for completeness.

A longest-chain protocol is carried out by a set of parties of two types: *honest* parties follow the protocol and *adversarial* parties may deviate arbitrarily. The execution timeline is divided into consecutive discrete short time intervals called *slots*. In each slot, each party evaluates a private lottery (implemented for example using a cryptographic hash function for PoW or a verifiable random function [5] for PoS) to determine whether she is eligible to act as a *slot leader* for that slot, which affords her the right to contribute to the ledger by creating block(s). We use a *characteristic string* to indicate a summary of the outcomes of the lottery in each slot.

More concretely, given an alphabet $\Sigma = \mathbb{N} \times \mathbb{N}$, a *characteristic string* $w = w_1 \ldots w_n \in \Sigma^n$ is a sequence of symbols over Σ. Intuitively, each symbol $w_i = (h_i, a_i) \in \Sigma$ indicates that h_i honest parties and a_i adversarial parties were eligible slot leaders for slot i, based on their private lotteries. For a characteristic string $w = w_1 \ldots w_n \in \Sigma^n$ where each $w_i = (h_i, a_i) \in \mathbb{N} \times \mathbb{N}$, we define $\#_{\mathsf{h}}(w) := \sum_{i=1}^{n} h_i$ and similarly $\#_{\mathsf{a}}(w) := \sum_{i=1}^{n} a_i$, i.e., the total number of honest and adversarial slot leaders over a sequence of slots corresponding to w. Moreover,

we sometimes make use of a similar quantity $\#_{[\mathsf{a}]}(w)$ that denotes the number of symbols in w with positive second coordinate, i.e.,

$$\#_{[\mathsf{a}]}(w) := |\{i \in \{1, \ldots, n\} \mid w_i \notin \mathbb{N} \times \{0\}\}| \ .$$

A longest-chain protocol calls for parties to exchange *blockchains*, each of which is an ordered sequence of blocks beginning with a distinguished "genesis block," known to all parties. When an honest party becomes a slot leader, she always creates a single block, and follows the *longest-chain rule* which dictates that she adds her block to the longest blockchain she has observed thus far; she also broadcasts the new block(-chain) to all other parties. When an adversarial party becomes a slot leader, what he is allowed to do differs between PoW and PoS. Intuitively, in PoW an adversarial success allows for creating a single block that extends an arbitrary chain chosen by the adversary, while in PoS an adversarial success can be used to create any number of blocks and hence extend any number of previously existing chains by one block. Naturally, the adversary is not forced to immediately propagate his blocks, and can distribute them strategically.

More formally, let $\mathcal{C}_t$ denote the collection of all blockchains created by the end of slot t and let $H(\mathcal{C}_t)$ denote the subset of all chains in $\mathcal{C}_t$ whose last block was created by an honest party. Set $\mathcal{C}_0 = \{G\}$, where G denotes the unique chain consisting solely of the genesis block. The genesis block is considered "honest"; thus $H(\mathcal{C}_0) = \mathcal{C}_0$. It is convenient to adopt the convention that $\mathcal{C}_{-t} = H(\mathcal{C}_{-t}) = \{G\}$ for any negative integer $-t < 0$. Then the protocol execution proceeds as follows. For each slot $t = 1, 2, \ldots$:

- Initiate $\mathcal{C}_t := \mathcal{C}_{t-1}$ and $H(\mathcal{C}_t) := H(\mathcal{C}_{t-1})$.
- Given $w_t = (h, a)$ the following modifications are applied:
  - The adversary *must* perform the following *honest iteration* exactly h times: select any collection of chains $\mathcal{V}$ for which $H(\mathcal{C}_{t-1-\Delta}) \subseteq \mathcal{V} \subseteq \mathcal{C}_t$. This is the "view" of the honest slot leader, who applies the longest chain rule to $\mathcal{V}$, selects the longest chain $L \in \mathcal{V}$ (resp. $L \in \mathcal{V} \cap \mathcal{C}_{t-1}$ in PoS) where ties are broken by the adversary, and adds a new block to create a new chain L', which is added to $\mathcal{C}_t$ and also $H(\mathcal{C}_t)$.
  - If $a > 0$, the adversary *may* perform the following *adversarial iteration* at most a times for PoW or an arbitrary number of times for PoS: select a single blockchain C from $\mathcal{C}_t$ (in PoS, it must be from $\mathcal{C}_{t-1}$) and add a block to create a new chain C', which is added to $\mathcal{C}_t$. $H(\mathcal{C}_t)$ remains unchanged.

Note that the synchrony assumption is reflected in the description of the honest iteration: the adversary is obligated to deliver all chains produced by honest parties that are Δ slots old, i.e., the set of chains in $H(\mathcal{C}_{t-1-\Delta})$.

Also note that the model grants the adversary to power to break ties in the longest-chain rule. Considering that the adversary selects both the view $\mathcal{V}$ of each honest party and is empowered to break ties, the structure of the resulting sequence of chains (that is, the directed acyclic graph naturally formed by the blocks) is determined entirely by the adversary and the characteristic string.

We make several additional remarks. First, we permit the adversary to have full view of the characteristic string during this process. In reality, a PoS adversary can only predict its own lottery successes, not those of honest parties, while in PoW, neither successes are predictable. Hence our modeling here only makes the adversary stronger. (Looking ahead, this strengthening only affects our upper bounds, as we determine lower bounds via concrete attacks that can be performed by a realistic adversary, see Sect. 5.) Second, we have placed an implicit constraint on the adversary: the only means of producing a new chain is to append a block (containing a proof of a slot leadership) to an existing chain. In practice, this constraint is guaranteed with cryptographic hash functions. Third, we assume that the distribution of slot leaders is impervious to adversarial tampering and, as in previous treatments, is fixed throughout the analysis. This is motivated by the fact that settlement in deployed protocols takes place at much smaller time scales than shifts in mining power or stake distributions. Lastly, the model does not reflect attacks exploiting rational behaviors of parties, such as selfish-mining attacks [10], beyond simply considering such parties corrupt.

2.2 Ledger Consensus

In the context of ledger consensus protocols (also referred to as blockchain [11] or state machine replication [22] protocols), one is usually interested in preserving two properties, *consistency* and *liveness*, formulated in [11,18,22]. Consistency means that once a block (or equivalently, a transaction within it) is considered *settled* by some honest party, then it is present in the currently held chains of all parties, and remains that way forever. In this work we consider the *block-based settlement* rule for longest-chain consensus, where a party considers a block settled if it appears a particular number of blocks deep in the longest chain currently known to that party. Block-based settlement is adopted in practice, and is generally preferable to time-based settlement, as argued in [13].

To describe consistency and liveness concisely under the longest-chain rule, we define the set of Δ-*dominant* chains $\mathcal{D}_t \subseteq \mathcal{C}_t$ in each time step t. The set $\mathcal{D}_t \subseteq \mathcal{C}_t$ is determined entirely by $\mathcal{C}_t$ and $H(\mathcal{C}_{t-1-\Delta})$: namely, $\mathcal{D}_t$ is the set of all chains in $\mathcal{C}_t$ that are at least as long as the longest chain in $H(\mathcal{C}_{t-1-\Delta})$. The intuition behind this definition is that, in a time slot t, it is in principle possible for the adversary to manipulate an honest party into adopting any Δ-dominant chain, as the adversary is only obligated to deliver those chains in $H(\mathcal{C}_{t-1-\Delta})$ and the chains in $\mathcal{D}_t$ are at least as long as those in $H(\mathcal{C}_{t-1-\Delta})$.

Consistency for block-based settlement; with parameter k. A block B that is k blocks deep in some chain in $\mathcal{D}_t$ is contained in every chain $C \in \mathcal{D}_{t'}$ for all $t' \geq t$.

The goal of this paper is to bound (from both above and below) the probability that consistency is violated as a function of the parameter k.

For completeness, we also mention the liveness property [12], though it is not the focus of this paper.

Liveness; with parameter u**.** For any two slots $t_1, t_2 > 0$ with $t_1 + u \leq t_2$, and any chain $C \in \mathcal{D}_{t_2}$, there is a time $t' \in \{t_1, \ldots, t_1 + u\}$ and a chain $C' \in H(\mathcal{C}_{t'}) \setminus H(\mathcal{C}_{t'-1})$ such that C' is a prefix of C.

3 Proof-of-Work Settlement

In this section we first showcase our approach in the more familiar PoW setting, where it provides tighter results than state-of-the-art settlement bounds.

3.1 Proof-of-Work Blocktrees

We formally capture the above protocol dynamics by the combinatorial notion of a *PoW tree*. It is a variant of the "fork" concept first considered for the proof-of-stake case in [2,7,17] and more recently also employed for PoW-analysis [1,12,13]. An example PoW tree is shown in Fig. 2, illustrating several of the concepts defined below.

Definition 1 (PoW tree). *Let $n \in \mathbb{N}$. A* PoW tree *for the string $w \in \Sigma^n$ is a directed, rooted tree $F = (V, E)$ with a pair of functions*

$$\mathsf{I}_{\#} : V \to \{0, \ldots, n\} \qquad \text{and} \qquad \mathsf{I}_{\mathsf{type}} : V \to \{\mathsf{h}, \mathsf{a}\}$$

satisfying the axioms below. Edges are directed "away from" the root so that there is a unique directed path from the root to any vertex. The value $\mathsf{I}_{\#}(v)$ is referred to as the label *of v. The value $\mathsf{I}_{\mathsf{type}}(v)$ is referred to as the* type *of the vertex: when $\mathsf{I}_{\mathsf{type}}(v) = \mathsf{h}$, we say that the vertex is* honest*; otherwise it is* adversarial*.*

(A1) the root $r \in V$ is honest and is the only vertex with label $\mathsf{I}_{\#}(r) = 0$;
(A2) for any pair of honest vertices v, w for which $\mathsf{I}_{\#}(v) + \Delta < \mathsf{I}_{\#}(w)$, $\mathsf{len}(v) < \mathsf{len}(w)$, where $\mathsf{len}()$ denotes the depth of the vertex;
(W3) the sequence of labels $\mathsf{I}_{\#}()$ along any directed path is non-decreasing;
(W4) if $w_i = (h_i, a_i)$ then there are exactly h_i honest vertices and at most a_i adversarial vertices in F with the label i.

We will refer to PoW trees simply as *trees* when the context is clear. Unless explicitly stated otherwise, throughout the paper we reserve the term "tree" for the above structure, as opposed to the underlying graph-theoretic notion.

A PoW tree abstracts a protocol execution with a simple but sufficiently descriptive discrete structure. Its vertices and edges stand for blocks and their connecting hash links (in reverse direction), respectively. The root represents the genesis block, and for each vertex v, $\mathsf{I}_{\#}(v)$ and $\mathsf{len}(v)$ denote the slot in which the corresponding block was created and the block's depth, respectively.

It is easy to see the correspondence between the above axioms and the constraints imposed in the protocol execution. In particular, (A1) corresponds to the trusted nature of the genesis block; (A2) reflects the fact that given sufficient time, as needed for block propagation in the network, an honest party will

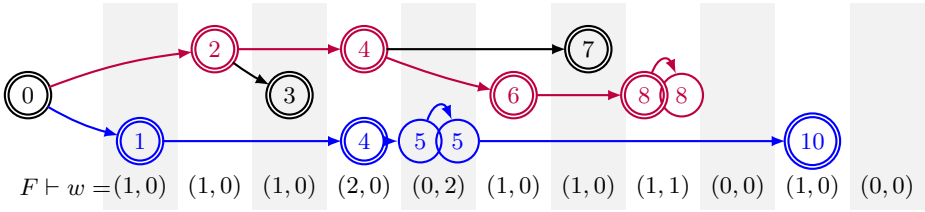

Fig. 2. A PoW tree F for the characteristic string w with $\Delta = 1$. Honest vertices are shown with double-struck boundaries, while adversarial vertices are simple circles. Vertices are labeled with $\mathsf{l}_{\#}(\cdot)$. The tree indicates a k-consistency violation for $k = 4$ —given by the red and blue chains—in a circumstance where the simple private-chain attack does not succeed: in particular, the tree constructs two alternate chains with disjoint suffixes of length 5, while only three adversarial proofs of work are discovered over this period. We remark that $F = F_{\lceil 1}$, since the last symbol of w is $(0, 0)$, and that $\overline{F_{\lceil 1}}$ is obtained by removing the adversarial vertex with label 8. Thus $\mathsf{len}(\overline{F_{\lceil 1}}) = 5$, this maximum length achieved by the blue chain. Note, then, that the two chains indicated in red and blue each have advantage equal to zero, and both are dominant. Considering that these chains share no vertices after the root, they witness $\beta_1(F) \geq 0$ for the tree F and hence $\beta_1(w) \geq 0$ for the characteristic string w. (Color figure online)

take into account the blocks produced by previous honest parties. Axiom (W3) reflects that the blocks' ordering in a chain must be consistent with the order of their creation and finally (W4) reflects that honest players produce exactly one block per PoW success, while the adversary might forgo a block-creation opportunity. Looking ahead, the labeling of the above axioms reflects that while (A1) and (A2) will apply universally to PoW and PoS alike, axioms (W3) and (W4) are specific to PoW, and will be replaced for PoS analysis by appropriate variations (S3) and (S4) in Sect. 4.

Definition 2 (PoW tree notation). *We write $F \vdash^{\boldsymbol{w}} w$ if F is a PoW tree for the characteristic string w. When the setting is clear from context or not germane to the discussion, we drop the superscript and simply write $F \vdash w$. If $F' \vdash w'$ for a prefix w' of w, we say that F' is a* subtree *of F if F contains F' as a consistently-labeled subgraph. A tree $F \vdash w$ is* public *if all leaves are honest. The trivial tree, consisting solely of a root vertex, is considered public. The* public core *of a tree F, denoted $\overline{F}$, is the maximal public subtree of F.*

An individual blockchain constructed during the protocol execution is represented by the notion of a *chain*, defined next.

Definition 3 (Chains). *A path in a tree F originating at the root is called a* chain *(note that a chain does not necessarily terminate at a leaf). As there is a one-to-one correspondence between directed paths from the root and vertices of a tree, we routinely overload notation so that it applies to both chains and vertices. Specifically, we let $\mathsf{len}(T)$ denote the* length *of the chain, equal to the*

number of edges on the path; recall that $\mathsf{len}(v)$ *also denotes the depth of a vertex. We sometimes emphasize the tree from which* v *is drawn by writing* $\mathsf{len}_F(v)$*. We further overload this notation by letting* $\mathsf{len}(F)$ *denote the length of the longest chain in a tree* F*. Likewise, we let* $\mathsf{l}_{\#}(\cdot)$ *apply to chains by defining* $\mathsf{l}_{\#}(T) := \mathsf{l}_{\#}(v)$*, where* v *is the terminal vertex on the chain* T*. We say that a chain is* honest *if the last vertex of the chain is honest. For a vertex* v *in a tree* F*, we denote by* $F(v)$ *the chain in* F *terminating in* v*.*

Definition 4 (Branches). *For an integer* $\ell \geq 1$ *and for two chains* T *and* T' *of a tree* F*, we write* $T \sim_\ell T'$ *if the two chains share a vertex with a label greater than or equal to* ℓ*. The set of all chains* $T' \in F$ *such that* $T \sim_\ell T'$ *is called the* branch *of* T *in* F *and denoted* $\mathsf{B}_F(T; \ell)$*; when* ℓ *can be inferred from context, we write* $\mathsf{B}_F(T)$*.*

Intuitively, $T \sim_\ell T'$ guarantees that the respective blockchains agree on the state of the ledger up to time slot ℓ. Looking ahead, the adversary can make two honest parties disagree on the state of the ledger up to time ℓ only if she makes them hold two blockchains chains $T \not\sim_\ell T'$.

Definition 5 (Tree trimming; dominance). *For a string* $w = w_1 \dots w_n$ *and some* $k \in \mathbb{N}$*, we let* $w_{\lceil k} = w_1 \dots w_{n-k}$ *denote the string obtained by removing the last* k *symbols. For a tree* $F \vdash w_1 \dots w_n$ *we let* $F_{\lceil k} \vdash w_{\lceil k}$ *denote the tree obtained by retaining only those vertices labeled from the set* $\{1, \dots, n-k\}$*. For convenience, we sometimes prefer to emphasize the* remaining *length of the string (resp. tree), and denote by* $w_{\lceil m}$ *and* $F_{\lceil m}$ *the* m*-symbol prefix of* w *and the corresponding tree, formally* $w_{\lceil m} := w_{\lceil n-m}$ *and* $F_{\lceil m} := F_{\lceil n-m}$*. We say that a chain* T *in* F *is* dominant *if* $\mathsf{len}(T) \geq \mathsf{len}(\overline{F_{\lceil \Delta}})$*.*

Observe that honest chains appearing in $F_{\lceil \Delta}$ are those that are necessarily visible to honest players at a round just beyond the last one described by the characteristic string. Correspondingly, the notion of a dominant chain matches the use of this term in Sect. 2.1.

Looking ahead, our approach is to analyze phases that end with Δ consecutive slots with no honest successes. Hence, at the end of each phase the characteristic string has the form wx with $x \in (\{0\} \times \mathbb{N})^\Delta$. We note the following fact.

Fact 1. *For a characteristic string of the form* wx*, where* $x \in (\{0\} \times \mathbb{N})^\Delta$*, any tree* $F \vdash wx$ *has the property that* $\mathsf{len}(\overline{F_{\lceil \Delta}}) = \mathsf{len}(\overline{F})$ *and hence dominance follows simply from* $\mathsf{len}(T) \geq \mathsf{len}(\overline{F})$*.*

Definition 6 (Honest depth h_Δ). *For* $x \in \{0,1\}^*$*, we define* $h_\Delta(x)$ *inductively so that* $h_\Delta(\epsilon) = 0$*,* $h_\Delta(x0) = h_\Delta(x)$*, and* $h_\Delta(x1) = h_\Delta(x_{\lceil \Delta}) + 1$*. We often overload* h_Δ *to apply to strings from* $\Sigma^* = (\mathbb{N} \times \mathbb{N})^*$*, in which case symbols with non-zero first coordinate (i.e., from* $((\mathbb{N} \setminus \{0\}) \times \mathbb{N})^*$ *are counted as* 1*s, while symbols from* $(\{0\} \times \mathbb{N})^*$ *are treated as* 0*s.*

The honest depth $h_\Delta(x)$ of a string x captures the minimum growth of honest blockchains over a period of slots corresponding to x. More concretely, it is the

minimum number of times during x that an honest slot leader must create a block at a higher depth because it is guaranteed to "see" an honest blockchain at one depth lower that was created at least Δ slots earlier.

3.2 PoW Characteristic Quantity: Margin (β_ℓ)

As shown in previous works [1,12,13], the core quantity useful for analyzing PoW longest-chain blockchains is *margin*, defined next.

Definition 7 (PoW Margin β_ℓ). *For a PoW tree $F \vdash^{\boldsymbol{w}} w$, we define the* advantage *of a chain $T \in F$ as*

$$\alpha_F(T) = \mathsf{len}(T) - \mathsf{len}(\overline{F_{\lceil \Delta}}) \,.$$

Observe that $\alpha_F(T) \geq 0$ if and only if T is dominant in F. For $\ell \geq 1$, we define the margin *of a tree F as*

$$\beta_\ell(F) = \max_{\substack{T_h \not\sim_\ell T_a \\ T_h \text{ is dominant}}} \alpha_F(T_a) \,,$$

this maximum extended over all pairs of chains (T_h, T_a) where T_h is dominant and $T_a \not\sim_\ell T_h$. We call the pair (T_h, T_a) the witness chains for F *if the above conditions are satisfied; i.e., T_h is dominant, $T_h \not\sim_\ell T_a$, and $\beta_\ell(F) = \alpha_F(T_a)$. Note that there might exist multiple such pairs in F, but under the condition $\ell \geq 1$ there will always exist at least one such pair, as the trivial chain T_0 containing only the root vertex satisfies $T_0 \not\sim_\ell T$ for any T and $\ell \geq 1$, in particular $T_0 \not\sim_\ell T_0$. For this reason, we will always consider β_ℓ only for $\ell \geq 1$.*

We overload the notation and let

$$\beta_\ell(w) = \max_{F \vdash^{\boldsymbol{w}}_{\sim} w} \beta_\ell(F) \,.$$

We call a PoW tree $F \vdash^{\boldsymbol{w}} w$ a witness tree for w *if $\beta_\ell(w) = \beta_\ell(F)$; again many witness trees may exist for a string w.*

There is a known tight connection between margin and settlement, captured by the following lemma and motivating our effort to upper-bound β_ℓ.

Lemma 1 ([12,13]). *Consider an execution of a PoW blockchain for L slots as described above, resulting in a characteristic string $w = w_1 \dots w_L$. Let B be a block produced in slot $\ell \in [L]$, and let $t > \ell$ be such that B is contained in some chain $C \in \mathcal{D}_t$. If for every $t' \in \{t, \dots, L\}$ we have $\beta_\ell(w_1 \dots w_{t'}) < 0$ then B is contained in every $C' \in \mathcal{D}_{t'}$ for all $t' \in \{t, \dots, L\}$.*

Note that if a tree $F \vdash^{\boldsymbol{w}} w$ has $\beta_\ell(F) < 0$ then all chains T of length at least $\mathsf{len}(\overline{F_{\lceil \Delta}})$ belong to the same branch, which we call the main branch.

Definition 8 (Main branch, PoW). *Let $w \in \Sigma^n$, $\ell \geq 1$, and $F \vdash^{\boldsymbol{w}} w$ such that $\beta_\ell(F) < 0$. The unique branch of F that contains all chains of length at least $\mathsf{len}(\overline{F_{\lceil \Delta}})$ (and possibly other chains) is called the* main branch *of F and denoted $\mathsf{M}^w(F)$.*

3.3 Main PoW Theorem

The goal of Sect. 3 is to provide recurrences that allow us to upper-bound the value of margin $\beta_\ell(w)$ for any PoW characteristic string w. As shown in Lemma 1, this allows us to upper-bound the probability of a settlement violation in any execution with leadership lottery outcomes captured by w.

We approach this challenge by splitting w into consecutive, non-overlapping substrings called *phases*, in a way that ensures the following property:

Phase property: any honest party producing a block in a particular phase is at that time necessarily aware of all honest blocks that have been produced in all previous phases.

To ensure this property, we determine phase boundaries in w so that two consecutive phases are separated by a Δ-long sequence of slots in which no honest successes occur. Notice that this clearly implies the phase property, as any honest block created in phase i will have been delivered to all honest parties within Δ slots, before the beginning of phase $i+1$.

More formally, we devise a recurrence determining the quantity $\beta_\ell(wsxt)$ based on the value $\beta_\ell(ws)$ and the suffix xt, where $w, x \in \Sigma^*$ are arbitrary characteristic strings, while $s, t \in (\{0\} \times \mathbb{N})^\Delta$ represent Δ-long periods with no honest successes. Together with the trivial initial condition $\beta_\ell(\varepsilon) = 0$ for the empty string ε, this gives us a phase-based inductive characterization of β_ℓ, where xt denotes the currently processed phase. Our main result in this section is the following theorem providing such a characterization.

Theorem 1 (The PoW Phase Recurrence). *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^\Delta$ be characteristic strings. We have:*

Margin recurrence. $\beta_\ell(\varepsilon) = 0$. *Furthermore,*

$$\beta_\ell(wsxt) \begin{cases} = \beta_\ell(ws) + \#_{\mathsf{a}}(xt) - h_\Delta(x) & \text{if } \beta_\ell(ws) < -\#_{\mathsf{a}}(xt) \\ & \text{or } \beta_\ell(ws) > h_\Delta(x), \\ \leq \min\{0, \beta_\ell(ws)\} + \#_{\mathsf{a}}(xt) & \text{otherwise.} \end{cases}$$

Crossing zero. *If $|ws| \geq \ell - 1$ and $\beta_\ell(ws) = 0$ then $\beta_\ell(ws(1,0)(0,0)^\Delta) = -1$.*

Hot, cold, and critical regions. We establish the recurrences above over a sequence of lemmas. These lemmas consider β_ℓ in one of the regions that we informally call *hot*, *cold*, and *critical*. A quantity is said to be in the hot region if its value is sufficiently above zero, such that the currently considered phase cannot bring it down to zero. On the other hand, it is said to be in the cold region if it is sufficiently negative so that it won't climb to zero within the current phase. Finally, it is said to be in a critical region if it is close to zero as detailed below.

The critical region corresponds to the situation of a "close race" discussed in Sect. 1. This is the most difficult situation to analyze as special behaviors of the considered quantities (in this case β_ℓ) manifest here: most notably, it is

possible in this region for a new honest success to make things worse for the honest players. For example, consider a situation where an honest player builds a new block B on a chain so that it is exactly one block longer than the best competing chain; now, an additional honest block produced by an honest player that has not seen B can be placed on the competing chain which "neutralizes" this one-block advantage. In contrast, in the hot and cold regions, this second honest block is merely wasted: it does not benefit the honest players but does not hurt either.

With the above discussion, Theorem 1 says that in both the hot and cold regions, β_ℓ exactly follows an ideal recurrence

$$\beta_\ell(wsxt) = \beta_\ell(ws) + \#_{\mathsf{a}}(xt) - h_\Delta(x) \tag{1}$$

where it increases by 1 for each adversarial success, and decreases by 1 whenever the pattern of honest successes enforces an increase in the honest depth h_Δ. In the critical region, $\beta_\ell(wsxt)$ can still be upper-bounded by both $\#_{\mathsf{a}}(xt)$ and $\beta_\ell(ws) + \#_{\mathsf{a}}(xt)$. Intuitively, this means that if $\beta_\ell(ws)$ is "close to zero" from the negative side, then the worst-case behavior observed in the subsequent phase xt is as if the ideal recurrence was applied but xt contained no honest successes, while if $\beta_\ell(ws)$ is "close to zero" from the positive side, $\beta_\ell(wsxt)$ is still upper-bounded by $\#_{\mathsf{a}}(xt)$. Finally, note that these rules by themselves would never permit β_ℓ to descend below zero; for this purpose we establish a separate statement that if $\beta_\ell(ws) = 0$, then a subsequent phase containing only a single success that is honest, brings margin into negative values.

We remark that the exact behavior near zero appears to be quite complicated, in part because there is no longer a clear optimal strategy for the adversary to neutralize honest successes. We identified the simplest and most common scenario, i.e., a single honest success followed by a Δ period of no success, that transitions the quantity from zero to negative. There might be other advanced patterns of honest successes that cannot be neutralized but we treat as thought they can. This is also why we give an upper bound rather than an exact recurrence in the critical region.

3.4 Existing Tools: Tree Compression and the PoW Restructuring Lemma

In our PoW arguments we make use of special honest vertices called *tight* that are, informally speaking, at the minimal depth that the preceding part of the tree allows without violating the axiom (A2). Here we define these vertices formally and summarize several useful properties they have. In particular, in Lemma 3 we show how a PoW tree that has a tight vertex at each possible depth (we call such trees *compressed*) allows for a complex restructuring operation that leads to a lower-bound on the margin of the underlying characteristic string.

Definition 9. *Let $F \vdash w \in \Sigma^n$. An honest vertex v of F is called* tight *if* $\mathsf{len}(v) = \mathsf{len}(\overline{F_{\lfloor \#(v) - \Delta - 1 \rfloor}}) + 1$. *The tree F is said to be* compressed *if, for every depth $0 \le d \le \mathsf{len}(\overline{F})$, there is a tight honest vertex v of depth d.*

We recall two lemmas established in [12]. The first asserts that witness trees may be assumed to be compressed without loss of generality. The second identifies and analyses a restructuring operation in compressed trees. Proofs of both lemmas, adapted to our notation but following those of previous work, appear in the full version [14].

Lemma 2 ([12]). *Let $w \in (\mathbb{N} \times \mathbb{N})^*$ and $s \in (\{0\} \times \mathbb{N})^\Delta$. Then there exists a witness tree $F \vdash^{\mathsf{W}} ws$ that is compressed.*

Lemma 3 (Restructuring lemma, [12]). *Let $\ell \geq 1$, let $w \in \Sigma^*$ be a characteristic string and $F \vdash^{\mathsf{W}} w$ be a compressed PoW tree for w; let $T_1 \not\sim_\ell T_2$ be arbitrary chains in F. For $i \in \{1, 2\}$, let v_i be an honest vertex on T_i and let A_i denote the set of all adversarial vertices on T_i deeper than v_i. If $\mathsf{l}_{\#}(v_1) \leq \mathsf{l}_{\#}(v_2)$ then*

$$\beta_\ell(w) \geq \alpha_F(v_1) + |A_1 \cup A_2| \,.$$

3.5 Outside of the Critical Region

We establish the ideal recurrence (1) outside of the critical region in a sequence of three lemmas: first, Lemma 4 shows that the recurrence gives a lower bound for β_ℓ, and then Lemmas 5 and 6 show that it is also an upper bound in the cold and the hot region, respectively.

Lemma 4 (Lower bound). *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^\Delta$ be characteristic strings. Then*

$$\beta_\ell(wsxt) \geq \beta_\ell(ws) + \#_{\mathsf{a}}(xt) - h_\Delta(x) \,.$$

Proof. Let F be a witness PoW tree $F \vdash^{\mathsf{W}} ws$, and let (T_h, T_a) be a pair of witness chains in F, i.e., $T_h \not\sim_\ell T_a$, T_h is dominant in F, and $\alpha_F(T_a) = \beta_\ell(F) = \beta_\ell(ws)$.

We construct a tree $F' \vdash^{\mathsf{W}} wsxt$ such that $\beta_\ell(F') = \beta_\ell(F) + \#_{\mathsf{a}}(xt) - h_\Delta(x)$. Namely, we add $\#_{\mathsf{a}}(xt) + \#_{\mathsf{h}}(x)$ new vertices to F in two steps. First, we extend T_a by a path consisting of $\#_{\mathsf{a}}(xt)$ adversarial vertices that we label consistently with xt to satisfy axiom (W4), call the resulting chain T'_a. Second, we also add $\#_{\mathsf{h}}(x)$ honest vertices that form a subtree rooted in the terminating vertex of T_h, where each of these honest vertices is always put at the minimal depth allowed by axiom (A2), and labeling them consistently with x to again satisfy axiom (W4). Let T'_h denote a chain terminating in some maximum-depth newly added honest vertex. The resulting tree (call it F') is indeed a PoW tree: it is easy to observe that all axioms of a PoW tree are satisfied by construction. Note that $\mathsf{len}_{F'}(T'_a) = \mathsf{len}_F(T_a) + \#_{\mathsf{a}}(xt)$, $\mathsf{len}_{F'}(T'_h) = \mathsf{len}_F(T_h) + h_\Delta(x)$, and we have $T'_h \not\sim_\ell T'_a$ as these chains share no new vertices. Finally, T'_h is clearly dominant, and hence the pair (T'_h, T'_a) witnesses $\beta_\ell(F') = \beta_\ell(F) + \#_{\mathsf{a}}(xt) - h_\Delta(x)$ as desired. □

Lemma 5 (Cold region). *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^\Delta$ be characteristic strings. If $\beta_\ell(ws) < -\#_{\mathsf{a}}(xt)$ then*

$$\beta_\ell(wsxt) \leq \beta_\ell(ws) + \#_{\mathsf{a}}(xt) - h_\Delta(x) \,.$$

The proof of Lemma 5 is an adaptation of the proof of Lemma 8 from [12] to our setting. Note that our new approach of processing the characteristic string by phases allows for a stronger statement: the ideal recurrence is shown to hold closer to the critical region. At the same time, the proof becomes simpler.

Proof. Let $w' := wsxt$ and let F' be a witness PoW tree $F' \vdash^{\mathsf{W}} w'$; let (T'_h, T'_a) be a pair of witness chains in F' such that $\mathsf{len}(T'_h) = \mathsf{len}(\overline{F'_{\lceil \Delta}})$. Furthermore, let $F := F'_{|ws|} \vdash^{\mathsf{W}} ws$ and define $T_h := (T'_h)_{|ws|}$ and $T_a := (T'_a)_{|ws|}$, i.e., T_h and T_a are the restrictions of T'_h and T'_a to vertices with labels at most $|ws|$; we have $T_h, T_a \in F$ by definition of F. Note that, as s, t contain no honest successes, we have $\overline{F_{\lceil \Delta}} = \overline{F}$ and $\overline{F'_{\lceil \Delta}} = \overline{F'}$, and

$$\mathsf{len}(\overline{F'_{\lceil \Delta}}) \geq \mathsf{len}(\overline{F_{\lceil \Delta}}) + h_\Delta(x) \; . \tag{2}$$

By our assumption of negative $\beta_\ell(ws)$, there is a well-defined main branch $\mathsf{M}^{\mathsf{W}}(F)$. We first establish that, intuitively speaking, any chains in F outside of $\mathsf{M}^{\mathsf{W}}(F)$ are, after ws, extended by adversarial vertices only.

Claim. Consider any chain $T \in F$ such that $T \notin \mathsf{M}^{\mathsf{W}}(F)$ and any $T' \in F'$ that extends T in F' so that $T = T'_{|ws|}$. Then the set of vertices $T' \setminus T$ contains no honest vertices.

To see this, observe that any honest vertex in F' with label greater than $|ws|$ must have depth at least $\mathsf{len}(\overline{F_{\lceil \Delta}}) + 1 = \mathsf{len}(\overline{F}) + 1$ by axiom (A2), hence all vertices in $T' \setminus T$ with depth at most $\mathsf{len}(\overline{F})$ must be adversarial. However, $\mathsf{len}(T) + \#_{\mathsf{a}}(xt) < \mathsf{len}(\overline{F})$. To see this, note that we have $\alpha_F(T) \leq \beta_\ell(ws)$ as $T \notin \mathsf{M}^{\mathsf{W}}(F)$ and hence again there exists some dominant chain in $\mathsf{M}^{\mathsf{W}}(F)$ that forms a witness pair with T. Moreover, $\beta_\ell(ws) < -\#_{\mathsf{a}}(xt)$ by assumption, and this together implies $\mathsf{len}(T) + \#_{\mathsf{a}}(xt) < \mathsf{len}(\overline{F})$ and hence $\mathsf{len}(T') < \mathsf{len}(\overline{F})$. This already shows that there are no honest vertices in $T' \setminus T$ and establishes Claim 3.5.

We now argue that $T_h \in \mathsf{M}^{\mathsf{W}}(F)$. Towards contradiction, assume that $T_h \notin \mathsf{M}^{\mathsf{W}}(F)$. Then Claim 3.5 applies to T_h and $T'_h \setminus T_h$ contains no honest vertices, hence

$$\mathsf{len}(T'_h) \leq \mathsf{len}(T_h) + \#_{\mathsf{a}}(xt) \; . \tag{3}$$

However, by assumption $\mathsf{len}(T_h) - \mathsf{len}(\overline{F}) = \alpha_F(T_h) \leq \beta_\ell(ws) < -\#_a(xt)$, where the first inequality holds as $T_h \notin \mathsf{M}^{\mathsf{W}}(F)$ and hence there exists some dominant chain in $\mathsf{M}^{\mathsf{W}}(F)$ that forms a witness pair with T_h. Hence $\mathsf{len}(T_h) < \mathsf{len}(\overline{F}) - \#_{\mathsf{a}}(xt)$, and using equations (3) and (2) gives us $\mathsf{len}(T'_h) < \mathsf{len}(\overline{F}) \leq \mathsf{len}(\overline{F'})$, a contradiction with the definition of T'_h. Therefore, $T_h \in \mathsf{M}^{\mathsf{W}}(F)$.

Since $T'_h \not\sim_\ell T'_a$, it also follows that $T_h \not\sim_\ell T_a$, and at most one of these chains belongs to $\mathsf{M}^{\mathsf{W}}(F)$, hence we have $T_a \notin \mathsf{M}^{\mathsf{W}}(F)$. By Claim 3.5, $T'_a \setminus T_a$ contains no honest vertices. Hence we have $\mathsf{len}(T'_a) \leq \mathsf{len}(T_a) + \#_{\mathsf{a}}(xt)$ and we can combine this with Eq. (2) to get

$$\begin{aligned} \beta_\ell(ws) \geq \alpha_F(T_a) = \mathsf{len}(T_a) - \mathsf{len}(\overline{F}) &\geq \mathsf{len}(T'_a) - \#_{\mathsf{a}}(xt) - \mathsf{len}(\overline{F'}) + h_\Delta(x) \\ &= \alpha_{F'}(T'_a) - \#_{\mathsf{a}}(xt) + h_\Delta(x) = \beta_\ell(w') - \#_{\mathsf{a}}(xt) + h_\Delta(x) \; , \end{aligned}$$

where the first inequality is again justified by $T_a \notin \mathsf{M}^{\mathsf{W}}(F)$. This concludes the proof of Lemma 5. □

Lemma 6 (Hot region). *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^{\Delta}$ be characteristic strings. If $\beta_\ell(ws) > h_\Delta(x)$ then*

$$\beta_\ell(wsxt) \leq \beta_\ell(ws) + \#_{\mathsf{a}}(xt) - h_\Delta(x) \,. \tag{4}$$

The proof of the above lemma employs as a crucial ingredient the tree compression concept and the restructuring lemma that we recalled in Sect. 3.4.

Proof of Lemma 6. As in the cold case, the proof begins with a witness tree F' for $wsxt$ and shows how to construct a tree $F^* \vdash ws$ for which $\beta_\ell(F^*) \geq \beta_\ell(wsxt) - \#_{\mathsf{a}}(xt) + h_\Delta(x)$; this completes the theorem as $\beta_\ell(ws) \geq \beta_\ell(F^*)$. To set down notation, define $F' \vdash wsxt$ to be a compressed witness tree with witness chains (T'_h, T'_a); we then consider the restriction $F \vdash ws$ of F' to the string ws and, in particular, the restrictions (T_h, T_a) of the witness chains (T'_h, T'_a) to F. To prepare for the main argument, we establish a few straightforward properties of these two trees. First, observe that the inequality

$$\mathsf{len}(\overline{F'_{\lceil \Delta}}) = \mathsf{len}(\overline{F'}) \geq \mathsf{len}(\overline{F}_{\lceil \Delta}) + h_\Delta(x) = \mathsf{len}(\overline{F}) + h_\Delta(x) \tag{5}$$

follows immediately from Fact 1, tree axiom (A2), the definition of honest height, and the fact that s and t contain no honest successes. We then establish that there are no honest vertices on T'_a with label exceeding $|ws|$; in other words, there are no honest vertices in $T'_a \setminus T_a$. Towards a contradiction, assume that there is an honest vertex in $T'_a \setminus T_a$ and let v'_{a} be such an honest vertex with maximum label (and hence maximum depth). Since $\mathsf{l}_{\#}(v'_{\mathsf{a}}) > |ws|$, all vertices u on T'_a with $\mathsf{len}(u) > \mathsf{len}(v'_{\mathsf{a}})$ also have $\mathsf{l}_{\#}(u) > \mathsf{l}_{\#}(v'_{\mathsf{a}}) > |ws|$ and, by maximality of v'_{a}, all these vertices are adversarial; hence there are at most $\#_{\mathsf{a}}(xt)$ subsequent vertices (on T'_h) by axiom (W4). However, as v'_{a} is honest we also have $\mathsf{len}(v'_{\mathsf{a}}) \leq \mathsf{len}(\overline{F'})$. Combining these, we conclude $\beta_\ell(wsxt) = \mathsf{len}(T'_a) - \mathsf{len}(\overline{F'}) \leq \mathsf{len}(T'_a) - \mathsf{len}(v'_{\mathsf{a}}) \leq \#_{\mathsf{a}}(xt)$. Combining this with the assumption $\beta_\ell(ws) > h_\Delta(x)$ yields a direct contradiction to Lemma 4, which asserts that $\beta_\ell(wsxt) \geq \beta_\ell(ws) + \#_{\mathsf{a}}(xt) - h_\Delta(x)$. We conclude that there are no honest vertices on $T'_a \setminus T_a$ and, in particular, that $\mathsf{len}(T'_a) - \mathsf{len}(T_a) \leq \#_{\mathsf{a}}(xt)$.

The last honest vertices on the chains T_h and T_a play a central role in the remainder of the analysis; these we denote v_h and v_a, respectively. We handle the two cases $\mathsf{l}_{\#}(v_h) \geq \mathsf{l}_{\#}(v_a)$ and $\mathsf{l}_{\#}(v_h) < \mathsf{l}_{\#}(v_a)$ separately, in either setting concluding the argument with an application of Lemma 3 to a vertex with minimal label.

The case $\mathsf{l}_{\#}(v_h) < \mathsf{l}_{\#}(v_a)$. We define the sets A_a and A'_a to consist of the adversarial vertices appearing after v_a on T_a in F and F', respectively; thus $A_a \subset A'_a$. We likewise define A_h and A'_h for the chain T_h and vertex v_h.

We first establish that

$$\mathsf{len}(v_a) \leq \mathsf{len}(v_h) + |A'_h| \,. \tag{6}$$

Recalling that T'_h is dominant, $\mathsf{len}(v_a) \le \mathsf{len}(\overline{F}) \le \mathsf{len}(\overline{F'}) = \mathsf{len}(\overline{F'_{\lceil \Delta}}) = \mathsf{len}(T'_h)$. In the case when all vertices of T'_h after v_h are adversarial, the inequality (6) follows immediately because $\mathsf{len}(T'_h) = \mathsf{len}(v_h) + |A'_h|$. Otherwise, there is a first honest vertex v'_h on T'_h that appears after v_h; by definition, this vertex does not lie in F and is labeled by an index in x. Considering that the quiet region s lies between the labels for v_a and v'_h, we must have $\mathsf{len}(v_a) < \mathsf{len}(v'_h)$ by axiom (A2). Combining this with the fact that $\mathsf{len}(v'_h) \le \mathsf{len}(v_h) + |A'_h| + 1$, the inequality (6) follows. We may then conclude that

$$\begin{aligned} \beta_\ell(wsxt) = \alpha_{F'}(T'_a) &= \mathsf{len}(T'_a) - \mathsf{len}(\overline{F'_{\lceil \Delta}}) = \mathsf{len}(T'_a) - \mathsf{len}(\overline{F'}) \\ &= \mathsf{len}(v_a) + |A'_a| - \mathsf{len}(\overline{F'}) \le \mathsf{len}(v_h) + |A'_a| + |A'_h| - \mathsf{len}(\overline{F'})\,. \end{aligned} \tag{7}$$

Now we invoke Lemma 3 with chains T_h, T_a and vertices v_h, v_a in F. By assumption $\mathsf{l}_{\#}(v_h) < \mathsf{l}_{\#}(v_a)$, and hence we obtain

$$\beta_\ell(ws) \ge \alpha_F(v_h) + |A_a \cup A_h| = \mathsf{len}(v_h) + |A_a| + |A_h| - \mathsf{len}(\overline{F}) \tag{8}$$

using the definition of α_F, Fact 1, and the observation that $\mathsf{l}_{\#}(v_h) < \mathsf{l}_{\#}(v_a)$ implies $v_h \ne v_a$ and together with the definition of v_h, v_a this means that $A_a \cap A_h = \emptyset$ and $|A_a \cup A_h| = |A_a| + |A_h|$. Combining (7) and (8), we conclude that

$$\begin{aligned} \beta_\ell(wsxt) - \beta_\ell(ws) &\le |A'_a| - |A_a| + |A'_h| - |A_h| - \big(\mathsf{len}(\overline{F'}) - \mathsf{len}(\overline{F})\big) \\ &\le \#_{\mathsf{a}}(xt) - h_\Delta(x) \end{aligned}$$

as desired. The last inequality follows from (5) and the fact that there are no more than $\#_{\mathsf{a}}(xt)$ adversarial vertices in F' that do not lie in F.

The case $\mathsf{l}_{\#}(v_h) \ge \mathsf{l}_{\#}(v_a)$. We remark that the tree F is compressed. To see this, note that any honest vertex v of F' labeled from the suffix xt must have height strictly larger than $\mathsf{len}(\overline{F_{\lceil \Delta}})$ by axiom (A2); on the other hand, in light of Fact 1 $\mathsf{len}(\overline{F_{\lceil \Delta}}) = \mathsf{len}(\overline{F})$ since ws ends with a quiet period and it follows that the removal of the honest vertices labeled by xs does not affect those of depth at most $\mathsf{len}(\overline{F})$. In particular, F still has an honest vertex of each relevant height and is compressed.

We now invoke Lemma 3 with the chains T_a, T_h and vertices v_a, v_h in F. Since $\mathsf{l}_{\#}(v_a) \le \mathsf{l}_{\#}(v_h)$, we obtain:

$$\begin{aligned} \beta_\ell(ws) &\ge \alpha_F(v_a) + |A_a| = \alpha_F(T_a) = \mathsf{len}_F(T_a) - \mathsf{len}(\overline{F_{\lceil \Delta}}) \\ &\ge (\mathsf{len}_{F'}(T'_a) - \#_{\mathsf{a}}(xt)) - \big(\mathsf{len}(\overline{F'}) - h_\Delta(x)\big) \qquad (9) \\ &= \alpha_{F'}(T'_a) - \#_{\mathsf{a}}(xt) + h_\Delta(x) = \beta_\ell(wsxt) - \#_{\mathsf{a}}(xt) + h_\Delta(x)\,, \end{aligned}$$

where the inequality in line (9) follows from (5). This concludes the proof. □

3.6 The Critical Region

Finally, it remains to tackle the behavior of β_ℓ in the critical region. We establish the two upper bounds mentioned in Sect. 3.3 in Lemmas 7 and 8, and show the crossing-zero property in Lemma 9.

Lemma 7. *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^{\Delta}$ be characteristic strings. Then*

$$\beta_\ell(wsxt) \leq \beta_\ell(ws) + \#_{\mathsf{a}}(xt) \ .$$

Proof. As before, let $w' := wsxt$ and let F' be a witness PoW tree $F' \vdash^{\mathbf{w}} w'$; let (T'_h, T'_a) be a pair of witness chains in F' such that $\mathsf{len}(T'_h) = \mathsf{len}(\overline{F'_{\lceil \Delta}}) = \mathsf{len}(\overline{F'})$ (cf. Fact 1). Furthermore, let $F := F'_{|ws|\rceil} \vdash^{\mathbf{w}} ws$ and define $T_h := (T'_h)_{|ws|\rceil}$ and $T_a := (T'_a)_{|ws|\rceil}$, i.e., T_h and T_a are the restrictions of T'_h and T'_a to vertices with labels at most $|ws|$; we have $T_h, T_a \in F$ by definition of F. Moreover, let T_H be a chain in F such that $\mathsf{len}(T_H) = \mathsf{len}(\overline{F})$.

If $T_H \not\sim_\ell T_a$, we have $\beta_\ell(ws) \geq \alpha_F(T_a)$. Looking at the set of vertices $T'_a \setminus T_a$ in F', let $\mathcal{H} \subseteq T'_a \setminus T_a$ denote the set of those vertices $v \in T'_a \setminus T_a$ that satisfy $\mathsf{len}(\overline{F}) < \mathsf{len}(v) \leq \mathsf{len}(\overline{F'})$. Intuitively, $\mathcal{H}$ covers the vertices in the extension $T'_a \setminus T_a$ that have depths in which F' might contain honest vertices with labels greater than $|ws|$. Observe that therefore $|\mathcal{H}| \leq \mathsf{len}(\overline{F'}) - \mathsf{len}(\overline{F})$ and all vertices in $(T'_a \setminus T_a) \setminus \mathcal{H}$ are adversarial. This gives us

$$\begin{aligned}
\beta_\ell(w') - \beta_\ell(ws) &\leq \big(\mathsf{len}(T'_a) - \mathsf{len}(\overline{F'})\big) - \big(\mathsf{len}(T_a) - \mathsf{len}(\overline{F})\big) \\
&= \big(\mathsf{len}(T'_a) - \mathsf{len}(T_a)\big) - \big(\mathsf{len}(\overline{F'}) - \mathsf{len}(\overline{F})\big) \\
&\leq \big(|\mathcal{H}| + \#_{\mathsf{a}}(xt)\big) - \big(\mathsf{len}(\overline{F'}) - \mathsf{len}(\overline{F})\big) \leq \#_{\mathsf{a}}(xt)
\end{aligned}$$

as desired.

On the other hand, if $T_H \sim_\ell T_a$ then we have $T_H \not\sim_\ell T_h$, and

$$\begin{aligned}
\beta_\ell(w') - \beta_\ell(ws) &\leq \big(\mathsf{len}(T'_a) - \mathsf{len}(\overline{F'})\big) - \big(\mathsf{len}(T_h) - \mathsf{len}(T_H)\big) \\
&\leq \big(\mathsf{len}(T'_a) - \mathsf{len}(\overline{F'})\big) + \big(\mathsf{len}(T_H) - \mathsf{len}(T_h)\big) \ .
\end{aligned}$$

Observe that if $\mathsf{len}(T'_a) - \mathsf{len}(\overline{F'}) > 0$, all vertices on T'_a with depth greater than $\mathsf{len}(\overline{F'})$ must be adversarial by definition of $\overline{F'}$. Similarly, if $\mathsf{len}(T_H) - \mathsf{len}(T_h) > 0$, then all vertices on T'_h with depth d satisfying $\mathsf{len}(T_h) \leq d \leq \mathsf{len}(\overline{F}) = \mathsf{len}(T_H)$ must be adversarial, as the minimum depth at which honest vertices labeled from x can appear is $\mathsf{len}(\overline{F}) + 1$ due to axiom (A2) and the fact that s contains no honest successes. Putting these two facts together, we get $\mathsf{len}(T'_a) - \mathsf{len}(\overline{F'}) + \mathsf{len}(T_H) - \mathsf{len}(T_h) \leq \#_{\mathsf{a}}(xt)$, concluding the proof also for this case. □

Lemma 8. *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^{\Delta}$ be characteristic strings. If $\beta_\ell(ws) \leq h_\Delta(x)$ then*

$$\beta_\ell(wsxt) \leq \#_{\mathsf{a}}(xt) \ .$$

Proof (sketch). The lemma can be established by an argument identical to the proof of Lemma 6, with a single exception.

Using the notation from that proof, in this case we do not prove that $T'_a \setminus T_a$ contains no honest vertices as before. Instead, we observe that if there actually *is* an honest vertex on $T'_a \setminus T_a$, then by definition of T_a this vertex has a label exceeding $|ws|$, and hence the deepest honest vertex in T'_a can only be followed

by at most $\#_{\mathsf{a}}(xt)$ adversarial vertices. This directly implies $\beta_\ell(wsxt) \le \#_{\mathsf{a}}(xt)$ and proves the lemma for this case.

Otherwise we again have no honest vertices on $T'_a \setminus T_a$, and the rest of the argument is identical to the proof of Lemma 6 as it never again invokes the assumption about $\beta_\ell(ws)$. The argument gives us $\beta_\ell(wsxt) \le \beta_\ell(ws) + \#_{\mathsf{a}}(xt) - h_\Delta(x)$, and since here we assume $\beta_\ell(ws) \le h_\Delta(x)$ we can conclude $\beta_\ell(wsxt) \le \#_{\mathsf{a}}(xt)$ as desired. □

Lemma 9. *Let $\ell \ge 1$, let $w, x \in (\mathbb{N}\times\mathbb{N})^*$ and $s, t \in (\{0\}\times\mathbb{N})^\Delta$ be characteristic strings. If $|ws| \ge \ell - 1$ and $\beta_\ell(ws) = 0$ then*

$$\beta_\ell(ws(1,0)(0,0)^\Delta) \le -1 \,.$$

Proof. Let $w' := ws(1,0)(0,0)^\Delta$ and towards a contradiction, assume $\beta_\ell(w') \ge 0$. By definition of β_ℓ, there exists a witness PoW tree $F' \vdash^{\mathsf{w}} w'$ and two chains T'_1, T'_2 in F' such that $T'_1 \not\sim_\ell T'_2$, $\alpha_{F'}(T'_1) = 0$, $\alpha_{F'}(T'_2) \ge 0$, and T'_1 terminates with the unique (and honest) vertex with $\mathsf{l}_{\#}(v'_{\mathsf{h}}) = |ws| + 1$ prescribed by w'; let us call this vertex v'_{h}. (Note that (T'_1, T'_2) are not necessarily witness chains as we don't ask for $\alpha_{F'}(T'_2) = \beta_\ell(F')$, this allows us to require that T'_1 terminates in v'_{h} without loss of generality.) Denote $F := F'_{|ws|\rceil} \vdash^{\mathsf{w}} ws$ and note that F is in fact obtained from F' by just removing v'_{h}. As $\#_{\mathsf{h}}(s) = 0$ and $|s| = \Delta$, by axiom (A2) we have $\mathsf{len}_{F'}(v'_{\mathsf{h}}) > \mathsf{len}(\overline{F})$ and hence $\mathsf{len}(\overline{F'_{\lceil\Delta}}) = \mathsf{len}(\overline{F'}) > \mathsf{len}(\overline{F}) = \mathsf{len}(\overline{F_{\lceil\Delta}})$. Let $T_1 := (T'_1)_{|ws|\rceil}$. Note that as $T'_1 \not\sim_\ell T'_2$ and $|ws(1,0)| \ge \ell$, we must have $v'_{\mathsf{h}} \notin T'_2$ and $T'_1 \ne T'_2$, hence T'_2 also exists in F and $T_1 \not\sim_\ell T'_2$ in F. As $\mathsf{len}(\overline{F_{\lceil\Delta}}) < \mathsf{len}(\overline{F'_{\lceil\Delta}})$, we have $\alpha_F(T_1) \ge 0$ and $\alpha_F(T'_2) > 0$, resulting in $\beta_\ell(ws) \ge \beta_\ell(F) > 0$, a contradiction. □

4 Proof-of-Stake Settlement

4.1 Proof-of-Stake Blocktrees

The execution of a longest-chain PoS protocol is in principle similar to the execution of its PoW counterpart, with two notable differences, described in passing already in Sect. 2.1. Most importantly, the effect of an adversarial lottery success is different: it allows the adversary to create an arbitrary number of blocks for the corresponding slot, while in PoW a single lottery success only leads to a single block. Second, a valid PoS chain may only contain at most one block from any given slot, while in PoW the adversary can in principle use multiple adversarial blocks from the same slot to extend the same chain.

To model this behavior, we consider the same alphabet $\Sigma = \mathbb{N} \times \mathbb{N}$ for characteristic strings also in the PoS case. However, the notion of a tree needs to be adapted to capture the above differences. The resulting notion of a PoS tree conceptually matches the 'fork' notion from previous PoS works [2,7].

Definition 10 (PoS tree). *A* PoS tree *is defined exactly as a PoW tree (cf. Definition 1), except that axioms (W3) and (W4) are replaced by the following axioms:*

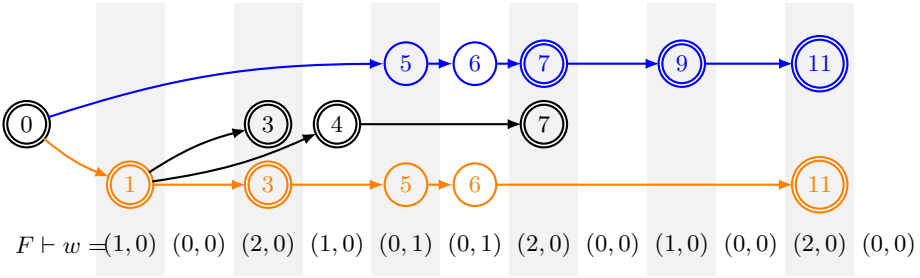

Fig. 3. A public PoS tree F for the characteristic string w with $\Delta = 1$, using the same graphical representation as Fig 2. The tree indicates a successful double spend attack given by the orange and blue chains and highlights a notable feature of the proof-of-stake setting: the adversary's ability to play multiple blocks in slots 5 and 6 permits a double spend attack in circumstances where there would be no attack in the proof-of-work case. We remark that $F = \overline{F} = F_{\lceil 1} = \overline{F_{\lceil 1}}$, since all leaves of F are honest and the last symbol of w is $(0, 0)$. Clearly $\mathsf{len}(\overline{F}_{\lceil 1}) = 5$. The two chains indicated in red and blue each have advantage equal to zero, and both are dominant. Considering that these chains share no vertices after the root, they witness $\mu_1(F) \geq 0$ for the tree F and hence for the characteristic string w. (Color figure online)

(S3) the sequence of labels $\mathsf{l}_{\#}()$ *along any directed path is strictly increasing;*
(S4) if $w_i = (h_i, a_i)$ *then there are exactly* h_i *honest vertices of* F *with the label* i *and if the number of adversarial vertices with label* i *is nonzero then* $a_i > 0$.

The two changes to tree axioms formally capture the two differences from the PoW setting, listed above. Note that the notation laid out in Definitions 2–6, as well as Fact 1, immediately apply also to the PoS case. An example PoS tree is depicted in Fig. 3.

4.2 PoS Characteristic Quantities: Reach (ρ) and Margin (μ_ℓ)

As established in an existing line of work on PoS protocols going back to [17], the dynamics of a PoS longest-chain protocol can be captured by a pair of quantities called *reach* and *margin*. Note that this is in contrast to the PoW case where a single quantity is sufficient (see Sect. 3), and represents the additional complexity in analyzing the PoS case. We recall the notions of reach and margin as generalized in [3], and adapt them to our notation. For consistency with these works, we refer to μ_ℓ as *margin*; no confusion should arise as it is always clear whether we consider the PoW or PoS margin.

Definition 11 (PoS reach, margin). *For a public PoS tree* $F \vdash^{\boldsymbol{s}} w$*, we define the* advantage *of a chain* $T \in F$ *exactly as in the PoW case in Definition 7. We define the* reserve *of a chain* T *in* F *to be the number of adversarial indices appearing in* w *after the last index in* T*; specifically, if* v *is the terminal vertex of* T*, we define*

$$\text{reserve}_F(T) := |\{i > \mathsf{l}_{\#}(v) \mid w_i = (h_i, a_i) \wedge a_i > 0\}| \ .$$

We then define

$$\text{reach}_F(T) := \alpha_F(T) + \text{reserve}_F(T) \ ,$$

$$\rho(F) := \max_{T \text{ in } F} \text{reach}_F(T) \qquad \textit{and} \qquad \rho(w) := \max_{\substack{F \vdash^{\mathbf{s}} w \\ F \text{ public}}} \rho(F) \ .$$

For a given w, we sometimes refer to a tree F and a chain T maximizing the above expressions as a witness tree *and a* witness chain*, respectively; note that these are not necessarily unique.*

For a public PoS tree $F \vdash^{\mathbf{s}} w$ we define the margin *of F, denoted $\mu_\ell(F)$, to be the "penultimate" reach taken over chains T_1, T_2 of F such that $T_1 \not\sim_\ell T_2$:*

$$\mu_\ell(F) := \max_{T_1 \not\sim_\ell T_2} \Big(\min\{\text{reach}_F(T_1), \text{reach}_F(T_2)\}\Big) \ .$$

There might exist multiple such pairs in F, but under the condition $\ell \geq 1$ there will always exist at least one such pair, as the trivial chain T_0 containing only the root vertex satisfies $T_0 \not\sim_\ell T$ for any T and $\ell \geq 1$, in particular $T_0 \not\sim_\ell T_0$. For this reason, we will always consider $\mu_\ell(\cdot)$ only for $\ell \geq 1$. We again overload the notation by defining

$$\mu_\ell(w) := \max_{\substack{F \vdash^{\mathbf{s}} w \\ F \text{ public}}} \mu_\ell(F) \ .$$

We use the terms witness tree and witness chains analogously also in the case of margin, it will be always clear from the context whether we are referring to witnesses with respect to reach or margin.

Intuitively, there is again a natural connection between margin and settlement: if w is a characteristic string capturing the execution of the PoS blockchain up to some current time t, and $\mu_\ell(w) < 0$ for some $\ell < t$, then any tree $F \vdash^{\mathbf{s}} w$ that resulted from the execution has $\mu_\ell(F) < 0$ and hence does not allow the adversary to make any honest party at time t adopt a blockchain that would not agree with its current chain up to the index ℓ. In other words, all chains with non-negative reach share their prefix up to slot ℓ, i.e., belong to the same branch. This connection was formally established for PoS in [7,17]; we summarize it for our setting in the following lemma. This will motivate our effort to upper-bound μ_ℓ.

Lemma 10 ([7,17]). *Consider an execution of a PoS blockchain for L slots as described above, resulting in a characteristic string $w = w_1 \dots w_L$. Let B be a block produced in slot $\ell \in [L]$, and let $t > \ell$ be such that B is contained in some chain $C \in \mathcal{D}_t$. If for every $t' \in \{t, \dots, L\}$ we have $\mu_\ell(w_1 \dots w_{t'}) < 0$ then B is contained in every $C' \in \mathcal{D}_{t'}$ for all $t' \in \{t, \dots, L\}$.*

Similarly as before, if a PoS-tree $F \vdash^{\mathbf{s}} w$ has $\mu_\ell(F) < 0$ then all chains T with $\text{reach}_F(T) \geq 0$ at least $\mathsf{len}(\overline{F_{\lceil \Delta}})$ belong to the same branch. This justifies the following definition.

Definition 12 (Main branch, PoS). *Let $w \in \Sigma^n$, $\ell \geq 1$, and $F \vdash^{\mathsf{S}} w$ such that $\mu_\ell(F) < 0$. The unique branch of F that contains all chains with non-negative reach (and possibly other chains) is called the* main *branch of F and denoted $\mathsf{M}^{\mathsf{S}}(F)$.*

4.3 Main PoS Theorem

The main result of this section is the following theorem, which is an analogue of Theorem 1 for the PoS case.

Theorem 2 (The PoS Phase Recurrences). *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^{\Delta}$ be characteristic strings. Then we have:*

Reach. $\rho(\varepsilon) = 0$. *Furthermore,*

$$\rho(wsxt) \begin{cases} = \rho(ws) + \#_{[\mathsf{a}]}(xt) - h_\Delta(x) & \text{if } \rho(ws) > h_\Delta(x), \\ \leq \#_{[\mathsf{a}]}(xt) & \text{otherwise.} \end{cases}$$

Margin. *If $|wsxt| < \ell$ then $\mu_\ell(wsxt) = \rho(wsxt)$, otherwise*

$$\mu_\ell(wsxt) \begin{cases} = \mu_\ell(ws) + \#_{[\mathsf{a}]}(xt) - h_\Delta(x) & \text{if } \mu_\ell(ws) < -\#_{[\mathsf{a}]}(xt), \\ \leq \rho(wsxt) & \text{otherwise.} \end{cases}$$

Crossing zero. *If $|ws| \geq \ell - 1$ and $\rho(ws) = \mu_\ell(ws) = 0$ then*

$$\mu_\ell(ws(1,0)(0,0)^{\Delta}) \leq -1 \,.$$

Theorem 2 describes the characteristic PoS quantities ρ and μ_ℓ in terms of phase-based recurrences. Similarly to the PoW case, the quantities behave differently in the three regions. Recall that a quantity is informally said to be in the hot region if it is sufficiently positive, such that the currently considered phase cannot bring it down to zero; it is said to be in the cold region if it is sufficiently negative so that it won't climb to zero within the current phase; and finally, it is said to be in the critical region if it is so close to zero that the effects of the special behavior the quantity exhibits around zero are manifested within this phase.

Informally speaking, Theorem 2 states that the reach quantity, as long as it remains within the hot region, exactly performs an "ideal recurrence"

$$\rho(wsxt) = \rho(ws) + \#_{[\mathsf{a}]}(xt) - h_\Delta(x) \,, \tag{10}$$

where it increases by 1 for each adversarially-successful slot, and decreases by 1 whenever the pattern of honest successes enforces an increase in the honest depth h_Δ. Whenever reach approaches the critical region (recall that reach is never negative by definition), we only upper-bound it with the quantity $\#_{[\mathsf{a}]}(xt)$—note that this is analogous to the outcome of the ideal recurrence in a hypothetical case where the honest successes first bring ρ to zero where the remaining honest

successes have no effect, while the remaining adversarial successes increase ρ back up to $\#_{[a]}(xt)$. As for margin, before slot ℓ it is identical to reach, and after slot ℓ it (again exactly) performs an analogue of the ideal recurrence (10) as long as it remains within the cold region, while outside of it we only make use of the trivial upper bound by ρ. Finally, we also establish a statement describing the crossing of zero, analogous to PoW.

4.4 Bounding Reach

The following lemma establishes the tightness of the ideal recurrence (10) for reach in the hot region.

Lemma 11 (Reach in the hot region). *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^{\Delta}$ be PoS characteristic strings. If $\rho(ws) > h_{\Delta}(x)$ then*

$$\rho(wsxt) = \rho(ws) + \#_{[a]}(xt) - h_{\Delta}(x) \ .$$

Proof. Denote $w' := wsxt$. We first prove a lower bound on $\rho(w')$. Towards that, consider a public witness tree $G \vdash^{\mathbf{s}} ws$ for reach in wx, and let U be the witness chain achieving $\mathrm{reach}_G(U) = \rho(ws)$. Let v_h be some maximum-depth honest vertex in G, i.e., $\mathsf{len}_G(v_h) = \mathsf{len}(G)$. Construct a labeled rooted tree G' from G by adding $\#_{\mathsf{h}}(x)$ honest vertices that form a subtree rooted in v_h, where each of these honest vertices is always put at the minimal depth allowed by axiom (A2), and labeling them consistently with x. Observe that by construction, G' is a valid public PoS-tree for w'. Using Fact 1 and the construction of G' we have $\mathsf{len}(\overline{G'_{\lceil \Delta}}) = \mathsf{len}(G') = \mathsf{len}(G) + h_{\Delta}(x)$, and hence

$$\begin{aligned} \rho(w') &\geq \mathrm{reach}_{G'}(U) = \alpha_{G'}(U) + \mathrm{reserve}_{G'}(U) \\ &= (\alpha_G(U) - h_{\Delta}(x)) + \big(\mathrm{reserve}_G(U) + \#_{[a]}(xt)\big) \\ &= \rho(ws) + \#_{[a]}(xt) - h_{\Delta}(x) > \#_{[a]}(xt) \ , \end{aligned} \tag{11}$$

where the last inequality follows by our assumption on $\rho(ws)$.

Towards an upper bound, let $F' \vdash^{\mathbf{s}} w'$ be a public witness tree for reach in w', and let T' be the witness chain for reach in F', i.e., $\mathrm{reach}_{F'}(T') = \rho(w')$. Let $F := \overline{F'_{|ws|}} \vdash^{\mathbf{s}} ws$ and let T be the restriction of T' to F. Using Fact 1 and the fact that F and F' are by definition public, we have $\overline{F_{\lceil \Delta}} = F$, $\overline{F'_{\lceil \Delta}} = F'$, and $\mathsf{len}(\overline{F'_{\lceil \Delta}}) \geq \mathsf{len}(\overline{F_{\lceil \Delta}}) + h_{\Delta}(x)$.

We now establish that $T = T'$. Indeed, if that is not the case, let v' be the terminating honest vertex of T'. Since $\mathsf{l}_{\#}(v') > |ws|$, it must be $\mathrm{reserve}(T') \leq \#_{[a]}(xt)$; and since v' is honest, we have $\rho(w') = \mathrm{reach}_{F'}(T') \leq \#_{[a]}(xt)$. This would be a contradiction with (11), proving that $T = T'$.

Given the above, we have

$$\begin{aligned} \rho(w') &= \mathrm{reach}_{F'}(T') = \alpha_{F'}(T') + \mathrm{reserve}_{F'}(T') \\ &\leq (\alpha_F(T) - h_{\Delta}(x)) + \big(\mathrm{reserve}_F(T) + \#_{[a]}(xt)\big) \\ &= \mathrm{reach}_F(T) + \#_{[a]}(xt) - h_{\Delta}(x) \leq \rho(ws) + \#_{[a]}(xt) - h_{\Delta}(x) \end{aligned}$$

as desired. □

It remains to prove the upper bound for reach in the critical region, this is done in the following lemma.

Lemma 12 (Reach approaching zero). *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^\Delta$ be PoS characteristic strings. If $\rho(ws) \leq h_\Delta(x)$ then*

$$\rho(wsxt) \leq \#_{[\mathsf{a}]}(xt) \ .$$

Proof. Let w', F', T', F, T be as in the proof of Lemma 11. We again have $\overline{F_{\lceil \Delta}} = F$, $\overline{F'_{\lceil \Delta}} = F'$, and $\mathsf{len}(\overline{F'_{\lceil \Delta}}) \geq \mathsf{len}(\overline{F_{\lceil \Delta}}) + h_\Delta(x)$.

Towards a contradiction, assume that $\rho(wsxt) = \mathrm{reach}_{F'}(T') > \#_{[\mathsf{a}]}(xt)$. As F' is public, clearly T' is honest without loss of generality, and $\mathsf{len}_{F'}(T') \leq \mathsf{len}(F')$, hence we have $\mathrm{reserve}_{F'}(T') > \#_{[\mathsf{a}]}(xt)$. However, this is only possible if $\mathrm{reserve}_{F'}(T')$ accounts also for some indices i (where $w'_i = (h_i, a_i)$ and $a_i > 0$) that satisfy $i < |ws|$, i.e., some adversarial vertices labeled from ws, and hence $T' = T$. However, given that $\mathsf{len}(\overline{F'_{\lceil \Delta}}) \geq \mathsf{len}(\overline{F_{\lceil \Delta}}) + h_\Delta(x)$, this means that $\mathrm{reach}_F(T) > h_\Delta(x)$ and therefore $\rho(w) > h_\Delta(x)$, contradicting our assumption and hence concluding the proof. □

4.5 Bounding Margin

Towards bounding the quantity $\mu_\ell(\cdot)$, first observe that its definition directly implies that $\mu_\ell(w) \leq \rho(w)$ for any $w \in \Sigma^*$. Moreover, for any w with $|w| < \ell$, we actually have $\mu_\ell(w) = \rho(w)$ as, recalling the definition of $\mu_\ell(F)$ and the relation $\not\sim_\ell$, notice that any chain T with $\mathsf{l}_\#(T) < \ell$ satisfies $T \not\sim_\ell T$, and hence the witness chains T_1, T_2 for $\mu_\ell(F)$ may satisfy $T_1 = T_2$.

We now proceed to prove a lower bound on μ_ℓ.

Lemma 13 (Margin lower bound). *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^\Delta$ be characteristic strings. If $\mu_\ell(ws) < -\#_{[\mathsf{a}]}(xt)$ then*

$$\mu_\ell(wsxt) \geq \mu_\ell(ws) + \#_{[\mathsf{a}]}(xt) - h_\Delta(x) \ .$$

The proof of the above lemma uses the same approach as the proof of Lemma 4 in the PoW case, we give it in the full version [14] for completeness. Note that in the PoS case, the construction given in the proof only works under the assumption $\mu_\ell(ws) < -\#_{[\mathsf{a}]}(xt)$, which is however exactly the region we are interested in.

We now turn to upper-bounding μ_ℓ in the specific case $\mu_\ell(ws) < -\#_{[\mathsf{a}]}(xt)$. Given that as observed above, for any ws satisfying $|ws| < \ell$ we have $\mu_\ell(ws) = \rho(ws) \geq 0$, this bound is only applicable after $|ws| \geq \ell$.

Lemma 14 (Margin in the cold region). *Let $\ell \geq 1$, let $w, x \in (\mathbb{N} \times \mathbb{N})^*$ and $s, t \in (\{0\} \times \mathbb{N})^\Delta$ be PoS characteristic strings. If $\mu_\ell(ws) < -\#_{[\mathsf{a}]}(xt)$ then*

$$\mu_\ell(wsxt) \leq \mu_\ell(ws) + \#_{[\mathsf{a}]}(xt) - h_\Delta(x) \ .$$

The proof of Lemma 14 is an adaptation of the proof of Lemma 5 to the PoS setting, we give it in the full version [14] for completeness.

The previous two lemmas together establish that margin follows an analogue of the ideal recurrence (10) in the cold region.

4.6 Crossing Zero

Finally, we show that if after slot ℓ both quantities are equal to zero, margin can descend to negative values. The proof of Lemma 15 uses essentially the same reasoning as that of Lemma 9, we provide it in the full version.

Lemma 15. *Let $\ell \geq 1$, let $w, x \in (\mathbb{N}\times\mathbb{N})^*$ and $s, t \in (\{0\}\times\mathbb{N})^{\Delta}$ be characteristic strings. If $|ws| \geq \ell - 1$ and $\rho(ws) = \mu_\ell(ws) = 0$ then*

$$\mu_\ell(ws(1,0)(0,0)^{\Delta}) \leq -1 \,.$$

4.7 A Practical PoS Adversary

In order to evaluate the strength of our settlement bounds, we describe and analyze a natural practical adversary in the proof-of-stake setting. It is analogous to the conventional "private-chain attack" adversary in the PoW setting.

In general, the adversary maintains two chains (L, S) and a "public depth" p, equal to the current depth of the deepest honest block. We collect this data together, writing $(\{L, S\}, p)$, and use l and s to denote the lengths of the chains L and S. We adopt the convention that L is the longer and S is the shorter of the two chains, with ties broken arbitrarily. The adversary will maintain the invariant that $l \geq \max(p, s)$ and that L and S diverge after ℓ. Then it is clear that $l - p$ is a lower bound for reach and that $s - p$ is a lower bound on margin. Given a current adversarial state $(\{L, S\}, p)$, we describe how the adversary responds to a new phase corresponding to a characteristic string x for which $\#_{\mathsf{a}}(x) = a$ and $h_{\Delta}(x) = h$.

In preparation for the full description, we set some terminology. Consider a chain C in this context (which is to say that C is one of S and L). We define the *adversarial extension* to be the chain C_a obtained by adding a path of $\#_{\mathsf{a}}(x)$ adversarial vertices to the end of C; this chain extension is consistent with x. If the length of C is at least p, we additionally define the *honest extension* C_h as follows: Define T_h to be a tree, rooted at the unique vertex of C of depth p, that contains one vertex for each honest success in x arranged so that each vertex is at the minimal depth dictated by Δ delay. The depth of this tree is $h_{\Delta}(x)$. Then define C_h, the honest extension of C, to be any path in this tree of maximal depth (thus having depth $p + h_{\Delta}(x)$). Note that this honest extension C_h is consistent with the characteristic string x of the new phase.

Prior to ℓ, the adversary maintains the invariant that $L = S$. A new phase with characteristic string x is fielded by constructing both the adversarial and honest extensions of L, called L_a and L_h, respectively, and assigning L' to be the longer of these. The resulting state is L' (and $S' = L'$). In the case where this phase includes the slot ℓ, the resulting state is $(\{L_a, L_h\}, p + h)$; observe that these do not share a vertex with label ℓ or more.

After ℓ, there are two cases depending on $(\{L, S\}, p)$. If $s < p$, no honest blocks can be immediately placed on S. In this case, define $S' = S_a$, the adversarial extension of S and define L' be the longer of L_a and L_h, the adversarial

and honest extensions of L. The resulting output state is $(\{L', S'\}, p + h)$. The second case arises when $s \geq p$: here we carry out the same procedure but reverse the roles of L and S: specifically, the honest spur is added to S rather than L, yielding two extensions of interest S_h and S_a. The longer of these is declared to be S'; L' is defined to be the simple adversarial extension of L. Note that while these rules are defined in terms of features of the entire phase, they can be carried out in an online fashion with no particular attention to placement of honest vertices (except that all blocks are delivered to honest parties with maximal delay). Note, furthermore, that the attack requires no tie breaking and can be thus carried out by an adversary that requires no capabilities beyond globally and uniformly delayed honest messages. (In contrast, it is not clear how to practically implement adaptive adversarial tie breaking.)

In terms of the recurrence relations (for μ_ℓ and ρ) that this yields, prior to ℓ we have $\mu_\ell = \rho$ by definition and $\rho' = \mu'_\ell = \max(0, \rho + a - h)$ by construction. After ℓ, reach continues to satisfy $\rho' = \max(0, \rho + a - h)$. If $\mu_\ell < 0$ it similarly satisfies $\mu'_\ell = \mu_\ell + a - h$. Otherwise, μ_ℓ is non-negative. We say that a configuration-input pair is "critical" if $\rho + a - h < 0$ (and $\mu_\ell \geq 0$). If the setting is critical, then margin satisfies $\mu'_\ell = \rho + a - h$. (Note that in this case, the two chains have switched roles.) Otherwise, set $\mu'_\ell = 0$ for convenience (as the exact value is not important to track).

5 Numerical Evaluation

In this section, we study explicit bounds provided by our analysis. We implement our analytical framework and make the code available at https://github.com/renling/LCanalysis/. We are interested in both PoW and PoS longest-chain consensus, and we pick one representative system for each. For PoW blockchains we study PoW-based Ethereum because its relatively short block interval presents a more challenging subject for analysis, while Bitcoin with its long block interval was already given fairly tight bounds [13,15]. For PoS blockchains we study Cardano, which implements the Ouroboros Praos protocol [7].

5.1 Modeling the Slot Leader Distribution

We assume the slot leader election is an ideal lottery. That is to say, the probability that any party (honest or adversary) becomes a leader in a slot is proportional to its hashing power (for PoW) or its stake[5] (for PoS), and this probability is independent of any other parties or any other slot. Thus, the total number of slot leaders in a given slot is given by a sum of Bernoulli random variables, one for each party. When there are sufficiently many parties, the sum of Bernoulli

[5] This is a slight simplification in the case of Ouroboros Praos, where the probability of a party that holds an s-fraction of stake (for $s \in [0, 1]$) becoming a slot leader is in fact $1-(1-f)^s$ for a constant f set to $1/20$ in Cardano. We adopt this simplification for the sake of broader applicability of our bounds.

random variables can be approximated by a Poisson random variable. More concretely, we model the number of honest leaders in a single slot as a Poisson random variable of parameter r_h, and the number of adversarial leaders in a single slot as a Poisson random variable of parameter r_a. Then, $\frac{r_h}{r_h+r_a}$ (resp. $\frac{r_a}{r_h+r_a}$) is the fraction of honest (resp. adversarial) hashing power in PoW, or stake in PoS. Furthermore, $\frac{1}{r_h+r_a}$ is the expected time it takes for one slot leader to appear, which is the target inter-block time. The inter-block time of Ethereum is roughly 13 s; the inter-block time of Cardano is 20 s. We can then derive r_h and r_a from the target block interval of the blockchain systems, and the assumed adversarial fraction.

Next, we need to make an assumption on the network propagation delay (recall that we denote it Δ_{r} when denominated in seconds). The 90th percentile block propagation time for Ethereum has been measured to be around 2 s [9], hence we will use 2 s as one example value of Δ_{r}. We will also give results for $\Delta_{\mathsf{r}} = 5$ s as a more conservative estimate. We did not find public propagation delay measurements for Cardano, but since Cardano and Ethereum have very similar block sizes, we use the same estimated values of Δ_{r} (i.e., 2 s and 5 s) for Cardano as well.

5.2 Symbol Distribution in a Phase

As our recurrences from Sects. 3 and 4 work at the phase granularity, the first step of the numeric evaluation is to compute the distribution of symbols in a phase. We will use PoS as the example in this subsection. The treatment for PoW is very similar.

Let xt be the characteristic string corresponding to a phase where $x \in (\mathbb{N} \times \mathbb{N})^*$ and $t \in (\{0\} \times \mathbb{N})^{\Delta}$. There are three quantities our recurrences need for each phase:

- $\#_{\mathsf{h}}(x)$, the total number of honest successes in the phase,
- $\#_{[\mathsf{a}]}(xt)$, the total number of slots with adversarial successes in the phase, and
- $h_{\Delta}(x)$, the honest depth of the phase.

The latter two quantities are directly used in the recurrences, and we will explain the role of $\#_{\mathsf{h}}(x)$ in Sect. 5.3.

We now explain how we can compute the joint distribution of the above three quantities for a given slot. By definition, whenever there is a Δ period with no honest successes, the phase ends. We will process one honest success at a time, and at each step, update the joint probability density functions (pdf) of the three quantities of interest. This way makes it easy to compute the distribution of $\#_{\mathsf{h}}(x)$. Each step has a probability of ending the phase and the i-th step gives the probability of $\#_{\mathsf{h}}(x) = i$. To compute the distributions of the other two quantities, we introduce and keep track of the distributions of two additional random variables representing elapsed times:

– S: elapsed time since the beginning of the phase, and
– S_{h_Δ}: elapsed time since the last increase of honest depth.

Let T_h be the interarrival time between the current honest success and the previous honest success. T_h as an interarrival time in a Poisson point process follows an exponential distribution. With each new honest success, the distribution of S is updated by a simple convolution of T_h and the original S (which yields the pdf of the sum of two random variables). The distribution of S_{h_Δ} can be similarly updated except that we always have $S_{h_\Delta} < \Delta$, so the post-Δ portion of the resulting pdf (after convolution) is reset and added to the its pdf at 0 (i.e., probability of $S_{h_\Delta} = 0$). This post-Δ portion of the resulting pdf is also the probability that the new honest success increments the honest depth $h_\Delta(x)$, allowing us to compute the distribution of $h_\Delta(x)$. From here, we can also compute the distribution of the latest inter-honest-success time T_h conditioned on whether or not $h_\Delta(x)$ is incremented. We can then compute the distribution of adversarial successes during this latest T_h, again conditioned on whether or not $h_\Delta(x)$ is incremented. Lastly, we can update the distribution of $\#_{[\mathsf{a}]}(xt)$ and the joint distribution of all three quantities by convolving it with the above conditional pdfs.

5.3 Evaluating the Recurrence

Once we have the characteristic string distribution within a phase, it is relatively straightforward to numerically evaluate the recurrences. We again focus on the PoS case. The PoW case is similar (in fact, simpler, because there is only one quantity involved in the PoW recurrence).

Initially, we must settle on a distribution of (μ_ℓ, ρ) at time ℓ, which corresponds to the moment the transaction of interest appears in a block). While this does depend on ℓ, the distribution converges quickly to a geometric distribution for reasonably large ℓ. For this reason, we will use the stationary distribution of the initial (μ_ℓ, ρ) as its distribution at time ℓ. Also observe that before ℓ, the margin μ_ℓ was equal to reach ρ, making ρ the only quantity of interest. Intuitively, this initial distribution of ρ represents the number of private blocks that the adversary has on top of the longest public honest block when the transaction of interest enters the ledger.

Next, we need to evolve the recurrence until settlement happens. Unfortunately, we do not know when exactly settlement happens as that depends on the adversarial strategy and the initial value of ρ. Therefore, we instead evolve the recurrence until the *earliest possible* time that a settlement error could occur. Observe that a settlement error can occur only after $2k - s$ lottery successes have occurred since ℓ, where s is value of ρ at time ℓ and k is the settlement depth. This is because two chains of length k must exist for the adversary to cause a settlement error. (For example, if $\rho = 2k$ at time ℓ, an adversary can immediately activate the settlement of "buried by k blocks" and violate consistency.) To do so, we need the distribution of the total number of successes in each phase, which is also why we need the quantity $\#_{\mathsf{h}}(x)$ in the joint distribution of a phase.

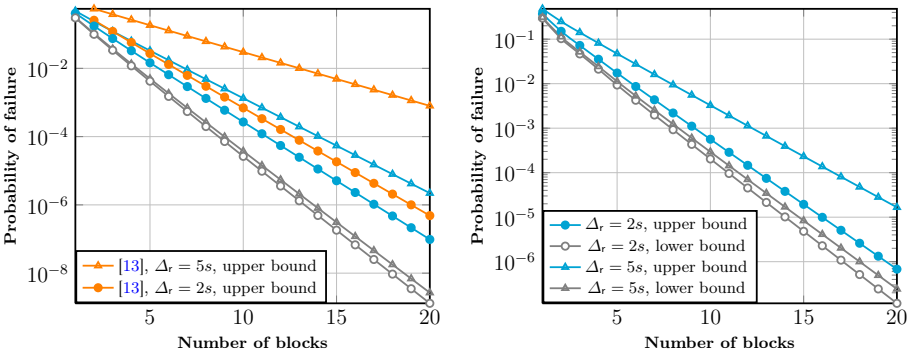

Fig. 4. Left: Ethereum (PoW) block-based settlement failure for a 10% adversary, results from [13] for comparison. **Right:** Cardano (PoS) block-based settlement failure for a 10% adversary. (The right-side legend applies to both figures.)

From here, we keep evolving the recurrence, but "freezing" any probability mass on positive values of the margin. We then evolve the system forward until the (exponentially decaying) contributions from further evolution are negligible.

The same approach can be used to numerically evaluate the concrete attack described in Sect. 4.7 to yield lower bounds on the settlement failure probability for PoS blockchains. The lower bounds for PoW blockchains are based on a simple private-mining attack.

5.4 Numerical Results

Figure 4 is a more detailed version of Fig. 1 from the introduction. It depicts our settlement bounds for Ethereum, compared to the best previous bounds for Ethereum [13]; as well as our new settlement bounds for Cardano. We provide more numerical results in the full version [14].

Our methods also enable a direct comparison between PoW and PoS blockchains in terms of settlement error and settlement delay. Figure 5 plots our settlement upper bound for Cardano and compares it against a hypothetical PoW blockchain with the same inter-block time as Cardano, and under the same network delay and adversarial ratio. We can see that given the same system parameters and at the same settlement depth, a PoS blockchain has a larger settlement error; equivalently, to obtain the same settlement error, a PoS blockchain needs a slightly higher settlement depth. This is an expected price being paid for avoiding the enormous energy consumption of PoW, and is caused by giving the adversary the extra power of creating as many blocks as it wishes using a single success. Our results allow this price to be precisely quantified for the first time.

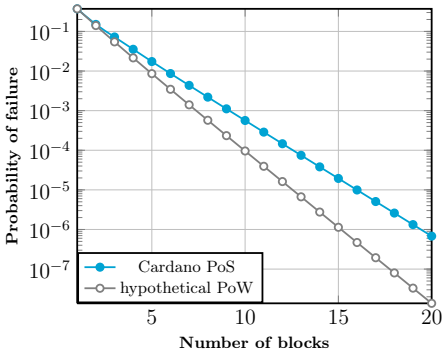

Fig. 5. Comparison of block-based settlement failure upper bounds for Cardano (PoS, 10% adversary, $\Delta_{\mathsf{r}} = 2s$) and a hypothetical PoW protocol with the same parametrization.

6 Conclusions: Practical Relevance

The goal of our work is to provide concrete settlement bounds with practical applicability to deployed longest-chain protocols. We provide the first such bounds for longest-chain PoS, and along the way also derive the best existing settlement bounds for PoW.

We remark that in specific PoS systems there may be additional security factors that affect settlement times. For example, while the lottery in protocols such as Ouroboros [17] cannot be biased by an adversary, another class of protocols including Ouroboros Praos [7] and Snow White [6] allow for so-called grinding of the randomness beacon. While the results of this paper describe the intrinsic aspects of longest-chain rule and apply to both of these protocol classes, for the latter one an additional term accounting for grinding must be considered. Fortunately, these two sources of settlement failure can be studied independently and combined in a straightforward fashion.

While the concrete results we quote consider the parametrizations of Ethereum and Cardano, our methods can be directly applied to compute these statistics for any other choice of block interval, block propagation delay Δ_{r}, and assumed share of adversarial power. In each specific case, the value Δ_{r} can be estimated based on measurements, such as those we reference for Ethereum. Finally, estimating the fraction of adversarially controlled stake ultimately comes down to each user's belief about the state of stake distribution across the set of users; nonetheless our results allow each individual user to choose their settlement rule based on their own beliefs about the system and their acceptable failure probability (perhaps depending on the transacted amount).

Acknowledgements. This work is funded in part by National Science Foundation award 2143058.

References

1. Badertscher, C., Gaži, P., Kiayias, A., Russell, A., Zikas, V.: Consensus redux: distributed ledgers in the face of adversarial supremacy. Cryptology ePrint Archive, Report 2020/1021 (2020). https://eprint.iacr.org/2020/1021
2. Badertscher, C., Gazi, P., Kiayias, A., Russell, A., Zikas, V.: Ouroboros genesis: composable proof-of-stake blockchains with dynamic availability. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018, pp. 913–930. ACM Press, October 2018. https://doi.org/10.1145/3243734.3243848
3. Blum, E., Kiayias, A., Moore, C., Quader, S., Russell, A.: The combinatorics of the longest-chain rule: linear consistency for proof-of-stake blockchains. In: Chawla, S. (ed.) 31st SODA, pp. 1135–1154. ACM-SIAM, January 2020. https://doi.org/10.1137/1.9781611975994.69
4. Bonneau, J., Meckler, I., Rao, V., Shapiro, E.: Coda: decentralized cryptocurrency at scale. Cryptology ePrint Archive, Report 2020/352 (2020). https://eprint.iacr.org/2020/352
5. Chen, J., Micali, S.: Algorand. arXiv preprint: arXiv:1607.01341 (2016)
6. Daian, P., Pass, R., Shi, E.: Snow White: robustly reconfigurable consensus and applications to provably secure proof of stake. In: Goldberg, I., Moore, T. (eds.) FC 2019. LNCS, vol. 11598, pp. 23–41. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32101-7_2
7. David, B., Gaži, P., Kiayias, A., Russell, A.: Ouroboros Praos: an adaptively-secure, semi-synchronous proof-of-stake blockchain. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part II. LNCS, vol. 10821, pp. 66–98. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78375-8_3
8. Dembo, A., et al.: Everything is a race and Nakamoto always wins. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 20, pp. 859–878. ACM Press, November 2020. https://doi.org/10.1145/3372297.3417290
9. Ethstats (2021). https://ethstats.net/
10. Eyal, I., Sirer, E.G.: Majority is not enough: Bitcoin mining is vulnerable. In: Christin, N., Safavi-Naini, R. (eds.) FC 2014. Lecture Notes in Computer Science(), vol. 8437, pp. 436–454. Springer, Berlin (2014). https://doi.org/10.1007/978-3-662-45472-5_28
11. Garay, J., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol: analysis and applications. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 281–310. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_10
12. Gazi, P., Kiayias, A., Russell, A.: Tight consistency bounds for bitcoin. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 20, pp. 819–838. ACM Press, November 2020. https://doi.org/10.1145/3372297.3423365
13. Gazi, P., Ren, L., Russell, A.: Practical settlement bounds for proof-of-work blockchains. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, pp. 1217–1230. Association for Computing Machinery, New York (2022). https://doi.org/10.1145/3548606.3559368
14. Gaži, P., Ren, L., Russell, A.: Practical settlement bounds for longest-chain consensus. Cryptology ePrint Archive, Paper 2022/1571 (2022). https://eprint.iacr.org/2022/1571
15. Guo, D., Ren, L.: Bitcoin's latency-security analysis made simple. In: Proceedings of the 4th ACM Conference on Advances in Financial Technologies (2022)

16. Kiayias, A., Quader, S., Russell, A.: Consistency of proof-of-stake blockchains with concurrent honest slot leaders. In: 40th IEEE International Conference on Distributed Computing Systems, ICDCS 2020, Singapore, November 29 - December 1, 2020, pp. 776–786. IEEE (2020). https://doi.org/10.1109/ICDCS47774.2020.00065
17. Kiayias, A., Russell, A., David, B., Oliynykov, R.: Ouroboros: a provably secure proof-of-stake blockchain protocol. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 357–388. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63688-7_12
18. Lamport, L.: The part-time parliament. In: Concurrency: the Works of Leslie Lamport, pp. 277–317 (2019)
19. Li, J., Guo, D.: On analysis of the bitcoin and prism backbone protocols in synchronous networks. In: 2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 17–24. IEEE (2019)
20. Li, J., Guo, D., Ren, L.: Close latency-security trade-off for the Nakamoto consensus. In: Proceedings of the 3rd ACM Conference on Advances in Financial Technologies, pp. 100–113 (2021)
21. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system (2008)
22. Schneider, F.B.: Implementing fault-tolerant services using the state machine approach: a tutorial. ACM Comput. Surv. (CSUR) **22**(4), 299–319 (1990)