

# Minotaur: Multi-Resource Blockchain Consensus

Matthias Fitzi\*  
IOHK  
matthias.fitzi@iohk.io

Xuechao Wang\*  
University of Illinois  
Urbana-Champaign  
xuechao2@illinois.edu

Sreeram Kannan  
University of Washington at Seattle  
ksreeram@uw.edu

Aggelos Kiayias  
University of Edinburgh and IOHK  
aggelos.kiayias@iohk.io

Nikos Leonardos  
University of Athens  
nikos.leonardos@gmail.com

Pramod Viswanath  
University of Illinois  
Urbana-Champaign  
pramodv@illinois.edu

Gerui Wang  
University of Illinois  
Urbana-Champaign  
gerui2@illinois.edu

## ABSTRACT

Resource-based consensus is the backbone of permissionless distributed ledger systems. The security of such protocols relies fundamentally on the level of resources actively engaged in the system. The variety of different resources (and related proof protocols, some times referred to as PoX in the literature) raises the fundamental question whether it is possible to utilize many of them in tandem and build *multi-resource* consensus protocols. The challenge in combining different resources is to achieve *fungibility* between them, in the sense that security would hold as long as the *cumulative* adversarial power across all resources is bounded.

In this work, we put forth Minotaur, a multi-resource blockchain consensus protocol that combines proof of work (PoW) and proof-of-stake (PoS), and we prove it *optimally* fungible. At the core of our design, Minotaur operates in epochs while continuously sampling the active computational power to provide a fair exchange between the two resources, work and stake. Further, we demonstrate the ability of Minotaur to handle a higher degree of work fluctuation as compared to the Bitcoin blockchain; we also generalize Minotaur to any number of resources.

We demonstrate the simplicity of Minotaur via implementing a full stack client in Rust (available open source [25]). We use the client to test the robustness of Minotaur to variable mining power and combined work/stake attacks and demonstrate concrete empirical evidence towards the suitability of Minotaur to serve as the consensus layer of a real-world blockchain.

## 1 INTRODUCTION

**Resource-based Consensus.** The fundamental feature of the decentralized computing paradigm of permissionless blockchains is that participation in the consensus protocol is enabled by proving possession of a resource. Bitcoin pioneered this idea based on proof-of-work (PoW), and its implied energy wastage inspired new designs based on alternative resources such as proof-of-stake (PoS), proof-of-space (PoSp) and proof-of-elapsed-time (PoD). These different resources cover the variegated and multidimensional forms of ‘capital’ that the participants bring. Each proof-of-X mechanism

is interesting on their merit in appropriate settings (computation in PoW, memory and storage in PoSp and time/delay in PoD, tokens/capital in PoS). Further, the mechanisms trigger different incentives even within the same resource format; focusing on PoW for example, we see that the hashcash algorithm [2] instantiated by the SHA256 function [29] implemented in Bitcoin has been replaced by others, including scrypt [32] and ethash [12]. Each of these different proof-of-X mechanisms have led to different, and isolated, blockchain ecosystems.

**Combining different resources.** Given the diversity of incentivization embodied by different resources, it is a natural question whether it is feasible to combine their features into a single blockchain design (e.g., hybrid PoW-PoS permissioned blockchains). The key challenge to combining multiple resources is in determining the *exchange rate*, i.e., to what extent the different resources are translated into power of authority in the system and to extract security from these resources in an optimal manner. Adapting the exchange rate dynamically to the participation levels in each resource type, resulting in a truly *fungible* notion of resource, is a basic and fundamental goal. By truly fungible, we mean that the security of the multi-resource consensus protocol is guaranteed as long as the honest players control a majority of the combined resources that consist of each resource type in the system — concretely,  $\sum_{i=1}^M \beta_i < M/2$  where  $M$  is the number of resources and  $\beta_i$  the adversarial power in the  $i$ -th resource (this property is appropriately generalized to any linear combination in Definition 2). Next we give some examples of *non-fungible* hybrid PoW/PoS protocol designs.

**First order hybrid PoW/PoS protocols.** Given the basic importance of incorporating multiple resources in a single blockchain design, there are several designs of hybrid PoW/PoS protocols in the literature [4, 10, 11, 19, 23]. These protocols, however, are either heuristic (lacking a formal security analysis) [4, 10, 23], or make assumptions that break fungibility (e.g. honest majority in stake [10, 11, 19] and/or static mining power [11]). Indeed, with sufficient assumptions the problem of a hybrid protocol is trivially resolved. For instance, if we assume an honest majority in stake (at all time), we can use a committee chosen randomly from the pool of stakeholders to assist a PoW ledger by regularly issuing


\*The first two authors contributed equally to this work.

checkpoints [19]. However, the security of this scheme is solely guaranteed by the stakeholders, and the trust is entirely supported by PoS (not from PoW).

**Natural solutions.** In fact, if we assume a static setting (where both the total mining power and total ‘active’ stake are fixed and known to the protocol designer), a simple solution is the following: PoW and PoS mining occur in parallel; and whichever PoW or PoS succeeds first, goes ahead and extends the longest chain. However, there does not appear to be any straightforward (or otherwise) approach to extend this simple idea to the non-static setting, since we can no longer normalize the stake and work in the system when the total mining power is changing and unknown (§4). Even a fungible combination of work (i.e., total amount of honest work is more than total Byzantine work) emanating from two different hash functions (thus allowing two PoW blockchains using the same longest-chain consensus protocol to co-exist) has been unsolved.

**Minotaur Protocol.** In this paper we present *Minotaur*, a blockchain protocol with proof of fungible work and stake. At its core, *Minotaur* is a longest chain protocol that bases its block-proposer schedule on our concept of *virtual stake* that fuses active *actual stake* and *work stake*. The work-stake distribution is determined per epoch by measuring the participants’ contributions in hashing power during a prior epoch. Work contribution is achieved by concurrent PoW mining of *PoW blocks* (similar to ‘endorser inputs’ [21] or ‘fruits’ [31]) to be eventually referenced by main-chain blocks. The work-stake distribution is then assigned proportionally to a participant’s share of endorser blocks referenced from the main chain during that epoch; and applying the fairness mechanisms in [17, 21, 31], this assignment indeed truly represents the PoW-resource distribution among the participants. Figure 1a illustrates the protocol (a detailed description is given in §5).

**Minotaur security is optimal.** We show that *Minotaur* is truly fungible between work and stake by showing it is secure when the sum of  $\omega$  times the proportion of adversarial hash power ( $\beta_w$ ) and  $1 - \omega$  times the proportion of the adversarial stake ( $\beta_s$ ) is smaller than  $1/2$ , for any  $\omega \in [0, 1]$ , a weighing parameter of PoW/PoS (see Figure 1b for  $\omega = 0.5$ ). Figure 1b also compares the security guarantee of *Minotaur* and a couple of non fungible PoW/PoS protocols in the literature (the checkpointed ledger [19], the 2-hop blockchain [11], and a few finality gadgets [6, 28, 33, 36]), with more details in Appendix A. We give a formal security analysis in §6. The new challenge that we have to tackle in our security analysis is to show a *fairness* guarantee in the work-stake conversion (e.g., a miner with 10% mining power should gain 10% work stake) in the *non-static* setting and in the presence of an *adversarial majority* of mining power. This requires a new understanding of how to bound the evolution of mining difficulty in the system compared to techniques used in previous works and the Bitcoin blockchain, and presents a significant barrier to surmount in our analysis. One immediate consequence is that *Minotaur* can tolerate a 51% mining adversary, as long as there is a supermajority in honest stake (and vice versa, a 51% stake adversary). We prove our security guarantee to be optimal, see §3, in the sense that otherwise the adversary will control the majority of the combined “resource” in the system and an analogy of Nakamoto’s private chain attack [26] could break the security (formal and detailed proof in §3). We also show that



**Figure 1: (a) Architecture of Minotaur consensus protocol. (b) Comparison of security regions achieved by different hybrid PoW/PoS protocols (more details in Appendix A).**

*Minotaur* is capable of tolerating fluctuations of work much more effectively than Bitcoin, cf. §6.4.

**Implementation.** We have implemented a prototype of a *Minotaur* client in about 6000 lines of Rust code [25] and provide experimental results to evaluate the concrete performance of *Minotaur* under different scenarios. We also implemented a Bitcoin client as a benchmark. We use that client to validate also experimentally that *Minotaur* can survive more drastic variations of network hash power, while Bitcoin is no longer live under the same scenario. We also evaluate practical security concerns, like selfish mining and spamming, on our *Minotaur* client. Our implementation shows that *Minotaur* performs well even under these attacks, and makes a stronger case for the practical viability and robustness of the system. More details on the implementation and the experiments can be found in §7.

**Launching considerations.** We point out that *Minotaur* provides also a secure and efficient method to bootstrap a PoW-based blockchains. In general, it is more challenging to securely launch PoW blockchains, as a new blockchain would presumably start off with a relatively smaller total computation power and be vulnerable to the so-called “51% attacks”. On the other hand, PoS blockchains are easy to launch with existing techniques, such as proof-of-burn [20], initial coin offering [24] and airdrop [1]. The key advantage of *Minotaur* is that it can launch in pure PoS mode (setting parameter

$\omega = 0$ ) and later transit into a pure PoW blockchain by gradually changing the weighing parameter  $\omega$  to a higher value (§8).

**Related work.** The idea of hybrid PoW/PoS block production was first mentioned in [23]. A first concrete construction was given in [4] under the label ‘proof of activity’, but without giving a rigorous security analysis. Their block-production mechanism essentially extends standard Bitcoin mining by having the mined block implicitly elect a set of stakeholders that are required to sign the block in order to validate it.

A similar construction was proposed in [11] and proven secure under a majority of adversarial hashing power—however, their security proof still implies an honest majority of stake. In particular, and contrary to their initial claims, the protocol is not proven secure under the condition that *any* minority of the combined resources is controlled by the adversary.<sup>1</sup>

The work of [11] was extended in [10], to adapt to dynamic participation of both, hashing power and stake, and in [37], to combine PoW with multiple resources (rather than just PoS)—both works lack a rigorous security analysis.

The application of PoW/PoS hybrid block production for the goal of protecting early-stage PoW systems against initial periods of adversarial dominance in hashing power was explored in [7]. They propose to start the system with hybrid block production (where each resource contributes to a certain fraction of the blocks) and then slowly fade out the stake contribution to eventually turn the system into pure PoW. Also this protocol is not proven secure.

There also exist another class of hybrid PoW/PoS protocols [6, 28, 33, 36], which uses a BFT protocol (PoS) to build a finality gadget/layer on the top of a Nakamoto-style longest chain protocol (PoW) to achieve important properties such as accountability and finality. However, these protocols require an honest majority (or even supermajority) on both the set of miners and the set of stakeholders hence they are not fungible.

## 2 PRELIMINARIES

### 2.1 Security model

**Time, slots, and synchrony.** We consider a setting where time is divided into discrete units called *slots*. Players are equipped with (roughly synchronized) clocks that indicate the current slot. Slots are indexed by integers, and further grouped into epochs with fixed size  $R$ , i.e., epoch  $e$  composes of slots  $\{(e-1)R+1, (e-1)R+2, \dots, eR\}$ . And we assume that the real time window that corresponds to each slot has the following two properties: (1) The current slot is determined by a publicly-known and monotonically increasing function of current time; (2) Each player has access to the current time and any discrepancies between parties’ local time are insignificant in comparison with the duration of a slot.

We describe our protocols in the now-standard  $\Delta$ -synchronous network model considered in [3, 8, 18, 30], where there is an (unknown) upper bound  $\Delta$  in the delay (measured in number of slots) that the adversary may inflict to the delivery of any message. Similar to [18, 30], the protocol execution proceeds in slots with inputs provided by an environment program denoted by  $\mathcal{Z}(1^\kappa)$  to parties that execute the protocol  $\Pi$ , where  $\kappa$  is a security parameter. The

network is modeled as a diffusion functionality similar to those in [18, 30]: it allows message ordering to be controlled by the adversary  $\mathcal{A}$ , i.e.,  $\mathcal{A}$  can inject messages for selective delivery but cannot change the contents of the honest parties’ messages nor prevent them from being delivered within  $\Delta$  slots of delay — a functionality parameter. Specially, for  $\Delta = 1$ , the network model is reduced to the so-called *lock-step synchronous model*, where messages are guaranteed to be delivered within one slot.

**Protocol player.** The protocol is executed by two types of players, PoW-miners (miners for short) and PoS-holders (stakeholders for short), who generate different types of blocks, PoW blocks and PoS blocks. Note that we allow for any relation among the sets of miners and stakeholders, including the possibility that all players play both roles, or the two types of players are disjoint.

**Random oracle.** For PoW mining, we abstract the hash function as a random-oracle functionality. It accepts queries of the form (compute,  $x$ ) and (verify,  $x, y$ ). For the first type of query, assuming  $x$  was never queried before, a value  $y$  is sampled from  $\{0, 1\}^\kappa$  and it is entered to a table  $T_H$ . If  $x$  was queried before, the pair  $(x, y)$  is recovered from  $T_H$ . In both cases, the value  $y$  is provided as an answer to the query. For the second type of query, a lookup operation is performed on the table. Honest miners are allowed to ask one query per slot of the type `compute` and unlimited queries of the type `verify`. The adversary  $\mathcal{A}$  is given a bounded number of `compute` queries per slot and also unlimited number of `verify` queries. The bound for the adversary is determined as follows. Whenever a corrupted miner is activated the bound is increased by 1; whenever a query is asked the bound is decreased by 1 (it does not matter which specific corrupted miner makes the query).

**Adversarial control of resources.** We assume a Byzantine adversary who can decide on the spot how many honest/corrupted miners are activated adaptively. In a slot  $r$ , the number of honest miners that are active in the protocol is denoted by  $h_r$  and the number of all active miners in slot  $r$  is denoted by  $n_r$ . Note that  $h_r$  can only be determined by examining the view of all honest miners and is not a quantity that is accessible to any of the honest miners individually. In order to obtain meaningful concentration bounds on the number of PoW blocks in a long enough window, we set a lower bound  $\alpha_0$  on the fraction of honest mining power, i.e.,  $h_r \geq \alpha_0 n_r$  for all  $r$ . Note that  $\alpha_0$  does not have to be  $1/2$  as required by Bitcoin, it can be a small positive constant. Sudden decreases of total mining power could hurt the liveness of the protocol due to too few PoW blocks mined in one epoch. Therefore, we need to restrict the fluctuation of the number of honest/adversarial queries in a certain limited fashion. Suppose  $\mathcal{Z}, \mathcal{A}$  with fixed coins produces a sequence of honest miners  $h_r$ , where  $r$  ranges over all slots of the entire execution, we define the following notation.

**DEFINITION 1 (FROM [16]).** For  $\gamma \in [1, \infty)$ , we call a sequence  $(x_r), r \in [LR]$ , as  $(\gamma, s)$ -respecting if for any set  $S \subseteq [0, LR]$  of at most  $s$  consecutive slots,

$$\max_{r \in S} x_r \leq \gamma \min_{r \in S} x_r,$$

We say that  $\mathcal{Z}$  is  $(\gamma, s)$ -respecting if for all  $\mathcal{A}$  and coins for  $\mathcal{Z}$  and  $\mathcal{A}$ , both the sequences of  $(h_r)$  and  $(n_r)$  are  $(\gamma, s)$ -respecting.

<sup>1</sup>This point was also confirmed by the last author of [11] in personal communication.

Finally, Minotaur achieves consensus via proof of fungible work and stake. The following is the formal definition of fungibility and our major assumptions on the adversarial power.

**DEFINITION 2 (FUNGIBILITY OF RESOURCES).** For a time window  $W$ , let  $\beta_s^W$  be the maximum fraction of adversarial stake in  $W$ , where the maximum is taken over all views of all honest players across all slots in  $W$ ; and let  $\beta_w^W$  be the maximum fraction of adversarial mining power over all slots in  $W$  (i.e.,  $\beta_w^W = 1 - \min_{r \in W} h_r/n_r$ ). For  $\theta \in [0, 1]$ , we say the adversary  $\mathcal{A}$  is  $(\theta, m, \omega)$ -bounded if for any time window  $W$  with at most  $m$  slots, we have  $\omega\beta_w^W + (1-\omega)\beta_s^W \leq \theta$ . We say a blockchain protocol achieves fungibility of work and stake if it is secure against such an adversary.

**ASSUMPTION 1 (INITIALIZATION).** During the initialization phase of the protocol, we assume:

- 1.1 The initial stake distribution has honest majority, i.e.,  $\beta_s^0 \leq 1/2 - \sigma$ , where  $\beta_s^0$  is the fraction of initial stake controlled by the adversary.
- 1.2 We have a good estimate of the initial honest mining power  $h_1$ . In particular, let  $\tilde{h}_1$  be the estimate of  $h_1$ , we have  $h_1/(1 + \delta)\gamma^2 \leq \tilde{h}_1 \leq \gamma^2 h_1/(1 - \delta)\alpha_0$ .

To keep the paper simple, the Minotaur protocol, as described, relies on Assumption 1, while we refer the reader to §8 for a discussion on how to modify the protocol such that this initial assumption (of an honest majority of stake) is no longer required.

**ASSUMPTION 2 (EXECUTION).** During the execution phase of the protocol, we assume:

- 2.1 Stake-work bound: The adversary  $\mathcal{A}$  is  $(1/2 - 2\sigma, 2R, \omega)$ -bounded.
- 2.2 Work fluctuation bound: The environment  $\mathcal{Z}$  is  $(\gamma, 2R)$ -respecting.
- 2.3 Work participation bound: For any slot  $r \in [LR]$ ,  $h_r \geq \alpha_0 n_r$ .

**Participation model.** Minotaur can be constructed based on different PoS longest chain protocols. Various versions of Minotaur will take different subsets of the following assumptions on the participation model:

- (P1) Honest stakeholders are always online<sup>2</sup>;
- (P2) Honest miners who mined PoW blocks in epoch  $e$  will stay online in epoch  $e + 2$ <sup>3</sup>
- (P3) In case an honest party joins after the beginning of the protocol, its initialization chain  $C$  provided by the environment should match an honest party's chain which was active in the previous slot.

## 2.2 Blockchain security properties

**NOTATION 1.** We denote by  $C^{\uparrow \ell}$  the chain resulting from “pruning” the blocks with timestamps within the last  $\ell$  slots. If  $C_1$  is a prefix of  $C_2$ , we write  $C_1 < C_2$ . The latest block in the chain  $C$  is called the head of the chain and is denoted by  $\text{head}(C)$ . We denote by  $C_1 \cap C_2$  the common prefix of chains  $C_1$  and  $C_2$ . Given a chain  $C$  and an interval  $S$  (or  $[r_1, r_2]$ ), let  $C(S)$  (or  $C[r_1, r_2]$ ) be the segment of  $C$  containing

<sup>2</sup>This assumption can be easily relaxed. Namely, it is sufficient for an honest stakeholder to come online at the beginning of each epoch, determine whether it belongs to the slot leader set for any slots within this epoch, and then come online and act on those slots while maintaining online presence at least every  $k$  slots. See Appendix H of [8] for more detail.

<sup>3</sup>This can be relaxed similarly as P1.

blocks with timestamps in  $S$  (or  $[r_1, r_2]$ ). We say that a chain  $C$  is held by or belongs to an honest node if it is one of the best chains in its view.

**DEFINITION 3 (COMMON PREFIX (CP)).** The common prefix property with parameter  $\ell_{\text{cp}} \in \mathbb{N}$  states that for any two honest nodes holding chains  $C_1, C_2$  at slots  $r_1, r_2$ , with  $r_1 \leq r_2$ , it holds that  $C_1^{\uparrow \ell_{\text{cp}}} < C_2$ .

**DEFINITION 4 (EXISTENTIAL CHAIN QUALITY ( $\exists\text{CQ}$ )).** The existential chain quality property with parameter  $\ell_{\text{cq}} \in \mathbb{N}$  states that for any chain  $C$  held by any honest party at slot  $r$  and any interval  $S \subseteq [0, r]$  with at least  $\ell_{\text{cq}}$  consecutive slots, there is at least one honestly generated block in  $C(S)$ .

Our goal is to generate a robust transaction ledger that satisfies persistence and liveness as defined in [21].

**DEFINITION 5.** A protocol  $\Pi$  maintains a robust public transaction ledger if it organizes the ledger as a blockchain of transactions and it satisfies the following two properties:

- (Persistence) Consider the confirmed ledger  $L_1$  on any node  $p_1$  at any slot  $r_1$ , and the confirmed ledger  $L_2$  on any node  $p_2$  at any slot  $r_2$  (here  $u_1(r_1)$  may or may not equal  $u_2(r_2)$ ). If  $r_1 + \Delta < r_2$ , then  $L_1$  is a prefix of  $L_2$ .
- (Liveness) Parameterized by  $u \in \mathbb{R}$ , if a transaction  $\text{tx}$  is received by all honest nodes for more than  $u$  slots, then all honest nodes will contain  $\text{tx}$  in the same place in the confirmed ledger.

## 3 IMPOSSIBILITY RESULT


Consider any protocol  $\Pi$  executed by two types of players, miners and stakeholders, and let  $D_\Pi$  be the region of  $(\beta_w, \beta_s)$  such that  $\Pi$  can generate a robust public transaction ledger under an adversary controlling a  $\beta_w$  fraction of mining power and a  $\beta_s$  fraction of stake (see Figure 1b).

**THEOREM 3.1.** Any two points,  $X_1 = (p_1, q_1)$  and  $X_2 = (p_2, q_2)$  such that  $p_1 + p_2 \geq 1$  and  $q_1 + q_2 \geq 1$ , cannot co-exist in  $D_\Pi$ .

**PROOF.** We show that if there exists a protocol  $\Pi$  secure under both both points  $X_1$  and  $X_2$ , then we can use  $\Pi$  to implement a robust transaction ledger with two players, one of them being Byzantine, which is impossible.

Let Alice and Bob be two players, one of them possibly being malicious. Let Alice control a  $p_1$  fraction of mining power and a  $1 - q_2$  fraction of stake. Let Bob control a  $1 - p_1$  fraction of mining power and a  $q_2$  fraction of stake. In the case where Alice is malicious, a  $p_1$  fraction of mining power and a  $1 - q_2 \leq q_1$  fraction of stake are malicious, which is dominated by point  $X_1$ . In the case where Bob is malicious, a  $1 - p_1 \leq p_2$  fraction of mining power and a  $q_2$  fraction of stake are malicious, which is dominated by point  $X_2$ . Assuming that  $\Pi$  implements a robust public transaction ledger implies that also the protocol among the 2 players does while one of them being malicious. However, by Theorem 1 in [15], tolerating  $f = 1$  malicious player requires at least  $2f + 1 = 3$  players. This is a contradiction.  $\square$

Figure 2 gives a few examples of possible *maximum* security regions that satisfy the constraint by Theorem 3.1. By maximum, we mean the region cannot be enlarged any more. All these regions



**Figure 2: A few examples of possible security regions satisfying the constraint by Theorem 3.1.**

are bounded by a *non-increasing* curve that is *symmetrical* about the point  $(1/2, 1/2)$ . Therefore, all these security regions have area  $1/2$ , which is an analogy of the  $1/2$  fault tolerance in the single-resource systems. In this work, we focus on achieving the security regions bounded by a linear curve (the yellow and red ones in Figure 2) and leave the achievability of other possible regions (e.g., the green and blue ones) to future work.

**COROLLARY 1.** *No protocol  $\Pi$  can generate a robust public transaction ledger under a  $(1/2, m, \cdot)$ -bounded adversary for any  $m$ .*

Corollary 1 thus implies that Assumption 2.1 is not only sufficient but also necessary.

## 4 BASELINE APPROACHES

In this section, we provide a couple of natural designs of hybrid PoW/PoS protocols, which fail to achieve the full goal as Minotaur. **Idea 1: Securing PoW chain via checkpointing.** The checkpointed ledger [19] employs an external set of parties or a committee chosen randomly from the pool of stakeholders to assist a PoW ledger by finalizing blocks shortly after their creation. The finalized blocks are called checkpoints and the final ledger is formed by the chain of checkpoints. This mechanism can also secure a PoW ledger in the presence of adversarial mining majorities, but the security is solely guaranteed by the external set or the stakeholders.

There also exist other checkpointing protocols [6, 28, 33, 36] achieving properties that are not possible with pure PoW protocols, such as accountability and finality. However, these protocols require an honest majority (or even supermajority) on both the set of miners and the set of stakeholders.

**Idea 2: Smooth interpolation among PoW and PoS blocks.** In the static setting (where both the total mining power and total active stake are fixed and known to the protocol designer), there is a simple protocol that can also achieve the regions defined by the red line in Figure 1b. In this protocol, PoW and PoS mining occur in parallel, following the longest chain rule. In the initialization phase, we tune the mining targets such that the mining rate (i.e., number of blocks produced per unit time) of PoW blocks is  $\omega f$  and

the mining rate of PoS blocks is  $(1 - \omega)f$ . In the execution phase, whichever miners or stakeholders succeeds first, it goes ahead and extends the longest chain.

Compared with a pure PoS protocol (e.g., Ouroboros Praos [3]), the adversary in this protocol has strictly smaller action space because it cannot equivocate with the PoW blocks. Therefore, the security of this protocol follows directly from the security of Ouroboros Praos via either the forkable string argument [3] or the Nakamoto block method developed in [9]. We remark that the existence of this simple protocol makes the 2-hop blockchain [11] less interesting as it only works in the static setting and does not even allow weighing PoW and PoS.

However, it is very hard to extend this simple idea to the non-static setting, particularly with variable mining power. A natural approach to support variable mining power is to have fixed-length epochs and adjust the mining target of PoW blocks every epoch as in Bitcoin. From the simple protocol described above, we learn that in order to guarantee security in the non-static setting, we need to make sure that the ratio of total mining rates of PoW blocks and PoS blocks is a constant  $\omega/(1 - \omega)$  in every epoch. However, this is impossible to achieve in the variable mining power setting because the adversary can always decide to hide or release its PoW blocks so that there is no way to estimate the total mining power accurately. Inaccurate PoW mining target adjustment could lead to a different weighing parameter  $\omega$  of the security region than the desired one. More importantly, the value of  $\omega$  can be easily manipulated by the adversary (via selectively publishing its PoW blocks), and is unknown to the honest players.

## 5 THE FULL PROTOCOL

In this section, we provide a detailed description of our protocol Minotaur. The protocol is built on an epoch-based PoS longest chain protocol (e.g., Ouroboros [21], Ouroboros Praos [8], Ouroboros Genesis [3]). Recall that in a epoch-based PoS protocol, time is divided into multiple epochs, each with a fixed number of slots. In each slot, one or multiple stakeholders are selected as block proposers by a PoS ‘lottery’ process. In this process, the probability of being selected is proportional to the relative stake a stakeholder has in the system, as reported by the blockchain itself.

In contrast to the original Ouroboros protocols, the block-proposer schedule accounts for both resources by means of *virtual stake*, which is a combination of the *actual stake* (representing stake as in the original protocol), and *work stake* representing the share in block-production rights attributed to the PoW resource.

Additionally, PoW miners participate by mining *endorser blocks* (along the lines of ‘endorser inputs’ [21] or ‘fruits’ [31]). These endorser blocks are not directly appended to the main chain, but are to be referenced by future main-chain blocks (i.e., PoS blocks scheduled by means of virtual stake). Each epoch is assigned a certain amount of work stake that is assigned to the PoW miners who succeed in mining PoW blocks; the work stake assigned to the respective miners is proportional to their contribution of PoW blocks referenced from the main chain during an epoch. PoS block production rights per epoch are then assigned by considering the sum of actual stake (contributed by tokens) and work stake (contributed by PoW blocks).

Note that the purpose of the PoW endorser blocks is to measure work fairly, not implying that they also need to carry the ledger transactions. Transaction inclusion is orthogonal to this aspect, and can still be implemented along the lines of Bitcoin (transaction inclusion in the main-chain blocks) or Input-Endorsers/Fruitchains (transaction inclusion in the endorser blocks), or variants thereof.

We now treat the underlying PoS protocol as a black-box and present the protocol in detail. The protocol runs in fixed-time epochs with  $R$  slots each.

- PoW mining: In slot  $sl$  of epoch  $e$ , a miner mines a PoW block if

$$H(pk, sl, h, mr, nonce) < T_e,$$

where  $pk$  is the public key of the miner,  $h$  is the hash of the last confirmed PoS block (according to the confirmation rule of the underlying PoS protocol),  $mr$  is the Merkle root of the payload,  $T_e$  is the mining target in epoch  $e$ . Like in PoW blockchains, miners try different values of  $nonce$  to solve this hash puzzle. Following the notation in [18], we define the *difficulty* of this PoW block to be  $1/T_e$ . Besides possible transactions, also the public key of the miner is included in the payload. A PoW block is called recent in current slot  $sl_{\text{now}}$  if it refers to a confirmed PoS block mined no earlier than slot  $sl_{\text{now}} - sl_{\text{re}}$ , where  $sl_{\text{re}}$  is called the *recency parameter*.

- PoS mining: in slot  $sl$  of epoch  $e$ , one or multiple stakeholders are selected to propose PoS blocks extending the best chain (according to the chain selection rule). The selection of PoS block proposers uses the same mechanism as in the underlying PoS protocol. But for each node, the probability of being selected is proportional to its relative “virtual” stake, instead of the relative actual stake. The virtual stake is the sum of actual stake and work stake. At the beginning of each epoch, we set the total work stake equal to  $\frac{\omega}{1-\omega}$  times the total actual stake in the system. And the work stake distribution used in epoch  $e$  is set to be the distribution of block difficulties from PoW blocks referred by PoS blocks in epoch  $e - 2$ .
- Adjust the PoW mining target: For epoch  $e$ ,  $T_e$  is adjusted according to  $D_{e-1}^{\text{total}}$  the total difficulty of PoW blocks referred by PoS blocks in slots  $[(e - 2) * R - k + 1, (e - 1) * R - k]$ , i.e.,  $T_e = f^{(w)}R/D_{e-1}^{\text{total}}$ , where  $f^{(w)}$  is a protocol parameter representing the expected PoW mining rate in number of blocks per slot. (Specially, for  $e = 2$ ,  $T_2 = f^{(w)}(R - k)/D_1^{\text{total}}$ .) Note that around the boundary of two epochs  $e - 1$  and epoch  $e$ , the adversary has the option to use targets from both epochs, as PoW blocks don’t have accurate timestamps.
- Chain selection rule: due to a long-range attack (see Appendix C), the longest chain rule would fail. We use the following chain selection rules, which have different security guarantees.
  - maxvalid-mc (from Ouroboros Praos [8], needs a trusted third party for bootstrapping, i.e., assumption (P3)): it prefers longer chains, unless the new chain forks more than  $k$  blocks relative to the currently held chain (in which case the new chain would be discarded).
  - maxvalid-bg (from Ouroboros Genesis [3], supports bootstrapping from genesis block): it prefers longer chains, if

the new chain  $C'$  forks less than  $k$  blocks relative to the currently held chain  $C$ . If  $C'$  did fork more than  $k$  blocks relative to  $C$ ,  $C'$  would still be preferred if it grows more quickly in the  $s$  slots following the slot associated with the last common block of  $C$  and  $C'$ .

## 6 SECURITY ANALYSIS

### 6.1 Main theorems

Under different security models and assumptions, we prove the security (in particular, *persistence* and *liveness* as defined in §2.2) of Minotaur constructed with three various versions of Ouroboros PoS protocols.

**THEOREM 6.1.** *Under Assumption 1, Assumption 2 and assumptions (P1)–(P3), when executed in a lock-step synchronous model (i.e.,  $\Delta$ -synchronous model with  $\Delta = 1$ ), Minotaur constructed with Ouroboros generates a transaction ledger that satisfies persistence and liveness with overwhelming probability.*

**THEOREM 6.2.** *Under Assumption 1, Assumption 2 and assumptions (P1)–(P3), when executed in a  $\Delta$ -synchronous model, Minotaur constructed with Ouroboros Praos generates a transaction ledger that satisfies persistence and liveness with overwhelming probability.*

**THEOREM 6.3.** *Under Assumption 1, Assumption 2 and assumptions (P1)–(P2), when executed in a  $\Delta$ -synchronous model, Minotaur constructed with Ouroboros Genesis generates a transaction ledger that satisfies persistence and liveness with overwhelming probability.*

### 6.2 Security with Ouroboros

We first provide a proof sketch for Theorem 6.1. All the parameters used in our analysis are listed in Table 1.

$h_r$	number of honest PoW queries in slot $r$
$n_r$	number of total PoW queries in slot $r$
$\alpha_0$	lower bound on the fraction of honest mining power
$sl_{\text{re}}$	recency parameter for PoW blocks
$\Delta$	network delay in slots
$\kappa$	security parameter; length of the hash function output
$R \in \mathbb{N}$	duration of an epoch in number of slots
$(\gamma, s)$	restriction on the fluctuation of the number of honest queries across slots (Definition 1)
$f^{(s)}$	expected PoS mining rate in number of blocks per slot
$f^{(w)}$	expected PoW mining rate in number of blocks per slot
$\epsilon$	quality of concentration of random variables
$\sigma$	PoW fairness parameter & advantage of honest parties (Assumption 1.1 and 2.1)
$\delta$	“goodness” parameter of an epoch/slot (Definition 6)
$\lambda$	typicality parameter of the execution
$\ell$	minimum number of slots for concentration bounds
$L$	total number of epochs in the execution
$\omega$	the weighing parameter

**Table 1: The parameters used in our analysis.**

**PROOF SKETCH.** The proof relies on two important building blocks:

- 1) **Single-epoch argument.** Given honest majority in the virtual stake (normalized work-stake + normalized actual-stake) used in one epoch, we prove the security properties (CP and  $\exists$ CQ) for this single epoch. For this, we use the so-called *forkable strings* technique developed in the Ouroboros papers [3, 8, 21].
- 2) **Fruitchains argument.** Given CP and  $\exists$ CQ of the PoS chain in an epoch, we prove *fairness* of PoW blocks in this epoch, i.e., a miner controlling a  $\phi$  fraction of the computational resources will contribute a  $\phi$  fraction of work. This is similar to the analysis on the variable difficulty Fruitchains [38].

Using these building blocks, we prove the security with maxvalid- $\mu$  inductively. Under the honest majority assumption in the initial stake distribution (Assumption 1.1), we can prove security of the PoS chain in epochs 1&2. Then applying the Fruitchains argument, we get fairness of PoW blocks in epochs 1&2. Combining with Assumption 2.1, i.e.,  $\beta_w + \beta_s < 1$  (for any long enough window), we have honest majority in the virtual stake distribution used in epoch 3&4. And the proof goes on till the last epoch (Figure 3).  $\square$

### 6.2.1 Single-epoch security.

**THEOREM 6.4 (THEOREM 3 FROM [8]).** *Let  $\kappa, R, \Delta \in \mathbb{N}$  and  $\sigma \in (0, 1)$ . Let  $\beta_v$  be the fraction of adversarial virtual stake satisfying*

$$\beta_v \leq 1/2 - \sigma$$

*for some positive constant  $\sigma$ . Then the probability that the adversary violates CP with parameter  $\ell_{cp} = \kappa$  and  $\exists$ CQ with parameter  $\ell_{cq} = \kappa$  throughout a period of  $R$  slots is no more than  $Re^{-\Omega(\kappa)}$ . The constant hidden by the  $\Omega(\cdot)$ -notation depends on  $\sigma$ .*

**6.2.2 Fruitchains argument.** Recall that  $h_r$  is the number of honest PoW queries in slot  $r$  and  $n_r$  is the number of total PoW queries in slot  $r$ . For a set of slots  $S$ , we define  $h(S) = \sum_{r \in S} h_r$  and  $n(S) = \sum_{r \in S} n_r$ . In order to obtain meaningful concentration bounds on the number of PoW blocks in one epoch, there is a lower bound  $\alpha_0$  on the fraction of honest mining power, i.e.,  $h_r \geq \alpha_0 n_r$  for all  $r$ .

In the analysis of this subsection, we assume the main chain (PoS chain) satisfies properties CP with parameter  $\ell_{cp} = \kappa$  and  $\exists$ CQ with parameter  $\ell_{cq} = \kappa$ . By the common prefix property, for a PoS chain  $C$  held by an honest node at slot  $r$ , the prefix  $C^{\lceil \kappa}$  are stabilized, so to mine a PoW block at slot  $r$  an honest miner will refer the tip of  $C^{\lceil \kappa}$  as the last confirmed PoS block. And we set the recency parameter  $sl_{re} = 3\kappa + \Delta$ , i.e., a PoW block  $B_w$  is recent w.r.t. a chain  $C$  at slot  $r$  if the confirmed PoS block referred by  $B_w$  is in  $C$  and has timestamp at least  $r - 3\kappa - \Delta$ . With this selection of the recency parameter, we can prove the following key property of the protocol: any PoW block mined by an honest miner will be incorporated into the stabilized chain (and thus never lost). We refer to this as the *Fruit Freshness Lemma*—PoW blocks stay “fresh” sufficiently long to be incorporated.

**LEMMA 6.5 (FRUIT FRESHNESS).** *Suppose the PoS chain satisfies properties CP with parameter  $\ell_{cp} = \kappa$  and  $\exists$ CQ with parameter  $\ell_{cq} = \kappa$ . Then, if  $sl_{re} = 3\kappa + \Delta$ , an honest PoW block mined at slot  $r$  will be included into the stabilized chain before slot  $r + r_{wait}$ , where  $r_{wait} = 2\kappa + \Delta$ .*

**PROOF.** Suppose an honest PoW block  $B_w$  is mined at slot  $r_0$  while the PoS chain  $C_0$  is adopted, then  $B_w$  will reference the tip of  $C_0^{\lceil \kappa}$  as the last confirmed PoS block. Further, by the  $\exists$ CQ property, the tip of  $C_0^{\lceil \kappa}$  has timestamp  $r_1 \geq r_0 - 2\kappa$ . By slot  $r_0 + \Delta$ , all honest nodes will receive  $B_w$ . Let  $r_2 = r_0 + 2\kappa + \Delta$  and  $C$  be any chain held by an honest node at slot  $r_2$ , then again by the  $\exists$ CQ property, there exists at least one honest block  $B_s$  on  $C$  whose timestamp  $r_3$  is in the interval  $[r_2 - 2\kappa, r_2 - \kappa)$ . We check that  $B_w$  is still recent at slot  $r_3$  as

$$r_3 - r_1 < (r_2 - \kappa) - (r_0 - 2\kappa) = 3\kappa + \Delta = sl_{re}.$$

As  $B_s$  is an honest block mined after  $r_0 + \Delta$ ,  $B_s$  or an ancestor of  $B_s$  must include  $B_w$ . And since  $B_s$  is stabilized in  $C$  at slot  $r_2$ , we have that  $r_{wait} = r_2 - r_0 = 2\kappa + \Delta$ .  $\square$

In a  $(\gamma, s)$ -respecting environment, we have the following useful proposition.

**PROPOSITION 1 (PROPOSITION 2 FROM [18]).** *In a  $(\gamma, s)$ -respecting environment, let  $U$  be a set of at most  $s$  consecutive slots and  $S \subseteq U$ . Then, for any  $h \in \{h_r : r \in U\}$  and any  $n \in \{n_r : r \in U\}$  we have*

$$\frac{h}{\gamma} \leq \frac{h(S)}{|S|} \leq \gamma h, \quad \frac{n}{\gamma} \leq \frac{n(S)}{|S|} \leq \gamma n,$$

$$h(U) \leq \left(1 + \frac{\gamma|U \setminus S|}{|S|}\right)h(S), \quad n(U) \leq \left(1 + \frac{\gamma|U \setminus S|}{|S|}\right)n(S).$$

Recall that  $T_e$  is the mining target in epoch  $e$  determined by the stabilized segment of the chain from epoch  $e - 1$ . In order to obtain meaningful concentration bounds on the number of PoW blocks, we need  $T_e$  to be “reasonable” for each epoch. Similar to [16, 18], we define a notation of “good” epochs as follows. By abuse of notation, we write the expected PoW block rate  $f^{(w)}$  simply as  $f$  in the analysis below.

**DEFINITION 6.** *Epoch  $e$  is good if  $(1 - \delta)\alpha_0 f / \gamma \leq ph^{(e)}T_e \leq (1 + \delta)\gamma f$ , where  $p = 1/2^\kappa$  and  $h^{(e)} = h_{(e-1)R+1}$ , i.e., the number of honest queries in the first slot of epoch  $e$ . A slot  $r$  is good if it is in a good epoch.*

**PROPOSITION 2.** *Under a  $(\gamma, 2R)$ -respecting environment, if slot  $r$  is a good slot in epoch  $e$ , then  $(1 - \delta)\alpha_0 f / \gamma^2 \leq ph_r T_e \leq (1 + \delta)\gamma^2 f$ .*

**PROOF.** By the definition of a good epoch, we have  $(1 - \delta)\alpha_0 f / \gamma \leq ph^{(e)}T_e \leq (1 + \delta)\gamma f$ . And under a  $(\gamma, 2R)$ -respecting environment and the fact that slot  $r$  is in epoch  $e$ , we know  $1/\gamma \leq h_r/h^{(e)} \leq \gamma$ . Therefore,  $(1 - \delta)\alpha_0 f / \gamma^2 \leq ph_r T_e \leq (1 + \delta)\gamma^2 f$ .  $\square$

Now, we prove *fairness* of PoW blocks in any large enough window, i.e., a miner controlling a  $\phi$  fraction of the computational resources will contribute a  $\phi$  fraction of work. We first give a form definition of fairness.

**DEFINITION 7.** *We say that  $\mathcal{H}$  is a  $\phi$ -fraction honest subset if miners in  $\mathcal{H}$  (that may change over time) are honest and  $n_r^{\mathcal{H}} > \phi n_r$  for any slot  $r$ , where  $n_r^{\mathcal{H}}$  is number of queries in  $\mathcal{H}$  at slot  $r$ .*

**DEFINITION 8 (FAIRNESS).** *We say that the protocol has (approximate)  $(W_0, \sigma)$ -fairness if, for any  $\phi, \alpha_0 \leq \phi < 1$ , any  $\phi$ -fraction honest subset  $\mathcal{H}$ , and any honest miner holding chain  $C$  at slot  $r$  and any interval  $S_0 \subseteq [0, r]$  with at least  $W_0$  consecutive slots, it holds that the PoW blocks included in  $C(S_0)$  mined by  $\mathcal{H}$  have total difficulty*




Figure 3: Proof sketch of §6.2.

at least  $(1 - \sigma)\phi d$ , where  $d$  is the total difficulty of all PoW blocks included in  $C(S_0)$ .

Define the random variable  $D_r$  equal to the sum of the difficulties of all PoW blocks computed by honest miners at slot  $r$ . And for fixed  $\phi$  and  $\phi$ -fraction honest subset  $\mathcal{H}$ , define the random variable  $D_r^{\mathcal{H}}$  equal to the sum of the difficulties of all PoW blocks computed by miners in  $\mathcal{H}$  at slot  $r$ . For a set of slots  $S$ , we define  $D(S) = \sum_{r \in S} D_r$ ,  $D^{\mathcal{H}}(S) = \sum_{r \in S} D_r^{\mathcal{H}}$ , and  $n^{\mathcal{H}}(S) = \sum_{r \in S} n_r^{\mathcal{H}}$ . For a set of  $J$  adversarial queries, define the random variable  $A(J)$ , as the sum of difficulties of all the PoW blocks created during queries in  $J$ .

Next we define the notion of typical executions, which will be shown to occur with overwhelming probability.

DEFINITION 9 (TYPICAL EXECUTION). *An execution  $E$  is typical if the following hold*

(a) *For any set  $S$  of at least  $\ell$  consecutive good slots,*

$$D(S) < (1 + \epsilon)ph(S).$$

(b) *For any set  $S$  of at least  $\ell$  consecutive good slots, let  $J$  be the set of adversarial queries in  $S$ . If we further know each query in  $J$  made at slot  $r$  with target  $T$  satisfies  $(1 - \delta)\alpha_0 f / \gamma^2 \leq pn_r T \leq (1 + \delta)\gamma^2 f / \alpha_0$ , then*

$$D(S) + A(J) < (1 + \epsilon)p(h(S) + |J|).$$

(c) *For any set  $S$  of at least  $\ell/\phi$  consecutive good slots,*

$$D^{\mathcal{H}}(S) > (1 - \epsilon)pn^{\mathcal{H}}(S).$$

To obtain meaningful concentration, we should be considering a sufficiently long slot sequence of at least

$$\ell \triangleq \frac{2(1 + \epsilon/3)}{\epsilon^2 \gamma^3 (1 - \delta) \alpha_0 f} \lambda,$$

where  $\lambda$  is called the typicality parameter of the execution.

For our analysis to go through, the protocol parameters should satisfy certain conditions which we now discuss. First, we will require that the number  $\ell$  defined above and the security parameter  $\kappa$  are appropriately small compared to  $R$ , the duration of an epoch.

$$R - 3\kappa - \Delta \geq \ell/\alpha_0 \geq \frac{\gamma}{\epsilon}(4\kappa + \Delta). \quad (\text{C1})$$

Note that (C1) implies  $R \geq (4\kappa + \Delta)/\epsilon$ . Second, the slack variables  $\epsilon$  and  $\delta$  should satisfy

$$4\epsilon \leq \delta \leq 1. \quad (\text{C2})$$

Next we bound the probability of an atypical execution (Lemma 6.6) and show that all epochs are good in a typical execution by induction (Lemma 6.7). Therefore, all epochs are good with overwhelming probability.

LEMMA 6.6. *For an execution  $\mathcal{E}$  of LR slots, in a  $(\gamma, 2R)$ -respecting environment, the probability of the event “ $\mathcal{E}$  is not typical” is bounded by  $\mathcal{O}(LR)e^{-\lambda}$ .*

LEMMA 6.7. *For a typical execution in a  $(\gamma, 2R)$ -respecting environment, all epochs are good.*

The proofs of Lemma 6.6 and Lemma 6.7 are in Appendix D.

THEOREM 6.8 (FAIRNESS). *For a typical execution in a  $(\gamma, 2R)$ -respecting environment, the protocol with recency parameter  $sl_{re} = 3\kappa + \Delta$  satisfies  $(W_0, \sigma)$ -fairness, where  $W_0 = \ell/\alpha_0 + 3\kappa + \Delta$  and  $\sigma = 4\epsilon$ .*

PROOF. Fix  $\phi$ , a  $\phi$ -fraction honest subset  $\mathcal{H}$ , and an honest node holding chain  $C$ . Let  $S_0 = \{u : r_1 \leq u \leq r_2\}$  be a window of at least  $W_0$  consecutive slots. Let  $C(S_0)$  be the segment of  $C$  containing PoS blocks with timestamps in  $S_0$ , let  $\mathcal{B}$  be all PoW blocks included in  $C(S_0)$ , and  $d$  be the total difficulty of all PoW blocks in  $\mathcal{B}$ . Facts 1-3 from the proof of Lemma 6.7 will be useful for proving the theorem. Let  $S_1 = \{u : r_1 - (4\kappa + \Delta) \leq u \leq r_2 + \kappa\}$ ,  $S_2 = \{u : r_1 \leq u \leq r_2 - (3\kappa + \Delta)\}$ , and  $J$  be the set of adversary queries associated with  $\mathcal{B}$  in  $S_1$ . Then by Fact 1 and Fact 2, we have that all PoW blocks in  $\mathcal{B}$  are mined in  $S_1$ ; by Fact 3, we have that all PoW blocks mined by  $\mathcal{H}$  in  $S_2$  are in  $\mathcal{B}$ . Similar to the arguments in Lemma 6.7, for each query in  $S_1$  made by either an honest node or the adversary at slot  $r$  in epoch  $e$ , the target  $T$  must satisfy  $(1 - \delta)\alpha_0 f / \gamma^2 \leq pn_r T \leq (1 + \delta)\gamma^2 f / \alpha_0$ .

Further note that, to prove  $\sigma$ -fairness, it suffices to show that

$$D^{\mathcal{H}}(S_2) \geq (1 - \sigma)\phi(D(S_1) + A(J)).$$

Under a typical execution, we have

$$D^{\mathcal{H}}(S_2) > (1 - \epsilon)pn^{\mathcal{H}}(S_2) \geq (1 - \epsilon)\phi pn(S_2),$$

and

$$D(S_1) + A(J) < (1 + \epsilon)p(h(S_1) + |J|) = (1 + \epsilon)pn(S_1).$$

By our choice of  $W_0$ , we have  $|S_2| \geq \ell/\alpha_0 \geq \ell/\phi$ . Furthermore, we may assume  $|S_2| \leq 2R$ . This is because we may partition  $S_2$  in parts such that each part has size between  $\ell/\alpha_0$  and  $2R$ , sum over all parts to obtain the desired bound. Then by Proposition 1, we



have

$$\begin{aligned} n(S_1) &\leq \left(1 + \frac{\gamma|S_1 \setminus S_2|}{|S_2|}\right)n(S_2) \\ &\leq \left(1 + \frac{\gamma(8\kappa + 2\Delta)}{\ell/\alpha_0}\right)n(S_2) \\ &\stackrel{(C1)}{\leq} (1 + 2\epsilon)n(S_2). \end{aligned}$$

Finally, by setting  $\sigma = 4\epsilon$ , we conclude the proof.  $\square$

**6.2.3 Lifting argument from single-epoch to multiple-epoch.** Theorem 6.4 gives bounds for CP and  $\exists\text{CQ}$  for a single-epoch run of the protocol with static stake distribution and perfect randomness. We now conclude our proof of Theorem 6.1 by showing that these blockchain properties hold throughout the whole lifetime of the system consisting of many epochs.

**THEOREM 6.9 (RESTATEMENT OF THEOREM 6.1).** *Fix parameters  $\kappa, \lambda, \ell, \gamma, \sigma, \epsilon, \delta, R$  and  $L$  satisfying conditions (C1) and (C2). Under Assumption 1, Assumption 2 and assumptions (P1)-(P3), when executed in a lock-step synchronous model (i.e.,  $\Delta$ -synchronous model with  $\Delta = 1$ ), Minotaur constructed with Ouroboros generates a transaction ledger that satisfies persistence and liveness throughout a period of  $L$  epochs (each with  $R$  slots) with probability  $1 - RL(e^{-\Omega(\kappa)} + e^{-\lambda})$ .*

This part of the analysis proceeds similarly as in Section 5 of [21] and hence we only sketch it in Appendix D.

### 6.3 Security with Ouroboros Praos/Genesis

By maxvalid-bg, an honest node prefers longer chains, if the new chain  $C'$  forks less than  $k$  blocks relative to the currently held chain  $C$ . If  $C'$  did fork more than  $k$  blocks relative to  $C$ ,  $C'$  would still be preferred if it grows more quickly in the  $s$  slots following the slot associated with the last common block of  $C$  and  $C'$ . In this section, we analyze the security of the protocol with maxvalid-bg, but drop assumption (P3) (i.e., allow bootstrapping from genesis).

**PROOF SKETCH.** This part is identical to the analysis in Ouroboros Genesis (Section 4.3-4.4 of [3]), so we only sketch it here. The proof proceeds in two steps:

- 1) For an honest node  $h_1$  that is always online, we show that when replacing maxvalid-mc with maxvalid-bg, the overall execution of the protocol in  $h_1$ 's view remains the same except with negligible probability. That is to show that with overwhelming probability, whenever  $h_1$  receives a new chain  $C'$  that forks more than  $k$  blocks from  $h_1$ 's local chain  $C$ ,  $h_1$  will always favor  $C$ , i.e.,  $C$  grows more quickly than  $C'$  in the first  $s$  slots right after the fork.
- 2) For a newly joined node  $h_2$ , we consider a "virtual" node  $h_2'$  that holds no stake, but was participating in the protocol since the beginning and was honest all the time. Then we show  $h_2$  will always adopt the same chains as  $h_2'$  after joining the network.

$\square$

### 6.4 Comparison to Bitcoin

In this paragraph, we observe that Minotaur (when executed in the pure PoW case,  $\omega = 1$ ) is more robust against fluctuations in

PoW participation than Bitcoin. In particular, we demonstrate that Bitcoin may not enjoy liveness if the number of parties is allowed to halve every two weeks—while Minotaur does. The theoretical results of this section are accompanied by experiments; see §7.2 and Figure 5.

Consider Bitcoin for the case that, initially, there are  $n$  honest parties and that the target is  $T$ . Using the model and notation of [16] with  $q = 1$ ,  $m$  is the number of blocks in an epoch and  $p = 1/2^\kappa$ . Thus, each slot is successful with probability  $f = 1 - (1 - pT)^n \approx pTn$ . We will show that for  $s = \lfloor m/pTn \rfloor$ , liveness may fail in a  $(2, s)$ -respecting environment. In such an environment, the adversary is allowed to halve the number of parties every  $s \approx m/f$  slots (note that  $m/f$  is the expected duration of an epoch, which for Bitcoin is two weeks).

The attack proceeds in stages of  $s$  slots, during which the adversary does not mine nor delays messages. Thus, in the beginning of each slot all parties have a chain of the same length. The adversary halves the number of parties at the beginning of each stage. Let  $N_i$  denote the number of honest parties in stage  $i$ , with  $N_0 = n$  and  $N_i = n/2^i$  for  $i > 0$ . The expected number of successful slots in stage  $i$  is

$$s \cdot [1 - (1 - pT)^{N_i}] < spTN_i = \frac{spTn}{2^i} \leq \frac{m}{2^i}.$$

Consider the stage  $j$  for which

$$k - 1 < \frac{m}{2^j} \leq 2(k - 1).$$

The expected number of blocks computed in the first  $j$  stages is

$$\frac{m}{2} + \frac{m}{2^2} + \dots + \frac{m}{2^j} = m - \frac{m}{2^j} < m - k + 1$$

and at most  $k - 1$  in stage  $j + 1$ . Recalling that the median of a binomial distribution with mean  $\mu$  is at most  $\lceil \mu \rceil$ , we obtain that with probability at least  $1/4$  at most  $m$  blocks have been computed by the end of stage  $j + 1$  and at most  $k - 1$  of them have been computed in the last stage. Assuming the liveness parameter  $u < s$ , any transaction provided to all honest parties for the first  $u$  slots of the final stage, will be in depth less than  $k$ . It follows that liveness does not hold in the final stage of the attack.

**PROPOSITION 6.10.** *Bitcoin's ledger does not satisfy liveness in a  $(2, s)$ -respecting environment, for any  $u < s$  and  $s < m/pTn$ .*

In contrast, Minotaur does not suffer from such fluctuation. Intuitively, this is because it inherits its security from the underlying proof-of-stake protocol. In particular, the epochs are of fixed duration, implying that target recalculation occurs regularly. This is evident from Condition C1, which allows much greater values for  $R$  and  $\gamma$ .

## 7 EXPERIMENTS

We implemented a prototype instantiation of a Minotaur client in Rust and the code can be found at [25]. We also implemented a Bitcoin client as a benchmark. In this section, we briefly describe the architecture of our implementation and present experimental results to evaluate the concrete performance of Minotaur under different scenarios.

## 7.1 Architecture

We implemented the Input-Endorsers/Fruitchains variant of Minotaur, where transactions are exclusively included in the PoW blocks. In this way, Minotaur can potentially achieve optimal throughput up to the capacity of the network as shown by the implementation of Prism [39], a pure PoW consensus protocol with similar blockchain structure. The system architecture is illustrated in Figure 4. Functionally it can be divided into the following three modules:

- (1) *Block Structure Manager*, which maintains the client’s view of the blockchain, and communicates with peers to exchange new blocks.
- (2) *Miner*, which assembles new PoW blocks.
- (3) *Staker*, which assembles new PoS blocks.

The goal of a Minotaur client is to maintain up-to-date information of the blockchain. They are stored in the following three data structures:

- (1) *Block Structure Database*, stores the graph structure of the blockchain (i.e., a directed acyclic graph (DAG) formed with PoW and PoS blocks).
- (2) *Memory Pool*, stores the set of transactions that have not been mined into any PoW block.
- (3) *PoW Block Pool*, stores the set of PoW blocks that have not been referenced by any PoS block.

At the core of the **Block Structure Manager** is an *event loop* which sends and receives network messages to/from peers, and a *worker thread pool* which handles those messages. When a new block arrives, the worker thread first checks its proof-of-work/proof-of-stake, and if valid, then proceeds to relay the block to peers that have not received it. Next, the worker thread checks whether all blocks referenced by the new block, e.g. its parent, are already present in the Block Structure Database. If not, it buffers the block in an in-memory queue and defers further processing until all the referenced blocks have been received. Finally, the worker validates the block (e.g., verifying transaction signatures), and inserts the block into the Block Structure Database. If the block is a PoW block, the Block Structure Manager checks the Memory Pool against the transactions in this new block and removes any duplicates or conflicts from the Memory Pool, and also puts the block into the PoW Block Pool.

The **Miner** module assembles new PoW blocks. It is implemented as a busy-spinning loop. At the start of each round, it polls the Block Structure Database and the Memory Pool to update the block it is mining. When a new PoW block is mined, it will be inserted into the Block Structure Database, then sent to peers by the Block Structure Manager. The Memory Pool and PoW Block Pool will also be updated accordingly.

The **Staker** module works similarly, but assembles new PoS blocks. At the start of each round, it polls the Block Structure Database and the PoW Block Pool to update the block it is assembling. When a new PoS block is generated, it will be inserted into the Block Structure Database, then sent to peers by the Block Structure Manager. The PoW Block Pool will also be updated accordingly.




Figure 4: Architecture of our Minotaur client implementation.

## 7.2 Performance under variable mining power

In §6.4, we have seen that Minotaur can survive more drastic variations of network hash power. We design the following experiment on our Minotaur/Bitcoin codebase to verify this argument.


Recall that both Bitcoin and Minotaur have epoch-based target adjustment rules, varying the difficulty target of block mining based on the median inter-block time from the previous epoch. However, the epoch length in Minotaur is defined as a fixed number of slots instead of a fixed number of blocks (e.g., 2016 blocks in Bitcoin). This definition makes sense in Minotaur because the main-chain blocks, which are PoS blocks, always have accurate timestamps (even if they are proposed by an adversary). In this experiment, we will see that this small change allows Minotaur to enjoy better liveness than Bitcoin when the total mining power is decreasing.

In our experiment, we set the epoch length in Bitcoin to be 400 blocks and the expected duration of an epoch is 2 minutes; while the epoch length of Minotaur is 2 minutes and the expected number of blocks in a epoch is 400 blocks. Then Bitcoin and Minotaur are experiencing the same mining power variation, i.e., starting with an insufficient mining rate (100 blocks per minute) and then halving the mining rate every epoch (2 minutes in the experiments). From Figure 5, we can see that the mining target never has a chance to be adjusted in Bitcoin as the length of the chain can never reach 400 blocks, which leads to a liveness failure; while the mining targets are adjusted every epoch (2 minutes) in Minotaur to make sure that the number of PoW blocks referenced by the main chain grows linearly.

## 7.3 Performance under attacks

In the following experiments, we evaluate how Minotaur performs in the presence of active attacks. Specifically, we consider two types of attacks: spamming and selfish mining attacks. Spamming attack aims to reduce network throughput, while selfish mining aims to compromise the PoW fairness of the protocol (i.e., reducing the work stake of honest miners).

**Spamming attack.** In a spamming attack, attackers flood the network with conflicting transactions. Among a bunch of conflicting transactions, at most one of them could be executed successfully. For example, in an account model, transactions with identical sender and nonce are conflicting transactions. Compared to Bitcoin, the Input-Endorsers/Fruitchains variant of Minotaur cannot ensure that no conflicting transactions enter the transaction ledger since PoW are mined *before* picked by main chain. Therefore, goodput




**Figure 5: Behavior of Bitcoin and Minotaur under decreasing mining power. When the total mining power is halved every two minutes, the Bitcoin chain stops growing as it can never reach its epoch size (400 blocks), while the number of PoW blocks referenced by the main chain grows linearly in Minotaur.**

(i.e., throughput of successfully executed transactions) of Minotaur transaction ledger will be comprised if conflicting transactions enter PoW blocks. To mitigate this attack, we propose a spam filter that operates as follows: a miner validates transactions with respect to the transaction ledger, PoW blocks that haven't been picked by main chain, and preceding pending transactions of this miner. In this way, if one of the conflicting transactions enters a PoW block and propagated throughout the network before other conflicting ones enter a PoW block, miners will filter out other conflicting transactions and increase goodput.

We implemented the spam filter and did experiments with four nodes in a line topology. Our metrics is *normalized spam*, which is the count of unsuccessfully executed transactions normalized by the count of transactions in ledger. In our experiments, we make transaction generators create conflicting transactions on purpose, and send them to two nodes sitting at the ends of the line topology. And other nodes receive non-conflicting transactions. Hence, without any spam precaution, normalized spam is 0.25. We set the block generation rate to be 0.44 block/s for PoS blocks and 1.39 block/s for PoW blocks. The results of our experiments with additional peer-to-peer latency is shown in Figure 6, and we observe that although normalized spam increases as additional latency increases, it is still lower than that without spam filter. The highest spam in our experiments is 0.063, far lower than 0.25. Noticed that in that experiment we apply an additional peer-to-peer latency of 0.3s, which is larger than most peer-to-peer latency in Bitcoin and Ethereum [22]. This means that our spam filter can reduce spam to minimal effect.

**Selfish mining attack.** It has been known that Bitcoin is vulnerable to the selfish mining attack [13, 27, 34], where a selfish miner withhold its mined blocks and release them later at an appropriate time to take the place of honest blocks in the longest chain. This attack hurts the fairness of the protocol, in the sense that the selfish



**Figure 6: Normalized spam in 4-node experiments with additional peer-to-peer latency. Spam is reduced mostly under various additional latencies.**

miner with  $\beta_w$  fraction of mining power can have more than  $\beta_w$  fraction of blocks in the main chain so that it will reap higher revenue. The Fruitchains protocol [30] was proposed as a solution to selfish mining. We observe that the selfish mining attack won't work on Minotaur since it has similar blockchain structure as Fruitchains. To verify this observation, we implemented the following selfish mining attacks on our Bitcoin/Minotaur code base:

- **On Bitcoin:** The selfish miner always mines on the block at the tip of the longest chain, whether the chain is private or public. Upon successful mining, the adversary maintains the block in private to release it at an appropriate time. In particular, when an honest miner publishes a block, the selfish miner will release a previously mined block at the same level (if it has one).
- **On Minotaur:** The selfish miner plays the same withholding-and-releasing strategy on its PoS blocks as above, while its PoS blocks contain PoW blocks mined by itself exclusively.


We assume honest nodes will choose the attacker's block with probability  $p$  whenever there is a tie. In our experiments, we set  $p > 0.5$  to simulate the case that the attacker has better network connection than honest nodes so its blocks are usually transmitted faster. Table 2 presents fractions of PoW blocks in the main chain of Bitcoin/Minotaur under various adversarial power (combinations of  $\beta_w$  and  $p$ ). On Minotaur, the attacker controls  $1/3$  of the virtual stake, i.e.,  $\omega\beta_w + (1 - \omega)\beta_s = 1/3$ , where  $\beta_s$  is the fraction of the actual stake controlled by the attacker. From Table 2, we can see that Bitcoin is indeed vulnerable to selfish mining, meaning that the attacker has more fraction of PoW blocks in the main chain (which means more block rewards in the context of Bitcoin) than its fraction of mining power, particularly when  $\beta_w \geq 0.5$ . Meanwhile, in Minotaur, the attacker can only have  $\beta_w$  fraction of PoW blocks in the main chain, no matter how much mining power it has, and how the network favors it. This property is crucial to the security of Minotaur as we have seen in §6.2.2.

## 8 DISCUSSION

**Variable weighing parameter.** In our security analysis, we assume a fixed weighing parameter  $\omega$  through the whole execution of the protocol. However, we point out that  $\omega$  can change over time as

Attacker's power		Fraction of PoW blocks				
$\beta_w$		0.75	0.67	0.50	0.33	0.25
$p = 1$	Bitcoin	0	0	0.007	0.513	0.669
	Minotaur	0.248	0.332	0.498	0.665	0.746
$p = 0.7$	Bitcoin	0	0.001	0.098	0.618	0.72
	Minotaur	0.248	0.333	0.499	0.666	0.750

**Table 2: Fractions of honest PoW blocks in the main chain under selfish mining attacks in Bitcoin/Minotaur. In these experiments, the attacker controls  $\beta_w$  fraction of mining power and the tie breaking rule favors the attacker's block with probability  $p$ .**



**Figure 7: An example of variable weighing parameter  $\omega$  in Minotaur. Suppose the value of  $\omega$  changes from  $1/2$  to  $2/5$  at the onset of epoch  $e_0 + 1$ . Then the adversarial stake/mining power needs to be restricted below the red line in epochs  $e \leq e_0 - 1$  and below the blue line in epochs  $e \geq e_0 - 1$ . In particular, the adversary is restricted by both lines in epoch  $e_0 - 1$ . When  $\omega$  changes smoothly over time, it is reasonable to assume exactly this.**

long as we put proper assumptions on the adversarial stake/mining power. Suppose  $\omega(e)$  is the weighing parameter of epoch  $e$ , i.e., the ratio of work stake and actual stake in epoch  $e$  is set to be  $\omega(e)/(1 - \omega(e))$ . Note that the function  $\omega(e)$  may be decided by the protocol designers and hard-coded in the genesis block, but the players can also reach an agreement (off-chain) to update it by doing a soft fork.

We give a brief security sketch, deferring the full analysis to future work. Recall that, in our protocol, the virtual stake of a player in epoch  $e$  composes of two parts: the actual stake drawn from the last PoS block in epoch  $e - 2$ ; and the work stake that is proportional to the amount of work it produced during epoch  $e - 2$ . Thus, to guarantee the security in epoch  $e$ , we need the adversary to be  $(1/2 - 2\sigma, 2R, \omega(e))$ -bounded (see Definition 2) in epoch  $e - 2$  and epoch  $e - 1$ . And similarly, the adversary needs to be  $(1/2 - 2\sigma, 2R, \omega(e + 1))$ -bounded in epoch  $e - 1$  and epoch  $e$ . Therefore, the adversarial stake/mining power must be restricted by both above bounds for epoch  $e - 1$ . As long as  $\omega(e + 1)$  does not differ too much from  $\omega(e)$ , this restriction on the adversary is reasonable to assume, and the weighing parameter can transit smoothly from  $\omega(e)$  to  $\omega(e + 1)$ . Figure 7 gives an example of how the assumptions shift when  $\omega(e)$  is updated at the onset of epoch  $e_0 + 1$ .

We point out that such flexible weighing between work and stake is very useful in practice. For example, in general, PoS blockchains are easy to launch with existing techniques, such as proof-of burn [20], initial coin offering [24] and airdrop [1]. Therefore, Minotaur can be launched as a pure PoS blockchain and later transit into a hybrid PoW/PoS one or a pure PoW blockchain. In addition, the security of Minotaur can be enhanced by assigning higher weight to the more decentralized resource.

**Removing the initial constraint.** As discussed in §2, we required an honest majority of stake for the two initial epochs (as there is no initial work-stake distribution). We note that we can remove this initial constraint by implementing the first two epochs using the smooth-interpolation approach (§4, Idea 2) where work blocks are mined based on 2-for-1 PoW [18] producing both, main-chain blocks (at a rate defined by  $\omega$ ) and endorser blocks (at a sufficient rate to guarantee fairness). This allows to start the protocol with any initial weighing by means of  $\omega$ . We defer a formal analysis to future work.

**Generalization to multiple resources.** Following the idea in [37], Minotaur can be extended to more than two resources. Note that, in contrast to their construction, this is achievable without fundamentally changing the structure of the protocol.

Assume  $M \geq 1$  different resources, and for each such resource, an independent lottery mechanism among the contributors to assign ‘successes’ proportionally to a party’s ratio of the total contributed resource.

Defining a fixed amount of virtual stake  $V$  and weights  $\omega_i > 0$ ,  $\sum_{i=1}^M \omega_i = 1$ , we have resource  $i$  control  $\omega_i \cdot V$  of the virtual stake.

During each epoch  $e$ , the different lotteries are run concurrently, wherein each success allows for the release of a respective block (to be eventually picked up by a block of the main chain), tied to the respective resource.

Main-chain block leadership for epoch  $e + 2$  is assigned to resource  $i$  based on relative virtual stake  $\omega_i$ , and further split among the contributors of that resource based on their production of respective ‘resource blocks’ during epoch  $e$ .

Naturally, this protocol tolerates


$$\sum_{i=1}^M \omega_i \beta_i < \frac{1}{2}, \quad (1)$$

where  $\beta_i \in [0, 1]$  is the fraction of resource  $i$  held by the adversary. A pictorial example for  $M = 3$  is given in Fig. 8.

Although the underlying blockchain protocol is PoS-based, stake (in the classical, non-virtual sense) is not required to be one of the  $M$  resources, i.e., the protocol can be run solely on virtual stake. Also, full adversarial control of some of the resources can be tolerated as long as Eq. (1) is satisfied.

Finally, note that the resources can also be of the same type, e.g., Minotaur can combine work emanating from different hash functions such as SHA256, scrypt, and ethash—answering an open question raised in §1.

**Fairness with respect to the combined resources.** Fairness with respect to PoW blocks (Definition 8) is a crucial property of Minotaur to guarantee fair assignment of work stake to the miners. However, this is not sufficient for aspects like fair reward sharing [31], since, for this purpose, fairness needs to be achieved for



**Figure 8: Tolerable resource corruption for  $\omega = (\frac{1}{4}, \frac{1}{3}, \frac{5}{12})$ . The volume enclosed by the gray surfaces is tolerated by the protocol:  $\sum_{i=1}^3 \omega_i \beta_i < \frac{1}{2}$ .**

overall block production with respect to the *combined* resources. We note that fairness with respect to the combined resources can be achieved by scheduling endorser blocks for every involved resource, *including stake*. A given epoch reward can now be shared by assigning an  $\omega_i$  fraction of the reward to the contributors of the  $i$ -th resource, and distributing each such fraction proportionally to the parties' contributions of endorser blocks for that resource.

## 9 ACKNOWLEDGES

This research is supported in part by the US National Science Foundation under grants CCF-1705007 and CNS-1718270 and the US Army Research Office under grant W911NF1810332.

## REFERENCES

- [1] Airdrop. <https://en.bitcoinwiki.org/wiki/Airdrop>.
- [2] BACK, A. Hashcash—a denial of service counter-measure.
- [3] BADERTSCHER, C., GAZI, P., KIAYIAS, A., RUSSELL, A., AND ZIKAS, V. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security* (2018), pp. 913–930.
- [4] BENTOV, I., LEE, C., MIZRAHI, A., AND ROSENFELD, M. Proof of activity: Extending bitcoin's proof of work via proof of stake [extended abstract] y. *ACM SIGMETRICS Performance Evaluation Review* 42, 3 (2014), 34–37.
- [5] BLUM, E., KIAYIAS, A., MOORE, C., QUADER, S., AND RUSSELL, A. The combinatorics of the longest-chain rule: Linear consistency for proof-of-stake blockchains. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020* (2020), S. Chawla, Ed., SIAM, pp. 1135–1154.
- [6] BUTERIN, V., AND GRIFFITH, V. Casper the friendly finality gadget. *arXiv preprint arXiv:1710.09437* (2017).
- [7] CHEN, L., XU, L., GAO, Z., LU, Y., AND SHI, W. Protecting early stage proof-of-work based public blockchain. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)* (2018), IEEE, pp. 122–127.
- [8] DAVID, B., GAZI, P., KIAYIAS, A., AND RUSSELL, A. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques* (2018), Springer, pp. 66–98.
- [9] DEMBO, A., KANNAN, S., TAS, E. N., TSE, D., VISWANATH, P., WANG, X., AND ZEITOUNI, O. Everything is a race and nakamoto always wins. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security* (2020), pp. 859–878.

- [10] DUONG, T., CHEPURNOY, A., FAN, L., AND ZHOU, H.-S. Twinscoin: A cryptocurrency via proof-of-work and proof-of-stake. In *Proceedings of the 2nd ACM Workshop on Blockchains, Cryptocurrencies, and Contracts* (2018), pp. 1–13.
- [11] DUONG, T., FAN, L., KATZ, J., THAI, P., AND ZHOU, H.-S. 2-hop blockchain: Combining proof-of-work and proof-of-stake securely. In *European Symposium on Research in Computer Security* (2020), Springer, pp. 697–712.
- [12] ETHEREUM. Ethash. <https://web.archive.org/web/20220105100244/https://eth.wiki/en/concepts/ethash/ethash>, 2022.
- [13] EYAL, I., AND SIRER, E. G. Majority is not enough: Bitcoin mining is vulnerable. In *International conference on financial cryptography and data security* (2014), Springer, pp. 436–454.
- [14] FELLER, W. *An Introduction to Probability Theory and Its Applications, Volume 2. An Introduction to Probability Theory and Its Applications*. Wiley, 1967.
- [15] GARAY, J., AND KIAYIAS, A. Sok: A consensus taxonomy in the blockchain era. In *Cryptographers' Track at the RSA Conference* (2020), Springer, pp. 284–318.
- [16] GARAY, J., KIAYIAS, A., AND LEONARDOS, N. The bitcoin backbone protocol with chains of variable difficulty. In *Annual International Cryptology Conference* (2017), Springer, pp. 291–323.
- [17] GARAY, J. A., KIAYIAS, A., AND LEONARDOS, N. The bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II* (2015), E. Oswald and M. Fischlin, Eds., vol. 9057 of *Lecture Notes in Computer Science*, Springer, pp. 281–310.
- [18] GARAY, J. A., KIAYIAS, A., AND LEONARDOS, N. Full analysis of nakamoto consensus in bounded-delay networks. *IACR Cryptol. ePrint Arch. 2020* (2020), 277.
- [19] KARAKOSTAS, D., AND KIAYIAS, A. Securing proof-of-work ledgers via checkpointing. In *2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)* (2021), IEEE, pp. 1–5.
- [20] KARANTIAS, K., KIAYIAS, A., AND ZINDROS, D. Proof-of-burn. In *International Conference on Financial Cryptography and Data Security* (2020), Springer, pp. 523–540.
- [21] KIAYIAS, A., RUSSELL, A., DAVID, B., AND OLYINYKOV, R. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual International Cryptology Conference* (2017), Springer, pp. 357–388.
- [22] KIM, S. K., MA, Z., MURALI, S., MASON, J., MILLER, A., AND BAILEY, M. Measuring ethereum network peers. In *Proceedings of the Internet Measurement Conference 2018* (2018), pp. 91–104.
- [23] KING, S., AND NADAL, S. Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. *self-published paper, August 19, 1* (2012).
- [24] LI, J., AND MANN, W. Initial coin offering and platform building. *SSRN Electronic Journal* (2018), 1–56.
- [25] MINOTAUR. Rust implementation of the minotaur consensus protocol. <https://anonymous.4open.science/r/Minotaur-D037>.
- [26] NAKAMOTO, S. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review* (2008), 21260.
- [27] NAYAK, K., KUMAR, S., MILLER, A., AND SHI, E. Stubborn mining: Generalizing selfish mining and combining with an eclipse attack. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)* (2016), IEEE, pp. 305–320.
- [28] NEU, J., TAS, E. N., AND TSE, D. Ebb-and-flow protocols: A resolution of the availability-finality dilemma. In *2021 IEEE Symposium on Security and Privacy (SP)* (2021), IEEE, pp. 446–465.
- [29] NIST. Fips 180-2. <http://csrc.nist.gov/encryption/tkhash.html>, 2002.
- [30] PASS, R., SEEMAN, L., AND SHELAT, A. Analysis of the blockchain protocol in asynchronous networks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques* (2017), Springer, pp. 643–673.
- [31] PASS, R., AND SHI, E. Fruitchains: A fair blockchain. In *Proceedings of the ACM Symposium on Principles of Distributed Computing* (2017), pp. 315–324.
- [32] PERCIVAL, C., AND JOSEFSSON, S. The scrypt password-based key derivation function. *IETF Draft URL: http://tools.ietf.org/html/josefsson-scrypt-kdf-00.txt* (accessed: 30.11.2012) (2016).
- [33] SANKAGIRI, S., WANG, X., KANNAN, S., AND VISWANATH, P. Blockchain cap theorem allows user-dependent adaptivity and finality. *arXiv preprint arXiv:2010.13711* (2020).
- [34] SAPIRSHTEIN, A., SOMPOLINSKY, Y., AND ZOHAR, A. Optimal selfish mining strategies in bitcoin. In *International Conference on Financial Cryptography and Data Security* (2016), Springer, pp. 515–532.
- [35] SHEVTSOVA, I. G. Moment-type estimates with an improved structure for the accuracy of the normal approximation to distributions of sums of independent symmetric random variables. *Theory of Probability & Its Applications* 57, 3 (2013), 468–496.
- [36] STEWART, A., AND KOKORIS-KOGIA, E. Grandpa: a byzantine finality gadget. *arXiv preprint arXiv:2007.01560* (2020).
- [37] THAI, P., NJILLA, L., DUONG, T., FAN, L., AND ZHOU, H.-S. A generic paradigm for blockchain design. In *Proceedings of the 15th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services* (2018), pp. 460–469.

- [38] WANG, X., MUPPIRALA, V. V., YANG, L., KANNAN, S., AND VISWANATH, P. Securing parallel-chain protocols under variable mining power. *arXiv preprint arXiv:2105.02927* (2021).
- [39] YANG, L., BAGARIA, V., WANG, G., ALIZADEH, M., TSE, D., FANTI, G., AND VISWANATH, P. Prism: Scaling bitcoin by 10,000 x. *arXiv preprint arXiv:1909.11261* (2019).

## APPENDIX

### A SECURITY REGIONS OF HYBRID PROTOCOLS IN FIGURE 1B

We give details of the security regions plotted in Figure 1b. Recall that  $\beta_w$  is the proportion of adversarial hash power and  $\beta_s$  is the proportion of the adversarial stake (or the proportion of Byzantine nodes if a permissioned BFT protocol is adopted).

**2-hop blockchain.** We derived the security region from the assumptions in Theorem 2&3 of [11]. When  $\Delta = 0$  (which gives the largest security region), the security proof of [11] assumes  $(1 - \beta_w)(1 - \beta_s) > \beta_s$ , i.e.,  $\beta_s < \frac{1 - \beta_w}{2 - \beta_w}$  implying  $\beta_s < 1/2$ .

**Checkpointed ledger.** [19] uses a synchronous BFT protocol (with 1/2 fault tolerance) to regularly issuing *checkpoints* on a PoW longest chain. The protocol is proven to be safe and live when  $\beta_s < 1/2$ .

**Finality gadgets.** [6, 28, 33, 36] use an asynchronous/partial synchronous BFT protocol (with 1/3 fault tolerance) to build a finality gadget/layer on the top of a PoW longest chain to achieve important properties such as finality (a.k.a deterministic safety under asynchrony) and accountability. These protocols are proven to be safe and live when  $\beta_s < 1/3$  and  $\beta_w < 1/2$ .

### B MATHEMATICAL FACTS

**THEOREM B.1 (FROM [18]).** *Let  $(X_1, X_2, \dots)$  be a martingale with respective the sequence  $(Y_1, Y_2, \dots)$ , if an event  $G$  implies  $X_k - X_{k-1} \leq b$  and  $V = \sum_k \text{var}[X_k - X_{k-1} | Y_1, \dots, Y_{k-1}] \leq v$ , then for non-negative  $n$  and  $t$*

$$P(X_n - X_0 \geq t, G) \leq e^{-\frac{t^2}{2v + \frac{2bt}{3}}}.$$

The following is known as the Berry-Esseen Theorem. See [14] as a standard reference and [35] for improvements with respect to the constant 1/2.

**THEOREM B.2.** *Let the  $X_i$  be independent variables with common distribution such that  $\mathbb{E}[X_i] = \mu$ ,  $\mathbb{V}[X_i] = \sigma^2 > 0$ ,  $\mathbb{E}[|X_i - \mu|^3] = \rho < \infty$ . If  $F_n$  is the distribution of  $(X_1 + \dots + X_n - \mu n) / \sqrt{n\sigma^2}$  and  $\Phi$  the standard normal, then*

$$|F_n(x) - \Phi(x)| \leq \frac{\rho}{2\sigma^3\sqrt{n}}, \quad \text{for all } x \text{ and } n.$$

### C A LONG RANGE ATTACK

We point out that Minotaur is insecure with the longest chain rule due to a long range attack. Let the weighing parameter  $\omega = 0.5$  and define  $f$  to be PoS block production rate by a stakeholder who controls all actual stake in the system. Suppose the adversary controls 0.8 fraction of stake and 0.1 fraction of mining power at some slot (not at the beginning of the execution), then after behaving honestly for some time, the adversary will control 0.8 fraction of actual stake and 0.1 fraction of work stake at the beginning of some

epoch  $e$ . Now, the adversary starts to grow a private chain  $C_1$ , while honest nodes grow a public chain  $C_2$ . Suppose  $C_1$  and  $C_2$  include the same set of transactions. Then at the beginning of epoch  $e + 2$ , the adversary will control 0.8 fraction of actual stake and all the work stake (after normalization) on its chain  $C_1$  because  $C_1$  only refers adversarial PoW blocks in epoch  $e$ . So the growth rate of  $C_1$  after epoch  $e + 2$  will be  $0.8f + f = 1.8f$ . Similarly, at the beginning of epoch  $e + 2$ , honest nodes will control 0.2 fraction of actual stake and all the work stake (after normalization) on the chain  $C_2$  and the growth rate of  $C_2$  after epoch  $e + 2$  will be  $0.2f + f = 1.2f$ . Therefore,  $C_1$  will catch up with  $C_2$  eventually if the longest chain rule is adopted. See Figure 9.

This long range attack is similar to the stake grinding attack on Ouroboros and it can be prevented by the new chain selection rules maxvalid-mc and maxvalid-bg. For maxvalid-mc, honest nodes won't accept  $C_1$  because it forks too long from  $C_2$ ; for maxvalid-bg, honest nodes won't accept  $C_1$  because it grows slower than  $C_2$  in epoch  $e$  (right after the fork).

### D PROOF FOR §6

**PROOF OF LEMMA 6.6.** The proof for (a) is the same as part (a) of Theorem 1 in [18]. For (b), by the condition, for each query in  $S$  made by either an honest node or the adversary at slot  $r$  in epoch  $e$ , the target  $T$  must satisfy  $(1 - \delta)\alpha_0 f / \gamma^2 \leq pn_r T \leq (1 + \delta)\gamma^2 f / \alpha_0$ . Therefore, the proof is also similar to part (a) of Theorem 1 in [18].

For (c), let the execution be partitioned into parts such that each part has at least  $\ell/\phi$  and at most  $s = 2R$  slots. We prove that the statement fails with a probability less than  $e^{-\lambda}$  for each part. Let  $J$  denote the queries made by  $\mathcal{H}$  in slots  $S$ . We have  $|J| = n^{\mathcal{H}}(S) = \phi n(S)$ . For  $k \in [|J|]$ , let  $Z_i$  be the difficulty of any block obtained from query  $j \in J$  and we write  $\mathcal{E}_{j-1}$  for the execution just before this query. Then

$$\begin{aligned} X_0 &= 0 \\ X_k &= \sum_{i \in [k]} Z_i - \sum_{i \in [k]} \mathbb{E}[Z_i | \mathcal{E}_{i-1}] \end{aligned}$$

is a martingale with respect to  $\mathcal{E}_0, \dots, \mathcal{E}_k$ . We have

$$\begin{aligned} X_k - X_{k-1} &= X_k - \mathbb{E}[X_k | \mathcal{E}_{k-1}] = Z_k - \mathbb{E}[Z_k | \mathcal{E}_{k-1}] \\ &\leq \frac{1}{T_k} = \frac{ph_k}{ph_k T_k} \leq \gamma^3 ph(S) / (1 - \delta)\alpha_0 f |S| \\ &\leq \gamma^3 pn(S) / (1 - \delta)\alpha_0 f |S| := b. \end{aligned}$$

Similarly

$$\begin{aligned} V &= \sum_k \text{var}[X_k - X_{k-1} | \mathcal{E}_{k-1}] \leq \sum_k \mathbb{E}[Z_k^2 | \mathcal{E}_{k-1}] \\ &= \sum_k p T_k \frac{1}{T_k^2} \leq \gamma^3 p^2 |J| h(S) / (1 - \delta)\alpha_0 f |S| \\ &\leq \gamma^3 p^2 |J| n(S) / (1 - \delta)\alpha_0 f |S| := v. \end{aligned}$$

Let the deviation  $t = \epsilon p |J| = \epsilon \phi p n(S)$ , then we have  $b = \frac{\gamma^3 t}{(1 - \delta)\alpha_0 \epsilon \phi f |S|}$  and  $v = \frac{\gamma^3 t^2}{(1 - \delta)\alpha_0 \epsilon^2 \phi f |S|}$ . Using the minimum value of  $|S|$  is  $\ell/\phi$  and applying Theorem B.1 to  $-X_{|J|}$ , we have

$$P[D(S) < (1 - \epsilon)pn^{\mathcal{H}}(S)] \leq \exp\left(-\frac{\epsilon t}{2b(1 + \epsilon/3)}\right) \leq \exp(-\lambda).$$

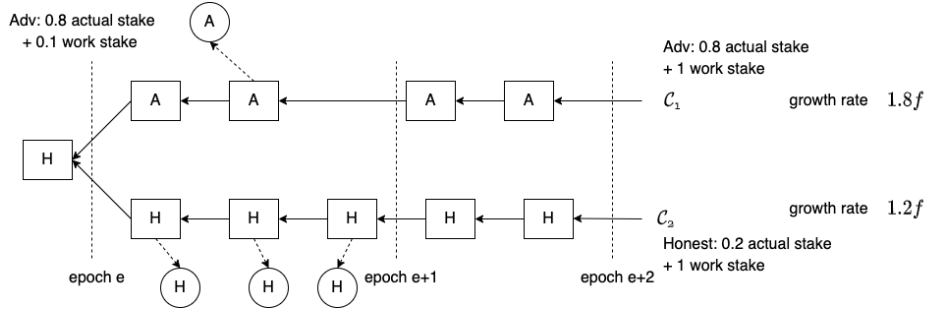


Figure 9: A long range attack on the longest chain rule.

This concludes the proof.  $\square$

**PROOF OF LEMMA 6.7.** We prove the lemma by induction. For epoch  $e = 1$ , it is trivial to just set  $T_1 = f/ph_1$  by Assumption 1.2. Now we assume all epochs are good until epoch  $e - 1$  ( $e \geq 2$ ), we will show epoch  $e$  is good.

Let  $S_0 = \{u : r_1 \leq u \leq r_2\}$  be the window that will be used to determine  $T_e$ , i.e.,  $r_1 = \max(0, (e - 2)R - \kappa) + 1$  and  $r_2 = (e - 1)R - \kappa$ . Let  $C(S_0)$  be the segment of  $C$  containing PoS blocks with timestamps in  $S_0$ , let  $\mathcal{B}$  be all PoW blocks included in  $C(S_0)$ , and  $d$  be the total difficulty of all PoW blocks in  $\mathcal{B}$ . Then we have the following facts:

- **Fact 1.** For any PoW block  $B \in \mathcal{B}$ ,  $B$  is mined after  $r_1 - 4\kappa - \Delta$ . Indeed by recency condition,  $B$  must refer to a confirmed PoS block  $B_s$  with timestamp at least  $r_1 - sl_{re}$ . By  $\exists CQ$ , the last honest ancestor block of  $B_s$  has timestamp at least  $r_1 - sl_{re} - \kappa$ . So  $B$  must be mined after  $r_1 - 4\kappa - \Delta$ .
- **Fact 2.** For any PoW block  $B \in \mathcal{B}$ ,  $B$  is mined before  $r_2 + \kappa$ . Indeed the PoS block (denoted as  $B_s$ ) including  $B$  has timestamp at most  $r_2$ , and again by  $\exists CQ$ , the first honest descendant block of  $B_s$  has timestamp at most  $r_2 + \kappa$ . So  $B$  must be mined before  $r_2 + \kappa$ .
- **Fact 3.** If a PoW block  $B$  is mined by an honest miner after  $r_1$  and before  $r_2 - 3\kappa - \Delta$ , then  $B \in \mathcal{B}$ . Indeed, by  $\exists CQ$ , the last honest block in  $C(S_0)$  has timestamp at least  $r_2 - \kappa$ . Hence by Lemma 6.5, all honest PoW blocks mined after  $r_1$  and before  $r_2 - \kappa - r_{wait}$  will be included into a PoS block in  $C(S_0)$ .

Let  $S_1 = \{u : r_1 - (4\kappa + \Delta) \leq u \leq r_2 + \kappa\}$ ,  $S_2 = \{u : r_1 \leq u \leq r_2 - (3\kappa + \Delta)\}$ , and  $J$  be the set of adversary queries associated with  $\mathcal{B}$  in  $S_1$ . Then by Fact 1 and Fact 2, we have all PoW blocks in  $\mathcal{B}$  are mined in  $S_1$ ; by Fact 3, we have all PoW blocks mined by honest nodes in  $S_2$  are in  $\mathcal{B}$ . Hence,  $D(S_2) \leq D_{e-1} \leq D(S_1) + A(J)$ . By Proposition 2, for each query in  $S_1$  made by an honest node at slot  $r$  in epoch  $e$ , the target  $T$  must be  $T_e$ , so we have  $(1 - \delta)\alpha_0 f / \gamma^2 \leq ph_r T \leq (1 + \delta)\gamma^2 f$ . For each query in  $J$  made by the adversary at slot  $r$  in epoch  $e$ , the target  $T$  may be either  $T_{e-1}$  or  $T_e$ , still we have  $(1 - \delta)\alpha_0 f / \gamma^2 \leq ph_r T \leq (1 + \delta)\gamma^2 f$  in both cases (under a  $(\gamma, 2R)$ -respecting environment). By the fact that  $\alpha_0 n_r \leq h_r \leq n_r$ , we have  $(1 - \delta)\alpha_0 f / \gamma^2 \leq pn_r T \leq (1 + \delta)\gamma^2 f / \alpha_0$ .

Under a typical execution, we have

$$D(S_2) > (1 - \epsilon)ph(S_2),$$

and

$$\begin{aligned} D(S_1) + A(J) &< (1 + \epsilon)p(h(S_1) + |J|) \\ &= (1 + \epsilon)pn(S_1) \leq (1 + \epsilon)ph(S_1) / \alpha_0. \end{aligned}$$

Therefore, by Proposition 1, we have

$$(1 - \epsilon)ph^{(e)}|S_2| / \gamma < D_{e-1}^{\text{total}} < (1 + \epsilon)\gamma ph^{(e)}|S_1| / \alpha_0.$$

By the difficulty adjustment rule, we have  $T_e = f|S_0| / D_{e-1}^{\text{total}}$ . Then  $T_e$  can be bounded as follows:

- **Lower bound:**

$$\begin{aligned} ph^{(e)}T_e &\geq \frac{|S_0|}{(1 + \epsilon)|S_1|} \frac{\alpha_0 f}{\gamma} \geq \frac{(1 - \epsilon)|S_0|}{|S_0| + 5\kappa + \Delta} \frac{\alpha_0 f}{\gamma} \\ &\geq \frac{(1 - \epsilon)(R - \kappa)}{R + 4\kappa + \Delta} \frac{\alpha_0 f}{\gamma} \stackrel{(C_1)}{\geq} \frac{(1 - \epsilon)(R - \epsilon R)}{R + \epsilon R} \frac{\alpha_0 f}{\gamma} \\ &\geq (1 - \epsilon)^3 \frac{\alpha_0 f}{\gamma} \geq (1 - 4\epsilon) \frac{\alpha_0 f}{\gamma} \stackrel{(C_2)}{\geq} (1 - \delta) \frac{\alpha_0 f}{\gamma}. \end{aligned}$$

- **Upper bound:**

$$\begin{aligned} ph^{(e)}T_e &\leq \frac{|S_0|}{(1 - \epsilon)|S_2|} \gamma f = \frac{|S_0|}{(1 - \epsilon)(|S_0| - 3\kappa - \Delta)} \gamma f \\ &\leq \frac{R - \kappa}{(1 - \epsilon)(R - 4\kappa - \Delta)} \gamma f \stackrel{(C_1)}{\leq} \frac{R}{(1 - \epsilon)(R - \epsilon R)} \gamma f \\ &= \frac{1}{(1 - \epsilon)^2} \gamma f \leq \frac{1}{1 - 2\epsilon} \gamma f \leq (1 + 4\epsilon)\gamma f \stackrel{(C_2)}{\leq} (1 + \delta)\gamma f. \end{aligned}$$

This concludes the proof.  $\square$

**PROOF SKETCH OF THEOREM 6.9.** When moving from the single-epoch setting to the multiple-epoch setting, two new aspects need to be considered.

- **Virtual stake distribution updates.** For epochs  $e = 1, 2$ , the virtual stake has the same distribution as the initial stake. Since we assume the initial stake has honest majority (Assumption 1.1), by Theorem 6.4, CP and  $\exists CQ$  are guaranteed in epochs 1&2. For epoch  $e \geq 3$ , the virtual stake of a node  $h$  composes of two parts, the actual stake recorded on the blockchain up to the last block of the epoch  $e - 2$  (by Ouroboros [21]) and the work stake decided by the amount of work  $h$  has contributed in epoch  $e - 2$ . Let  $S_s$  and  $S_w$  be the total actual stake and total work stake in epoch  $e$ . Recall that we set  $\omega_{S_s} = (1 - \omega)S_w$  in Minotaur, i.e., the total virtual stake is  $S_v = S_s / (1 - \omega)$ . Denote that the adversary controls  $\beta_s S_s$  actual stake and let  $\beta_w$  be maximum

fraction of adversarial mining power in epoch  $e - 2$ , i.e.,  $\beta_w = 1 - \min_{r \in [(e-3)R+1, (e-2)R]} (h_r/n_r)$ . By Assumption 2.1, we have  $\omega\beta_w + (1 - \omega)\beta_s \leq 1/2 - 2\sigma$ . And by Theorem 6.8, honest nodes control at least  $(1 - \sigma)(1 - \beta_w)S_w$  work stake for epoch  $e$ . Therefore, the virtual stake controlled by the adversary is at most

$$\begin{aligned} & \beta_s S_s + (1 - (1 - \sigma)(1 - \beta_w))S_w \\ &= ((1 - \omega)\beta_s + \omega(\beta_w + \sigma(1 - \beta_w)))S_v \\ &\leq (1/2 - \sigma)S_v. \end{aligned}$$

Therefore, by an induction argument, we can guarantee an honest majority in the virtual stake for all epochs. The analysis critically relies on the fact that the CP property is immutable: specifically, when all honest parties agree on a common prefix  $C^{(t)}$  at some slot  $t$  and, as maxvalid-mc can only revise the last  $k$  blocks of a currently adopted chain,  $C^{(t)}$  will be a prefix of all future chains held by the honest parties. Check the proof of Theorem 5.3 in [21] for details.

- **Randomness updates.** Every epoch needs new public randomness to be used for sampling slot leaders from the above virtual stake distribution. In Ouroboros [21], elected slot leaders (one per slot) from epoch  $e - 1$  runs a publicly verifiable secret sharing (PVSS) protocol to generate the randomness for epoch  $e$ . The core idea is the following: given that we have guaranteed that an honest majority among elected leaders in epoch  $e$  will hold with very high probability, we have that the PVSS protocol suitably simulates a beacon with the relaxation that the output may become known to the adversary before it is known to the honest parties. However, as long as the distribution of virtual stake is determined prior to this leakage, the sampling of leaders in epoch  $e$  will still be unbiased.

At last, it is not hard to see that, the CP property is equivalent to the persistence of the ledger, while CP together with  $\exists\text{CQ}$  implies liveness.  $\square$

## E DETAILS ON THE OUROBOROS PROTOCOL FAMILY

We give a summary of the different PoS protocols we explicitly base our generic hybridization construction on in the main part of the paper, Ouroboros Classic [21], Praos [8], and Genesis [3]. In order of this sequence, each version of the protocol gives stronger security guarantees.

We first give quick summaries of the respective protocol guarantees and their underlying assumptions. Finally, in Appendix E.1, we give a more detailed description of Ouroboros Praos as the reference protocol; and sketch how the other variants differ from it.

*Ouroboros Classic.* Classic is secure against a minority of adversarially controlled stake under the following assumptions:

- Network. The communication network is synchronous.
- Corruption. The adversary is ‘moderately’ adaptive (participant corruption only takes effect after a certain delay).
- Stake shift. There is an upper bound on the stake shift, i.e., the stake distribution among the stake holders does not change too fast.

Offline tolerance. The protocol participants only go offline for short periods of time.

*Ouroboros Praos assumptions:* Praos is secure against a minority of adversarially controlled stake under the following assumptions:

- Network. The communication network is semi-synchronous, i.e., that the network delay is bounded by some delay  $\Delta$  not known to the participants.
- Corruption. The adversary is fully adaptive.
- Stake shift. As in Classic.

*Ouroboros Genesis assumptions:* Classic is secure against a minority of adversarially controlled stake among all participants who are active in the system—under the following assumptions:

- Network. As in Praos.
- Corruption. As in Praos.
- Offline tolerance. The protocol participants can join later during any stage of the protocol, or go offline for extended periods of time during participation.

### E.1 Ouroboros Praos

The protocol proceeds in epochs of  $R$  slots, each slot representing a given ‘unit of time’, say, 1 second of the protocol run-time. For ease of exposition, let the genesis block represent epoch 0 of the protocol. We now describe how the protocol operates per epoch  $e > 0$ .

*Slot-leader election.* During each slot  $sl_j$ , a slot-leader election is held among the stakeholders, and a winning stakeholder is allowed to publish a new block associated with this slot.

The slot-leader election during epoch  $e$  is based on the stake distribution at the end of epoch  $\max(0, e - 2)$ , i.e., the stake distribution that results after the processing of the last block of epoch  $e - 2$  (or the stake distribution from the genesis block).

Let  $\alpha_i$  be the relative stake held by stakeholder  $P_i$  (holding  $\alpha_i \cdot S$  of the total stake  $S$ ) at the end of epoch  $\max(e - 2, 0)$ . Per slot, the probability  $p_i$  for stakeholder  $P_i$  to be a block leader is defined as

$$p_i = \phi_f(\alpha_i) \triangleq 1 - (1 - f)^{\alpha_i} \quad (2)$$

for some appropriate *active-slots coefficient*  $f$  (the probability that, among the total stake, at least one slot leader is elected during any given slot).

Slot leadership is pseudo-randomly assigned based on the *epoch nonce*  $\eta$ , a seed calculated in epoch  $e - 1$  (as described further below). To become a slot leader for slot  $sl_m$ , the stakeholder  $P_i$  evaluates a verifiable random function (VRF) (bound to a public key registered by  $P_i$ )  $(y, \pi) = \text{VRF}_i(\eta, sl_m)$ . Slot leadership is satisfied iff

$$y \stackrel{!}{<} T \triangleq 2^{\ell_{\text{VRF}}} \cdot \phi_f(\alpha_i),$$

where  $\text{VRF}_i$  produces outputs of  $\ell_{\text{VRF}}$  bits, and  $T$  is a threshold to enforce the desired probabilities.

*Block production.* Besides a hash of its predecessor and the payload, a block contains

- the slot number  $m$ ;
- the above proof of leadership  $(y, \pi)$  (such that  $y < T$ );
- an additional, independent, VRF output  $(y_\rho, \pi_\rho)$  contributing to the epoch-nonce generation; and



- a signature on  $(m, (y, \pi), (y_\rho, \pi_\rho))$  by  $P_i$  of a key-evolving signature scheme (KES).<sup>4</sup>

*Block settlement.* A block is considered settled if it sits at least  $k$  blocks deep in a node's main chain where  $k$  is the *prefix parameter* of the protocol.

*Chain selection rule:* maxvalid-mc. Upon the arrival of a new block, each node chooses, as their main chain, the longest chain that does not fork from the previous main chain by more than  $k$  blocks where  $k$  is the prefix parameter.

*Epoch-nonce generation.* The epoch nonce  $\eta$  for epoch  $e > 1$  is computed as a hash of all VRF outputs  $y_\rho$  included in the blocks of the main chain up to  $2/3$  of epoch  $e - 1$ . The epoch nonce for epoch  $e = 1$  is included in the genesis block.

## E.2 Differences in Classic and Genesis

*Classic.* The main difference to Praos is that the epoch randomness is obtained from a coin-flip protocol based on a publicly verifiable secret sharing protocol—with the disadvantage that the leader schedule is public and thus allows for adaptive corruption.

*Genesis.* Genesis differs from Praos by applying a different chain selection rule to allow protocol participants to be offline for extended periods of time. The modified chain-selection rule maxvalid-bg works as follows: A longer chain is adopted if

- it forks from the current main chain by at most  $k$  blocks; or
- if forks by more than  $k$  blocks but contains more blocks in the  $s$  slots following the last common block of the main chain and the longer chain.

## F ANALYSIS TIGHTNESS

The analysis of our main security theorem, Theorem 6.1, draws from [21]. A more careful and improved analysis appears in [5], from which we may obtain the refined bound  $1 - RL(e^{-\Omega(\delta^3 \kappa)} + e^{-\lambda})$  for the statement of Theorem 6.1. In this section we analyze a private chain attack and show a corresponding lower bound. It reveals that  $\delta^2 \kappa$  needs to be bounded below by a constant and, consequently, the dependency on  $\delta^3$  cannot be improved much further.

**PROPOSITION 3.** *Common prefix does not hold against a  $(1/2 - \delta, m, \omega)$ -bounded adversary, for  $\delta \leq 1/\sqrt{\ell_{\text{cp}}}$  and sufficiently large  $\ell_{\text{cp}}$ .*

**PROOF.** The adversary follows the protocol using his mining power to produce PoW blocks. The attack begins when the virtual stake of the adversary is  $\omega\beta_w + (1 - \omega)\beta_s \geq \frac{1}{2} - \delta$ .

Define the random variable  $X_i$  taking values in  $\{-1, 0, 1\}$  according to who produced blocks of height  $i$ . If it was the adversary only, then  $X_i = 1$ ; if both the adversary and honest parties, then  $X_i = 0$ ; if honest parties only, then  $X_i = -1$ . Let  $p = 1 - (1 - f)^{1/2 + \delta}$  and  $q = 1 - (1 - f)^{1/2 - \delta}$ . We have

$$\begin{aligned} \Pr[X_i = -1] &= p(1 - q)/f, \\ \Pr[X_i = 1] &= q(1 - p)/f, \quad \text{and} \\ \Pr[X_i = 0] &= pq/f. \end{aligned}$$

Note that the adversary will create a fork of length  $k$ , if  $X_1 + \dots + X_k \geq 0$ . We will use the Berry-Esseen bound to lower bound the probability he succeeds. To that end we compute

$$\begin{aligned} \mu &= \mathbb{E}[X_i] = (q - p)/f = [(1 - f)^{1/2 + \delta} - (1 - f)^{1/2 - \delta}]/f \\ &= [(1 - f)^{1 + \delta} - (1 - f)^{1 - \delta}]/[f(1 - f)^{1/2}] \\ &> [(1 - f)^{1 + \delta} - (1 - f)^{1 - \delta}]/f > -2\delta, \\ \sigma^2 &= \mathbb{E}[X_i^2] - \mu^2 = 1 - pq/f - \mu^2 > 1 - f - 4\delta^2, \\ \rho &= \mathbb{E}[|X_i - \mu|^3] < (1 + 2\delta)^3. \end{aligned}$$

Observe that for  $\delta \leq 1/3$  and  $f \leq 1/9$ .

Let  $F_k = (X_1 + \dots + X_k - \mu k)/\sqrt{k\sigma^2}$ . By the Berry-Essen Theorem we have that, for  $\delta \leq 1/\sqrt{k}$ ,

$$\Pr[F_k \geq -\mu k/\sqrt{k\sigma^2}] \geq \Phi^c(3) - O(k^{-1/2}).$$

□

Although it might be possible to strengthen the above attack and analysis to obtain slight improvements, it remains an open question whether the exponential drop in the probability of security failure is in the order of  $\delta^3 \kappa$  or  $\delta^2 \kappa$ .

<sup>4</sup>The stakeholders update the private keys of their KES instance after every slot.