

Blockchain and Consensus from Proofs of Work without Random Oracles

Juan A. Garay
Texas A&M University
garay@cse.tamu.edu

Aggelos Kiayias*
University of Edinburgh
akiayias@inf.ed.ac.uk

Giorgos Panagiotakos
University of Edinburgh
giorgos.pan@ed.ac.uk

March 22, 2018

Abstract

One of the most impactful applications of “proofs of work” (POW) currently is in the design of blockchain protocols such as Bitcoin. Yet, despite the wide recognition of POWs as the fundamental cryptographic tool in this context, there is no known cryptographic formulation that implies the security of the Bitcoin blockchain protocol. Indeed, all previous works formally arguing the security of the Bitcoin protocol relied on direct proofs in the random oracle model, thus circumventing the difficulty of isolating the required properties of the core POW primitive.

In this work we fill this gap by providing a formulation of the POW primitive that implies the security of the Bitcoin blockchain protocol in the standard model. Our primitive entails a number of properties that parallel an efficient non-interactive proof system: completeness and fast verification, security against malicious provers (termed “hardness against tampering and chosen message attacks”) and efficiency and security for honest provers (the latter captured as almost k -wise independence of the proving algorithm running time). Interestingly, our formulation is incomparable with previous formulations of POWs that applied the primitive to contexts other than the blockchain and highlights the importance of *run-time independence* as a property for POWs suitable for blockchain protocols.

Using our primitive and standard properties of the underlying hash function, we establish the security of the Bitcoin backbone protocol [Eurocrypt 2015] without relying on random oracles. We then tackle the question of constructing a consensus protocol based on POW. We illustrate that previously known solutions essentially relied on the random oracle and propose a new blockchain-based consensus protocol provably secure under the same assumptions as above. This yields the first consensus protocol for honest majority reducible to a POW primitive without random oracles.

*Research partly supported by Horizon 2020 project PANORAMIX, No. 653497.

Contents

1	Introduction	3
2	Preliminaries	6
3	Proofs of Work	8
4	The Bitcoin Backbone from POW	10
4.1	Bitcoin backbone model and definitions	10
4.2	Security properties of the blockchain	12
4.3	Robust public transaction ledgers	12
4.4	The modified Bitcoin backbone protocol	13
4.5	Security proof	14
4.6	Comparison with the random-oracle analysis	23
5	Consensus from Proofs of Work without Random Oracles	27
A	Hash Functions and Unpredictability Preservation	33
B	Proof of the modified bitcoin backbone protocol	33

1 Introduction

For over 20 years the term “proof of work” (POW) has been used liberally in the cryptography and security literature to describe a number of cryptographic primitives that were applied in a variety of settings, including spam mitigation [19], sybil attacks [18], denial of service protection [27, 6], and other applications. A POW scheme can be typified as a primitive that entails two algorithms: a proving algorithm and a verification algorithm, an abstraction which, despite minor differences, can be used to describe most schemes in the literature. The proving algorithm will receive as input a parameter specifying the intended “difficulty” of the POW as well as additional information such as a context or message with respect to which the POW is to be computed. (Intuitively, the higher level of difficulty selected, the more effort [in terms of computational steps] would be required.)

The main properties that almost all previous works share are: (i) *amortization resistance*, which guarantees that the adversary cannot compute POWs without doing work proportional to the computational resources available to it; (ii) the *existence of trapdoors*, which when appropriately used, bypasses the previous restriction, and may even allow producing a POW with effort independent of the given difficulty level; (iii) *POW sampling*, which prescribes that one can randomly sample POWs; and (iv) *fast verification*, which requires that the time it takes to verify the correctness of a POW is a lot less than the time it takes to compute it.

However, the most impactful application of POWs to date is in the design of “permissionless” blockchain protocols [31], where the use of POWs enables the construction of distributed consensus protocols that are not based on a pre-existing public-key infrastructure or authenticated channels. At a high level, the way POWs help in this context is by slowing down message passing for all parties indiscriminately, thus generating opportunities for honest parties to converge to a unique view under the assumption that the aggregate computational power of honest parties sufficiently exceeds that of the adversary. Now, while this intuition matches the more rigorous analyses of the Bitcoin protocol that have been carried out so far [22, 32, 23, 7], these works have refrained from actually formally defining the underlying POW algorithms as a stand-alone cryptographic primitive and relied instead on the random oracle (RO) model [12] to prove *directly* the properties of the blockchain protocol, as opposed to following a reduction approach to the underlying cryptographic primitives. The same is true for other provably secure POW-based distributed protocols [4, 28, 24]. In fact, to our knowledge, no prior work has identified the set of properties for the POW primitive that are sufficient to imply the security of a blockchain protocol.

Our contributions. In this work we fill the above gap by putting forth a formalization of a notion of POW that we prove to be sufficient for reducing in a black-box manner the security of the Bitcoin backbone protocol [22] to it. This paves the way for establishing the security of Bitcoin in the standard model, as opposed to the random oracle model, assuming our POW primitive can be realized under a simple computational assumption. Interestingly, the set of POW properties we identify is novel and the properties not been considered in previous works that attempted to formalise the POW primitive.

Formalization of the POW primitive. Our syntax of the POW primitive entails a proving and verification algorithm, **Prove** and **Verify**, respectively. **Prove** is invoked on an input (r, msg, h) , where r is a random key that guarantees freshness, msg is a message that will be associated with the POW, and h is the hardness level. Expectedly, **Verify** is invoked on input (r, msg, h, w) where w is (possibly) an output of **Prove**. There are five properties that we put forth:

- **Completeness:** As in the case of an interactive proof, we require that proofs produced by **Prove** should be acceptable by the **Verify** algorithm.

- **t -Verifiability:** This property mandates that the POW can be verified in t steps by the honest verifier.
- **α -Efficiency:** This property bounds the computational effort of an honest prover to produce a POW. Specifically, it requires the event that **Prove** terminates in a certain number of steps to be lower bounded by α , a function of t .
- **Hardness against Tampering and Chosen-Message Attacks (H-TCMA):** This property deals with security against malicious provers, and is akin to the property of existential unforgeability under chosen-message attacks of digital signatures. Specifically, it captures the fact that producing a sequence of POWs, for chosen messages, does not provide an advantage to an adversary in terms of running time. Furthermore, this should hold against an adversary able to tamper with the keys (that’s the ‘T’ in the acronym), and even in the presence of a **Prove** oracle (excluding of course any proofs produced by the oracle from the adversary’s output).
- **Run-time m -wise independence:** This final property deals with security for honest provers. It captures the setting where honest provers are potentially invoked on adversarial inputs and ensures that their running time enjoys some degree of independence. Specifically, we call a POW *run-time m -wise independent* if in the course of an adversarial program execution that makes function calls to **Prove**, the random variables defined as the running time of each **Prove** invocation, is a set of almost m -wise independent random variables (cf. [2]).

Security of the Bitcoin backbone protocol assuming POW. We then turn to the analysis of the Bitcoin backbone protocol [22] given our POW primitive. We consider a fixed number of parties in the standard model assuming our POW, following the approach of [22]. We recall the three basic properties of the blockchain data structure, (strong) common prefix, chain quality and chain growth, and show how the Bitcoin backbone protocol should be modified in order to incorporate our POW primitive. Besides the POW primitive, proving the security of the backbone protocol requires standard properties of the underlying hash function that “glues” the blocks together, namely, collision and preimage resistance.

We first prove that using the H-TCMA property and assuming the adversarial hashing power is suitably bounded, it is unlikely in any sufficiently long time window for the adversary to exceed the number of POWs of the honest parties. Then, using the α -Efficiency and (β, ϵ) -HTCMA properties in conjunction with run-time m -wise independence, where $m = \Theta(n)$, the number of parties, we establish that summations of running times of successive **Prove** invocations have the variance needed to ensure that “uniquely successful rounds” (rounds where exactly one of the honest parties produces a POW) happen with high density in any sufficiently long time window. (Recall that a high rate of uniquely successful rounds was fundamental in the security proof of [22].) Using these last two core results, and under suitable constraints for the basic POW parameters $\alpha, \beta, \epsilon, h$ and number of parties n , we prove the security of the Bitcoin backbone protocol in the standard model assuming just our POW primitive and the security of the underlying hash function.

The above analysis also highlights the distinctive nature of our notion of POWs that is geared to the setting of blockchain protocols. Importantly, using a POW that is merely efficiently verifiable, and even H-TCMA secure but without the run-time m -wise independence property, it is easy to construct an attack against the blockchain protocol. Specifically, without m -wise independence, it is conceivable to have a POW that always takes the same number of steps to be solved. Using such a POW, the honest parties will find a solution in the same round and thus, an adversary, exploiting the network scheduling, can maintain an indefinite fork. On the other hand, using m -wise independence this symmetry between honest parties can be broken for a suitable choice of m , and honest parties are capable of converging to a single blockchain following the longest chain rule (note that we assume a static number of parties). We note that run-time m -wise independence is trivial in

the random oracle model and as such it was never highlighted in any previous work on POW-based blockchain protocols.

Finally, and as a “sanity check,” we show that the Bitcoin POW scheme in the random oracle model is secure according to our definitions, and that according to the security parameters we obtain for the scheme, the security guarantees we get from our analysis of the Bitcoin backbone are similar to those proved in [22, 32].

Consensus from POW. We next tackle the problem of designing a consensus protocol [33, 30] for an honest majority of parties that can reduce to our POW primitive. The only known consensus protocol in this setting (where no *private* setup assumption is provided) is given in [22] where it is shown how using a technique called “2-for-1 POW” two POW-based protocols can be run concurrently and create a blockchain where the number of honest party contributions is proportional to their actual number; this property is then shown to imply consensus. We stress that this is not the case for the Bitcoin blockchain protocol, where attacks such as block withholding (as used in the “selfish mining” attack [21]) enable the adversary to create blockchains that misrepresent the number of honest parties.

We observe that this protocol cannot work in our setting as it cannot be reduced to our POW primitive in a black-box way. Indeed, 2-for-1 POW non-trivially relies on the random oracle model and the fact that each witness for a POW in the RO setting can be rearranged in a certain way so as to obtain a test for a witness for *another* POW in a way that is independent from the first solution. Such property is achievable in the RO model but not black-box-reducible to our POW primitive. We propose a new blockchain protocol that circumvents this problem and preserves the required property, namely, that honest parties’ contributions to the blockchain reflect their actual number. The core idea of our protocol is as follows: When mining a block the parties include their input as well as the headers of “orphan” blocks that exist in forks stemming off their main chain and have not been included so far. The header of a block contains the hash of the previous block in the chain, the witness to the POW, application specific data (such as the input to the consensus protocol), and a hash that may contain headers of other blocks. Using this mechanism, we show how it is possible to reconstruct the whole tree of block headers from the blockchain contents and thus in this way preserve all block headers produced by the honest parties. This ensures that the resulting ledger will reflect the number of parties and hence a consensus protocol may now be easily reduced to this blockchain protocol.

Prior and related work. Dwork and Naor [20] first considered POWs under the term “pricing functions,” as a means of protection against spam e-mail. The main properties discussed in their work are amortization resistance, “moderate hardness” and the existence of trapdoors (“shortcuts” in their terms). Three different constructions were presented there. The first one is based on the difficulty of extracting square roots modulo a prime p and does not have a trapdoor; the second construction has a trapdoor and is based on the hardness of forging signatures (i.e., hardness of factoring in the specific instantiation), while the last construction is based on “broken” signature schemes, where moderately hard algorithms that can forge signatures exist, but the signing key is protected from these attacks.

In a different direction, Juels and Jakobsson [26] and Back [5, 6] use POWs to construct electronic payment systems. In [26] the authors consider the following properties: amortization resistance, fast verification, and some special composability property which states that generating a POW for some scheme may help in generating a POW for another scheme. As acknowledged by the authors themselves, the definitions they provide are only sketches. In [5, 6] another set of closely related properties is considered, but again the approach is not rigorous. Notably, the author mentions the concept of “trapdoor-freeness,” i.e., that the party which generates the initial parameters of the

scheme should not be able to also generate a trapdoor regarding these parameters. This property comes in sharp contrast with the POW sampling property, signifying a conflict on what actually are the *essential* properties of a POW.

More recently, Bitansky *et al.* [15] construct time-lock puzzles as well as POW schemes from randomized encodings. Since the focus of their work is time-lock puzzles, the properties of POW schemes—amortization resistance, fast verification, and POW sampling—are only briefly investigated, although they do instantiate a POW scheme based on randomized encodings and the existence of non-amortizing languages in the worst case.

Another interesting approach is that of Ball *et al.* [9], who construct a POW scheme that can be used to do “useful” work. That is, an instance of a problem is selected and a proof of work on this instance along with a solution can be constructed, while at the same time preserving amortization resistance. While the authors mention that their POW scheme can be used for Bitcoin, no formal proof is provided that Bitcoin security reduces to the primitive they construct.

Taking a further step in this same direction, i.e., proving that Bitcoin’s security can be reduced to an assumption on the underlying POW scheme, Poelstra [34] proposes that one could base the security of Bitcoin on so-called “dynamic membership multi-party signatures.” Unfortunately, this work does not offer a formal treatment of the primitive nor a reduction of the security of the Bitcoin protocol to it.

Finally, in [3], Alwen and Tackmann study moderately hard functions (MoHF), providing simulation based definitions for what they call “non-interactive proofs of effort” (niPoE), which—as explicitly acknowledged by the authors—cannot be used to analyze Bitcoin. The main impediment is that the adversary can only invoke the same MoHF only once per protocol session, while for the Bitcoin protocol multiple invocations of the same MoHF should be allowed.

We note that we perform our analysis in the static setting, i.e., where the POW difficulty remains the same throughout the execution as in [22]. We refer to [23] for an analysis of Bitcoin in the dynamic setting and leave as an open question the extension of our results to this setting.

Organization of the paper. The basic computational model, definitions and cryptographic building blocks are presented in Section 2. Formal definition of a POW and its security properties are presented in Section 3. Section 4 is dedicated to Bitcoin: Showing how the Bitcoin backbone protocol should be modified to incorporate our POW primitive and reducing its security to it (Sections 4.4 and 4.5, respectively), and analyzing Bitcoin’s POW implementation according to our definition (Section 4.6). The new consensus protocol for an honest majority based on POW is presented in Section 5.

2 Preliminaries

We let λ denote the security parameter, and use $x \stackrel{\$}{\leftarrow} X$ to denote choosing a value uniformly at random from set X . In this paper we will follow the concrete approach [11, 13, 25, 14] to security evaluation rather than the asymptotic one. We will use functions t, ϵ , whose range is \mathbb{N}, \mathbb{R} respectively and have possibly many different arguments, to denote concrete bounds on the running time (number of steps) and probability of adversarial success of an algorithm in some given computational model, respectively. When we speak about running time this will include the execution time plus the length of the code (cf. [14]; note also that we will be considering uniform machines). We will always assume that t is a polynomial on the security parameter λ , although we will sometimes omit this dependency for brevity.

Instead of using interactive Turing machines (ITMs) as the underlying model of distributed computation, we will use (interactive) RAMs. The reason is that we need a model where subroutine

access and simulation do not incur a significant overhead. ITMs are not suitable for this purpose, since one needs to account for the additional steps to go back-and-forth all the way to the place where the subroutine is stored. A similar choice was made by Garay *et al.* [25]; refer to [25] for details on using interactive RAMs in a UC-like framework, as well as to Section 4.1. Given a RAM M , we will denote by $\mathbf{Steps}_M(x)$ the random variable that corresponds to the number of steps of M given input x . We will say that M is t -bounded if it holds that $\Pr[\mathbf{Steps}_M(x) \leq t(|x|)] = 1$.

Finally, we remark that in our analyses there will be asymptotic terms of the form $\mathbf{negl}(\lambda)$ and concrete terms; throughout the paper, we will assume that λ is large enough to render the asymptotic terms insignificant compared to the concrete terms.

Cryptographic primitives and building blocks. We will make use of the following notions of security for cryptographic hash functions (see, e.g., [35] for a thorough review):

Definition 1. Let $\mathcal{H} = \{H_\lambda : K_\lambda \times M_\lambda \rightarrow Y_\lambda\}_{\lambda \in \mathbb{N}}$ be a hash-function family, and \mathcal{A} be a PPT adversary. Then \mathcal{H} is

– *Collision resistant* if and only if for any $\lambda \in \mathbb{N}$ and corresponding $H : K \times M \rightarrow Y$ in \mathcal{H}

$$\Pr[k \xleftarrow{\$} K; (m, m') \leftarrow \mathcal{A}(1^\lambda, k); (m \neq m') \wedge (H(k, m) = H(k, m'))] \leq \mathbf{negl}(\lambda);$$

– *everywhere preimage resistant* if and only if for any $\lambda \in \mathbb{N}$ and corresponding $H : K \times M \rightarrow Y$ in \mathcal{H}

$$\Pr[(y, s) \leftarrow \mathcal{A}(1^\lambda); k \xleftarrow{\$} K; m \leftarrow \mathcal{A}(1^\lambda, k, s) : H(k, m) = y] \leq \mathbf{negl}(\lambda).$$

As we will see, POW schemes critically depend on a key that ensures the proof is “fresh”. Nevertheless, in many settings it is useful to allow the adversary to tamper this key with a certain function f . Given that arbitrary tampering may eliminate freshness, we have to suitably restrict the tampering function class. The following class of functions will be useful in our security analysis.

Definition 2. Let $\mathcal{F} = \{f_{i,\lambda}\}_{i,\lambda \in \mathbb{N}}$ be a function family such that $f_{i,\lambda} : \{0,1\}^{m(\lambda)} \rightarrow \{0,1\}^\lambda$, for some function $p : \mathbb{N} \rightarrow \mathbb{N}$, for all $i, \lambda \in \mathbb{N}$. We say that \mathcal{F} is a *unpredictability-preserving* function family if and only if for any uniform PPT RAM machine \mathcal{A} it holds that:

$$\Pr[(y, s) \leftarrow \mathcal{A}(1^\lambda); x \xleftarrow{\$} U_{m(\lambda)}; f \leftarrow \mathcal{A}(1^\lambda, y, s, x) : f \in \mathcal{F} \wedge f(x) = y] \leq \mathbf{negl}(\lambda).$$

Observation 3. The identity function $\{\mathbf{id}_\lambda(\cdot)\}_{\lambda \in \mathbb{N}}$ and an everywhere preimage resistant hash function family (Definition 1; see also Appendix A) are unpredictability-preserving.

We will also need some classical notions related to random variables:

Definition 4. The *statistical distance* of two random variables X, Y with range U , denoted $\Delta(X, Y)$, is defined as

$$\Delta(X, Y) = \frac{1}{2} \sum_{u \in U} |\Pr[X = u] - \Pr[Y = u]|.$$

For $\epsilon > 0$, we say that X, Y are ϵ -close when $\Delta(X, Y) \leq \epsilon$.

Definition 5. A set of random variables X_1, \dots, X_n is *k-wise independent* if, for any subset $I \subseteq [n]$, where $|I| \leq k$, and for any values $x_i, i \in I$:

$$\Pr[\bigwedge_{i \in I} X_i = x_i] = \prod_{i \in I} \Pr[X_i = x_i].$$

3 Proofs of Work

At a high level, a proof-of-work (POW) scheme is a protocol that enables one party to convince others that she has invested some computational power during some specific time interval and with respect to a specific piece of information, which we will generically refer to as ‘message.’ In this section we formalize this notion and present its desired security properties.

POW syntax. We formalize the POW notion as follows. Given a specific security parameter λ , let R be the key space, M the message space and W the witness (or POW) space. With foresight, the role of the key is to provide freshness for the POW computation and the role of the message is to allow information to be embedded into the POW; see Remark 1 regarding the significance of these input elements.

Definition 6. A POW scheme consists of two algorithms $\text{POW} = (\text{Prove}, \text{Verify})$ where:

- $\text{Prove}(r, \text{msg}, h)$ is a randomized algorithm that takes as input a key $r \in R$, a message $\text{msg} \in M$ and hardness parameter $h \in \mathbb{N}$, and returns a witness $w \in W$. Sometimes informally we may refer to w as the POW.
- $\text{Verify}(r, \text{msg}, h, w)$ is a deterministic algorithm that takes as input a key $r \in R$, message $\text{msg} \in M$, hardness parameter $h \in \mathbb{N}$ and a witness $w \in W$, and returns **true** or **false** to indicate the validity of the proof.

Remark 1. In contrast with related literature on the subject [15, 8], we use in our syntax both a key and a message argument. By explicitly separating the key from the message, we aim at decoupling two requirements of any POW scheme: (i) the key is a way to guarantee that the computational work is “fresh” i.e., executed during a specific time interval (say, from the time the key became known to the prover), and (ii) the message is a way to allow arbitrary information to be encoded in the computed POW. Regarding the relevance of these inputs parameters, observe the following. If there is no way to tell when a POW was created (i.e., there is no way to guarantee freshness), then the adversary may precompute an arbitrary number of POWs before the protocol starts and try to trick other parties by pretending they all started at the same time. On the other hand, if there is no way to encode information in the POW, the usefulness of the primitive would be rather limited. Even though POW’s may be guaranteed to be fresh, if there is no context related to their computation, it will be easy for an adversary to copy POW’s across sessions or to different possibly adversarial contexts and thus “steal” the computational effort of honest participants.

Security properties. Next, we present a number of security properties that we will require POW schemes to satisfy. We start with the completeness property.

Definition 7. We say that a POW scheme is *complete* if for every $r \in R, h \in \mathbb{N}$, and $\text{msg} \in M$:

$$\Pr [w \leftarrow \text{Prove}(r, \text{msg}, h) : w \neq \perp \Rightarrow \text{Verify}(r, \text{msg}, h, w) = \text{true}] = 1.$$

Next we require that verification time is bounded by a parameter t .

Definition 8. We say that a POW scheme is *t-verifiable*, if the **Verify** algorithm takes time at most t (on all inputs).

Next we need to capture how the adversary can subvert an honest verifier and an honest prover in the context of POWs. In the first case, the adversary’s objective is to compute a number of POWs a lot faster than an honest prover while in the second case, it is to make the honest prover take too much time to generate a POW. It is useful to notice the parallel to the properties of soundness and

zero-knowledge for interactive proofs, As in this case, the two properties we are after here will be complementary.

We deal with malicious provers first. We put forth an attack that we will use to express a class of adversaries that attempt create POW's faster than expected. Intuitively this constitutes an attack against an honest verifier that may be trying to gauge a certain measure using the number of POWs. We call this *hardness against tampering and chosen message attack* (H-TCMA). The game defining the attack is shown in Figure 1. In order to capture the interactions that the adversary may have with different parties, we allow the adversary to have access to a **Prove** oracle, essentially allowing him to simulate any scenario where he interacts with other parties that use the same POW scheme. Every time the oracle is queried, we assume that it runs the **Prove** procedure for some uniformly sampled random tape. A subtle point in the modeling of security in the presence of such oracle is that the **Prove** oracle should also “leak” the time it took for a query to be processed. In an actual execution while interacting with honest parties that are producing POWs, time is a side-channel that may influence the adversarial strategy; in order to preserve this dependency on this side-channel we will require from the **Prove** oracle to leak this information. In addition, we require that the keys used by the adversary to construct proofs should be fresh, i.e., we want to avoid situations where the adversary outputs POWs that he has precomputed a long time ago. We model this by requiring that the keys must be generated by some oracle $\mathcal{R}()$ that samples values from a uniformly random distribution and is *external* to the adversary. Hence, the adversary will only know the value after he has queried $\mathcal{R}()$, thus providing us with a point of reference before which this value cannot be predicted. Furthermore we allow the adversary to tamper the POW keys by manipulating them via applying tampering functions that belong to a function family \mathcal{F} .

More formally, the adversary will have oracle access to \mathcal{R} , which is a sampling oracle of the uniform distribution. We use $\text{Ans}_{\mathcal{R}}$ to denote the set of responses of oracle \mathcal{R} . The adversary will also have access to $\mathcal{P}(\cdot, \cdot, \cdot)$, a proof-of-work oracle that on input (f, r, msg) , where $f \in \mathcal{F}$, r belongs to $\text{Ans}_{\mathcal{R}}$, and msg is some message, returns the pair (w, t) where w is the output of $\text{Prove}(f(r), msg, h)$ and t is the number of steps taken by the **Prove** algorithm on these parameters. Function $\text{Asked}(r', msg, w)$ is true if w was a response of \mathcal{P} to some query (f, r, msg) , where $f(r) = r'$.

$\text{Exp}_{\mathcal{A}, \mathcal{F}}^{\text{H-TCMA}}(1^\lambda, h, k)$
<ol style="list-style-type: none"> 1. $\{(r_1, f_1, msg_1, w_1), \dots, (r_k, f_k, msg_k, w_k)\} \leftarrow \mathcal{A}^{\mathcal{P}(\cdot, \cdot, \cdot), \mathcal{R}()}(1^\lambda, h, k);$ 2. $\text{output } \bigwedge_{i=1}^k (\text{Verify}(f_i(r_i), msg_i, w_i) \wedge \neg \text{Asked}(f_i(r_i), msg_i, w_i) \wedge r_i \in \text{Ans}_{\mathcal{R}} \wedge f_i \in \mathcal{F}).$

Figure 1: *The Hardness against Tampering and Chosen-Message Attack experiment for a POW scheme. $\mathcal{R}(\cdot)$ is a sampling oracle of the uniform distribution. \mathcal{F} is the class of tampering functions.*

We are now ready to formulate the security property of Hardness against Tampering and Chosen Message Attacks (H-TCMA). It is parameterized by two values β, ϵ , and informally, it states that no adversary \mathcal{A} exists that can participate in the experiment of Figure 1 so that \mathcal{A} takes at most t steps after querying \mathcal{R} for the first time and it produces $k \geq \beta \cdot t$ POWs with probability better than ϵ . Note that in total we allow any polynomial number of steps to \mathcal{A} , i.e., the adversary is allowed to execute a precomputation stage that permits it to obtain an arbitrary number of POWs. In the definition below, we allow β to depend on the hardness level h , and ϵ on β, t and $q_{\mathcal{P}}$, the number of queries the adversary makes to the proving oracle.

Definition 9. We say that a POW scheme has (β, ϵ) -*Hardness against Tampering and Chosen-Message Attacks* (H-TCMA), with tampering function class \mathcal{F} , if for every polynomial $t_{\text{pre}}(\cdot), t(\cdot)$, for any $h \in \mathbb{N}$, and every adversary \mathcal{A} , where \mathcal{A} is t_{pre} -bounded and starting from the query

on \mathcal{R} takes t steps and makes at most $q_{\mathcal{P}}$ queries to oracle \mathcal{P} , the probability of \mathcal{A} winning in $\text{Exp}_{\mathcal{A}, \mathcal{F}}^{\text{H-TCMA}}(1^\lambda, h, \lceil \beta(h) \cdot t \rceil)$ (Figure 1) is less than $\epsilon(\beta(h), t, q_{\mathcal{P}})$.

Next, we consider the setting of subverting the honest prover. Given that our primitive is non-interactive, subverting an honest prover, amounts to finding a certain set of keys over which the honest prover algorithm fails to produce proofs of work sufficiently fast and regularly. In more detail a POW scheme is α -efficient when the probability that the prover computes a POW in t steps is at least α .

Definition 10. We say that a POW scheme is α -efficient if for any $t \in \mathbb{N}$, $r \in R$, $msg \in M$ and $h \in \mathbb{N}$ it holds that:

$$\Pr [\text{Steps}_{\text{Prove}}(r, msg, h) < t] \geq \alpha(h, t)$$

In the same setting we require honest proving time to have at least some limited independence. This property, in combination with the efficiency and H-TCMA properties will prove crucial in ensuring that when multiple provers work together the distribution of the number of them who succeed in producing a POW has some “good” properties like variance and concentration.

Definition 11. Let M^{Prove} be a polynomially bounded RAM machine, \mathcal{E} denote the ensemble of random variables corresponding to the executions of M^{Prove} and \mathcal{I} be an ordered set of calls made to the **Prove** function. We denote by the random variable X_i the number of steps taken by the i -th function call in \mathcal{I} , where X_i is defined only in the executions where $|\mathcal{I}| \geq i$. We say that a POW scheme is *run-time m -wise independent*, if for any $I \subseteq \mathcal{I}$, there exists some distribution $(Y_i)_{i \in I}$ that is $(\text{negl}(\lambda))$ -close to $(X_i)_{i \in I}$ and¹ $\{Y_i\}_{i \in I}$ is m -wise independent.

4 The Bitcoin Backbone from POW

As mentioned in Section 1, existing analyses of Bitcoin have relied on the random oracle model to prove the properties of the blockchain protocol directly. In this section we take a reduction approach to the underlying cryptographic primitive—POW, as defined in Section 3—to prove the security of the Bitcoin backbone protocol [22] in the standard model. We start with some additional details about the model and pertinent preliminary definitions. We then continue with the security proof of the Bitcoin backbone protocol in the black-box POW setting. We conclude the section showing how Bitcoin’s POW scheme is secure in the random oracle model according to our definitions, and that the security guarantees we get from our analysis are similar to those proved in [22, 32].

4.1 Bitcoin backbone model and definitions

In [22] a security model was proposed for the analysis of the Bitcoin backbone protocol. Here we overview the basics, substituting IRAMs for ITMs for the reasons explained in Section 2. The execution of a protocol Π is driven by an “environment” program \mathcal{Z} that may spawn multiple instances running the protocol Π . The programs in question can be thought of as “interactive RAMs” communicating through registers in a well-defined manner, with instances and their spawning at the discretion of a control program which is also an IRAM and is denoted by C . In particular, the control program C forces the environment to perform a “round-robin” participant execution sequence for a fixed set of parties.

¹Depending on the subset \mathcal{I} , we take in account only the executions where the random variable family $\{X_i\}_{i \in I}$ is well defined.

Specifically, the execution driven by \mathcal{Z} is defined with respect to a protocol Π , an adversary \mathcal{A} (also an IRAM) and a set of parties P_1, \dots, P_n ; these are hardcoded in the control program C . The protocol Π is defined in a “hybrid” setting and has access to one “ideal functionality,” called the *diffusion channel* (see below). It is used as subroutine by the programs involved in the execution (the IRAMs of Π and \mathcal{A}) and is accessible by all parties once they are spawned.

Initially, the environment \mathcal{Z} is restricted by C to spawn the adversary \mathcal{A} . Each time the adversary is activated, it may communicate with C via messages of the form $(\text{Corrupt}, P_i)$. The control program C will register party P_i as corrupted, only provided that the environment has previously given an input of the form $(\text{Corrupt}, P_i)$ to \mathcal{A} and that the number of corrupted parties is less or equal t , a bound that is also hardcoded in C . The first party to be spawned running protocol Π is restricted by C to be party P_1 . After a party P_i is activated, the environment is restricted to activate party P_{i+1} , except when P_n is activated in which case the next party to be activated is always the adversary \mathcal{A} . Note that when a corrupted party P_i is activated the adversary \mathcal{A} is activated instead.

Initially, the diffusion functionality sets a variable *round* to be 1. It also maintains a $\text{Receive}()$ string (register) defined for each party P_i . A party is allowed at any moment to fetch the contents of its personal $\text{Receive}()$ string. Moreover, when the functionality receives an instruction to diffuse a message m from party P_i it marks the party as complete for the current round; note that m is allowed to be empty. At any moment, the adversary \mathcal{A} is allowed to receive the contents of all messages for the round and specify the contents of the $\text{Receive}()$ string for each party P_i . The adversary has to specify when it is complete for the current round. When all parties are complete for the current round, the functionality inspects the contents of all $\text{Receive}()$ strings and includes any messages that were diffused by the parties in the current round but not contributed by the adversary to the $\text{Receive}()$ tapes. The variable *round* is then incremented.

Next, we consider the complications in the modeling due to the analysis of Bitcoin in the concrete security setting. Both in [22] and [32] a modified version of the standard simulation-based paradigm of [16] is followed, where there exist both a malicious environment and a malicious adversary. In addition, the POW scheme is modeled in a non black-box way using a random oracle (RO), and the computational power of the adversary is then bounded by limiting the number of queries it can make to the RO per round. Since in this work the POW scheme is modeled in a black-box way, an alternative approach to bound the adversary’s power is needed.

A naïve first approach is to only bound the computational power of \mathcal{A} . Unfortunately this will not work for several reasons. Firstly, nothing stops the environment from aiding the adversary, i.e., computing POWs, and then communicating with it through their communication channel or some other subliminal channel; secondly, even if we bound the *total* number of steps of \mathcal{A} , it is not clear how to bound the steps it is taking per round in the model of [16], which we build on. Furthermore, another issue arising is that if the adversary is able to send, say, θ messages in each round, it can force each honest party to take $\theta \cdot t_{\text{ver}}$ extra steps per round.

In order to capture these considerations we are going to define a predicate on executions and prove our properties in disjunction with this predicate.

Definition 12. Let $(t_{\mathcal{A}}, \theta)$ -good be a predicate defined on executions in the hybrid setting described above. Then E is $(t_{\mathcal{A}}, \theta)$ -good, where E is one such execution, if

- the total number of steps taken by \mathcal{A} and \mathcal{Z} per round is no more than $t_{\mathcal{A}}$;²
- the adversary sends at most θ messages per round.

²The adversary cannot use the running time of honest parties that it has corrupted; it is activated instead of them during their turn. Also, note that it is possible to compute this number by counting the number of configurations that \mathcal{A} or \mathcal{Z} are activated per round.

For the rest of the paper, we assume that the environment is restricted so that with overwhelming probability the execution is good for fixed parameters $t_{\mathcal{A}}$ and θ , and that there exists some polynomial $p(\cdot)$ such that the number of rounds of all executions is upper bounded by $p(\lambda)$.

4.2 Security properties of the blockchain

A number of desired basic properties for the blockchain were introduced in [22, 29]. At a high level, the first property, called *common prefix*, has to do with the existence, as well as persistence in time, of a common prefix of blocks among the chains of honest players. Here we will consider a stronger variant of the property, presented in [29, 32], which allows for the black-box proof of application-level properties (such as the *persistence* of transactions entered in a public transaction ledger built on top of the Bitcoin backbone—cf. Appendix 4.3). We will use $\mathcal{C} \preceq \mathcal{C}'$ to denote that some chain \mathcal{C} is a prefix of some other chain \mathcal{C}' , and $\mathcal{C}^{\lceil k}$ to denote the chain resulting from removing the last k blocks of \mathcal{C} .

Definition 13 ((Strong) Common Prefix). The *strong common prefix property* Q_{cp} with parameter $k \in \mathbb{N}$ states that the chains $\mathcal{C}_1, \mathcal{C}_2$ reported by two, not necessarily distinct honest parties P_1, P_2 , at rounds r_1, r_2 , with $r_1 \leq r_2$, satisfy $\mathcal{C}_1^{\lceil k} \preceq \mathcal{C}_2$.

The next property relates to the proportion of honest blocks in any portion of some honest player’s chain.

Definition 14 (Chain Quality). The *chain quality property* Q_{cq} with parameters $\mu \in \mathbb{R}$ and $k, k_0 \in \mathbb{N}$ states that for any honest party P with chain \mathcal{C} in $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}(\kappa, z)$, it holds that for any k consecutive blocks of \mathcal{C} , excluding the first k_0 blocks, the ratio of adversarial blocks is at most μ .

Further, in the derivations in [22] an important lemma was established relating to the rate at which the chains of honest players were increasing as the Bitcoin backbone protocol was run. This was explicitly considered in [29] as a property under the name *chain growth*. Similarly to the variant of the common prefix property above, this property along with chain quality were shown sufficient for the black-box proof of application-level properties (in this case, transaction ledger *liveness*; see Appendix 4.3).

Definition 15 (Chain Growth). The chain growth property Q_{cg} with parameters $\tau \in \mathcal{R}$ (the “chain speed” coefficient) and $s, r_0 \in \mathbb{N}$ states that for any round $r > r_0$, where honest party P has chain \mathcal{C}_1 at round r and chain \mathcal{C}_2 at round $r + s$ in $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}(\kappa, z)$, it holds that $|\mathcal{C}_2| - |\mathcal{C}_1| \geq \tau \cdot s$.

4.3 Robust public transaction ledgers

A *public transaction ledger* is defined with respect to a set of valid ledgers \mathcal{L} and a set of valid transactions \mathcal{T} , each one possessing an efficient membership test. A ledger $\mathbf{x} \in \mathcal{L}$ is a vector of sequences of transactions $\text{tx} \in \mathcal{T}$. Each transaction tx may be associated with one or more *accounts*. Ledgers correspond to chains in the backbone protocol. In the protocol execution there also exists an oracle Txgen that generates valid transactions. Note, that it is possible for the adversary to create two transactions that are conflicting; valid ledgers must not contain conflicting transaction. We will assume that the oracle is unambiguous, i.e., that the adversary cannot create transactions that come in ‘conflict’ with the transactions generated by the oracle. A transaction is called *neutral* if there does not exist any transactions that comes in conflict with it.

In order to turn the backbone protocol into a protocol realizing a public transaction ledger suitable definitions were given for functions $V(\cdot), R(\cdot), I(\cdot)$ in [22]. Namely, $V(\langle x_1, \dots, x_m \rangle)$ is true

if its input is a valid ledger. Function $R(\mathcal{C})$ returns the contents of the chain if they constitute a valid ledger, otherwise it is undefined. Finally, $I(st, \mathcal{C}, round, \text{INPUT}(), \text{RECEIVE}())$ returns the largest subsequence of transactions in the input and receive tapes that constitute a valid ledger, with respect to the contents of the chain the party already has, together with a randomly generated neutral transaction. We denote the instantiation of our protocol with these functions by Π_{PL} . For more details we refer to [22].

Definition 16. A protocol Π implements a *robust public transaction ledger* in the q -bounded synchronous setting if it satisfies the following two properties:

- **Persistence:** Parameterized by $k \in \mathbb{N}$ (the “depth” parameter), if in a certain round an honest player reports a ledger that contains a transaction tx in a block more than k blocks away from the end of the ledger, then tx will always be reported in the same position in the ledger by any honest player from this round on.
- **Liveness:** Parameterized by $u, k \in \mathbb{N}$ (the “wait time” and “depth” parameters, resp.), provided that a transaction either (i) issued by Tngen , or (ii) is neutral, is given as input to all honest players continuously for u consecutive rounds, then there exists an honest party who will report this transaction at a block more than k blocks from the end of the ledger.

4.4 The modified Bitcoin backbone protocol

In this section we describe the Bitcoin backbone protocol in the black-box POW setting. First, we introduce some notation needed to understand the description of the algorithms. We will use the terms *block* and *chain* to refer to tuples of the form $\langle s, x, w \rangle$ and sequences of such tuples, respectively. Each block contains a seed, data, and a witness denoted by s, x, w , respectively. We say that a chain is *valid* if and only if for any two consecutive blocks $\langle s_i, x_i, w_i \rangle, \langle s_{i+1}, x_{i+1}, w_{i+1} \rangle$ it holds that $H(s_i, G(x_i), w_i) = s_{i+1}$. We call this value the *hash of block B_i* and extend the definition of H to blocks so that it is equal to $H(B_i)$.

Algorithm 1 The *proof of work* function, parameterized by q, D and hash functions $H(\cdot), G(\cdot)$. The input is (x, \mathcal{C}) .

```

1: function pow( $x, \mathcal{C}$ )
2:   if  $\mathcal{C} = \varepsilon$  then                                     ▷ Determine proof of work instance
3:      $s \leftarrow 0$ 
4:   else
5:      $s \leftarrow H(\text{head}(\mathcal{C}))$ 
6:   end if
7:    $B \leftarrow \varepsilon$ 
8:    $w \leftarrow \text{Prove}(s, x, h)$                              ▷ Run the prover of the POW scheme.
9:   if  $w \neq \perp$  then
10:     $B \leftarrow \langle s, x, w \rangle$ 
11:  end if
12:   $\mathcal{C} \leftarrow \mathcal{C}B$                                          ▷ Extend chain
13:  return  $\mathcal{C}$ 
14: end function

```

The modifications with respect to the original protocol in [22] are as follows: In Algorithm 1 (proof of work function) the **Prove** function of the underlying POW scheme is invoked for a limited

Algorithm 2 The *chain validation predicate*, parameterized by q, D , the hash functions $G(\cdot), H(\cdot)$, and the *input validation predicate* $V(\cdot)$. The input is \mathcal{C} .

```

1: function validate( $\mathcal{C}$ )
2:    $b \leftarrow V(\mathbf{x}_{\mathcal{C}}) \wedge (\mathcal{C} \neq \varepsilon)$ 
3:   if  $b = \text{True}$  then                                      $\triangleright$  The chain is non-empty and meaningful w.r.t.  $V(\cdot)$ 
4:      $\langle s, x, w \rangle \leftarrow \text{head}(\mathcal{C})$ 
5:      $s' \leftarrow H(\text{head}(\mathcal{C}))$ 
6:     repeat
7:        $\langle s, x, w \rangle \leftarrow \text{head}(\mathcal{C})$ 
8:       if  $\text{Verify}(s, x, h, w) \wedge (H(\text{head}(\mathcal{C})) = s')$  then
9:          $s' \leftarrow s$                                       $\triangleright$  Retain hash value
10:         $\mathcal{C} \leftarrow \mathcal{C}^{\uparrow 1}$                               $\triangleright$  Remove the head from  $\mathcal{C}$ 
11:       else
12:          $b \leftarrow \text{False}$ 
13:       end if
14:     until  $(\mathcal{C} = \varepsilon) \vee (b = \text{False})$ 
15:   end if
16:   return ( $b$ )
17: end function

```

number of steps so that the total number of steps of the invoking party does not exceed the $t_{\mathcal{H}}$ bound per round; in Algorithm 2 (chain validation predicate) the `Verify` predicate is replaced with a call to the `Verify` algorithm of the POW scheme; and in Algorithm 3 we assume that the honest parties start the execution with a “genesis” block. We require the message (content) of this block to be sufficiently random, in order to ensure that the adversary cannot precompute blocks related to it. Differently stated, we assume the existence of a *common random string* (CRS), that becomes available to all parties at the start of the execution. We call the resulting protocol $\Pi_{\text{PL}}^{\text{POW}}$ (‘PL’ for “public ledger”).

4.5 Security proof

Next, we prove that $\Pi_{\text{PL}}^{\text{POW}}$ implements a robust public transaction ledger (Definition 16). First, we introduce some additional notation.

For each round j , we define the Boolean random variables X_j and Y_j as follows. Let $X_j = 1$ if and only if j was a *successful round*, i.e., at least one honest party computed a POW at round j , and let $Y_j = 1$ if and only if j was a *uniquely successful round*, i.e., exactly one honest party computed a POW at round j . With respect to a set of rounds S and some block B , let $Z_B(S)$ denote the number of blocks broadcast by the adversary during S that have B as their ancestor. Also, let $X(S) = \sum_{j \in S} X_j$ and define $Y(S)$ similarly.

Let t_{bb} (bb for backbone) be an upper bound on the number of steps needed to run the code of an honest party in one round, besides the `Prove` and `Verify` calls. By carefully analyzing the backbone protocol one can extract an upper bound on this value.³ To aid our presentation, we will use $t'_{\mathcal{A}}$ and $t'_{\mathcal{H}}$ to denote the following quantities which roughly correspond to the time needed by a

³Note that t_{bb} depends on the running time of three external functions: $V(\cdot), I(\cdot)$ and $R(\cdot)$. For example, in Bitcoin these functions include the verification of digital signatures, which would require doing modular exponentiations. In any case t_{bb} is at least linear in λ .

Algorithm 3 The Bitcoin backbone protocol, parameterized by the *input contribution function* $I(\cdot)$ and the *chain reading function* $R(\cdot)$.

```

1:  $\mathcal{C} \leftarrow \text{Gen}$ 
2:  $st \leftarrow \varepsilon$ 
3:  $round \leftarrow 0$ 
4: while TRUE do
5:    $\tilde{\mathcal{C}} \leftarrow \text{maxvalid}(\mathcal{C}, \text{any chain } \mathcal{C}' \text{ found in RECEIVE}())$ 
6:    $\langle st, x \rangle \leftarrow I(st, \tilde{\mathcal{C}}, round, \text{INPUT}(), \text{RECEIVE}())$  ▷ Determine the  $x$ -value.
7:    $\mathcal{C}_{\text{new}} \leftarrow \text{pow}(x, \tilde{\mathcal{C}})$ 
8:   if  $\mathcal{C} \neq \mathcal{C}_{\text{new}}$  then
9:      $\mathcal{C} \leftarrow \mathcal{C}_{\text{new}}$ 
10:    BROADCAST( $\mathcal{C}$ )
11:  end if
12:   $round \leftarrow round + 1$ 
13:  if INPUT() contains READ then
14:    write  $R(\mathbf{x}_{\mathcal{C}})$  to OUTPUT()
15:  end if
16: end while

```

Algorithm 4 The function that finds the “best” chain, parameterized by function $\text{max}(\cdot)$. The input is $\{\mathcal{C}_1, \dots, \mathcal{C}_k\}$.

```

1: function maxvalid( $\mathcal{C}_1, \dots, \mathcal{C}_k$ )
2:    $temp \leftarrow \varepsilon$ 
3:   for  $i = 1$  to  $k$  do
4:     if validate( $\mathcal{C}_i$ ) then
5:        $temp \leftarrow \text{max}(\mathcal{C}, temp)$ 
6:     end if
7:   end for
8:   return  $temp$ 
9: end function

```

RAM machine to simulate one round in the execution of the Bitcoin protocol, without taking into account calls made to the **Prove** routine by the honest parties, and to the minimum number of steps spent by an honest party on running the **Prove** routine in a round, respectively:

$$t'_{\mathcal{A}} = t_{\mathcal{A}} + n \cdot t_{\text{bb}} + \theta t_{\text{ver}} \quad \text{and} \quad t'_{\mathcal{H}} = t_{\mathcal{H}} - t_{\text{bb}} - \theta t_{\text{ver}}.$$

It holds that at least $n - t$ parties will run the **Prove** routine for at least $t'_{\mathcal{H}}$ steps at every round.

Next, we focus on the hash functions used by the Bitcoin protocol, and the necessary security assumptions to avoid cycles in the blockchains or the possibility that the adversary will predict the value of some “future” block. First, note that in the actual implementation of Bitcoin an unkeyed hash function is used, namely, a double invocation of SHA-256. In previous analyses of the protocol this was modeled by a random oracle. We choose to model it as a keyed hash function family

$$\mathcal{H} = \{H_K(\cdot) : \{0, 1\}^{\lceil \log(|R|) \rceil + \lceil \log(|W|) \rceil + \lambda} \rightarrow \{0, 1\}^\lambda\}_{K \in \mathcal{K}}$$

that is *collision* and *everywhere preimage resistant* (Definition 1); the CRS we have already assumed can serve as the key of our hash function. In our analysis we will also need G to be collision resistant and take inputs of variable length. It is well known (see, e.g., [17]) that given a fixed-length collision-resistant hash function family, we can construct an arbitrary-length collision-resistant hash function family. To aid readability, we will sometimes omit the key of both functions (as we already do in the description of the protocol). Finally, we require the input contribution function $I(\cdot)$ of our protocol to satisfy the *input entropy assumption* made in [22], which states that two independent invocations of $I(\cdot)$ with arbitrary parameters output the same values only with negligible probability.




Figure 2: *The hash structure of the blocks in the Bitcoin protocol.*

Figure 2 depicts a chain of blocks maintained by the Bitcoin protocol. Notice that the repeated hash function application is similar to the basic cascade construction [10], denoted here by $H^*(\cdot, \cdot)$, when the fixed-length hash function family used is \mathcal{H} . In this construction, it holds that for any chain \mathcal{C} there always exists a message M such that the hash of the last block, i.e., $H(\text{tail}(\mathcal{C}))$, is equal to $H^*(H(\text{head}(\mathcal{C})), M)$. In Appendix A, we prove that $\{H^*(K, \cdot) : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda\}_{K \in \mathcal{K}}$ is collision resistant, and $\{H^*(\cdot, G(I(\cdot; K)))\}_{K \in \mathcal{K}}$ is everywhere preimage resistant. These two properties imply the next lemma, which has to do with the way the blocks are connected. (Proof in Appendix ??.)

Definition 17. An *insertion* occurs when, given a chain \mathcal{C} with two consecutive blocks B and B_0 , a block B^* created after B_0 is such that B, B^*, B_0 form three consecutive blocks of a valid chain. A *copy* occurs if the same block exists in two different positions. A *prediction* occurs when a block extends one honest block which was computed at a later round.

Lemma 18. *Assuming that \mathcal{H} is collision and everywhere preimage resistant, no insertions, no copies and no predictions occur with probability $1 - \text{negl}(\lambda)$.*

Proof. Insertions and copies imply that one block extends two distinct blocks, i.e. there exist blocks B_1, B_2, B_3 such that $s_3 = H(B_1) = H(B_2)$. This in turn implies that there exist M_1, M_2 such that $H^*(H(\text{Gen}), M_1) = H^*(H(\text{Gen}), M_2)$, where Gen is the genesis block and $M_1 \neq M_2$. This constitutes a break of the collision resistance property of H^* , and thus it only happens with negligible probability on λ .

Next, for the sake of contradiction assume that prediction happens with non-negligible probability. We can use this adversary that breaks prediction to break the ePre property. The adversary would just output at random the seed of one of the blocks he has mined (since they are polynomial in number at most), say B_2 , as y in the ePre game. By our assumption with non negligible probability some honest block B_1 will be mined afterwards such that $y = H(B_1)$, hence we can use the parameters of this block to generate an output to win the ePre game. \square

As mentioned in Section 2 (and proved in Appendix A), the everywhere preimage resistance of H^* implies that the function family $\mathcal{F} = \{H^*(\cdot, M)\}_{M \in \mathcal{M}}$ is unpredictability preserving. For the

λ	security parameter
n	number of parties
$t_{\mathcal{H}}$	number of steps per round per honest party
$t_{\mathcal{A}}$	total number of adversarial steps per round
θ	upper bound on the number of messages sent by the adversary per round
β	upper bound on POW computation rate per step
γ	lower bound on the rate of uniquely successful rounds
f	lower bound on the rate of successful rounds
δ	advantage from the honest majority assumption
σ	quality of concentration of random variables in typical executions
k	number of blocks for the common-prefix property
ℓ	number of blocks for the chain-quality property
η	parameter determining block to round translation
ε	probability that an execution is typical

Table 1: The parameters in our analysis.

rest of this section we will assume that the underlying POW scheme used in the Bitcoin backbone protocol with respect to the unpredictability-preserving tampering function class \mathcal{F} is:

- Complete;
- (β, ϵ) -H-TCMA secure;
- α -efficient;
- run-time m -wise independent;
- t_{ver} -verifiable,

where $\epsilon \in \text{negl}(t)$ and $m \geq n \log \lambda$.

Next, we prove that the adversary cannot mine blocks that extend an honest block created recently at a very high rate with probability better than that of breaking the H-TCMA property. For a summary of our notation we refer to Table 1.

Lemma 19. *For any set of consecutive rounds S , it holds that the probability that there exists some honest block B mined in S such that $Z_B(S) > \beta t'_{\mathcal{A}} |S|$, is at most $\epsilon(\beta, t'_{\mathcal{A}} \cdot |S|, n \cdot |S|)$.*

Proof. Let $S = \{r' | r \leq r' < r + s\}$ and let E be the event where in $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{t, n}$ the adversary has mined at least $\beta t'_{\mathcal{A}} s$ blocks that descend some honest block B until round $r + s$. For the sake of contradiction, assume that the lemma does not hold, and thus the probability that the execution is $(t_{\mathcal{A}}, \theta)$ -good and E holds is greater than $\epsilon(\beta, t'_{\mathcal{A}} s, ns)$. Using \mathcal{A} , we will construct an adversary \mathcal{A}' that wins the H-TCMA game with probability greater than that. \mathcal{A}' is going to run internally \mathcal{A} and \mathcal{Z} , while at the same time perfectly simulating the view of honest parties using the two oracles that he has in his disposal on the H-TCMA game. This way, the view of \mathcal{A}, \mathcal{Z} will be indistinguishable both in the real and the simulated runs, and thus the probability that E happens will be the same in both cases.

We are going to describe the two stages of \mathcal{A}' separately, i.e. before and after querying \mathcal{R} for the first time. First, \mathcal{A}' perfectly simulates honest parties up to round $r - 1$ and at the same time runs \mathcal{A} and \mathcal{Z} in a black-box way. He can do this since he has polynomial time on λ on his disposal. Note, that up until this point in the eyes of \mathcal{A} and \mathcal{Z} the simulated execution is indistinguishable compared to the real one. If r is the first round, \mathcal{A}' will do nothing.

For the second stage, \mathcal{A}' is again going to simulate honest parties behavior, from round r until round $r + s$, but in a different way. Instead of running the **Prove** algorithm for each non-corrupted

honest party at every round, it makes a query to the \mathcal{P} oracle with the respective parameters. Then, it checks if the honest party succeeded in making a POW in this round by comparing the number of steps needed to make this POW to the number of proving steps available to the party at this round. Hence, honest parties have to do n queries to the proving oracle per round. The adversary can also send up to θ messages per round to honest parties which they have to verify, thus inducing an additional $\theta \cdot t_{\text{ver}}$ overhead in the simulation. Note that \mathcal{A}' has to run the verification procedure only once per message.

Moreover, for the reduction to work we also want all honest blocks created after round r to be related to the outputs of the randomness oracle \mathcal{R} through some function in \mathcal{F} . In our scenario, \mathcal{R} will generate uniformly random strings of the same length as the randomness needed by the input contribution function I . At every round starting from r and for each honest party, \mathcal{A}' will query \mathcal{R} and use the response to run I . If an honest block is mined using this response, then any block descending it will be related to the randomness used through some function in \mathcal{F} . More precisely, if $\langle s, I(m; r_i), w \rangle$ is the honest block mined using the response r_i of oracle \mathcal{R} , then any block descending it is of the form $\langle H^*(H(\langle s, I(m; r_i), w \rangle), M), x', w' \rangle$ for some M, x', w' . Hence, it is related to r_i through some function in \mathcal{F} as required by the H-TCMA security definition.

Since \mathcal{A} and \mathcal{Z} cannot distinguish between the bitcoin execution and the one we described above, E will occur with probability at least $\epsilon(\beta, t'_A s, ns)$, i.e. \mathcal{A} will compute at least $\beta t'_A s$ blocks starting from round r and up to round $r + s$ that descend some honest block B mined during these rounds. Note, that these blocks are also valid POWs, that are related through \mathcal{F} to the responses of \mathcal{R} . Hence, \mathcal{A}' will win the H-TCMA game with probability greater than $\epsilon(\beta, t'_A s, ns)$, while being $s \cdot (t_A + \theta \cdot t_{\text{ver}} + t_{\text{bb}} \cdot n) = s \cdot t'_A$ -bounded and having made at most ns queries to the proving oracle, which is a contradiction to our initial assumption. A sketch of the reduction is given at Figure 3.

Note that we can do exactly the same reduction without using the oracle to simulate the proving procedure of the honest parties. The total running time of the second stage of \mathcal{A}' is at least $s \cdot (t'_A + nt'_{\mathcal{H}})$ -bounded and hence the probability he can win is $\epsilon(\beta, s \cdot (t'_A + nt'_{\mathcal{H}}), 0)$ □

Moreover, if in the previous proof instead of using the proving oracle provided in the H-TCMA game to simulate honest parties' work, we run their code as it is, we can derive the following bound on the total number of blocks produced by both honest and malicious parties during a certain number of rounds.

Corollary 20. *For any set of consecutive rounds S , the probability that there exists some honest block B mined in S such that $Z_B(S) + X(S) > \beta(t'_A + nt'_{\mathcal{H}}) \cdot |S|$ is less than $\epsilon(\beta, |S| \cdot (t'_A + nt'_{\mathcal{H}}), 0)$.*

Next, we prove lower bounds on the rate of successful and uniquely successful rounds. Our proof crucially depends on the independence of the POW scheme. More specifically, the property is needed to ensure that the sum of the Bernoulli random variables of the events that a round is uniquely successful concentrate around the mean, with the implication that we can lower-bound the rate of uniquely successful rounds with good probability. Note that the model we described at the beginning of this section is captured by the class of models mentioned in the definition of the independence property.

Lemma 21. *For any set of consecutive rounds S and for any $\sigma \in (0, 1)$ it holds that:*

- *The probability that less than $\gamma \cdot |S|$ uniquely successful rounds occur in S is at most $\epsilon_s(\gamma, |S|)$;*
- *the probability that less than $f \cdot |S|$ successful rounds occur in S is at most $\epsilon_s(f, |S|)$; and*
- *$\gamma \leq f$,*

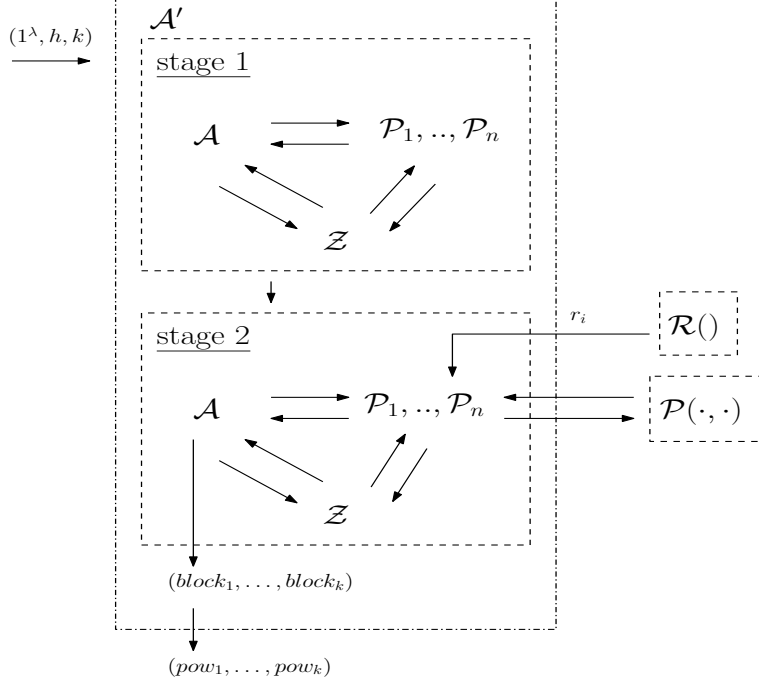


Figure 3: The reduction from the Bitcoin backbone to the H-TCMA game of Lemma 19.

for $\gamma = (1 - \sigma)(n - t)\alpha(h, t'_{\mathcal{H}})(1 - \epsilon(\beta, t'_{\mathcal{H}}, 0))^{n-t-1}$, $f = (1 - \sigma)(1 - (1 - \alpha(h, t'_{\mathcal{H}}))^{n-t})$, and $\epsilon_s(x, |S|) = \left(\frac{(1-\sigma)^2 \log \lambda}{4|S|\sigma^2 x^2}\right)^{\lfloor \log \lambda/2 \rfloor}$.

Proof. For some fixed execution we will denote by the array $\mathbf{T}_{S \times n} = (t_{i,j}) \in \mathbb{N}^{|S| \times n}$ the number of steps each honest party takes running the **Prove** routine, for each round in the set S . It holds that at most t elements of each column are zero, i.e. corrupted, and the rest are lower bounded by $t'_{\mathcal{H}}$ and upper bounded by $t_{\mathcal{H}}$. W.l.o.g let $S = \{1, \dots, s\}$.

Since this lemma talks about the steps taken by the **Prove** function, we are going to use the almost m -wise independence property of the POW scheme, and do all the analysis on the m -wise independent random variable defined by this property. For the rest of this proof, unless explicitly stated, assume that the $\text{Steps}_{\text{Prove}}(r, m, h)$ random variable refers to its idealized m -wise independent version. We first buildup some notation to help in our analysis. For some array $(r_{i,j}) \in R^{s \times n}$, $(\text{msg}_{i,j}) \in M^{s \times n}$ and for $h \in \mathbb{N}$ let:

- random variable $P_{i,j} = 1$ if $\text{Steps}_{\text{Prove}}(r_{i,j}, \text{msg}_{i,j}, h) \leq t_{i,j}$, and 0 otherwise;
- random variable $Y_i = 1$ if $\sum_{j=1}^n P_{i,j} = 1$ and 0 otherwise.
- random variable $Y = \sum_{i \in [s]} Y_i$.

It easily follows from the efficiency property that $\Pr[P_{i,j} = 1] \geq \alpha(h, t_{i,j})$.

Claim 1. *The random variable families $(P_{i,j})_{i \in [s], j \in [n]}$ and $(Y_i)_{i \in [s]}$ are m and $\lfloor m/n \rfloor$ -wise independent respectively.*

Proof of Claim. First, notice that the m -wise independence of the scheme implies m -wise independence of $(P_{i,j})$. We will show this for two random variables and the extension to m variables will

be obvious. Let $P_1, P_2 \in (P_{i,j})$ and $x_1, x_2 \in \{0, 1\}$, then

$$\begin{aligned}
\Pr[P_1 = x_1 \wedge P_2 = x_2] &= \Pr[\text{Steps}_{\text{Prove}}(r_1, m_1, h) \in S_1 \wedge \text{Steps}_{\text{Prove}}(r_2, m_2, h) \in S_2] \\
&= \sum_{(s_1, s_2) \in S_1 \times S_2} \Pr[\text{Steps}_{\text{Prove}}(r_1, m_1, h) = s_1 \wedge \text{Steps}_{\text{Prove}}(r_2, m_2, h) = s_2] \\
&= \sum_{(s_1, s_2) \in S_1 \times S_2} \Pr[\text{Steps}_{\text{Prove}}(r_1, m_1, h) = s_1] \cdot \Pr[\text{Steps}_{\text{Prove}}(r_2, m_2, h) = s_2] \\
&= \sum_{s_1 \in S_1} \Pr[\text{Steps}_{\text{Prove}}(r_1, m_1, h) = s_1] \cdot \sum_{s_2 \in S_2} \Pr[\text{Steps}_{\text{Prove}}(r_2, m_2, h) = s_2] \\
&= \Pr[\text{Steps}_{\text{Prove}}(r_1, m_1, h) \in S_1] \cdot \Pr[\text{Steps}_{\text{Prove}}(r_2, m_2, h) \in S_2] \\
&= \Pr[P_1 = x_1] \cdot \Pr[P_2 = x_2]
\end{aligned}$$

where S_1, S_2 are either $[0, t]$ or (t, ∞) depending on x_1, x_2 , and r_1, m_1, r_2, m_2 are the parameters of the random processes. We use the m -wise independence property on the third line, hence as long as the number of processes is not bigger than m the claim will follow.

Next, we prove the second point of the claim. Again, w.l.o.g we only show it for 2 random variables, Y_1, Y_2 and the extension to k/n is obvious. Let $y_1, y_2 \in \{0, 1\}$, then

$$\begin{aligned}
\Pr[Y_1 = y_1 \wedge Y_2 = y_2] &= \Pr[\sum_{j \in [n]} P_{1,j} \in S_1 \wedge \sum_{j \in [n]} P_{2,j} \in S_2] \\
&= \sum_{(s_1, s_2) \in S_1 \times S_2} \Pr[\sum_{j \in [n]} P_{1,j} = s_1 \wedge \sum_{j \in [n]} P_{2,j} = s_2] \\
&= \sum_{(s_1, s_2) \in S_1 \times S_2} \Pr[\sum_{j \in [n]} P_{1,j} = s_1] \cdot \Pr[\sum_{j \in [n]} P_{2,j} = s_2] \\
&= \sum_{s_1 \in S_1} \Pr[\sum_{j \in [n]} P_{1,j} = s_1] \cdot \sum_{s_2 \in S_2} \Pr[\sum_{j \in [n]} P_{2,j} = s_2] \\
&= \Pr[Y_1 = y_1] \cdot \Pr[Y_2 = y_2]
\end{aligned}$$

where S_1, S_2 are $\{1\}$ or $\{0, 2, 3, \dots\}$ depending on y_1, y_2 . Since each variable of the family (Y_i) is an n -term sum of variables from the family ($P_{i,j}$), independence drops by a factor of n for these variables. \dashv

Claim 2. *It holds that for any $i \in S$: $\mathbb{E}[Y_i] \geq \gamma$*

Proof of Claim. Since $m > n$ it follows that

$$\begin{aligned}
\mathbb{E}[Y_i] &= \Pr[Y_i = 1] = \Pr[\sum_{j \in [n]} P_{i,j} = 1] \\
&= \sum_{j \in [n]} \Pr[P_{i,j} = 1] \cdot \prod_{m \in [n] \setminus \{j\}} \Pr[P_{i,m} = 0] \\
&\geq \sum_{j \in [n]} \alpha(h, t_{i,j}) \prod_{m \in [n] \setminus \{j\}} (1 - \epsilon(\beta, t_{i,m}, 0)) \\
&\geq (n - t) \alpha(h, t'_{\mathcal{H}}) (1 - \epsilon(\beta, t'_{\mathcal{H}}, 0))^{n-t-1} = \gamma
\end{aligned}$$

The inequalities follow from the efficiency and H-TCMA properties. Note that $\beta t_{i,j}$ is greater than 0, hence the adversary has to compute at least one POW in the H-TCMA experiment. Also note,

that in order for $\mathbb{E}[Y_i]$ to be big, α must be as big as possible and ϵ must be as small as possible. Hence, as we will see later, our goal would be to calibrate h so that it maximizes the ratio γ/β . \dashv

The lemma follows from a direct application of the following claim about the concentration of m -wise independent r.v.'s, found as a problem statement in [36], and improved here.

Claim 3. *If Y is the sum of t k -wise independent r.v.'s each of which is confined to the interval $[0, 1]$ with $\mu = \mathbb{E}[Y]$ and $\epsilon > 0$, then*

$$\Pr[|Y - \mu| \geq t\epsilon] \leq \left(\frac{k}{t\epsilon^2}\right)^{\lfloor k/2 \rfloor}$$

Proof of Claim. Assume k is even. First, by the linearity of expectation it holds that:

$$\mathbb{E}[(Y - \mathbb{E}[Y])^k] = \sum_{S \subseteq [t]^k} \mathbb{E}\left[\prod_{i \in S} (Y_i - \mathbb{E}[Y_i])\right]$$

Due to the k -wise independence property, all terms in the sum where an index occurs only one time are zero, since $\mathbb{E}[Y_i - \mathbb{E}[Y_i]] = \mathbb{E}[Y_i] - \mathbb{E}[Y_i] = 0$. Hence, each index should show up twice at S . The number of possible pairs of indexes is upper bounded by $k^{k/2}$ and the number of possible assignments for a specific set of pairs is $t^{k/2}$. Thus, there are at most $(tk)^{k/2}$ non-zero terms in the sum. Moreover, each term is less than 1. Therefore

$$\mathbb{E}[(Y - \mathbb{E}[Y])^k] \leq (tk)^{k/2}$$

By Markov's inequality we get the desired result:

$$\Pr[Y - \mathbb{E}[Y] \geq t\epsilon] \leq \frac{\mathbb{E}[(Y - \mathbb{E}[Y])^k]}{(t\epsilon)^k} \leq \frac{(tk)^{k/2}}{(t\epsilon)^k} \leq \left(\frac{k}{t\epsilon^2}\right)^{k/2}$$

\dashv

By the linearity of expectation we have that $(1 - \delta)E[Y] \geq \gamma|S|$. For $\epsilon = \delta\gamma/(1 - \delta)$, $t = |S|$, $k = \log \lambda$ we get that:

$$\begin{aligned} \Pr[Y \leq \gamma|S|] &\leq \Pr[Y \leq (1 - \delta)E[Y]] \\ &\leq \Pr[|Y - E[Y]| \geq \delta E[Y]] \\ &\leq \Pr[|Y - E[Y]| \geq \delta\gamma|S|/(1 - \delta)] \\ &\leq \left(\frac{(1 - \delta)^2 \log \lambda}{4|S|\delta^2\gamma^2}\right)^{\lfloor \log \lambda / 2 \rfloor} \end{aligned}$$

The same fact follows with only negligible difference in probability for our scheme due to the $(\text{negl}(\lambda))$ -closeness property and since Y is a function of the distribution referred by the property.

Moreover, we can calculate in a similar way a bound for $X(S)$. First, note that $(1 - \delta)\mathbb{E}[X_i] \geq f$. By applying the concentration bound for $\epsilon = \delta f/(1 - \delta)$, $t = |S|$ we get:

$$\begin{aligned} \Pr[X \leq f|S|] &\leq \Pr[X \leq (1 - \delta)E[X]] \\ &\leq \Pr[|X - E[X]| \geq \delta E[X]] \\ &\leq \Pr[|X - E[X]| \geq \delta|S|f/(1 - \delta)] \\ &\leq \left(\frac{(1 - \delta)^2 \log \lambda}{4|S|\delta^2 f^2}\right)^{\lfloor \log \lambda / 2 \rfloor} \end{aligned}$$

We note that our results also hold when some of the honest parties are corrupted, i.e. the number of steps they execute is 0. Finally, it is easy to see that $\gamma \leq f$ since they correspond to the probabilities of the following two events: throwing $n - t$ coins that are $(1 - \alpha)$ -biased and having exactly one come heads, and at least one come heads respectively. Obviously, the second event includes the first one, and the result follows. \square

We are now ready to define the set of *typical executions* for this setting. This strategy was also followed in [22]. However, here we will need to adapt the definition due to the difficulties associated with performing a black-box reduction to a POW scheme.

Definition 22 (Typical execution). An execution is η -*typical* if and only if for any set S of consecutive rounds with $|S| \geq \eta\lambda$, the following hold:

1. $Y(S) \geq \gamma|S|$ and $X(S) \geq f|S|$;
2. for any block B mined by an honest party during S , $Z_B(S) \leq \beta t'_A \cdot |S|$ and $Z_B(S) + X(S) \leq \beta(t'_A + nt'_H) \cdot |S|$; and
3. no insertions, no copies and no predictions occurred.

Theorem 23. *An execution is not typical with probability at most*

$$\varepsilon = \sum_{|S| \geq \eta\lambda} (\epsilon(\beta, t'_A |S|, n|S|) + \epsilon_s(\gamma, |S|) + \epsilon_s(f, |S|) + \epsilon(\beta, |S| \cdot (t'_A + nt'_H), 0)) \leq \text{negl}(\lambda).$$

Proof. By applying the union bound on Lemmas 19 and 21 and Corollary 20 it holds that the probability that there exists any set of consecutive rounds S , where $|S| > \eta\lambda$, and either $Y(S) < \gamma|S|$ or $X(S) < f|S|$ or for some honest block mined during S , $Z_B(S) > \beta t'_A |S|$ or $Z_B(S) + X(S) > \beta(t'_A + nt'_H) \cdot |S|$, is at most $\sum_{|S| \geq \eta\lambda} (\epsilon(\beta, t'_A |S|, n|S|) + \epsilon_s(\gamma, |S|) + \epsilon_s(\zeta, |S|) + \epsilon(\beta, |S| \cdot (t'_A + nt'_H), 0))$.

$$\begin{aligned} \varepsilon &\leq \sum_{|S| \geq \eta\lambda} (\epsilon(\beta, t'_A |S|, n|S|) + \epsilon_s(\gamma, |S|) + \epsilon_s(f, |S|) + \epsilon(\beta, |S| \cdot (t'_A + nt'_H), 0)) \\ &\leq \sum_{|S| \geq \eta\lambda} [\text{negl}(t'_A |S|) + \text{negl}((t'_A + nt'_H) |S|) \\ &\quad + \left(\frac{(1 - \sigma)^2 \log \lambda}{4|S|\sigma^2\gamma^2} \right)^{\lceil \log \lambda/2 \rceil} + \left(\frac{(1 - \sigma)^2 \log \lambda}{4|S|\sigma^2 f^2} \right)^{\lceil \log \lambda/2 \rceil}] \\ &\leq \sum_{|S| \geq \eta\lambda} [\text{negl}(\lambda) + \text{negl}(\lambda) + \text{negl}(\lambda) + \text{negl}(\lambda)] \leq \text{negl}(\lambda) \end{aligned}$$

By taking the negation of this event, properties 1 and 2 of a typical execution follow with the desired probability. Property 3 follows from Lemma 18. \square

Our proof strategy is going to be based on the fact that the rate of uniquely successful rounds exceeds the rate at which the adversary produces blocks. In previous works the main security assumption was that the total running time of honest parties per round exceeds that of the adversary. Interestingly, in our approach the running time of the adversary and the running time of honest parties do not have the same value, i.e., the adversary may use a superior proving algorithm. To take this into account we introduce the following Honest Majority Assumption, which also depends

on the security parameters of the POW scheme used. Note that γ , as defined in Lemma 21, depends on the parameters of the efficiency and H-TCMA properties.

Honest Majority Assumption. It holds that $\gamma \geq (1 + \delta)\beta \cdot t'_A$, for some $\delta \in (0, 1)$.

t'_A is directly related to the computational power of the adversary. As protocol designers, our goal is to be able to tolerate t'_A 's that are as big as possible. The Honest Majority Assumption implies that to achieve that we need to maximize the ratio of γ over $\beta t'_A$, i.e., the ratio of uniquely successful rounds over the number of blocks the adversary mines. The next lemma exactly highlights this implication.

Lemma 24. *For any set S of at least $\eta\lambda$ rounds in a η -typical execution and for any block B mined by an honest party during S , it holds that $Z_B(S) \leq (1 - \frac{\delta}{2})Y(S)$.*

Proof.

$$Z_B(S) \leq \beta t'_A \cdot |S| \leq \frac{1}{1 + \delta} \gamma |S| < (1 - \frac{\delta}{2})Y(S)$$

The first and the last inequality follow from a typical execution. The one in the middle from the Honest Majority Assumption. \square

We can now use the machinery built in [22] to prove the common prefix, chain quality and chain growth properties, with only minor changes. Refer to Appendix B for the details. Using these properties we prove that the modified Bitcoin backbone protocol implements a robust transaction ledger.

Theorem 25. *Assuming the existence of a common random string, a hash function that is collision and everywhere preimage resistant, a POW scheme that satisfies $(n \log \lambda)$ -wise independence and $(\beta, \text{negl}(\lambda))$ -H-TCMA, and model parameters $\{n, t, h, t_H, t_A, \theta\}$ that comply with the Honest Majority Assumption, protocol $\Pi_{\text{PL}}^{\text{POW}}$ implements a robust public transaction ledger with parameters $u = 4\eta\lambda$ and $k = \eta\lambda(\gamma + \beta n t'_H)$ and with overwhelming probability in λ .*

4.6 Comparison with the random-oracle analysis

In this subsection, and as a sanity check, we compare the results we got from our black-box analysis to those from the random oracle analysis of [22]. To do this, we first we show that the POW scheme used in the Bitcoin protocol (call it **BPOW**) is secure in the random oracle model according to our definitions. Then, based on the security parameters we obtain for **BPOW**, we show that the security guarantees we get from our analysis of the Bitcoin backbone protocol are similar to those proved in [22, 32].

In a nutshell, Bitcoin's **Prove** algorithm tries to find a block header with a small hash. The main components of the header are as follows: (i) the hash of the header of the previous block, (ii) the hash of the root of the Merkle tree of the transactions that are going to be added to Bitcoin's ledger, including the *randomly* created coinbase transaction, and (iii) a counter. The algorithm works by first fetching available transactions from the network, then computing a random public key that will be used for the coinbase transaction, and then iteratively incrementing the counter and calculating the hash of the header of the block until a small value is found. Casting this in our terms, the key is the hash of the previous block, which by itself depends on Bitcoin's genesis block, and the transactions received by the network as well as the coinbase transaction constitute the message. It is important to note that it is not possible to consider the key to be the coinbase transaction, as there is no guarantee it has any entropy when produced by an adversarial prover. To abstract

the randomization of the proving procedure, which in the actual implementation is captured by the coinbase transaction, we hash msg together with a randomly generated string. This should be part of the witness in our POW syntax since it is produced by the proving process and is necessary for verification. Similarly, the counter is also part of the witness produced by the proving process. BPOW, a simplified version of the scheme described above with the transaction semantics omitted for simplicity, is presented in Figure 4.

Remark 2. In the Bitcoin implementation, the hash of the root of the Merkle tree of the transactions is not “salted.” This means that if we consider the adversary to be non-uniform, she could get collisions for free in her advice string and use them to compute two POWs at the cost of one. This would be problematic for our H-TCMA security game. Thus, in order to strengthen the security of the scheme, we choose to also include the key in the hash of the message.

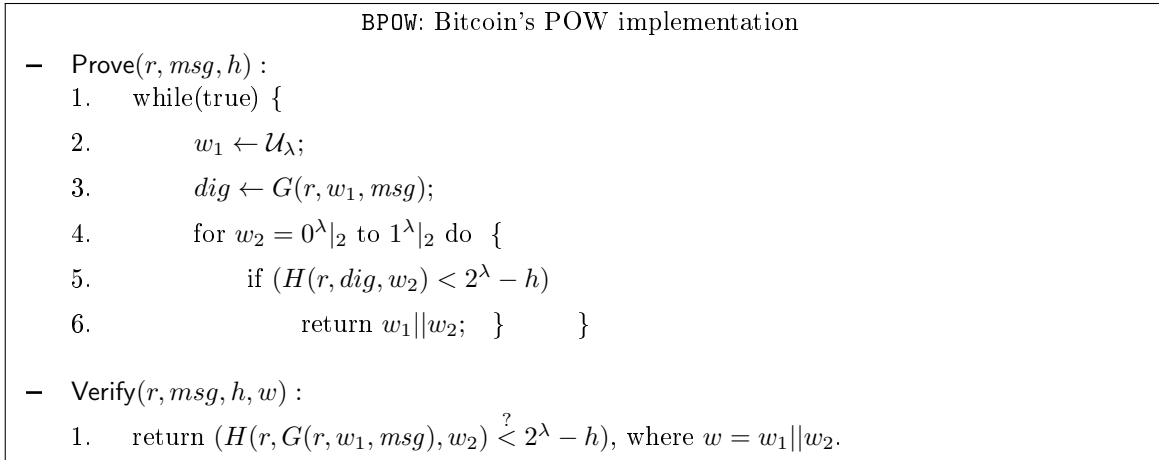


Figure 4: Bitcoin’s POW scheme. H and G are hash functions instantiated with SHA-256.

We will assume that both H and G are idealized hash functions, i.e., our analysis is in the random oracle.

Theorem 26. For any $\sigma \in (0, 1)$, and for any unpredictability-preserving tampering function class \mathcal{F} , BPOW is

- complete,
- $O(\lambda)$ -verifiable,
- $((1 + \sigma)(1 - \frac{h}{2^\lambda}), \epsilon(\beta(h), t, q_{\mathcal{P}}))$ -H-TCMA secure;
- $1 - (\frac{h}{2^\lambda})^t$ -efficient;
- run-time m -wise independent, for any $m > 0$ and

$$\epsilon(\beta(h), t, q_{\mathcal{P}}) = \begin{cases} 1 - (\frac{h}{2^\lambda})^t & , \text{ if } (1 + \sigma)(1 - \frac{h}{2^\lambda})t \leq 1 \\ \exp(-\beta(h, t)(1 + \log(q_{\mathcal{P}}))\sigma^2 t / 6) & , \text{ otherwise.} \end{cases}$$

Proof. Let $p_h = 1 - \frac{h}{2^\lambda}$ be the probability that a query to the random oracle returns a value less than $2^\lambda - h$, and let $q_{\mathcal{H}}$ be the number of queries the adversary makes to the RO. We consider each property in turn.

Completeness. The completeness property is trivially satisfied by the scheme.

H-TCMA security. Let $k = \lceil \beta(h)t \rceil$. First, we show that for any adversary \mathcal{A} there exists an adversary \mathcal{A}' that succeeds in winning $\text{Exp}_{\mathcal{A}', \mathcal{F}}^{\text{H-CMA}}$ (Figure 1) with almost the same time complexity

and probability that \mathcal{A} wins, without using the proving oracle \mathcal{P} . Hence, we will not have to take \mathcal{P} into account in our analysis. \mathcal{A}' is going to run \mathcal{A} internally, and all calls made by \mathcal{A} to \mathcal{P} are going to be simulated, i.e., assuming \mathcal{A} queries \mathcal{P} with values (f, r, m) , \mathcal{A}' will respond with some number t sampled from the time distribution of \mathcal{P} (t can be efficiently sampled from a geometric distribution, since queries are i.i.d Bernoulli trials) and some random witness $w = (w_1, w_2)$, where $t > w_2$. \mathcal{A}' is also going to store this query in some efficient structure that allows for search in logarithmic time. Any calls made by the adversary afterwards to the RO that are related to (f, r, m) will be answered accordingly; if \mathcal{A}_2 queries the RO with some string $H(f(r), G(f(r), w_1, msg), w'_2)$, where $w' = w_2$, then \mathcal{A}' will respond with the same value he responded on the initial query to \mathcal{P} , otherwise if $w'_2 < w_2$, he responds by $2^\lambda - h + (y \bmod h)$, where y is the output of the real RO in this query. Hence, the view of \mathcal{A} will be the same in both executions and he will output k valid POWs with respect to the simulated view with the same probability that he wins in the real experiment. It could be the case that the output of \mathcal{A} contains a POW related to the queries asked to (the simulated) oracle \mathcal{P} , and thus it does not correspond to a winning output for \mathcal{A}' ; i.e., there exists a query $\langle f, r, m, (w_1, w_2) \rangle$ and a POW on the output $\langle f', r', m', (w'_1, w'_2) \rangle$ such that $G(f(r), w_1, m_1) = G(f'(r'), w'_1, m'_1)$ and $m \neq m'$ or $w_1 \neq w'_1$. This implies that the adversary has found a collision in G , which only happens with negligible probability in λ . Hence, \mathcal{A}' will provide the same output as \mathcal{A} and with the same probability (minus some negligible term in λ) it will win the experiment. Moreover, the overhead incurred to \mathcal{A}' running time will be only logarithmic on qp i.e. \mathcal{A}' can simulate the t steps taken by \mathcal{A} after his first call to \mathcal{R} in time $t \cdot (1 + \log(qp))$; he has to maintain a heap of the queries made to \mathcal{P} and search it each time the RO is queried.

Let A be the event where \mathcal{A} asks t queries the RO after querying \mathcal{R} , and receives more than $\lceil k \rceil$ responses that have value less than $2^\lambda - h$. Let random variable X be equal to the number of these responses that are less than $2^\lambda - h$. Since the queries are i.i.d. Bernoulli random variables with probability of success p_h , we can use the Chernoff bound to bound the probability of A . By setting k equal to $(1 + \delta)p_h t$ it follows that, for any $\delta \in (0, 1)$:

$$\Pr[A] = \Pr[X > k] = \Pr[X > (1 + \delta)p_h t] = \Pr[X > (1 + \delta)\mathbb{E}[X]] \leq e^{-\frac{\mathbb{E}[X]\delta^2}{3}} \leq e^{-\frac{k\delta^2}{6}}.$$

For the special case where $k \leq 1$ we can derive a stricter bound for the probability of A as follows:

$$\Pr[A] = \Pr[X \geq 1] = 1 - (1 - p_h)^t$$

Let B be the event where \mathcal{A} can find an $f \in \mathcal{F}$ where for some query $H(y, x, w)$ he made to the random oracle before querying \mathcal{R} for the first time and some $r \in \text{Ans}_{\mathcal{R}}$, it holds that $f(r) = y$. Notice that he does not know the responses of \mathcal{R} when he picks y . Assume that B happens with non-negligible probability. Then we can use \mathcal{A} to break the unpredictability-preserving property of \mathcal{F} . \mathcal{A} will randomly output the y part of one of the queries he has made before querying \mathcal{R} , and then he will try to find a function f matching this query. The probability of this event is non-negligible which is a contradiction. Thus, $\Pr[B]$ must be negligible in λ .

Let C be the event where the adversary outputs two POWs that correspond to the same query in RO. This implies that the adversary can find a collision on G . In time $L = t + t_{\text{pre}}$ polynomial in λ , the probability that \mathcal{A} finds a collision is $\binom{L}{2} 2^{-\lambda+1} = e^{-\Omega(\lambda)} = \text{negl}(\lambda)$. Finally, let D be the event where the adversary outputs a valid POW that he has not asked the RO. This event occurs with probability p_h which is $\text{negl}(\lambda)$.

If A, B, C, D do not occur, it is implied that \mathcal{A}' will lose in the experiment. Thus, by an

application of the union bound we get:

$$\begin{aligned}
\Pr[\text{Exp}_{\mathcal{A}, \mathcal{F}}^{\text{H-TCMA}}(1^\lambda, h, k) = 1] &\leq 1 - \Pr[\text{Exp}_{\mathcal{A}, \mathcal{F}}^{\text{H-TCMA}}(1^\lambda, h, k) = 0] \\
&\leq 1 - \Pr[\neg(A \vee B \vee C \vee D)] \\
&\leq 1 - (1 - \Pr[A \vee B \vee C \vee D]) \\
&\leq \Pr[A \vee B \vee C \vee D] \\
&\leq \begin{cases} 1 - (1 - p_h)^t + \text{negl}(\lambda), & \text{when } (1 - \frac{h}{2^\lambda})t \leq 1 \\ e^{-\frac{k(1+\log(qp))\delta^2}{6}} + \text{negl}(\lambda), & \text{otherwise} \end{cases}
\end{aligned}$$

Verifiability. Assuming H and G take constant time, verification takes time $c_{\text{ver}}\lambda$, for some small constant c_{ver} which can be easily computed by careful inspection of the verification protocol.

Efficiency. For any $t \in \mathbb{N}$, $r \in R$, $msg \in M$ and $h \in \mathbb{N}$ it holds that:

$$\Pr[\text{Steps}_{\text{Prove}}(r, msg, h) < t] = 1 - (1 - p_h)^t$$

Independence. Let $\{Y_i\}_{i \in I}$ be the same as $\{X_i\}_{i \in I}$ with the only difference that the random oracle is replaced with a random function, i.e., every time **Prove** is called and the oracle is queried it generates a random output. Obviously the random variables in $\{Y_i\}_{i \in I}$ are mutually independent, since their output only depends on their own local coins.

Regarding the second property, let E be the event that all w_1 sampled are different among all the invocations of **Prove**, and that no collisions occurs in G . Note, that in any polynomially bounded execution this event happens with overwhelming probability in λ . Moreover, in any execution that E occurs, $\{X_i\}_{i \in I}$ and $\{Y_i\}_{i \in I}$ are equal. Therefore, if $p(\cdot)$ is a polynomial that upper bounds the duration of the execution, it holds that for any $z \in [p(\lambda)]^{|I|}$

$$\begin{aligned}
&\Pr[\{X_i\}_{i \in I} = z] - \Pr[\{Y_i\}_{i \in I} = z] = \\
&= \Pr[\{X_i\}_{i \in I} = z | E] \Pr[E] + \Pr[\{X_i\}_{i \in I} = z | \neg E] \Pr[\neg E] \\
&\quad - \Pr[\{Y_i\}_{i \in I} = z | E] \Pr[E] - \Pr[\{Y_i\}_{i \in I} = z | \neg E] \Pr[\neg E] \\
&\leq (\Pr[\{X_i\}_{i \in I} = z | \neg E] - \Pr[\{Y_i\}_{i \in I} = z | \neg E]) \Pr[\neg E]
\end{aligned}$$

Hence, it follows that the two distributions are $\text{negl}(\lambda)$ -close:

$$\begin{aligned}
2\Delta[\{X_i\}_{i \in I}, \{Y_i\}_{i \in I}] &= \sum_z |\Pr[\{X_i\}_{i \in I} = z] - \Pr[\{Y_i\}_{i \in I} = z]| \\
&\leq \sum_z |(\Pr[\{X_i\}_{i \in I} = z | \neg E] - \Pr[\{Y_i\}_{i \in I} = z | \neg E]) \Pr[\neg E]| \\
&\leq \Pr[\neg E] \sum_z |(\Pr[\{X_i\}_{i \in I} = z | \neg E] - \Pr[\{Y_i\}_{i \in I} = z | \neg E])| \\
&\leq \text{negl}(\lambda) (\sum_z \Pr[\{X_i\}_{i \in I} = z | \neg E] + \sum_z \Pr[\{Y_i\}_{i \in I} = z | \neg E]) \leq \text{negl}(\lambda)
\end{aligned}$$

The last inequality follows from the fact that each of the sums should be less or equal to 1, as the events described are disjoint and their union covers the entire sample space. \square

Since parameter ϵ of the H-TCMA property of **BPOW** is negligible in t and the scheme is m -wise independent for any $m > 0$, we can use Theorem 26 and obtain meaningful bounds for the γ, f

quantities introduced in the previous subsection.⁴ These quantities are important since γ determines how powerful the adversary our system can handle can be, and f is related to how fast blocks are produced. Both of them also appear in [22] and are (expressed in the language of our model) equal to

$$\gamma_{\text{RO}} = (1 - \sigma)(n - t) \left(1 - \frac{h}{2^\lambda}\right) t'_{\mathcal{H}} \left(\frac{h}{2^\lambda}\right)^{((n-t)-1)t'_{\mathcal{H}}}, f_{\text{RO}} = (1 - \sigma) \left(1 - \left(\frac{h}{2^\lambda}\right)^{(n-t)t'_{\mathcal{H}}}\right)$$

respectively. The next lemma shows that our bounds are *greater or equal* to the ones calculated there. Moreover, our analysis also reveals other factors that affect security such as the verification time, the number of messages the adversary sends to the honest parties, as well as the cost of reducing the security to that of the POW primitive. The proof appears in Appendix ?? . Let $\Pi_{\text{PL}}^{\text{BPOW}}$ denote the Bitcoin protocol using BPOW. Then:

Lemma 27. *For $\Pi_{\text{PL}}^{\text{BPOW}}$, it holds that $\gamma \geq \gamma_{\text{RO}}$ and $f \geq f_{\text{RO}}$.*

Proof. For γ we have:

$$\begin{aligned} \gamma &\geq (1 - \delta')(n - t) \left(1 - \left(\frac{h}{2^\lambda}\right)^{t'_{\mathcal{H}}}\right) \left(1 - \left(1 - \left(\frac{h}{2^\lambda}\right)^{t'_{\mathcal{H}}}\right)\right)^{n-t-1} \\ &= (1 - \delta') \cdot (n - t) \left(1 - \left(\frac{h}{2^\lambda}\right)^{t'_{\mathcal{H}}}\right) \left(\frac{h}{2^\lambda}\right)^{(n-t-1)t'_{\mathcal{H}}} \\ &\geq (1 - \delta')(n - t) \left(1 - \frac{h}{2^\lambda}\right) t'_{\mathcal{H}} \left(\frac{h}{2^\lambda}\right)^{((n-t)-1)t'_{\mathcal{H}}} \end{aligned}$$

The first inequality follows from the honest majority assumption. Also, w.l.o.g assume that $t'_{\mathcal{A}} \geq t'_{\mathcal{H}}$, i.e. the adversary has at least the computational power of one honest party. Then:

$$(1 + \delta') \left(1 - \frac{h}{2^\lambda}\right) t'_H = \beta t'_{\mathcal{H}} \leq \beta t'_{\mathcal{A}} < \gamma \leq 1$$

On the other hand, for f we have that:

$$f = (1 - \delta) \left(1 - \left(1 - \left(1 - \left(\frac{h}{2^\lambda}\right)^{t'_{\mathcal{H}}}\right)\right)^{n-t}\right) = (1 - \delta) \left(1 - \left(\frac{h}{2^\lambda}\right)^{(n-t)t'_{\mathcal{H}}}\right)$$

□

Hence, $\Pi_{\text{PL}}^{\text{BPOW}}$ implements a robust transaction ledger with overwhelming probability in λ and with bounds comparable to those in [22].

5 Consensus from Proofs of Work without Random Oracles

In this section we show how to achieve consensus (a.k.a. Byzantine agreement [33, 30]) under exactly the same assumptions used for proving the security of the Bitcoin backbone protocol in Section 4. In [22], consensus is achieved under the Honest Majority Assumption by using the POW construction in a *non-black-box* way, through a mining technique called “2-for-1 POWs.” In more

⁴Earlier, in order to make our presentation easier to follow, we chose to allow the BPOW Prove routine to run until it finds a solution. However, in the Bitcoin backbone protocol it should be stopped after running a specific number of steps. We can implement this modification using a step counter, without affecting any of the security properties of the scheme.

detail, the technique shows how miners can compute POWs for two different POW schemes at the cost of one, while at the same time ensuring that their resources cannot be used in favor of one of the two schemes. This cannot be directly translated to the POW security properties we have defined, as they all apply to the execution of a *single Prove* subroutine. Hence, we would have to introduce “2-for-1 POWs” as an extra property. In this section we show how blockchain-based consensus can be achieved by only using the security properties we have so far, directly, and without the extra non-black-box machinery used in [22]. The protocol is based on the Bitcoin backbone protocol, and formally specified by providing adequate definitions for the I, R, V functions.

We first give a definition of the consensus problem. There are n parties, $t < n$ of which might be corrupted, taking an initial input $x \in V$ (without loss of generality, we can assume $V = \{0, 1\}$).

Definition 28. A protocol Π solves the consensus problem provided it satisfies the following two properties:

- **Agreement.** There is a round after which all honest parties output the same value.
- **Validity.** If all the honest parties have the same input, then they all output this value.

Next, we define some notation and terminology that will be used in the remainder of the section. We will use the terms “input” and “vote” interchangeably, referring to the parties’ input in the consensus problem. We will use $header(\langle s, x, w \rangle)$ to denote the “compressed” version of block $\langle s, x, w \rangle$, equal to $\langle s, G(x) || vote, w \rangle$. Note that, as defined, the header of any block is of a fixed size. We also extend the definition of our hash function H on headers of blocks. The hash of the header of some block B will be equal to the hash of B , i.e., $H(header(B)) = H(B)$ (note that the header of B provides all the information need to calculate the hash of B).

We now present a high-level description of the protocol. The basic idea is that during block mining, parties are going to include in their blocks not only their own votes, but also headers of other blocks that they have seen and that are not part of their chain. Then, after a predetermined number of rounds, the parties will count the votes “referenced” in a prefix of their chain, including the votes found in the headers of the blocks referenced. In this way, they can take advantage of the robust transaction ledger built in Section 4. The common prefix property implies that the honest parties will all agree on *which* votes should be counted, while the chain quality and chain growth properties guarantee that the majority of the counted votes come from honest parties. A main technical challenge is to be able to add the block references without making the honest parties’ chains grow too big, and at the same time to ensure that the number of honest votes exceeds the adversarial ones.

To overcome this challenge, we are going to change POW generation (Algorithm 1, line 8) and POW verification (Algorithm 2, line 8) so that they are run on the header of the block, i.e., $Prove(s, G(x) || vote, h)$ and $Verify(s, G(x) || vote, h, w)$, respectively. This way we are able to verify the validity of a block as a POW and determine the block’s vote by only knowing its header. This is exactly the properties we need for the consensus application.

Moreover, we should be able to tell whether the referenced blocks are “fresh”; that is, the adversary should not be able to reference blocks that it has precomputed and are not related to the genesis block. We achieve this by requiring blockchains to have a special structure in order to be considered valid by the content validation predicate $V(\cdot)$. A chain will be *valid* when the referenced blocks on every prefix of the chain form a tree that has the genesis block at its root. In order to check this efficiently, we require that the reference list of each block is ordered, so that each entry extends some block header found in previous entries of the same or parent blocks. The detailed description of the content validation predicate is given in Algorithm 5.

In line 2 we use an AVL tree [1] to efficiently check for membership in the hash tree. (Any other data structure supporting efficient updates and search would also work.) In line 5 the referenced

Algorithm 5 The *content validation predicate*. The input is the contents of the blocks of some chain.

```

1: function  $V(\langle x_1, \dots, x_m \rangle)$ 
2:    $D \leftarrow \text{new AVL}()$  ▷ Create a new (empty) AVL tree.
3:    $D.add(H(\text{Gen}))$  ▷ Add the hash of the genesis block on the tree.
4:   for  $i = 1, \dots, m$  do
5:      $queue \leftarrow \text{references}(x_i)$  ▷ Add all block references in a queue.
6:     while  $queue \neq \emptyset$  do
7:        $\langle s, G(x) || \text{vote}, w \rangle \leftarrow queue.top()$ 
8:       if  $((D.exists(s)) \wedge \text{Verify}(s, G(x) || \text{vote}, h, w))$  then
9:          $D.add(H(\langle s, G(x) || \text{vote}, w \rangle))$  ▷ Add new entry on the tree.
10:         $queue.pop()$ 
11:      else
12:        return False ▷ If not, the chain is invalid.
13:      end if
14:    end while
15:  end for
16:  return True
17: end function

```

Content validation predicate $V(\cdot)$	The validation function $V(\cdot)$ checks that the block headers included in the chain create a a block tree of valid POW's that starts from the genesis block.
Chain reading function $R(\cdot)$ (parametrized by M)	$R(\cdot)$ outputs the majority of the votes found in the block headers of the first M blocks of the selected chain.
Input contribution function $I(\cdot)$	The input function $I(\cdot)$ maintains state of which blocks have been received and outputs the input value x that contains (i) the headers of all valid blocks that extend the genesis block and are not mentioned in the chain that the party is currently extending, and (ii) the party's input (i.e., 0 or 1).

Figure 5: *Consensus from POWs: The $\Pi_{\text{BA}}^{\text{POW}}$ protocol.*

blocks are extracted and pushed into a queue. We note that during this process it is also checked that the contents of the block have a correct format, i.e., a vote field and list of block headers, and that the order in which the headers are organized in the block is the same as the order in which they are popped from the queue.

Finally, after L rounds (to be determined), the algorithm is going to output the majority of the votes found in a prefix of the selected chain, of a predetermined length M . We call the resulting protocol $\Pi_{\text{BA}}^{\text{POW}}$ (“BA” for Byzantine agreement, to be consistent with [22]). A description of the consensus protocol (specifically, the V, R, I functions) is presented in Figure 5, and an example in Figure 6.

Next, we formally prove that our protocol achieves consensus. Note also that all parties will terminate the protocol simultaneously.

Theorem 29. *Assuming the existence of a common random string, a collision- and everywhere-preimage-resistant hash function, a POW scheme satisfying run-time $(n \log \lambda)$ -wise independence*




Figure 6: The data structure maintained by $\Pi_{\text{BA}}^{\text{POW}}$. Block B references block C , block F references blocks D and G , and block D references some invalid blocks. This is not a problem, since (i) any chain that contains D will not be selected by any party, and (ii) D 's vote is correctly counted since D is a descendant of Gen .

and $(\beta, \text{negl}(\lambda))$ - H -TCMA, and model parameters $\{n, t, h, t_{\mathcal{H}}, t_{\mathcal{A}}, \theta\}$ that comply with the Honest Majority Assumption, protocol $\Pi_{\text{BA}}^{\text{POW}}$ solves consensus in $O(\lambda)$ rounds with overwhelming probability.

Proof. We are going to show that protocol $\Pi_{\text{BA}}^{\text{POW}}$, parameterized with $k \geq \eta\lambda(\gamma + \beta nt'_{\mathcal{H}})$, $L \geq 2k \frac{1}{1 - (1 - \frac{\delta}{2}) \frac{\gamma^2}{f}}$ and $M = (\gamma \cdot L - k)$ solves consensus with overwhelming probability in λ .

First, note that for any chain \mathcal{C} where the head of the chain is mined by an honest party it holds that $V(\mathcal{C})$ is equal to True. This follows from the modifications we did on the protocol and the fact that the adversary will only include headers on his list of references that satisfy the requirements of predicate $V(\cdot)$. Moreover, since the output of the predicate only depends on the chain that is being validated, if one honest party accepts a chain as valid, all honest parties accept.

Next, we prove Agreement. Assume that an execution is typical. Due to the chain growth property, after L rounds the chain of honest parties will have length at least $\gamma \cdot L$ blocks, and due to the common prefix property they will all agree on the first $\gamma \cdot L - k$ blocks. Hence, all honest parties will decide on their output value based on the “votes” mentioned in each block header that is referenced in these blocks, and they will all agree on the same value.

Regarding Validity, we are going to show that the majority of the counted “votes,” i.e., from blocks and block headers found in blocks $B_1, \dots, B_{\gamma \cdot L - k}$ of the selected chain, have been mined by honest parties. By the chain quality property, at least one block from $B_{\gamma \cdot L - 2k}, \dots, B_{\gamma \cdot L - k}$ is honest. Moreover, by chain growth the blocks preceding this honest block must have been computed up to round $\frac{L - 2k}{\gamma}$. Hence, all honest blocks mined up to this round will contribute votes to the agreement process. On the other hand, the last block in the sequence of $\gamma \cdot L - k$ blocks was computed at most at round L , due to chain growth and common prefix. Hence, it remains to show that for $S_1 = \{1, \dots, \frac{L - 2k}{\gamma}\}$ and $S_2 = \{1, \dots, L\}$ it holds that $Z(S_2) < X(S_1)$:

$$Z(S_2) \leq \beta t'_{\mathcal{A}} L < \frac{1}{1 + \delta} \gamma L \leq (1 - \frac{\delta}{2}) \gamma L \leq f \frac{L - 2k}{\gamma} \leq X(S_1).$$

The first and the last inequalities hold due to the fact that the execution is typical. The fourth inequality follows from our assumption about L . The theorem follows since the majority of the referenced blocks in the chain agreed upon, have been mined by honest parties.

Concluding, notice that the total size of any chain is bounded by the total number of blocks mined, since each block's header is mentioned at most once in a single chain. Hence, in r rounds of a typical execution a chain has size at most $O(r \cdot \beta(t'_{\mathcal{A}} + nt'_{\mathcal{H}}) \cdot \lambda)$ bits. \square

References

- [1] G. Adelson-Velskii and E. Landis. An algorithm for the organization of information.
- [2] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple constructions of almost k-wise independent random variables. In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume II*, pages 544–553. IEEE Computer Society, 1990.
- [3] J. Alwen and B. Tackmann. Moderately hard functions: Definition, instantiations, and applications. In Y. Kalai and L. Reyzin, editors, *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part I*, volume 10677 of *Lecture Notes in Computer Science*, pages 493–526. Springer, 2017.
- [4] M. Andrychowicz and S. Dziembowski. Distributed cryptography based on the proofs of work. Cryptology ePrint Archive, Report 2014/796, 2014. <http://eprint.iacr.org/>.
- [5] A. Back. Hashcash-amortizable publicly auditable cost functions. *Early draft of paper*, 2000.
- [6] A. Back. Hashcash—a denial of service counter-measure, 2002.
- [7] C. Badertscher, U. Maurer, D. Tschudi, and V. Zikas. Bitcoin as a transaction ledger: A composable treatment. In J. Katz and H. Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 324–356. Springer, 2017.
- [8] F. Baldimtsi, A. Kiayias, T. Zacharias, and B. Zhang. Indistinguishable proofs of work or knowledge. Cryptology ePrint Archive, Report 2015/1230, 2015. <http://eprint.iacr.org/2015/1230>.
- [9] M. Ball, A. Rosen, M. Sabin, and P. N. Vasudevan. Proofs of useful work. *IACR Cryptology ePrint Archive*, 2017:203, 2017.
- [10] M. Bellare, D. J. Bernstein, and S. Tessaro. Hash-function based prfs: Amac and its multi-user security. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 566–595. Springer, 2016.
- [11] M. Bellare, A. Desai, E. Jorjani, and P. Rogaway. A concrete security treatment of symmetric encryption. In *Foundations of Computer Science, 1997. Proceedings., 38th Annual Symposium on*, pages 394–403. IEEE, 1997.
- [12] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *CCS '93, Proceedings of the 1st ACM Conference on Computer and Communications Security, Fairfax, Virginia, USA, November 3-5, 1993.*, pages 62–73, 1993.
- [13] M. Bellare and P. Rogaway. The exact security of digital signatures-how to sign with rsa and rabin. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 399–416. Springer, 1996.
- [14] D. J. Bernstein and T. Lange. Non-uniform cracks in the concrete: the power of free precomputation. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 321–340. Springer, 2013.
- [15] N. Bitansky, S. Goldwasser, A. Jain, O. Paneth, V. Vaikuntanathan, and B. Waters. Time-lock puzzles from randomized encodings. Cryptology ePrint Archive, Report 2015/514, 2015. <http://eprint.iacr.org/2015/514>.
- [16] R. Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
- [17] I. B. Damgård. A design principle for hash functions. In *Conference on the Theory and Application of Cryptology*, pages 416–427. Springer, 1989.

- [18] J. R. Douceur. The sybil attack. In P. Druschel, M. F. Kaashoek, and A. I. T. Rowstron, editors, *Peer-to-Peer Systems, First International Workshop, IPTPS 2002, Cambridge, MA, USA, March 7-8, 2002, Revised Papers*, volume 2429 of *Lecture Notes in Computer Science*, pages 251–260. Springer, 2002.
- [19] C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In E. F. Brickell, editor, *CRYPTO*, volume 740 of *Lecture Notes in Computer Science*, pages 139–147. Springer, 1992.
- [20] C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In *Proceedings of the 12th Annual International Cryptology Conference on Advances in Cryptology, CRYPTO '92*, pages 139–147, London, UK, UK, 1993. Springer-Verlag.
- [21] I. Eyal and E. G. Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography*, 2014.
- [22] J. A. Garay, A. Kiayias, and N. Leonardos. 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*, pages 281–310, 2015.
- [23] J. A. Garay, A. Kiayias, and N. Leonardos. The bitcoin backbone protocol with chains of variable difficulty. *IACR Cryptology ePrint Archive*, 2016:1048, 2016.
- [24] J. A. Garay, A. Kiayias, N. Leonardos, and G. Panagiotakos. Bootstrapping the blockchain directly, with applications to consensus and fast pki setup. *Cryptology ePrint Archive*, Report 2016/991, 2016. <http://eprint.iacr.org/2016/991>. To appear in *PKC '18*.
- [25] J. A. Garay, P. MacKenzie, M. Prabhakaran, and K. Yang. Resource fairness and composability of cryptographic protocols. *Journal of cryptology*, 24(4):615–658, 2011.
- [26] M. Jakobsson and A. Juels. Proofs of work and bread pudding protocols. In *Proceedings of the IFIP TC6/TC11 Joint Working Conference on Secure Information Networks: Communications and Multimedia Security, CMS '99*, pages 258–272, Deventer, The Netherlands, The Netherlands, 1999. Kluwer, B.V.
- [27] A. Juels and J. G. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *NDSS*. The Internet Society, 1999.
- [28] J. Katz, A. Miller, and E. Shi. Pseudonymous secure computation from time-lock puzzles. *IACR Cryptology ePrint Archive*, 2014:857, 2014.
- [29] A. Kiayias and G. Panagiotakos. Speed-security tradeoffs in blockchain protocols. Technical report, IACR: Cryptology ePrint Archive, 2015.
- [30] L. Lamport, R. E. Shostak, and M. C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- [31] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <http://bitcoin.org/bitcoin.pdf>, 2008.
- [32] R. Pass, L. Seeman, and A. Shelat. Analysis of the blockchain protocol in asynchronous networks. In J. Coron and J. B. Nielsen, editors, *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part II*, volume 10211 of *Lecture Notes in Computer Science*, pages 643–673, 2017.
- [33] M. C. Pease, R. E. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.
- [34] A. Poelstra. On stake and consensus (2015). URL <https://download.wpsoftware.net/bitcoin/pos.pdf>.
- [35] P. Rogaway and T. Shrimpton. Cryptographic hash-function basics: Definitions, implications, and separations for preimage resistance, second-preimage resistance, and collision resistance. In *International Workshop on Fast Software Encryption*, pages 371–388. Springer, 2004.
- [36] S. P. Vadhan et al. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012.

A Hash Functions and Unpredictability Preservation




Figure 7: Bitcoin's block hash structure.

In [10], Bellare *et al.* describe the notion of the basic cascade as follows:

$$H^*(K, X) : Y \leftarrow K; \text{ for } i = 1 \text{ to } n \text{ do } Y \leftarrow H(Y, X[i]); \text{ return } Y$$

where in practical applications H is an unkeyed compression function. For bitcoin in particular H would be a double invocation of SHA-256. We choose to analyze H, G as keyed hash function families that are collision and preimage resistant, from now on $\{H(K, \cdot)\}_{K \in \mathcal{K}}, \{G(K, \cdot)\}_{K \in \mathcal{K}}$. We can think that the key of these two functions is part of the CRS. Furthermore, we assume that I has some entropy (similar to the input entropy assumption in [22]).

Lemma 30. *Assuming that $\{H(K, \cdot)\}_{K \in \mathcal{K}}$ is collision and everywhere preimage resistant, $\{H^*(K, \cdot)\}_{K \in \mathcal{K}}$ is collision resistant and $\{H^*(\cdot, G(I(\cdot; K))||\cdot)\}_{K \in \mathcal{K}}$ is every-where preimage resistant.*

Proof. If the adversary is able to find a collision in H^* , it follows that he can find a collision in H , which happens with negligible probability in λ . Regarding the ePre property, note that the input entropy assumption states that the probability that two invocations of I return the same value for arbitrary inputs and honestly generated randomness, i.e. K , is negligible. Hence, an adversary that can break the ePre property for H^* , also breaks the same property for H by outputting the parameters of the last block in the chain. \square

Having a function family that is ePre resistant implies that we can construct a function family that is unpredictability preserving, since the two notions are closely related.

Lemma 31. *If an arbitrary function family $\{H(K, \cdot)\}_{K \in \mathcal{K}}$ is an ePre resistant hash function family, then $\mathcal{F} = \{H(\cdot, M)\}_{M \in \mathcal{M}}$ is an unpredictability preserving function family.*

Proof. Assume \mathcal{F} was not unpredictability preserving. Then there exists some adversary \mathcal{A} that breaks the property with non-negligible probability. We can use him exactly as he is to break the ePre property with non-negligible probability. The function that \mathcal{A} outputs in the second stage of the game corresponds to a message that breaks the ePre property. Hence, the lemma follows. \square

B Proof of the modified bitcoin backbone protocol

The notion of a typical execution is at the core of the proof of security of Bitcoin in [22]. Here, we describe the minor changes one has to do after proving the typical execution theorem with respect to the analysis of [22], in order to prove the security of the protocol in our model. We only give brief proof sketches of lemmas or theorems from [22] that are exactly the same for our own setting.

λ :	security parameter
n :	number of parties
$t_{\mathcal{H}}$:	number of steps per round per honest party
$t_{\mathcal{A}}$:	total number of adversarial steps per round
θ :	upper bound on the number of messages sent by the adversary per round
β :	upper bound on POW computation rate per step
γ :	lower bound on the rate of uniquely successful rounds
f :	lower bound on the rate of successful rounds
δ :	advantage from the honest majority assumption
σ :	quality of concentration of random variables in typical executions
k :	number of blocks for the common-prefix property
ℓ :	number of blocks for the chain-quality property
η :	parameter determining block to round translation
ε :	probability that an execution is typical

Table 2: The parameters in our analysis.

Lemma 32. (*Chain-Growth Lemma*). *Suppose that at round r an honest party has a chain of length ℓ . Then, by round $s \geq r$, every honest party has adopted a chain of length at least $\ell + \sum_{i=r}^{s-1} X_i$.*

Proof. The main idea of the proof of this lemma is that, after each successful round at least one honest party will have received a chain that is at least one block longer than the chain it had, and all parties pick only chains that are longer than the ones they had. \square

Theorem 33. (*Chain-Growth*). *In a typical execution the chain-growth property holds with parameters $\tau = \gamma$ and $s \geq \eta\lambda$.*

Proof. Let S be any set of at least s consecutive rounds. Then, since the execution is η typical: $X(S) \geq \gamma \cdot |S| \geq \tau \cdot |S|$. By Lemma 32, each honest player's chain will have grown by that amount of blocks at the end of this round interval. Hence, the chain growth property follows. \square

Lemma 34. *In a typical execution let B be some honest block. Any $k \geq \eta\lambda(\gamma + \beta nt'_{\mathcal{H}})$ consecutive blocks descending B in a chain in this execution have been computed in at least $\eta\lambda$ rounds, starting from the round that B was computed.*

Proof. Assume there is a set of rounds S' , such that $|S'| < \eta\lambda$ and more than $\eta\lambda(\gamma + \beta nt'_{\mathcal{H}})$ blocks that descend block B have been computed. Then, there is a set of rounds S , where $|S| \geq \eta\lambda$ such that $X(S) + Z_B(S) \geq |S|(\gamma + \beta nt'_{\mathcal{H}}) \geq |S|\beta(t'_{\mathcal{A}} + nt'_{\mathcal{H}})$. This contradicts the typicality of the execution, hence the lemma follows. \square

Lemma 35. (*Common-prefix Lemma*). *Assume a typical execution and consider two chains \mathcal{C}_1 and \mathcal{C}_2 such that $\text{len}(\mathcal{C}_2) \geq \text{len}(\mathcal{C}_1)$. If \mathcal{C}_1 is adopted by an honest party at round r , and \mathcal{C}_2 is either adopted by an honest party or diffused at round r , then $\mathcal{C}_1^{[k]} \leq \mathcal{C}_2$ and $\mathcal{C}_2^{[k]} \leq \mathcal{C}_1$, for $k \geq \eta\lambda(\gamma + \beta nt'_{\mathcal{H}})$.*

Proof. In Lemma 19, instead of bounding the number of blocks mined by the adversary in a set of rounds, we bound the number of blocks mined by the adversary with the additional condition that these blocks extend some specific honest block. If we also use the previous lemma, the proof is exactly the same as in [22]. Note, that all adversarial blocks in the matching between uniquely successful rounds and adversarial blocks are descendants of the last honest block in the common prefix of \mathcal{C}_1 and \mathcal{C}_2 . \square

Theorem 36. (*Common-prefix*). *In a typical execution the common-prefix property holds with parameter $k \geq \eta\lambda \cdot (\gamma + \beta nt'_{\mathcal{H}})$.*

Proof. The main idea of the proof is that if there exists a deep enough fork between two chains, then the previously proved lemma cannot hold. Hence, the theorem follows. \square

Theorem 37. (*Chain-Quality*). *In a typical execution the chain-quality property holds with parameter $\mu < 1 - \delta/2$ and $\ell \geq \eta\lambda(\gamma + \beta nt'_{\mathcal{H}})$.*

Proof. The main idea of the proof is the following: a large enough number of consecutive blocks will have been mined in a set rounds that satisfies the properties of Definition 22. Hence, the number of blocks that belong to the adversary will be upper bounded, and all other blocks will have been mined by honest parties. \square

Finally, the Persistence and Liveness properties follow from the three basic properties, albeit with different parameters than in [22].

Lemma 38. (*Persistence*). *It holds that Π_{PL} with $k = \eta\lambda(\gamma + \beta nt'_{\mathcal{H}})$ satisfies Persistence with probability at least $1 - \varepsilon$.*

Proof. The main idea is that if persistence is violated, then the common-prefix property will also be violated. Hence, if the execution is typical the lemma follows. \square

Lemma 39. (*Liveness*). *It holds that Π_{PL} with $u = 4\eta\lambda$ rounds and $k = \eta\lambda(\gamma + \beta nt'_{\mathcal{H}})$ satisfies Liveness with probability at least $1 - \varepsilon$.*

Proof. The main idea here is that after u rounds at least $2k$ successful rounds will have occurred. Thus, by the chain growth lemma the chain of each honest party will have grown by $2k$ blocks, and by the chain quality property at least one of these blocks that is deep enough in the chain is honest. \square