

# Non-interactive proofs of proof-of-work

Aggelos Kiayias<sup>1</sup>, Andrew Miller<sup>2</sup>, and Dionysis Zindros<sup>3,\*</sup>

<sup>1</sup> University of Edinburgh, IOHK

<sup>2</sup> University of Illinois

<sup>3</sup> National and Kapodistrian University of Athens

**Abstract.** Blockchain protocols such as bitcoin provide decentralized consensus mechanisms based on proof-of-work (PoW). In this work we introduce and instantiate a new primitive for blockchain protocols called Noninteractive-Proofs-of-Proof-of-Work (NIPoPoWs) which can be adapted into existing proof-of-work-based cryptocurrencies. Unlike a traditional blockchain client which must verify the entire linearly-growing chain of PoWs, clients based on NIPoPoWs require resources only logarithmic in the length of the blockchain. NIPoPoWs solve two important open questions for PoW based consensus protocols. Specifically the problem of constructing efficient transaction verification clients, sometimes called “simple payment verification” or SPV, and the problem of constructing efficient sidechain proofs. We provide a formal model for NIPoPoWs. We prove our construction is secure in the backbone model and show that the communication is succinct in size. We provide simulations and experimental data to measure concrete communication efficiency and security. Finally, we provide two ways that our NIPoPoWs can be adopted by existing blockchain protocols, first via a soft fork, and second via a new update mechanism that we term a “velvet fork” and enables to harness some of the performance benefits of NIPoPoWs even with a minority upgrade.

## 1 Introduction

Cryptocurrencies such as Bitcoin [1][2] and Ethereum [3] are peer-to-peer networks that maintain a globally consistent transaction ledger, using a consensus protocol based on proof-of-work (PoW) puzzles [4] [5]. Worker nodes called “miners” expend computational work in order to reach agreement on the state of the network. Clients on the network, such as mobile phone apps, must verify these proofs-of-work in order to determine the correct view of the network’s state, as necessary to transmit and receive payments.

In this work we introduce, analyze and instantiate a new primitive, Noninteractive Proofs of Proof-of-Work (NIPoPoWs), which can be adapted into existing cryptocurrencies to support more efficient clients. A traditional blockchain client must verify the entire chain of proofs-of-work, which grows linearly over time. On the contrary, clients based on NIPoPoWs require resources only logarithmic in the length of the blockchain.

### 1.1 Motivation

There are three trends that motivate the need for more efficient clients, specifically clients that can verify transactions without having to process the whole blockchain. First, while Bitcoin remains by far the largest and most widely used cryptocurrency, the ecosystem has become significantly more diverse. Today there are hundreds of competitors, commonly called “altcoins”, and Bitcoin’s dominance in the market is at an all time low [6].<sup>4</sup> We envision that our protocol will form the basis of an efficient multi-blockchain client, which could efficiently support payments using hundreds of separate cryptocurrencies.

---

\* Research supported by ERC project CODAMODA, project #259152

<sup>4</sup> There are numerous problems with using market capitalization as a measure of economic size, such as defining what it means to quantify the quantity “in circulation”; further quantifying this is outside the scope of our paper.

Second, cryptocurrencies are increasingly used as components of larger systems (the cryptocurrency system is typically called a “blockchain” when its currency feature is not salient). As one example, a recent system called Catena [7] uses the Bitcoin blockchain as an authenticated log; the browser embeds a Bitcoin client that uses the blockchain to validate HTTPS certificates. The bandwidth costs that are tolerable for a dedicated Bitcoin wallet are likely unacceptable in such embedded contexts. Our techniques can directly reduce the cost of Catena and similar systems.

Finally, there has been significant interest in the development of “cross-chain” applications, i.e. logical transactions that span multiple separate cryptocurrencies. Simple cross-chain transactions are feasible today: the most well-known is the atomic exchange [8] e.g. a trade of Bitcoins for Ether. However, more sophisticated applications could be enabled by our protocol, which would allow one cryptocurrency to embed a client of a separate cryptocurrency. This concept initially popularized by a proposal by Back et al. [9] particularly can be used to avoid a difficult upgrade process: a new blockchain with additional features, such as experimental opcodes, can be backed by deposits in the original Bitcoin currency, obviating the need to bootstrap new currencies from scratch. This has been heralded as a key enabling scalability of the Bitcoin ecosystem; however, the sidechains proposal is incomplete, as it does not provide a protocol that satisfies a natural security definition, i.e. where attacks succeed with only negligible probability. Our work is the first to answer the open problem posed by Back et al.

These examples illustrate that our solution is a key component for two important pillars needed for next-generation blockchains: *interoperability* and *scalability*.

## 1.2 Our contributions

In summary, we make the following contributions in this paper.

Our main technical contribution is the introduction of a new cryptographic primitive called Non-Interactive Proofs of Proof of Work (NIPoPoW).

We present a formal model and a provably secure instantiation of NIPoPoWs in the model of [10]. Our contribution builds on previous work by Kiayias et al. who posed interactive proofs of proof-of-work [11], which, in turn, are based on previous discussion of such concepts in the bitcoin forums [12]. We show that the construction of [11] does not really solve the PoPoW problem because it is insecure and we illustrate this with a concrete attack. We prove that the attack can be realized with overwhelming probability. We then prove that our construction is in fact a secure solution to the problem. Furthermore, our solution has the additional property of being non-interactive, as the previous construction could be forced into interaction by an adversarial prover is present. Similar to previous work, we prove that these proofs are optimistically succinct which means that in honest conditions, the proofs are logarithmic in size. Our definition fills the gap in terms of security modeling and design that existed in previous proposals, especially the notion of cumulative “Dynamic Member Multisignature” [9].

We provide concrete parameterization and empirical analysis focusing to show the potential savings of our approach versus existing clients. Using actual data from the bitcoin blockchain we quantify the actual savings of NIPoPoWs over the previous known techniques of constructing efficient SPV verifiers.

We describe two practical deployment paths for our NIPoPoWs that existing cryptocurrencies can adopt to support our efficient clients. Specifically, we first outline how NiPoPoW’s can be deployed using either a “soft fork” or a “hard fork” upgrade procedure, both of which have been successfully used by existing cryptocurrencies [13]. We then present a less disruptive update mechanism that we term a “velvet fork” and which may

be of independent interest. In a velvet fork update the clients that have “forked” from the original implementation continue to be fully compatible with unupdated clients. It follows that in a velvet fork, the blockchain system can remain supported by a diverse software codebase indefinitely, while it can still enjoy, at least in proportion, some of the (efficiency in our case) benefits of the update without any of the security downsides.

## 2 Related work

*Lightweight “SPV” Clients.* Nakamoto’s original Bitcoin whitepaper [1] anticipated the rising costs of an evergrowing blockchain, and proposed a protocol for lightweight clients, called “Simplified Payment Verification” (SPV). Unlike “full nodes” which process and validate the entire blockchain (including all transactions and signatures), SPV clients only process the proof-of-work chain and transactions directly pertaining to them. When initialized for the first time, an SPV client downloads (and stores) the entire chain blockchain, which is thus linear in the chain size  $|\mathcal{C}|$ , and therefore grows unbounded over time, although it has low constant factors, i.e.  $80|\mathcal{C}|$  bytes. When an SPV client receives a payment, it requests a proof-of-inclusion (i.e., a Merkle tree branch proof) in order to confirm that the transaction is included in one of the blocks whose header is already stored by the client. Most widely used clients, especially mobile wallets based on BitcoinJ, implement SPV.

As an alternative to downloading all block headers at first startup, the SPV client software could embed a hardcoded checkpoint (perhaps chosen conservatively, e.g. many months in the past), blocks prior to which are ignored.<sup>5</sup> The BitcoinJ software provides mechanisms for users to build “Checkpoint Files,” and at least some applications, such as Multibit,<sup>6</sup> include checkpoints. Although this approach is very efficient, it introduces additional trust assumptions on software maintainers.

*Proofs of proof-of-work.* Kiayias et al. (KLS) [11] proposed a variation of the SPV protocol that avoids the need to transmit the entire chain of proof-of-work. Instead, in their scheme only a sublinear-sized sample of the proofs-of-work are sent, i.e. a “proof of proof-of-work” or PoPoW. They define security in a setting with a verifier (modeling the client) and multiple provers (modeling peers in the network), at least one of which is assumed to be honest. The KLS protocol is interactive; the verifier may need to interact with the provers over a number of rounds logarithmic in the length of the proof-of-work chain. Our protocol is inspired by KLS, but improves on it in several ways. First, our protocol is provably secure in contrast to that construction for which we describe a concrete attack that works for a high level, albeit still below 50%, of adversarial power. Second, our protocol is non-interactive, completing in one communication round regardless of adversarial prover’s behavior. Last, we also provide concrete implementation recommendations including security parameter values, proof size estimates in bytes, specific optimizations to reduce proof size as well as smooth upgrade recommendations.

In an influential whitepaper [9], Back et al. proposed the use of efficient non-interactive SPV proofs as a way of enabling applications spanning multiple cryptocurrencies. One such application would be a virtual asset in a new cryptocurrency, called a “sidechain,” that is backed by a deposit of value in an existing cryptocurrency such as Bitcoin. Several subsequent efforts by Poelstra [14] and Friedenbach [15] explored this design space, however a security definition and suitable construction has remained an open problem. In particular,

<sup>5</sup> See <https://bitcoinj.github.io/speeding-up-chain-sync>

<sup>6</sup> <https://multibit.org/help/hd0.3/how-spv-works.html>

prior proposals have been vulnerable to “high variance” attacks, i.e. attacks that fail on average, but nonetheless succeed with non-negligible probability. In Section 10, we explain how our protocol solves this open problem and can fulfill the sidechains vision.

### 3 Model and Definitions

We define our problem in the setting of a blockchain client which is completely stateless beyond the knowledge of a common reference string, the Genesis block. Without loss of generality, this single reference block could be any stable checkpoint block, either hard-coded in the client or obtained through a previous interaction with the network. However, importantly, the client does not maintain a blockchain. The challenge is, then, to build a client that is able to communicate with the network and, without downloading the whole chain headers, is convinced in a secure manner that a transaction has taken place. We term such a client a *verifier* and we term *provers* the full nodes it connects to.

The intuition behind these proofs is quite simple: Let  $T$  be the mining target for the network. Then because block ids are generated uniformly at random, half of the blocks will have an id less than  $T/2$ . By having the prover present to the verifier just these blocks which capture more proof-of-work than their neighbor counterparts, the verifier is convinced that the proof-of-work of the omitted blocks also took place. The same logic can be applied to any level  $\mu$ , namely blocks with id less than  $T2^{-\mu}$ .

#### 3.1 Backbone algorithms

We assume the verifier will always connect to at least one honest prover. We now move on to formally define these entities in terms of a mathematical model which can be used for our analysis.

In order to study the syntax and properties of NiPoPoWs we abstract their operation in a way that is akin to non-interactive zero-knowledge proofs (NIZK) [16]. In such protocols, a prover shares some context with a verifier (called the common reference string). For a given predicate, the prover computes the proof string and transmits it to the verifier. The verifier accepts or rejects the string. Given that in our setting the predicates will be referring to an active evolving chain we should define them in the presence of a blockchain protocol. Our theoretical abstraction is based on the “backbone” model of [10] that formalizes the operation of the Bitcoin protocol.

The backbone approach models proof-of-work discovery attempts by using a random oracle [17]. The model uses previous formulations of secure multiparty computation, specifically Canetti’s formulation of “real world” execution [18] [19] [20]. For following the full formal security analyses, we kindly invite our reader to familiarize themselves with the backbone model of GKL[10] and the respective quantities.

We introduce the Prover and the Verifier in the Backbone model as follows. The Prover will extend the functionality of the miner node. The Verifier will be a completely new entity that can be spawned by the environment and process a NiPoPoW produced by the miner nodes running the Prover code.

The predicates of interest in our context are predicates on the active blockchain. Some of the predicates are more suitable for succinct proofs than others. We focus our attention in what we call *reliable* predicates; that is, predicates which are monotonic and stable notions that we define below. These predicates have the exceptional property that all honest miners share their view of them in a way that is updated in a predictable manner, with a truth-value that persists as the blockchain grows. Furthermore, we are interested in *succinct* predicates, predicates which can be proven using a succinct protocol, a protocol

which requires proofs that are short in terms of the total blockchain size. Finally, we formalize the notion of non-interactivity.

The honest entities that live on the blockchain network are of three kinds: First, miners, who try to mine new blocks on top of the longest known blockchain and broadcast their blocks as soon as they are discovered. Second, full nodes, which maintain the longest blockchain but without mining. Full nodes also act as the provers in the network, producing proofs by calling a Prove function which we leave undefined for now, as it is part of the proof protocol. And third, verifiers or stateless clients, which connect to provers and ask for proofs in regards to which blockchain is the largest. The verifiers attempt to determine the value of a predicate on these chains.

A non-interactive proof of proof-of-work protocol is a game between a prover and a verifier parameterized by predicate  $Q$  which the prover tries to convince the verifier for. The predicate  $Q$  is a function of a chain  $\mathcal{C}$ , which the prover claims to be the longest chain.

When asked to do so by the environment, two provers generate proofs  $\pi_A, \pi_B$  claiming potentially different truth values for the predicate  $Q$  based on their claimed local longest chains. The verifier receives these proofs and accepts one of the two proofs, determining the truth value of the predicate. In this setting, one of the two proofs could be adversarial, as defined in the security section.

We define a *blockchain proof protocol* for a predicate  $Q$  as a pair  $(P, V)$  where  $P$  is called the *prover* and  $V$  is called the *verifier*.  $P$  is a PPT algorithm that is spawned by a full node when they wish to produce a proof, accepts as input a full chain  $\mathcal{C}$  and produces a proof  $\pi$  as its output.  $V$  is a PPT algorithm which is spawned by the environment at some round, receives a pair of proofs  $(\pi_A, \pi_B)$  from both an honest party and the adversary and returns its decision  $d \in \{T, F, \perp\}$  before the next round and terminates. The honest parties produce proofs for  $V$  using  $P$ .

### 3.2 Blockchain addressing

Blockchains, or *chains*, are finite sequences of blocks such that they obey the *blockchain property* that for every block in the chain there exists a pointer to its previous block in the chain. A chain  $\mathcal{C}$  is *anchored* if its first block is the designated *Genesis* block, a known reference string. We denote the genesis block as  $Gen$ .

For chain addressing we use the Python range notation of brackets  $[ ]$ . A positive number in a bracket indicates the indexed block in the chain, starting from zero. For example, if  $\mathcal{C}$  is anchored, then  $\mathcal{C}[0] = Gen$ . A negative number indicates a block indexed from the end. For example  $\mathcal{C}[-1]$  indicates the tip of the chain. A range of two numbers  $\mathcal{C}[i : j]$  indicates a subarray starting on index  $i$  (inclusive) and ending on index  $j$  (exclusive), which can be negative. Omitting one of the range indices takes it to the beginning or end of the chain respectively. So, for example  $\mathcal{C}[1 : ]$  indicates a chain with the first block omitted and  $\mathcal{C}[: -k]$  indicates a chain with the last  $k$  blocks omitted, the “stable” part of the chain according to [10].

Given two chains  $\mathcal{C}_1$  and  $\mathcal{C}_2$  we write  $\mathcal{C}_1\mathcal{C}_2$  to indicate the concatenation of these two chains into a new chain. Then the  $\mathcal{C}_2[0]$  must point back to  $\mathcal{C}_1[-1]$ . Similarly with a block  $B$  we write  $\mathcal{C}_1B$  to indicate the concatenation of block  $B$  to the tip of  $\mathcal{C}_1$ , implying that block  $B$  points back to  $\mathcal{C}_1[-1]$  to preserve the blockchain property.

We will also use block-based addressing with the range notation of curly braces  $\{ \}$ . Given a chain  $\mathcal{C}$  and two blocks  $A, Z$  which are in  $\mathcal{C}$ , we denote  $\mathcal{C}\{A : Z\}$  the subarray of the chain starting from block  $A$  (inclusive) and ending in block  $Z$  (exclusive). Again we can omit blocks from the range to take the chain to the beginning or end respectively.

Hence, for example,  $\mathcal{C}\{A : \}$  indicates that the subarray of the chain from block  $A$  and until the end of the chain is to be taken.

Some helper functions need to be defined. The *id* function calculates the id of a block given its data. This is done by applying the hashfunctions  $H$  and  $G$  in a manner similar to [10]:  $id = H(ctr, G(s, x, interlink))$ . However, note that block  $B$  maintains an *interlink* data structure instead of a previous block pointer. We will explain this further in section 4. Note also that, while we spell out the transaction vector  $x$ , this in practice contains only the root of the Merkle tree of the transactions [21].

The id of a block must always be less than or equal to the mining target  $T$  to satisfy the proof-of-work[22] condition:  $id \leq T$ . However, some blocks will achieve a lower id. If  $id \leq \frac{T}{2^\mu}$  we say that the block is of level  $\mu$ . Clearly all blocks are blocks of level 0. We term blocks that have a level above 0 *superblocks* of their respective level. The level of a block is given as  $\mu = \lfloor \log(T) - \log(id(B)) \rfloor$  and denoted  $level(B)$ . By convention, we set the id of Genesis to be 0 and hence it has  $\mu = \infty$ .

In this paper, we will extend blocks to contain multiple pointers to previous blocks. Hence, certain blocks can be omitted from a chain, obtaining a subchain. However, the subchain must still maintain the blockchain property that each block must contain a pointer to its previous block in the sequence (except for the first block in the sequence). Subarrays of chains obtained by addressing with  $[ ]$  or  $\{ \}$  are, naturally, subchains, as the only omit blocks from the beginning or the end of a chain. It will also become apparent that it is useful to omit certain blocks from the middle of a chain, while maintaining the blockchain property.

While chains are sequences and not sets, we will use some set notation for convenience. The notation  $B \in \mathcal{C}$  will be used to indicate that a block  $B$  is part of a chain  $\mathcal{C}$ , while  $B \notin \mathcal{C}$  will be used to indicate that it isn't. We will use  $\emptyset$  to denote the empty blockchain. Finally, given a chain  $\mathcal{C}$  and two subchains  $\mathcal{C}_1$  and  $\mathcal{C}_2$  of it, we will write  $\mathcal{C}_1 \subseteq \mathcal{C}_2$  to indicate that  $\mathcal{C}_1$  is a (potentially improper) subchain of  $\mathcal{C}_2$ , i.e. that for all blocks  $B$  of  $\mathcal{C}_1$ , we have that  $B \in \mathcal{C}_2$ , while  $\mathcal{C}_1 \not\subseteq \mathcal{C}_2$  will be used to indicate that it isn't. Because both  $\mathcal{C}_1 \subseteq \mathcal{C}$  and  $\mathcal{C}_2 \subseteq \mathcal{C}$ , we can define  $\mathcal{C}_1 \cup \mathcal{C}_2$  to mean the chain obtained by sorting the blocks contained in both  $\mathcal{C}_1$  and  $\mathcal{C}_2$  into a sequence, provided that we can ensure the blockchain property is maintained. Hence, for example, if blocks  $(B_1, B_2, B_3)$  form a blockchain  $\mathcal{C}$  such that  $B_2$  points to  $B_1$  and  $B_3$  points to  $B_2$ , but also  $B_3$  points to  $B_1$ , then letting  $\mathcal{C}_1 = (B_1, B_3)$  and  $\mathcal{C}_2 = (B_2)$ , we note that both  $\mathcal{C}_1$  and  $\mathcal{C}_2$  are blockchains and that  $\mathcal{C}_1 \cup \mathcal{C}_2$  is a blockchain. Additionally, take two anchored chains  $\mathcal{C}_1$  and  $\mathcal{C}_2$ , and note that  $\mathcal{C}_1$  and  $\mathcal{C}_2$  belong to the same blocktree but can potentially have forked. Then we use the notation  $\mathcal{C}_1 \cap \mathcal{C}_2$  to mean the anchored chain  $\{B : B \in \mathcal{C}_1 \wedge B \in \mathcal{C}_2\}$  provided it satisfies the blockchain property. We also define the lowest common ancestor of these chains  $LCA(\mathcal{C}_1, \mathcal{C}_2) = (\mathcal{C}_1 \cap \mathcal{C}_2)[-1]$ . We will also freely use set builder notation  $\{B \in \mathcal{C} : p(B)\}$  to obtain a subchain  $\mathcal{C}'$  from a chain  $\mathcal{C}$  by filtering with a block predicate  $p(B)$ , provided the blockchain property is maintained. Note that in all these cases, because we are talking about blockchains and not blocktrees, any set of blocks has a unique order obtained by topologically sorting the blocks based on the order indicated by the underlying chain and so can be put into a unique sequence. For two chains  $\mathcal{C}_1$  and  $\mathcal{C}_2$ , if  $\mathcal{C}_1[0] = \mathcal{C}_2[0]$  and  $\mathcal{C}_1[-1] = \mathcal{C}_2[-1]$ , we will say that the two chains *span* the same range of blocks.

Given a chain  $\mathcal{C}$  and a level  $\mu$ , we define the *upchain operator*  $\mathcal{C}^\uparrow^\mu$  to mean the chain  $\{B \in \mathcal{C} : level(B) \geq \mu\}$ . We will call such a chain that contains only  $\mu$ -superblocks a  $\mu$ -*superchain*. Given a chain  $\mathcal{C}'$  and a chain  $\mathcal{C}$  such that  $\mathcal{C}' \subseteq \mathcal{C}$ , we define the *downchain operator*  $\mathcal{C}' \downarrow^{\mathcal{C}}$  as the chain  $\mathcal{C}[\mathcal{C}'[0] : \mathcal{C}'[-1]]$  and we will call  $\mathcal{C}$  the *underlying chain* of  $\mathcal{C}'$ . If the underlying chain is implied by context, for example, if the chain  $\mathcal{C}'$  was obtained by

applying the upchain operator on a known underlying chain  $\mathcal{C}$ , we will simply write  $\mathcal{C}'\downarrow$ . Note as well that the  $\mathcal{C}\uparrow$  operator is absolute, in the sense that  $(\mathcal{C}\uparrow^\mu)^{\mu+i} = \mathcal{C}\uparrow^{\mu+i}$ .

If a set of consecutive rounds  $S$  is given, we define

$$\mathcal{C}^S = \{B \in \mathcal{C} : B \text{ was generated during } S\}$$

### 3.3 Desired properties

We now define two desired properties of a non-interactive proof protocol, in particular *succinctness* and *security*.

**Definition 1.** (*Security*) A blockchain proof protocol is secure if for all environments and for all PPT adversaries  $\mathcal{A}$  the output of  $V$  on round  $r$  is the same as the evaluation of  $Q(\mathcal{C})$  on all honest parties' chains  $\mathcal{C}$  as long as  $Q(\mathcal{C})$  is the same for all honest parties after round  $r$ .

**Definition 2.** (*Succinctness*) A blockchain proof protocol is succinct if the maximum proof size  $|\pi|$  across all honest provers for a given round  $r$  is  $O(\text{polylog}(r))$ .

Based on these definitions, it is easy to see that it is trivial to construct a secure but not succinct protocol: The prover provides the chain  $\mathcal{C}$  itself as a proof and the verifier simply compares the proofs received by length: Since the longest proof is literally the longest chain, the protocol is trivially secure but not succinct. The challenge we will solve over the next sections is to provide a non-interactive protocol which at the same time achieves security and succinctness over a large class of predicates.

### 3.4 Provable chain predicates

Not all predicates can be proved. In this section we characterize the class of blockchain predicates that will be of interest to us for constructing proofs. A useful property of predicates in our context is *monotonicity*. Based on predicate monotonicity, we show that a predicate has *liveness* and *persistence*. Formalizing the security of our blockchain proofs will be preconditioned on such blockchain predicates. Given that we are in a decentralized setting it can be the case that for certain sensible blockchain predicate, there will be honest nodes that have not reached a conclusion about its value. For this reason we allow our predicates to have an undefined value  $\perp$  before they take on a truth value of 0 or 1. We define a number of predicate properties and prove a simple but useful relation between them.

**Definition 3.** (*Predicate monotonicity*) A chain predicate  $Q(\mathcal{C})$  is monotonic if for all chains  $\mathcal{C}$  and for all blocks  $B$  we have that:  $Q(\mathcal{C}) \neq \perp \Rightarrow Q(\mathcal{C}) = Q(\mathcal{C}B)$ .

(*Predicate switching position*) A chain predicate  $Q(\mathcal{C}) \neq \perp$  has a switching position  $i \in \mathbb{N}$  in the chain  $\mathcal{C}$  defined as the minimum  $i$  such that:  $Q(\mathcal{C}[i]) \neq \perp$ .

(*Predicate stability*) Parameterized by  $k \in \mathbb{N}$  and a chain  $\mathcal{C}$ , a chain predicate  $Q$  has stability if its switching position  $i$  satisfies  $i < |\mathcal{C}| - k$ .

(*Predicate persistence*) Parameterized by  $k \in \mathbb{N}$  and a chain  $\mathcal{C}$ , a chain predicate  $Q$  has persistence if the following is true: When in a certain round the predicate is  $k$ -stable in some honest party's chain, then whenever it is  $k$ -stable in any honest party's chain, the truth value of the predicate is the same. Furthermore, the change of truth value from  $\perp$  to its final truth value happened at the same blockchain depth.

**Theorem 1.** *If a chain predicate is monotonic, then it satisfies persistence with parameter  $k = 2\eta\kappa f$  with probability at least  $1 - e^{-\Omega(\kappa)}$ .*

*Proof.* Assume a typical execution as defined in [10] and let  $\mathcal{C}_1$  be the chain of some honest player  $P_1$  at round  $r_1$ . Assume that the chain predicate  $Q(\mathcal{C})$  is  $k$ -stable for  $P_1$  at round  $r_1$ , i.e.  $Q(\mathcal{C}[: -k]) \neq \perp$ . The chain of any honest player  $P_2$  will be at least as long from the next round on and from common prefix we have that  $\mathcal{C}_1[: -k] \preceq \mathcal{C}_2$ . Therefore the predicate will be  $k$ -stable, have the same truth value, and have the same switching position for all honest parties.  $\square$


## 4 Consensus layer support

### 4.1 The interlink pointers data structure

We will now start constructing a concrete blockchain proof protocol. Before we're able to define the prover and the verifier algorithms, we need to propose certain modifications to the consensus layer.

We now observe that some blocks will achieve a much lower id. Observe then that in expectation, half of the blocks will be of level 1, 1/4 of the blocks will be of level 2, 1/8 will be of level 3 and  $1/2^\mu$  blocks will be of level  $\mu$ . In expectation, the number of superblock levels of a chain  $\mathcal{C}$  will be  $\log(\mathcal{C})$ .

Figure 1 illustrates the blockchain superblocks starting from level 1 and going up to level 4 in case these blocks are distributed exactly according to expectation. Note that each level contains half the blocks of the level below.



**Fig. 1.** The hierarchical blockchain. Higher levels have achieved a lower target (higher difficulty) during mining.

The *interlink* data structure is proposed to be included in each block, replacing the existing pointer to the previous block with a list of pointers to a small number of previous blocks. For each superblock level  $\mu$ , this data structure contains a pointer to the most recent preceding block of level  $\mu$ . The algorithm for this construction is shown in Algorithm 1 and is borrowed from [11]. Genesis is of infinite level and hence a pointer to it is included in every block at the first available index within the interlink data structure. The interlink data structure turns the blockchain into a skiplist-like data structure. Observe that the number of pointers that need to be included per block, in expectation, is  $\log(\mathcal{C})$ .

The `updateInterlink` algorithm accepts a block  $B'$ , which already has an interlink data structure defined on it. The function then evaluates the interlink data structure which needs to be included as part of the next block. The algorithm proceeds by copying the existing interlink data structure and then modifying some of its entries from level 0 to the level of block  $B'$  to point to the block  $B'$ . Observe that the pointers stored are block ids and hence block data needs to be looked up based on blockid. To this end, the full node maintains a *blockbyid* dictionary which, queried with the block id, returns the block data.



Finally, the full node also maintains a *depth* data structure which, given a block, returns its distance from the Genesis block.

---

**Algorithm 1** updateInterlink

---

```

1: function updateInterlink( $B'$ )
2:   interlink  $\leftarrow B'.\text{interlink}$ 
3:   id  $\leftarrow \text{blockid}(B')$ 
4:   for  $\mu = 0$  to  $\text{level}(B')$  do
5:     interlink[ $\mu$ ]  $\leftarrow$  id
6:   end for
7:   return interlink
8: end function

```

---

## 4.2 Superchain quality

We now formally describe the distribution of superblocks within a blockchain.

**Definition 4 (Locally good superchain).** *A superchain  $\mathcal{C}'$  of level  $\mu$  with underlying chain  $\mathcal{C}$  is said to be  $\mu$ -locally-good with respect to security parameter  $\delta$ , written  $\text{local-good}_\delta(\mathcal{C}', \mathcal{C}, \mu)$ , if  $|\mathcal{C}'| > (1 - \delta)2^{-\mu}|\mathcal{C}|$ .*

**Lemma 1 (Local goodness).** *Assume chain  $\mathcal{C}$  contains only honestly-generated blocks and has been adopted by an honest party in an execution with random network scheduling. Then for all levels  $\mu$ , for all constant  $\delta > 0$ , all continuous subchains  $\mathcal{C}' = \mathcal{C}[i : j]$  with  $|\mathcal{C}'| \geq m$  are locally good,  $\text{local-good}_\delta(\mathcal{C}', \mathcal{C}, \mu)$ , with overwhelming probability in  $m$ .*

*Proof. (Sketch)* Observing that for each honestly generated block the probability of being a  $\mu$ -superblock for any level  $\mu$  follows an independent Bernoulli distribution, we can apply a Chernoff bound to show that the number of superblocks within a chain will be close to its expectation, which is what is required for local goodness.  $\square$

**Definition 5 (Chain superquality).** *The  $(\delta, m)$  superquality property  $Q_{scq}^\mu$  of a chain  $\mathcal{C}$  pertaining to level  $\mu$  with security parameters  $\delta \in \mathbb{R}$  and  $m \in \mathbb{N}$  states that for all  $m' \geq m$ , it holds that  $\text{local-good}_\delta(\mathcal{C}^{\uparrow\mu}[-m' :], \mathcal{C}^{\uparrow\mu}[-m' :]\downarrow, \mu)$ . That is, all suffixes that are sufficiently large are locally good.*

**Lemma 2 (Superquality).** *For all levels  $\mu$ , for all constant  $\delta > 0$ , a chain  $\mathcal{C}$  containing only honestly-generated blocks adopted by an honest party in an execution with random network scheduling has  $(\delta, m)$ -superquality at level  $\mu$  with overwhelming probability in  $m$ .*

*Proof.* Let  $\mathcal{C}' = \mathcal{C}^{\uparrow\mu}$  and let  $\mathcal{C}^* = \mathcal{C}'[-m' :]$  for some  $m' \geq m$ . Then let  $B \in \mathcal{C}^*\downarrow$  and let  $X_B$  be the random variable equal to 1 if  $\text{level}(B) \geq \mu$  and 0 otherwise.  $\{X_B : B \in \mathcal{C}^*\downarrow\}$  are mutually independent Bernoulli random variables with expectation  $E(X_B) = 2^{-\mu}|\mathcal{C}^*\downarrow|$ . Let  $X = \sum_{B \in \mathcal{C}^*\downarrow} X_B$ . Then  $X$  follows a Binomial distribution with parameters  $(m', 2^{-\mu})$  and note that  $|\mathcal{C}^*\downarrow| = X$ . Then  $\mathbb{E}(|\mathcal{C}^*\downarrow|) = 2^{-\mu}|\mathcal{C}^*|$ . Applying a Chernoff bound on  $|\mathcal{C}^*\downarrow|$  we obtain that:

$$\Pr[|\mathcal{C}^*\downarrow| \leq (1 - \delta)2^{-\mu}|\mathcal{C}^*\downarrow|] \leq \exp(-\delta^2 2^{-\mu-1}|\mathcal{C}^*|)$$

$\square$

**Definition 6 (Multilevel quality).** A  $\mu$ -superchain  $\mathcal{C}'$  is said to have multilevel quality, written  $\text{multi-good}_{\delta, k_1}(\mathcal{C}, \mathcal{C}', \mu)$  with respect to an underlying chain  $\mathcal{C} = \mathcal{C}' \downarrow$  with security parameters  $k_1, \delta$  if for all  $\mu' < \mu$  it holds that for any  $\mathcal{C}^* \subseteq \mathcal{C}$ , if  $|\mathcal{C}^* \uparrow^{\mu'}| = k_1$ , then  $|\mathcal{C}^* \uparrow^\mu| \geq (1 - \delta)2^{\mu - \mu'} |\mathcal{C}^* \uparrow^{\mu'}|$

Note also that for  $\delta < 0.5$ , multilevel quality implies that  $|\mathcal{C}^* \uparrow^\mu| \geq 1$ .

**Lemma 3 (Multilevel quality).** For all levels  $\mu$ , for all constant  $0 < \delta \leq 0.5$ , a chain  $\mathcal{C}$  containing only honestly-generated blocks adopted by an honest party in an execution with random scheduling has  $(\delta, k_1)$  multilevel quality at level  $\mu$  with overwhelming probability in  $k_1$ .

*Proof. (Sketch)* This proof is identical to the one of Lemma 1.  $\square$

**Definition 7 (Good superchain).** A  $\mu$ -superchain  $\mathcal{C}'$  is said to be good, written  $\text{good}_{\delta, k_1}(\mathcal{C}, \mathcal{C}', \mu)$ , with respect to an underlying chain  $\mathcal{C} = \mathcal{C}' \downarrow$  if it has both superquality and multilevel quality with the same parameters  $(\delta, m)$ .

**Corollary 1 (Optimistic superchain distribution).** For all levels  $\mu$ , for all constant  $0 < \delta < 0.5$ , a chain  $\mathcal{C}$  containing only honestly-generated blocks adopted by an honest party in an execution with random scheduling is  $(\delta, m)$ -good at level  $\mu$  with overwhelming probability in  $m$ .

*Proof.* This is a direct consequence of Lemma 2 and Lemma 3.  $\square$

## 5 An attack against proofs of proof-of-work

In this section, we revisit the construction for interactive proofs of proof-of-work from [11] and its security. Their construction is the starting point and basis for our non-interactive proofs of proof-of-work. We show that the construction is susceptible to a double-spending attack even in the case it controls a minority of the parties (hashing power). We recommend that the interested reader skims through the KLS [11] work to fully understand the scheme and the attacks that follow.

We first show that a powerful attacker can break chain superquality with non-negligible probability and we construct a concrete attack based on it. Even though maintaining chain superquality is not in the original security model, we show how the property affects the security of the underlying proofs by presenting a concrete attack against their scheme that allows an adversary to double spend with non-negligible probability assuming sufficient hashing power (that is still below 50%).

### 5.1 The interactive proofs of proof-of-work protocol

The interactive proof of proof-of-work of [11] operates as follows. The core of the verification process is distinguishing between two candidate proofs  $(\pi_A, \chi_A)$  and  $(\pi_B, \chi_B)$ . In case  $\pi_A = \pi_B$ , the decision can be drawn immediately (without interaction) based on  $\chi_A, \chi_B$ . Otherwise, the verifier queries the two provers for their claimed anchored superchains  $\pi_A \uparrow^\mu$  and  $\pi_B \uparrow^\mu$  at a certain level  $\mu$ . The verifier starts querying at the highest possible level  $\mu$  and successively descends until level  $\mu$  is sufficiently low such that the lowest common ancestor block  $b = \text{LCA}(\pi_A \uparrow^\mu, \pi_B \uparrow^\mu)$  of the two provers'  $\mu$ -superchains is  $m$  blocks deep for at least one of the provers. More specifically, the querying stops at such  $\mu$  when  $\max(|\pi_A \uparrow^\mu \{b : \cdot\}|, |\pi_B \uparrow^\mu \{b : \cdot\}|) \geq m$  holds. Subsequently, the winner is designated as the prover with the most blocks after  $b$  at that level. More precisely, the winner is  $A$  if  $|\pi_A \uparrow^\mu \{b : \cdot\}| \geq |\pi_B \uparrow^\mu \{b : \cdot\}|$  and otherwise it is  $B$ . The communication overhead is reduced by only requesting blocks after the purported LCA.

## 5.2 Attacking chain superquality

We present next a general attack against chain superquality that works against any blockchain protocol that incorporates the interlink data structure, cf. Section 4.1. Observe that in expectation, in a completely honest setting with random network scheduling, the honestly adopted chain will be  $\mu$ -locally-good for sufficiently long subchains everywhere as proven in Lemma 1.

Let us now construct an adversary to break this property at level  $\mu$ . Suppose the adversary  $\mathcal{A}$  controls a portion of the hashing power equal to  $t/n$ .  $\mathcal{A}$  then works as follows. Assume she wants to attack the honest party  $B$  in order to have him adopt a chain  $\mathcal{C}_B$  which has a bad distribution of superblocks, i.e. such that local goodness is violated in some sufficiently long subchain. She continuously determines the current chain  $\mathcal{C}_B$  adopted by  $B$  (as honest chain adoption is deterministic, this can be easily done by inspecting the messages exchanged by the honest parties). The adversary starts playing after  $|\mathcal{C}_B| \geq 2$ . If  $\text{level}(\mathcal{C}_B[-1]) < \mu$ , then  $\mathcal{A}$  remains idle. However, if  $\text{level}(\mathcal{C}_B[-1]) \geq \mu$ , then  $\mathcal{A}$  attempts to mine an adversarial block  $b$  on top of  $\mathcal{C}_B[-2]$ . If she is successful, then she attempts to mine another block  $b'$  on top of  $b$ . If she is successful, she broadcasts both blocks  $b, b'$ . The adversarial mining continues until party  $B$  adopts a new chain, which can be due to two reasons: Either the adversary managed to successfully mine  $b, b'$  on top of  $\mathcal{C}_B[-2]$  and succeeded in having  $B$  adopt it; or one of the honest parties (potentially  $B$ ) was able to mine a block which was subsequently adopted by  $B$ . In either case, the adversary continues with the strategy by inspecting  $\mathcal{C}[-1]$  and acting accordingly.

Assume now that an honestly-generated  $\mu$ -superblock has been adopted by  $B$  at position  $\mathcal{C}_B[i]$  at some round  $r$ . Let us now examine the probability that  $\mathcal{C}_B[i]$  will remain a  $\mu$ -superblock in the long run. Suppose  $r' > r$  is the first round after  $r$  during which any block is generated. It is clear that  $\mathcal{A}$  will succeed in this attack with non-negligible probability and cause  $B$  to abandon the  $\mu$ -superblock from their adopted chain. Therefore, there will be some  $\delta$  such that the adversary will be able to cause such variance with non-negligible probability in  $m$ . This suffices to show that superquality is harmed by this attack. To avoid repeating ourselves, the full understanding of the probabilities for this attack will become clear in the next section where the double-spending attack is explored in detail.

## 5.3 From chain superquality attacks to double-spending

Extending the above attack, we now modify the superquality attacker into an attacker,  $\mathcal{A}$ , that can cause a double spending attack in the proof of proof-of-work construction. As before,  $\mathcal{A}$  targets the proofs generated by the honest party  $B$  by violating  $\mu$ -superquality in  $B$ 's adopted chain.  $\mathcal{A}$  begins by remaining idle until a certain chosen block  $b$ . After block  $b$  is produced,  $\mathcal{A}$  starts mining a secret chain which forks off from  $b$  akin to a selfish mining attacker [23]. The adversary performs a normal spending transaction on the honestly adopted blockchain and has it confirmed in the block immediately following block  $b$ . They also produce a double spending transaction which they secretly confirm in their secret chain in the block immediately following block  $b$ .

$\mathcal{A}$  keeps extending their own secret chain as usual. However, whenever a  $\mu$ -superblock is adopted by  $B$ , the adversary temporarily pauses mining in their secret chain and devotes her mining power to harm the  $\mu$ -superquality of  $B$ 's adopted chain. Intuitively, for large enough  $\mu$ , the time spent trying to harm superquality will be limited, because the probability of a  $\mu$ -superblock occurring will be small. Therefore, the adversary's superchain quality will be larger than the honest parties' superchain quality (which will be harmed

by the adversary) and therefore, even though the adversary’s 0-level blockchain will be shorter than the honest parties’ 0-level blockchain, the adversary’s  $\mu$ -superchain will be longer than the honest parties’  $\mu$ -superchain.

We now proceed to calculate the attack probabilities precisely and in formal detail. We simplify the above attack to ease the formal probabilistic analysis. The attack is parameterized by two parameters  $r$  and  $\mu$  which are picked by the adversary.  $\mu$  will be the superblock level at which the adversary will attempt to produce a proof longer than the honest proof. The modified attack works precisely as follows: Without loss of generality, we fix block  $b$  to be the Genesis block. The adversary always mines on the secret chain which forks off from genesis, unless a *superblock generation event* occurs. If a superblock generation event occurs, then the adversary pauses mining on the secret chain and attempts a *block suppression attack* on the honest chain. The adversary devotes exactly  $r$  rounds to this suppression attack and then resumes mining on their secret chain. Our goal is to show that, despite this simplification (of fixing  $r$ ) which is harmful to the adversary, the probability of a successful attack is non-negligible for certain (reasonable) values of the protocol parameters.

A superblock generation event is detected by the adversary by monitoring the network. Whenever an honest party diffuses an honestly-generated  $\mu$ -superblock at the end of a given round  $r_1$ , then the superblock generation event will have occurred and the adversary will start devoting their mining power to block suppression starting from the next round.

A block suppression attack works as follows. Let  $B$  be the honestly generated  $\mu$ -superblock which was diffused at the end of the previous round. If the round was not uniquely successful, let  $B$  be any of the diffused honestly-generated  $\mu$ -superblocks. Let  $B$  be the tip of an honest chain  $\mathcal{C}_B$ . The adversary first mines on top of  $\mathcal{C}_B[-2]$ . If she is successful in mining a block  $B'$ , then she continues extending the chain ending at  $B'$  (to mine  $B''$  and so on). The value  $r$  is fixed, and so the adversary devotes exactly  $r$  rounds to this whole process; the adversary will keep mining on top of  $\mathcal{C}_B[-2]$  (or one of the adversarially-generated extensions of it) for exactly  $r$  rounds, regardless of whether  $B'$  or  $B''$  have been found. At the same time, the honest parties will be mining on top of  $B$  (or a competing block in the case of a non-uniquely successful round). Again, further successful block diffusion by the honest parties shall not affect that the adversary is going to spend exactly  $r$  rounds for suppression.

Having laid out the attack scenario precisely, we are ready to prove that it will succeed with non-negligible probability. In fact, as we will see, perhaps surprisingly, it will succeed with overwhelming probability for the right choice of protocol values.

**Theorem 2.** *There exist parameters  $p, n, t, q, \mu, \delta$ , with  $t \leq (1 - \delta)(n - t)$ , and a double spending attack against KLS PoPoW that succeeds with overwhelming probability.*

*Proof.* Recall that in the backbone notation  $n$  denotes the total number of parties,  $t$  denotes the number of adversarial parties,  $q$  denotes the number of the random oracle queries allowed per party per round and  $p$  is the probability that one random oracle query will be successful and remember that  $p = T/2^\kappa$  where  $T$  is the mining target and  $\kappa$  is the security parameter (or hash function bit count). Then  $f$  denotes the probability that a given round is successful and we have that  $f = 1 - (1 - p)^{q(n-t)}$ . Recall further that a requirement of the backbone protocol is that the honest majority assumption is satisfied, that is that  $t \leq (1 - \delta)(n - t)$  where  $\delta \geq 2f + 3\epsilon$ , where  $\epsilon \in (0, 1)$  is an arbitrary small constant describing the quality of the concentration of the random variables.

Denote  $\alpha_A$  the secret chain generated by the adversary and  $\alpha_B$  the honest chain belonging to any honest party. We will show that for certain protocol values we have that  $\Pr[|\alpha_A \uparrow^\mu| \geq |\alpha_B \uparrow^\mu|]$  is overwhelming.

Assume that, to the adversary's harm and to simplify the analysis, the adversary plays at beginning of every round and does not perform adversarial scheduling. At the beginning of the round when it is the adversary's turn to play, she has access to the blocks diffused during the previous round by the honest parties.

First, observe that at the beginning of each round, the adversary finds herself in one of two different situations: Either she has been forced into an  $r$ -round-long period of suppression, or she is not in that period. If she is within that period, she blindly performs the suppression attack without regard for the state of the world. If she is not within that period, then she must initially observe the blocks diffused at the end of the previous round by the honest parties. Call these rounds during which the diffused data must be examined by the adversary *decision rounds*. Let there be  $\omega$  decision rounds in total. In each such decision round, it is possible that the adversary discovers a diffused  $\mu$ -superblock and therefore decides that a suppression attack must be performed starting with the current round. Call these rounds during which this discovery is made by the adversary *migration rounds*. Let there be  $y$  migration rounds in total. The adversary devotes the migration round to performing the suppression attack as well as  $r - 1$  non-migration rounds after the migration round. Call these rounds, including the migration round, *suppression rounds*. In the rest of the decision rounds, the adversary will not find any  $\mu$ -superblocks diffused. Call these *secret chain rounds*; these are rounds where the adversary devotes her queries to mining on the secret chain. Let there be  $x$  secret chain rounds. If the adversary devotes  $\omega$  decision rounds to the attack in total, then clearly we have that  $\omega = x + y$ . If the total number of rounds during which the attack is running is  $s$  then we also have that  $s = x + ry$ , because for each migration round there are  $r - 1$  non-decision rounds that follow.

We will analyze the honest and adversarial superchain lengths with respect to  $\omega$ , which roughly corresponds to time (because note that  $\omega \geq s/r$ , and so  $\omega$  is proportional to the number of rounds).

Let us calculate the probability  $p_{SB}$  ("superblock probability") that a decision round ends up being a migration round. Ignoring the negligible event that there will be random oracle collisions, we have that  $p_{SB} = (n - t)qp2^{-\mu}$ .

Given this, note that the decision taken at the beginning of each decision round follows independent Bernoulli distributions with probability  $p_{SB}$ . Denote  $z_i$  the indicator random variable indicating whether the decision round was a migration round. Therefore we can readily calculate the expectations for the random variables  $x$  and  $y$ , as  $x = \omega - y$ ,  $y = \sum_{i=1}^{\omega} z_i$ . We have  $E[x] = (1 - p_{SB})\omega$  and  $E[y] = p_{SB}\omega$ . Applying a Chernoff bound to the random variables  $x$  and  $y$ , we observe that they will attain values close to their mean for large  $\omega$  and in particular  $\Pr[y \geq (1 + \delta)E[y]] \leq \exp(-\frac{\delta^2}{3}E[y])$  and similarly  $\Pr[x \leq (1 - \delta)E[x]] \leq \exp(-\frac{\delta^2}{2}E[x])$ , which are negligible in  $\omega$ .

Given that there will be  $x$  secret chain rounds, we observe that the random variable indicating the length of the secret adversarial superchain follows the binomial distribution with  $xtq$  repetitions and probability  $p2^{-\mu}$ . We can now calculate the expected secret chain length as  $E[|\alpha_{\mathcal{A}} \uparrow^{\mu}|] = xtqp2^{-\mu}$ . Observe that in this probability we have given the adversary the intelligence to continue using her random oracle queries during a round even after a block has been found during a round and not to wait for the next round. Applying a Chernoff bound, we obtain that  $\Pr[|\alpha_{\mathcal{A}} \uparrow^{\mu}| \leq (1 - \delta)E[|\alpha_{\mathcal{A}} \uparrow^{\mu}|]] \leq \exp(-\frac{\delta^2}{2}E[|\alpha_{\mathcal{A}} \uparrow^{\mu}|])$ , which is negligible in  $\omega$  (because we know that with overwhelming probability  $x > (1 - \delta)(1 - p_{SB})\omega$ ).

It remains to calculate the behavior of the honest superchain.

Suppose that a migration round occurs during which at least one superblock  $B$  is diffused. We will now calculate the probability  $p_{sup}$  that the adversary is able to suppress

that block after  $r$  rounds by performing the suppression attack and cause all honest parties to adopt a chain not containing  $B$ .

One way for this to occur is if the adversary has generated exactly 2 shallow blocks (blocks which are not  $\mu$ -superblocks) after exactly  $r$  rounds and the honest parties having generated exactly 0 blocks after exactly  $r$  rounds. This provides a lower bound for the probability, which is sufficient for our purposes. Call ADV-WIN the event where the adversary has generated exactly 2 shallow blocks after exactly  $r$  rounds since the diffusion of  $B$  and call HON-LOSE the event where the honest parties have generated exactly 0 blocks after exactly  $r$  rounds since the diffusion of  $B$ .

The number of blocks generated by the adversary after the diffusion of  $B$  follows the binomial distribution with  $r$  repetitions and probability  $p_{LB}$ , where  $p_{LB}$  denotes the probability that the adversary is able to produce a shallow block (“low block probability”) during a single round. We have that  $p_{LB} = tqp(1 - 2^{-\mu})$ . To evaluate  $\Pr[\text{ADV-WIN}]$ , we evaluate the binomial distribution for 2 successes to obtain  $\Pr[\text{ADV-WIN}] = \frac{r(r-1)}{2} p_{LB}^2 (1 - p_{LB})^{r-2}$ . The number of blocks generated by the honest parties after the diffusion of  $B$  follows the binomial distribution with  $r$  repetitions and probability  $f$ . To evaluate  $\Pr[\text{HON-LOSE}]$ , we evaluate the binomial distribution for 0 successes to obtain  $\Pr[\text{HON-LOSE}] = (1 - f)^r$ . Note that this is an upper bound in the probability, in particular because there can be multiple blocks during a non-uniquely successful round during which a  $\mu$ -superblock was generated.

Then observe that the two events ADV-WIN and HON-LOSE are independent and therefore  $p_{sup} = \Pr[\text{ADV-WIN}] \Pr[\text{HON-LOSE}] = \frac{r(r-1)}{2} p_{LB}^2 (1 - p_{LB})^{r-2} (1 - f)^r$ .

Now that we have evaluated  $p_{sup}$ , we will calculate the honest chain length in two chunks: The superblocks generated and adopted by the honest parties during secret chain rounds,  $\mathcal{C}_1$ , and the superblocks generated and adopted by the honest parties during suppression rounds,  $\mathcal{C}_2$  (and note that these sets of blocks are not blockchains on their own).

$|\mathcal{C}_1|$  is a random variable following the binomial distribution with  $s(n - t)q$  repetitions and probability  $p2^{-\mu}(1 - p_{sup})$ . In the evaluation of this distribution, we give the honest parties the liberty to belong to a mining pool and share mining information within a round, an assumption which only makes matters for the adversary worse. We can now calculate the expected length of  $\mathcal{C}_1$  to find  $E[|\mathcal{C}_1|] = s(n - t)qp2^{-\mu}(1 - p_{sup})$ . Applying a Chernoff bound, we find that  $\Pr[|\mathcal{C}_1| \geq (1 + \delta)E[|\mathcal{C}_1|]] \leq \exp(-\frac{\delta^2}{3}E[|\mathcal{C}_1|])$ , which is negligible in  $s$ .

Finally, some additional  $\mu$ -superblocks could have been generated by the honest parties while the adversary is spending  $r$  rounds attempting to suppress a previous  $\mu$ -superblock. These  $\mu$ -superblocks will be adopted in the case the adversary fails to suppress the previous  $\mu$ -superblock. As the adversary does not devote any decision rounds to these new  $\mu$ -superblocks, they will never be suppressed if the previous  $\mu$ -superblock is not suppressed. We collect these in the set  $\mathcal{C}_2$ . To calculate  $|\mathcal{C}_2|$ , observe that the number of unsuppressed  $\mu$ -superblocks which caused an adversarial suppression period is  $|\mathcal{C}_1|$ . For each of these blocks, the honest parties spend  $r$  rounds attempting to form further  $\mu$ -superblocks on top. The total number of such attempts is  $r|\mathcal{C}_1|$ . Therefore, the number of further honestly generated  $\mu$ -superblocks attained during the  $|\mathcal{C}_1|$  different  $r$ -round periods follows a binomial distribution with  $|\mathcal{C}_1|rq(n - t)$  repetitions and probability  $p2^{-\mu}$ . Here we allow the honest parties to use repeated queries within a round even after a shallow success and to work in a pool to obtain an upper bound for the expectation. Therefore  $E[|\mathcal{C}_2|] = |\mathcal{C}_1|rq(n - t)p2^{-\mu}$  and applying a Chernoff bound we obtain that  $\Pr[|\mathcal{C}_2| \geq (1 + \delta)E[|\mathcal{C}_2|]] \leq \exp(-\frac{\delta^2}{3}E[|\mathcal{C}_2|])$ , which is negligible in  $s$  and has a quadratic error term. We deduce that  $|\mathcal{C}_2|$  will have a very small length compared to the rest of the honest chain, as it is a vanishing term in  $\mu$ .

Concluding the calculation of the adversarial superchain, we get  $E[|\alpha_B \uparrow^\mu|] = E[|\mathcal{C}_1|] + E[|\mathcal{C}_2|]$ .

Finally, it remains to show that there exist values  $p, n, t, q, r, \mu, \delta$  such that  $E[|\alpha_A \uparrow^\mu|] \geq (1 + \delta)E[|\alpha_B \uparrow^\mu|]$ .

Using the values  $p = 10^{-5}, q = 1, n = 1000, t = 489, \mu = 25, r = 200$ , we observe that the honest majority assumption is preserved. Replacing these values into the expectations formulae above, we obtain  $E[|\alpha_A \uparrow^\mu|] \approx 1.457 * 10^{-10} * \omega$  and  $E[|\alpha_B \uparrow^\mu|] \approx 1.424 * 10^{-10} * \omega$ , which result to a constant gap  $\delta$ . Because the adversarial chain grows linearly in  $\omega$ , the adversary only has to wait sufficient rounds for obtaining  $m$  blocks to create a valid proof. Therefore, for these values, the adversary will be able to generate a convincing PoPoW at some level  $\mu$  which is longer than the honest parties' proof, even when the adversary does not have a longer underlying blockchain.  $\square$

**Remark.** It is worth isolating the mistake in the security proof from the interactive construction paper [11]. Suppose player  $B$  is honest and player  $A$  is adversarial and suppose  $b$ , the LCA block, was honestly generated and suppose that the superchain comparison happens at level  $\mu$ . Their proof then argues that there will have been more honestly- than adversarially-generated  $\mu$ -superblocks after block  $b$ . Nevertheless, we observe that the mere fact that there have been more honestly- than adversarially-generated  $\mu$ -superblocks after  $b$  does not imply that  $|\bar{\pi}_A[\mu]\{b : \}| \leq |\bar{\pi}_B[\mu]\{b : \}|$ . The reason is that some of these superblocks could belong to blocktree forks that have been abandoned by  $B$ . Thus, the security conclusion does not follow.

Regardless, their basic argument and construction is what we will use as a basis for constructing a system that is both provably secure and succinct under the same assumptions, albeit requiring a more complicated proof structure to obtain security.

## 6 Blockchain suffix proofs

We now provide a concrete non-interactive PoPoW construction which allows proving certain predicates  $Q$  of the chain  $\mathcal{C}$ . The construction will be done in two steps: First, we will build a construction which is limited to a class of predicates that is easy to prove. Later, we will extend this class of predicates by augmenting the construction.

Among the predicates which are stable, we now limit ourselves to the class of predicates which are functions of only the chain suffix  $\mathcal{C}[-k - 1 : ]$ . We call these predicates  $(k + 1)$ -*suffix sensitive*. As we will illustrate, these predicates are significantly easier to prove, but at the same time allow us to use the construction as a scaffold for the more generic case.

**Definition 8 (Suffix sensitivity).** *A chain predicate  $Q$  is called  $k$ -suffix sensitive if for all chains  $\mathcal{C}, \mathcal{C}'$  with  $|\mathcal{C}| \geq k$  and  $|\mathcal{C}'| \geq k$  such that  $\mathcal{C}[-k : ] = \mathcal{C}'[-k : ]$  we have that  $Q(\mathcal{C}) = Q(\mathcal{C}')$ .*

Suffix-sensitive predicates allow proving that a recent transaction occurred in the longest chain and in a block buried under at most  $k + 1$  blocks. Because the predicates we are interested in are stable, they must also pertain to properties that are independent of the last  $k$  blocks. Therefore, combining the two, the predicate can only describe something that is visible at the exact block  $\mathcal{C}[-k - 1]$  and that will also remain true as the chain grows in perpetuity (due to monotonicity). These predicates may seem contrived for blockchains like Bitcoin. However, they make sense for blockchains such as Ethereum which maintain state [24] that is propagated and committed to in every block. The existence of a stable monotonic state bit flip, for example the creation of an indestructible smart contract, can therefore be described using such predicates.

## 6.1 Construction

We will now describe the details of the NIPoPoW construction by presenting a specific prover and verifier. We present a generic form of the verifier first and the prover afterwards. The generic form of the verifier will work with any practical suffix proof protocol (we make this more exact below). Therefore, it makes sense to describe the generic verifier first before we talk about the specific instantiation of our protocol. The generic verifier is given access to call a proof comparison operator  $\leq_m$  which is protocol-specific. We therefore begin the description of our protocol by first illustrating the generic verifier. Next, we describe the prover specific to our protocol. Finally, we show the instantiation of the  $\leq_m$  operator, which plugs into the generic verifier to make a concrete verifier for our protocol.

### The generic verifier

The Verify function of our NIPoPoW construction is described in Algorithm 2. The verifier algorithm is parameterized by a chain predicate  $Q$  and security parameters  $k$  and  $m$ . The parameter  $k$  pertains to the amount of proof-of-work needed to bury a block so that it is believed to remain stable; in bitcoin’s case, we often set  $k = 6$ . The parameter  $m$  is a security parameter pertaining to the prefix of the proof, which connects the Genesis block to the  $k$ -sized suffix. The verifier receives several proofs by different provers in a collection of proofs  $\mathcal{P}$ , at least one of which must be generated by an honest prover. Iterating over these proofs, it extracts the best one by comparing these proofs.

Each of these proofs is a chain. For honest provers, these chains are subchains of the currently adopted chain. The proofs consist of two parts,  $\pi$  and  $\chi$ , such that  $\pi\chi$  is a valid chain.  $\chi$  is called the proof suffix, for which we require that  $|\chi| = k$ , while  $\pi$  is called the proof prefix. For honest provers,  $\chi$  will be the last  $k$  blocks of their adopted chain, while  $\pi$  will consist of a selected subset of blocks from the rest of their chain preceding  $\chi$ . The exact method of choice of this subset will become clear when we describe the concrete prover in the next section.

The verifier compares the various proofs provided to it by calling the  $\geq_m$  comparison operator, which is defined by the protocol. We will get to the operator’s definition after we have examined the prover algorithm. All proofs are checked for validity before comparison by checking that  $|\chi| = k$  and subsequently using the *validChain* function. This function simply checks that  $\pi\chi$  is an anchored blockchain by ensuring that the first block in the proof is the common reference string (Genesis) and that each block in the proof contains a pointer to the previous block within the proof.

---

**Algorithm 2** The Verify algorithm for the NIPoPoW protocol

---

```
1: function Verify $_{m,k}^Q(\mathcal{P})$ 
2:    $\tilde{\pi} \leftarrow (\text{Gen})$  ▷ Trivial anchored blockchain
3:   for  $(\pi, \chi) \in \mathcal{P}$  do ▷ Examine each proof  $(\pi, \chi)$  in  $\mathcal{P}$ 
4:     if validChain( $\pi\chi$ ) and  $|\chi| = k$  and  $\pi \geq_m \tilde{\pi}$  then
5:        $\tilde{\pi} \leftarrow \pi$ 
6:        $\tilde{\chi} \leftarrow \chi$  ▷ Update current best
7:     end if
8:   end for
9:   return  $Q(\tilde{\chi})$ 
10: end function
```

---

At each loop iteration, the verifier compares the next candidate proof prefix  $\pi$  against the currently best known proof prefix  $\tilde{\pi}$  by calling  $\pi \geq_m \tilde{\pi}$ . If the candidate prefix is



better than the currently best known proof prefix, then the currently known best prefix is updated by setting  $\tilde{\pi} \leftarrow \pi$ . When the best known prefix is updated, the suffix  $\tilde{\chi}$  associated with the best known prefix is also updated to match the suffix  $\chi$  of the candidate proof by setting  $\tilde{\chi} \leftarrow \chi$ . While  $\tilde{\chi}$  is needed for the final predicate evaluation, it is not used as part of any comparison, as it has the same size  $k$  for all proofs. The best known proof prefix is initially set to  $(Gen)$ , the trivial anchored chain containing only the genesis block. Any well-formed proof compares favourably against the trivial chain.

After the end of the for loop, the verifier will have determined the best proof  $(\tilde{\pi}, \tilde{\chi})$ . We will later prove that this proof will necessarily belong to an honest prover with overwhelming probability. Since the proof has been generated by an honest prover, it is associated with an underlying honestly adopted chain  $\mathcal{C}$ . The verifier then evaluates the predicate  $Q$  on the underlying chain associated with the proof. Because the predicate is suffix-sensitive, its value does not depend on the prefix, and so to evaluate the predicate on the underlying blockchain it suffices to evaluate it on any sequence with the same suffix, namely  $\tilde{\chi}$ . More concretely, from the definition of suffix sensitivity, we have that  $Q(\mathcal{C}) = Q(\tilde{\pi}\tilde{\chi}) = Q(\tilde{\chi})$ .

This generic form of the verifier will work for any protocol for suffix proofs, as long as the underlying operator  $\leq_m$  is defined in a manner that is transitive and returns the same result if invoked multiple times with the same inputs with overwhelming probability.

### The concrete prover

The NIPoPoW honest prover construction is shown in Algorithm 3. The honest prover is supplied with an honestly adopted chain  $\mathcal{C}$  and security parameters  $m, k, \delta$  and returns a proof  $\pi\chi$ , which is a chain. The proof consists of a suffix  $\chi$  and a prefix  $\pi$ . The suffix  $\chi$  is simply the last  $k$  blocks of  $\mathcal{C}$ . The prefix  $\pi$  is constructed by selecting various blocks and adding them to  $\pi$ . which consists of a number of blocks for every level  $\mu$ . At the highest possible level at which at least  $m$  blocks, all these blocks are presented. Then, inductively, for every superchain of level  $\mu$  that is included in the proof, the suffix of length  $m$  is taken. Then the underlying superchain of level  $\mu - 1$  covering the same span as that suffix is also included, until level 0 is reached. This underlying superchain will have  $2m$  blocks in expectation and always at least  $m$  blocks.

---

#### Algorithm 3 The Prove algorithm for the NIPoPoW protocol

---

```

1: function Provem,k,δ( $\mathcal{C}$ )
2:    $B \leftarrow \mathcal{C}[0]$  ▷ Genesis
3:   for  $\mu = |\mathcal{C}[-k].\text{interlink}|$  down to 0 do
4:      $\alpha \leftarrow \mathcal{C}[: -k]\{B : \}^\mu$ 
5:      $\pi \leftarrow \pi \cup \alpha$ 
6:     if goodδ,m( $\mathcal{C}, \alpha, \mu$ ) then
7:        $B \leftarrow \alpha[-m]$ 
8:     end if
9:   end for
10:   $\chi \leftarrow \mathcal{C}[-k : ]$ 
11:  return  $\pi\chi$ 
12: end function

```

---


The algorithm returns a chain  $\pi\chi$ . In this chain,  $\chi$  is the suffix of the blockchain containing the most recent  $k$  blocks.  $\pi$  is a subchain of the underlying blockchain with the last  $k$  blocks removed,  $\mathcal{C}[: -k]$ . In each iteration of the for loop, blocks of level  $\mu$  are

considered, starting from the top-most level  $|\mathcal{C}[-k].\text{interlink}|$  and descending down to level 0.

When we take a  $\mu$ -superchain and are interested in its last  $m$  blocks, we fill the same range of blocks with blocks from the superchain of level  $\mu - 1$  below. All the blocks of level  $\mu$  which are within this last  $m$  blocks range will also be superblocks of level  $\mu - 1$  and so we do not want to keep them in the proof structure twice. Note furthermore that no check is necessary to make sure the top-most level has at least  $m$  blocks, even though the verifier requires this. The reason is the following: Assume the blockchain has at least  $m$  blocks in total. Then when a superchain of level  $\mu$  has less than  $m$  blocks in total, these blocks will all be necessarily included into the proof by a lower-level superchain  $\mu - i$  for some  $i > 0$ . Therefore, it does not hurt to add them to the set  $\pi$  earlier.

Figure 2 contains an example proof constructed for parameters  $m = k = 3$ . The top superchain level which contains at least  $m$  blocks is level  $\mu = 3$ . For the  $m$ -sized suffix of that level, 5 blocks of superblock level 2 are included for support spanning the same range. For the last 3 blocks of the level 2 superchain, blocks of level 1 are included for support.

**Fig. 2.** Non-interactive proof-of-work prefix  $\pi$  for  $m = 3$  on a chain good everywhere.



### The concrete verifier

Finally, the  $\geq_m$  operator which performs the actual comparison of proofs is presented in Algorithm 4. It takes two proofs as parameters and returns true if the first proof is winning, otherwise it return false to indicate the second proof is winning. Initially the algorithm parses the sets of blocks into chains by arranging them in topological order. It rejects proofs that are not valid chains. It then computes the lowest common ancestor block  $b$  between the two proofs by calling the LCA function. Finally, it finds the highest superblock level in which at least one of the proofs has at least  $m$  blocks. If the other one doesn't, then the one who does wins. Otherwise, the one with most blocks wins. The number of blocks is counted using the findInChain function, which returns the index of a block within the respective superchain.

---

**Algorithm 4** The algorithm implementation for the  $\geq$  operator to compare two proofs in the NIPoPoW protocol parameterized with security parameter  $m$ . Returns True if the underlying chain of player  $A$  is deemed longer than the underlying chain of player  $B$

---

```

1: function best-argm( $\pi, b$ )
2:    $M \leftarrow \{\mu : |\pi \uparrow^\mu \{b : \}| \geq m\} \cup \{0\}$ 
3:   return  $\max_{\mu \in M} \{2^\mu |\pi \uparrow^\mu \{b : \}|\}$ 
4: end function
5: operator  $\pi_A \geq_m \pi_B$ 
6:    $b \leftarrow \text{LCA}(\pi_A, \pi_B)$ 
7:   return best-argm( $\pi_A, b$ )  $\geq$  best-argm( $\pi_B, b$ )
8: end operator

```

---

## 6.2 Security

We will call a query to the random oracle  $\mu$ -successful if the random oracle returns a value  $h$  such that  $h \leq 2^{-\mu}T$ .

In order to prove our construction secure, we first define the boolean random variables  $X_r^\mu$ ,  $Y_r^\mu$  and  $Z_r^\mu$ , akin to the random variables  $X_r$ ,  $Y_r$  and  $Z_r$  of the backbone setting[10]. For each round  $r$  and for each query index  $j$  and for each  $k$  adversarial party index (out of  $t$ ), define these random variables as follows. If at round  $i$  an honest party obtains a PoW with  $id < 2^{-\mu}T$ , then  $X_r^\mu = 1$ , otherwise  $X_r^\mu = 0$ . If at round  $r$  exactly one honest party obtains a PoW with  $id < 2^{-\mu}T$ , then  $Y_r^\mu = 1$ , otherwise  $Y_r^\mu = 0$ . Regarding the adversary, if at round  $r$  the  $j$ -th query of the  $k$ -th corrupted party is  $\mu$ -successful, then  $Z_{ijk}^\mu = 1$ , otherwise  $Z_{ijk}^\mu = 0$ . Further define  $Z_r^\mu = \sum_{k=1}^t \sum_{j=1}^q Z_{ijk}^\mu$ , where  $q$  denotes the total number of queries. For a set of rounds  $S$ , let  $X^\mu(S) = \sum_{r \in S} X_r$  and similarly define  $Y^\mu(S)$  and  $Z^\mu(S)$ .

We now extend the definition of a *typical execution* from the backbone setting[10].

**Definition 9.** (*Typical execution*) An execution of the protocol is  $(\epsilon, \eta)$ -typical if:

**Block counts don't deviate.** For all  $\mu \geq 0$  it holds that for any set  $S$  of consecutive rounds with  $|S| \geq 2^\mu \eta \kappa$ , the following hold:

- $(1 - \epsilon)E[X^\mu(S)] < X^\mu(S) < (1 + \epsilon)E[X^\mu(S)]$  and  $(1 - \epsilon)E[Y^\mu(S)] < Y^\mu(S)$ .
- $Z^\mu(S) < (1 + \epsilon)E[Z^\mu(S)]$ .

**Round count doesn't deviate.** Let  $S$  be a set of consecutive rounds such that  $Z^\mu(S) \geq k$  for some security parameter  $k$ . Then we have that  $|S| \geq (1 - \epsilon)2^\mu \frac{k}{pqt}$  with overwhelming probability in  $k$ .

**Chain regularity.** No insertions, no copies, and no predictions [10] have occurred.

**Theorem 3 (Typicality).** Executions are  $(\epsilon, \eta)$ -typical with overwhelming probability in  $\kappa$ .

*Proof.* **Block counts and regularity.** For the blocks count and regularity, we refer the reader to [10] for the full proof.

**Round count.** First, observe that  $Z_{ijk}^\mu \sim \text{Bern}(2^{-\mu}p)$  and these are jointly independent. Therefore  $Z_S^\mu \sim \text{Bin}(tq|S|, 2^{-\mu}p)$  and  $|S| \sim \text{NB}(Z_S, 2^{-\mu}p)$ . So  $\mathbb{E}(|S|) = 2^\mu \frac{Z_S}{pqt}$ . Applying a tail bound to the negative binomial distribution, we obtain that  $\Pr[|S| < (1 - \epsilon)\mathbb{E}(|S|)] \in \Omega(\epsilon^2 m)$ .  $\square$

The following lemma is at the heart of the security proof that will follow.

**Lemma 4.** Suppose  $S$  is a set of consecutive rounds  $r_1 \dots r_2$  and  $\mathcal{C}_B$  is a chain adopted by an honest party at round  $r_2$  of a typical execution. Let  $\mathcal{C}_B^S = \{b \in \mathcal{C}_B : b \text{ was generated during } S\}$ . Let  $\mu_A, \mu_B \in \mathbb{N}$ . Suppose  $\mathcal{C}_B^S \uparrow^{\mu_B}$  is good. Suppose  $\mathcal{C}'_A$  is a  $\mu_A$ -superchain containing only adversarially generated blocks generated during  $S$  and suppose that  $|\mathcal{C}'_A| \geq k$ . Then:

$$2^{\mu_A} |\mathcal{C}'_A| < \frac{1}{3} 2^{\mu_B} |\mathcal{C}_B^S \uparrow^{\mu_B}|$$

*Proof.* From  $|\mathcal{C}'_A| \geq k$  we know that  $Z_S \geq k$ . Applying Theorem 3, we conclude that  $|S| \geq (1 - \epsilon') 2^{\mu_A} \frac{1}{pqt} |\mathcal{C}'_A|$ .

Applying the chain growth theorem [10] we obtain that  $|\mathcal{C}_B^S| \geq (1 - \epsilon)f|S|$ . But from the goodness of  $\mathcal{C}_B^S \uparrow^{\mu_B}$ , we know that  $|\mathcal{C}_B^S \uparrow^{\mu_B}| \geq (1 - \delta) 2^{-\mu_B} |\mathcal{C}_B^S|$ . Therefore  $|\mathcal{C}_B^S \uparrow^{\mu_B}| \geq 2^{-\mu_B} (1 - \delta)(1 - \epsilon)f(1 - \epsilon') 2^{\mu_A} \frac{1}{pqt} |\mathcal{C}'_A|$  and so:

$$2^{\mu_A} |\mathcal{C}'_{\mathcal{A}}| < \frac{pqt}{(1-\delta)(1-\epsilon')(1-\epsilon)f} 2^{\mu_B} |\mathcal{C}_B^S \uparrow^{\mu_B}|$$

□

**Definition 10 (Adequate level of honest proof).** Let  $\pi$  be an honestly generated proof constructed upon some adopted chain  $\mathcal{C}$  and let  $b \in \pi$ .

Then  $\mu'$  is defined as follows:

$$\mu' = \max\{\mu : |\pi\{b : \} \uparrow^\mu| \geq \max(m+1, (1-\delta)2^{-\mu} |\pi\{b : \} \uparrow^\mu \downarrow|)\}$$

We call  $\mu'$  the adequate level of proof  $\pi$  with respect to block  $b$  with security parameters  $\delta$  and  $m$ . Note that the adequate level of a proof is a function of both the proof  $\pi$  and the chosen block  $b$ .

**Lemma 5.** Let  $\pi$  be some honest proof generated with security parameters  $\delta, m$ . Let  $\mathcal{C}$  be the underlying chain,  $b \in \mathcal{C}$  be any block and  $\mu'$  be the adequate level of the proof with respect to  $b$  and the same security parameters.

Then

$$\mathcal{C}\{b : \} \uparrow^{\mu'} = \pi\{b : \} \uparrow^{\mu'}$$

*Proof.*  $\pi\{b : \} \uparrow^{\mu'} \subseteq \mathcal{C}\{b : \} \uparrow^{\mu'}$  is trivial.

For the converse, we will show that for all  $\mu^* > \mu'$ , we have that in the iteration of the Prove for loop with  $\mu = \mu^*$ , the block stored in variable  $B$  precedes block  $b$  in  $\mathcal{C}$ .

Suppose  $\mu = \mu^*$  is the first for loop iteration during which the property is violated. Clearly this cannot be the first iteration, as there  $B = \mathcal{C}[0]$  and Genesis precedes all blocks, including itself. By the induction hypothesis we see that during the iteration  $\mu = \mu^* + 1$ , block  $B$  preceded block  $b$ . But from the definition of  $\mu'$  we know that  $\mu'$  is the highest level for which  $|\pi\{b : \} \uparrow^{\mu'} [1 :]| \geq \max(m, (1-\delta)2^{-\mu'} |\pi\{b : \} \uparrow^{\mu'} [1 :] \downarrow|)$ .

Hence, this property cannot hold for  $\mu^* > \mu'$  and therefore  $|\pi_B\{b : \} \uparrow^{\mu^*} [1 :]| < m$  or  $\neg \text{local-good}_\delta(\pi\{b : \} \uparrow^{\mu^*} [1 :], \mathcal{C}, \mu^*)$ .

In case local-good is violated, variable  $B$  remains unmodified and the induction step holds. If local-good is not violated, then  $|\pi\{b : \} \uparrow^{\mu^*} [1 :]| < m$  and so  $\pi \uparrow^{\mu^*} [-m]$  precedes  $b$ , and so we are done. □

**Lemma 6.** Suppose the verifier evaluates  $\pi_A \geq \pi_B$  in a protocol interaction where  $B$  is honest and assume during the comparison that the compared level of the honest party is  $\mu_B$ . Let  $b = \text{LCA}(\pi_A, \pi_B)$  and let  $\mu'_B$  be the adequate level of  $\pi_B$  with respect to  $b$ . Then  $\mu'_B \geq \mu_B$ .

*Proof.* Because  $\mu_B$  is the compared level of the honest party we have:

$$2^{\mu_B} |\mathcal{C}\{b : \} \uparrow^{\mu_B}| > 2^{\mu_B} |\mathcal{C}\{b : \}|$$

We will prove the claim by contradiction. Suppose  $\mu'_B < \mu_B$ . Because by definition  $\mu'_B$  is the maximum level such that  $|\pi_B\{b : \} \uparrow^{\mu'} [1 :]| \geq \max(m, (1-\delta)2^{-\mu'} |\pi_B\{b : \} \uparrow^{\mu'} [1 :] \downarrow|)$ , therefore  $\mu_B$  does not satisfy this condition. But we know that  $|\pi_B\{b : \} \uparrow^{\mu_B} [1 :]| \geq m$  because  $\mu_B$  was selected by the Verifier. So therefore,

$$2^{\mu_B} |\mathcal{C}_B\{b : \} \uparrow^{\mu_B}| < (1-\delta) |\mathcal{C}\{b : \}|$$

By the fact that  $\mu'_B$  satisfies goodness, we have that:

$$2^{\mu'_B} |\mathcal{C}_B\{b : \} \uparrow^{\mu'_B}| > (1-\delta) |\mathcal{C}\{b : \}|$$

From the last two equations, we obtain that:

$$(1 - \delta)|\mathcal{C}\{b : \}| > 2^{\mu'_B} |\mathcal{C}\{b : \} \uparrow^{\mu'_B} |$$

But the last two equations contradict, hence the claim holds.  $\square$

**Theorem 4.** *The non-interactive proofs-of-proof-of-work construction for  $k$ -stable suffix-sensitive predicates is secure with overwhelming probability in  $\kappa$ .*

*Proof.* We will prove security by contradiction. Let  $m = k_1 + k_2 + k_3$  and let  $k_1, k_2, k_3$  be polynomial functions of the security parameter  $\kappa$ . Let  $Q$  be a  $k$ -stable suffix sensitive chain predicate. Assume the NIPoPoWs construction with parameters  $m, k$  is not secure with respect to  $Q$ . Then, during an execution at some round  $r_3$ ,  $Q(\mathcal{C})$  is defined and has the same value for all honest miners. Assume the execution is typical. The verifier  $V$  communicates with at least two provers,  $\mathcal{A}$ , the adversary, and  $B$ , an honest prover. The verifier receives the proofs  $\pi_{\mathcal{A}}, \pi_B$  from the adversary and the honest prover respectively. Because  $B$  is honest,  $\pi_B$  will be a proof constructed based on an underlying blockchain  $\mathcal{C}_B$  (with  $\pi_B \subseteq \mathcal{C}_B$ ) which is the adopted blockchain of  $B$  during round  $r_3$  at which  $\pi_B$  was generated. Furthermore,  $\pi_{\mathcal{A}}$  will have been generated at some round  $r'_3 \leq r_3$ .

The verifier then outputs  $\neg Q(\mathcal{C}_B)$ , and so  $\text{Verify}_{m,k}^Q = \neg Q(\mathcal{C}_B)$ . Thus it is necessary that  $\pi_{\mathcal{A}} \geq \pi_B$ , otherwise, because  $Q$  is suffix sensitive,  $\text{Verify}^Q$  would have returned  $Q(\mathcal{C}_B)$ . We will now show that  $\pi_{\mathcal{A}} \geq \pi_B$  is a negligible event.

Let  $b = \text{LCA}(\pi_{\mathcal{A}}, \pi_B)$  and let  $b^*$  be the most recent honestly generated block in  $\mathcal{C}_B$  preceding  $b$  (and note that  $b^*$  necessarily exists because Genesis was honestly generated). Let the levels of comparison decided by the verifier be  $\mu_{\mathcal{A}}$  and  $\mu_B$  respectively.

Let  $\mu'_B$  be the adequate level of proof  $\pi_B$  with respect to block  $b$ . Call  $\alpha_{\mathcal{A}} = \pi_{\mathcal{A}} \uparrow^{\mu_{\mathcal{A}}} \{b : \}$ ,  $\alpha'_B = \pi_B \uparrow^{\mu'_B} \{b : \}$ .

We will proceed in the proof by showing three successive claims: First,  $\alpha_{\mathcal{A}}$  and  $\alpha'_B \downarrow$  are mostly disjoint. Second,  $\alpha_{\mathcal{A}}$  contains mostly adversarially-generated blocks. And third, the adversary is able to produce this  $\alpha_{\mathcal{A}}$  with negligible probability.

**Claim 1a:** If  $\mu'_B \leq \mu_{\mathcal{A}}$  then  $\alpha_{\mathcal{A}}[1 : ]$  and  $\alpha_B[1 : ] \downarrow$  are completely disjoint.

Applying Lemma 5 to  $\mathcal{C}_B \{b : \} \uparrow^{\mu'_B}$  we see that  $\mathcal{C}_B \{b : \} \uparrow^{\mu'_B} = \pi_B \uparrow^{\mu'_B} \{b : \}$  and so indeed  $\pi_B \uparrow^{\mu'_B} \{b : \}[1 : ] \cap \pi_{\mathcal{A}} \uparrow^{\mu_{\mathcal{A}}} \{b : \}[1 : ] = \emptyset$ .

**Claim 1b:** If  $\mu_{\mathcal{A}} < \mu'_B$  then  $|\alpha_{\mathcal{A}}[1 : ] \cap \alpha_B \downarrow [1 : ]| \leq 2^{\mu'_B - \mu_{\mathcal{A}}} k_1$ .

First, observe that, because the adversary is winning, therefore  $|\alpha_{\mathcal{A}}| > 2^{\mu'_B - \mu_{\mathcal{A}}} m$ .


Suppose for contradiction that  $|\alpha_{\mathcal{A}}[1 : ] \cap \alpha_B \downarrow [1 : ]| > 2^{\mu'_B - \mu_{\mathcal{A}}} k_1$ . This means that there are indices  $1 \leq i < j$  such that  $|\mathcal{C}_B \uparrow^{\mu_{\mathcal{A}}} [i : j]| > 2^{\mu'_B - \mu_{\mathcal{A}}} k_1$  but  $|\mathcal{C}_B \uparrow^{\mu_{\mathcal{A}}} [i : j] \downarrow \uparrow^{\mu'_B}| = 0$ . But this contradicts the goodness of  $\mathcal{C}_B \uparrow^{\mu'_B}$ . Therefore there are more than  $2^{\mu'_B - \mu_{\mathcal{A}}} (k_2 + k_3)$  blocks in  $\alpha_{\mathcal{A}}$  that are not in  $\alpha_B$ , and clearly also more than  $k_2 + k_3$  blocks.

From Claim 1a and Claim 1b, we conclude that there are at least  $k_2 + k_3$  blocks after block  $b$  in  $\alpha_{\mathcal{A}}$  which do not exist in  $\alpha_B$ . We now set  $b_2 = \text{LCA}(\mathcal{C}_B, \alpha_{\mathcal{A}})$ .

**Claim 2:** At least  $k_3$  superblocks of  $\alpha_{\mathcal{A}}$  are adversarially generated.

We will show this by showing that  $\alpha_{\mathcal{A}}[k_2 + 1 : ]$  contains no honestly mined blocks. By contradiction, assume that the block  $\alpha_{\mathcal{A}}[i]$  for some  $i \geq k_1 + k_2 + 1$  was honestly generated. But this means that an honest party had adopted the chain  $\alpha_{\mathcal{A}}[i - 1]$  at some round  $r_2 \leq r_3$ . Because of the way the honest parties adopt chains, this means that the superchain  $\alpha_{\mathcal{A}}[: i - 1]$  has an underlying properly constructed 0-level anchored chain  $\mathcal{C}_{\mathcal{A}}$  such that  $\mathcal{C}_{\mathcal{A}} \subseteq \alpha_{\mathcal{A}}[: i - 1]$ . Let  $j$  be the index of block  $b_2$  within  $\mathcal{C}_{\mathcal{A}}$ . As  $\alpha_{\mathcal{A}} \subseteq \mathcal{C}_{\mathcal{A}}$ , observe that  $|\mathcal{C}_{\mathcal{A}}[j + 1 : ]| > i - 1 \geq k_2 + k_1$ . Therefore  $\mathcal{C}_{\mathcal{A}}[: -(k_2 + k_1)] \not\subseteq \mathcal{C}_B$ . But  $\mathcal{C}_{\mathcal{A}}$  was adopted by an honest party at round  $r_2$  which is prior to round  $r_3$  during which  $\mathcal{C}_B$  was adopted

**Fig. 3.** Two competing proofs at different levels.



by an honest party. This contradicts the Common Prefix [10] property with parameter  $k_2$ . It follows that with overwhelming probability in  $k_2$ , the  $k_3 = m - k_2 - k_1$  last blocks of the adversarial proof have been adversarially mined.

**Claim 3:**  $\mathcal{A}$  is able to produce a  $\alpha_{\mathcal{A}}$  that wins against  $\alpha_B$  with negligible probability.

Let  $b'$  be the latest honestly generated block in  $\alpha_{\mathcal{A}}$ , or  $b^*$  if no honest block exists in  $\alpha_{\mathcal{A}}$ . Let  $r_1$  be the round during which  $b^*$  was generated. Let  $j$  be the index of block  $b^*$ . Consider now the set  $S$  of consecutive rounds  $r_1 \dots r_3$ . Every block in  $\alpha_{\mathcal{A}}[-k_3 : ]$  has been adversarially generated during  $S$  and  $|\alpha_{\mathcal{A}}[-k_3 : ]| = k_3$ .  $\mathcal{C}_B$  is a chain adopted by an honest party at round  $r_3$  and filtering the blocks by the rounds during which they were generated to obtain  $\mathcal{C}_B^S$ , we see that  $\mathcal{C}_B^S = \mathcal{C}_B\{b^* : \}$ . But chain  $\mathcal{C}_B^S \uparrow \mu'_B$  is good with respect to  $\mathcal{C}_B^S$ . Applying Lemma 4, we obtain that with overwhelming probability

$$2^{\mu_{\mathcal{A}}} |\alpha_{\mathcal{A}}\{b' : \}| < \frac{1}{3} 2^{\mu'_B} |\mathcal{C}_B^S \uparrow \mu'_B|$$

But  $|\alpha_B| \geq |\mathcal{C}_B^S \uparrow \mu'_B|$  and  $|\alpha_{\mathcal{A}}\{b' : \}| \geq |\alpha_{\mathcal{A}}| - k_2$ , therefore:

$$2^{\mu_{\mathcal{A}}} |\alpha_{\mathcal{A}}| - k_2 < \frac{1}{3} 2^{\mu'_B} |\alpha_B|$$

But  $|\alpha_{\mathcal{A}}| - k_2 \geq k_3$ , therefore  $\frac{1}{3} 2^{\mu'_B} |\alpha_B| > k_3$  and so  $2^{\mu'_B} |\alpha_B| > 3k_3$

Taking  $k_2 = k_3$ , we obtain:

$$2^{\mu_{\mathcal{A}}} |\alpha_{\mathcal{A}}| < \frac{1}{3} 3k_3 + k_3 = 2k_3 < 2^{\mu'_B} |\alpha_B|$$

But this contradicts the fact that  $\pi_{\mathcal{A}} \geq \pi_B$ , and so the claim is proven.

Therefore we have proven that  $2^{\mu'_B} |\pi_B \uparrow \mu'_B| > 2^{\mu_{\mathcal{A}}} |\pi_{\mathcal{A}}^{\mu_{\mathcal{A}}}|$ .

From the definition of  $\mu_B$ , we know that  $2^{\mu_B} |\pi_B \uparrow \mu_B| \geq 2^{\mu'_B} |\pi_B \uparrow \mu'_B|$ , and therefore we conclude that  $2^{\mu_B} |\pi_B \uparrow \mu_B| > 2^{\mu_{\mathcal{A}}} |\pi_{\mathcal{A}} \uparrow \mu_{\mathcal{A}}|$ , which completes the proof.  $\square$

*Remark 1 (Variance attacks).* The critical issue addressed by this security proof is to avoid Bahack-style attack [25] where the adversary constructs “lucky” high-difficulty superblocks without filling in the underlying proof-of-work in the lower levels. Observe that, while setting  $m = 1$  “preserves” the proof-of-work in the sense that expectations remain the same, the probability of an adversarial attack becomes approximately proportional to the adversary power if the adversary follows a suitable strategy (for a description of such a strategy, see the parametrization section). With higher values of  $m$ , the probability of an adversarial attack drops exponentially, even though they maintain constant computational power, and hence satisfy a strong notion of security.

### 6.3 Succinctness

We now prove that our construction produces succinct proofs.

Observe that full succinctness in the standard honest majority model is impossible to prove for our construction. To see why, recall that an adversary with sufficiently large mining power can significantly harm superquality as described in Subsection 5.2. By causing this harm in superquality of a sufficiently low level  $\mu$ , for example  $\mu = 3$ , the adversary can cause the honest prover to digress in their proofs from high-level superchains down to low-level superchains, causing a linear proof to be produced. For instance, if superquality is harmed at  $\mu = 3$ , the prover will need to digress down to level  $\mu = 2$  and include the whole 2-superchain, which, in expectation, will be of size  $|\mathcal{C}|/2$ .

Having established security in the general case of the standard honest majority model, we therefore concentrate our succinctness claims to the particular case where there is no adversary. This case, called the “optimistic succinctness” case following the lead of [11], models the system as follows. First, the adversary has no mining power. In other words, we set  $t = 0$ . This establishes that all honestly adopted chains contain only honestly generated blocks. Second, the adversary is unable to perform adversarial network scheduling. Instead, network scheduling is random. What this means is that messages diffused by honest parties are reordered in a random permutation before they are delivered to other honest parties, but the adversary has no saying as to this order. And last, the adversary is unable to produce any proofs; that is, all proofs consumed by the verifier are honestly-generated. We will now explore why proofs produced in these circumstances are succinct.

**Theorem 5 (Number of levels).** *The number of superblock levels which have at least  $m$  blocks are at most  $\log(|S|)$ , where  $S$  is the set of all blocks produced, with overwhelming probability in  $m$ .*

*Proof.* Let  $S$  be the set of all blocks successfully produced by the honest parties or the adversary. Because each block id is generated by the random oracle, the probability that it is less than  $T2^{-\mu}$  is  $2^{-\mu}$ . These are independent Bernoulli trials. For each block  $B \in S$ , define  $X_B^\mu \in \{0, 1\}$  to be the random variable indicating whether the block belongs to superblock level  $\mu$  and let  $D_\mu$  indicate their sum, which is a Binomial distribution with parameters  $(|S|, 2^{-\mu})$  and expected value  $E[D_\mu] = |S|2^{-\mu}$ .

For level  $\mu$  to exist in any valid proof, at least  $m$  blocks of level  $\mu$  must have been produced by the honest parties or the adversary. We will now show that  $m$  blocks of level  $\mu = \log(|S|)$  are produced with negligible probability in  $m$ .

As all of the  $X^\mu$  are independent, we can apply a Binomial Chernoff bound to the probability of the sum. Therefore we have

$\Pr[D_\mu \geq (1 + \Delta)E[D_\mu]] \leq \exp(-\frac{\Delta^2}{2+\Delta}E[D_\mu])$ . But for this  $\mu$  we have that  $E[D_\mu] = 1$ . Therefore  $\Pr[D_\mu \geq 1 + \Delta] \leq \exp(-\frac{\Delta^2}{2+\Delta})$ . Requiring  $1 + \Delta = m$ , we get  $\Pr[D_\mu \geq m] \leq \exp(-\frac{(m-1)^2}{m+1})$ , which is negligible in  $m$ .  $\square$

The above theorem establishes that the number of superchains is small. What remains to be shown is that few blocks will be included at each superchain level.

**Theorem 6 (Large upchain expansion).** *Let  $\mathcal{C}$  be an honestly generated chain and let  $\mathcal{C}' = \mathcal{C}^{\uparrow^{\mu-1}} [i : i + \ell]$  with  $\ell \geq 4m$ . Then:*

$$|\mathcal{C}'^{\uparrow^{\mu}}| \geq m$$

*with overwhelming probability in  $m$ .*

*Proof.* Assume the  $(\mu - 1)$ -level superchain has  $4m$  blocks. Because each block of level  $\mu - 1$  was generated as a query to the random oracle, it constitutes an independent Bernoulli trial and the number of blocks in level  $\mu$ , namely  $\pi^{\uparrow^{\mu}}$ , is a Binomial distribution with parameters  $(4m, 1/2)$ . Clearly  $\Pr[|\pi^{\uparrow^{\mu}}| = m] \leq \Pr[|\pi^{\uparrow^{\mu}}| \leq m]$ . Observing that  $E[\pi^{\uparrow^{\mu}}] = 2m$  and applying a Chernoff bound, we get  $\Pr[|\pi^{\uparrow^{\mu}}| \leq (1 - \frac{1}{2})2m] \leq \exp(-\frac{(1/2)^2}{2}2m)$  which is negligible in  $m$ .

This probability bounds the probability of fewer than  $m$  blocks occurring in the  $\mu$  level restriction of  $(\mu - 1)$ -level superchains of more than  $4m$  blocks.  $\square$

**Corollary 2 (Small downchain support).** *Assume an honestly generated chain  $\mathcal{C}$  and let  $\mathcal{C}' = \mathcal{C}^{\uparrow^{\mu}} [i : i + m]$ . Then:*

$$|\mathcal{C}'^{\downarrow^{\mu-1}}| \leq 4m \tag{1}$$

*with overwhelming probability in  $m$ .*

*Proof.* Assume the  $(\mu - 1)$ -level superchain had at least  $4m$  blocks. Then by Theorem 6 it follows that more than  $m$  blocks exist in level  $\mu$  with overwhelming probability in  $m$ , which is a contradiction.  $\square$

This last theorem establishes the fact that the support of blocks needed under the  $m$ -sized chain suffix to move from one level to the one below is small. Based on this, we can establish our theorem on succinctness:

**Theorem 7 (Optimistic succinctness).** *Non-interactive proofs-of-proof-of-work produced by honest provers on honestly generated chains with random scheduling are succinct with the number of blocks bounded by  $4m \log(|\mathcal{C}|)$ , with overwhelming probability in  $m$ .*

*Proof.* Assume  $\mathcal{C}$  is the honest parties chain. From Theorem 5, the number of levels in the NIPoPoW is at most  $\log(|\mathcal{C}|)$  with overwhelming probability in  $m$ .

First, observe that the count of blocks in the highest level will be less than  $4m$  from Theorem 6; otherwise a higher superblock level would exist.

From Corollary 1, we know that at all levels  $\mu$  the chain will be good. Therefore, for each  $\mu$  superchain  $\mathcal{C}$  the supporting  $(\mu - 1)$ -superchain will only need to span the  $m$ -long suffix of the  $\mu$ -superchain above.

Then for the  $m$ -long suffix of each superchain of level  $\mu$ , the supporting superchain of level  $\mu - 1$  will have at most  $4m$  blocks from Corollary 2.

Therefore the size of the proof is  $4m \log(|\mathcal{C}|)$ , which is succinct.  $\square$

*Remark 2 (Denial of Service Attacks).* Examine the non-optimistic case where the adversary is allowed to provide proofs, yet, as in the optimistic case, they have no mining power or adversarial scheduling capabilities. It is then still possible for that adversary to produce



large incorrect proofs that use up the resources of a verifier with logarithmic time. Therefore, if such attacks are of concern, the verifier can relax the security guarantees to achieve succinctness in this model as follows: Whenever a large proof is provided by a prover, the verifier terminates early with output  $\perp$  regardless of which proof is better, indicating that they are unable to reach a conclusion. It is important to note that a verifier knowing ahead of time that a proof will be large cannot conclude that the prover providing it is malicious, as an attack on superquality of the blockchain could be taking place, requiring an honest prover to provide long proofs.

## 7 Blockchain Infix proofs

### 7.1 Construction

While the previous proofs we constructed were proving predicates on the blocks and transactions pertaining to the chain suffix of the last  $k$  blocks, similar to the kind of proofs possible with [11], we note that the most useful class of proofs would allow proving more general predicates that can depend on multiple blocks, sometimes even blocks buried deep within the blockchain.

We will now extend our prover to support this generalization of predicates. The generalized prover for *infix proofs* allows proving any predicate  $Q(\mathcal{C})$  that depends on a number of blocks that can appear anywhere within the chain. These blocks constitute a *subchain*  $\mathcal{C}'$  which is a subsequence of the original blockchain. For our proofs to work, the blocks contained within the subchain have to be  $k$ -stable; that is, they cannot be one of the most recent  $k$  blocks.

The proofs are able to prove any predicate of the general class of predicates that depend on both the inclusion of the blocks of the subchain as well as their relative order within the subchain. This allows proving powerful statements such as, for example, whether a transaction took place at any point in history, or even “following the money” and proving a claim that a certain coin was moved between a series of addresses in a particular order. While, as illustrated in the previous section, it is possible to build an SPV client which is aware of the latest blocks using suffix proofs, a client which needs to learn about a transaction buried deep in the blockchain cannot expect suffix proofs to be efficient. Hence, the infix proofs construction is necessary.

The manner to construct these proofs is shown in Algorithm 6. The infix prover accepts two parameters: The chain  $\mathcal{C}$  which is the full blockchain and  $\mathcal{C}'$  which is a subchain of the blockchain whose blocks are of interest for the predicate in question. This prover simply calls the previous suffix prover to produce a proof as usual. Then, having the prefix  $\pi$  and suffix  $\chi$  of the suffix proof in hand, the infix prover adds a few auxiliary blocks to the prefix  $\pi$ . The prover ensures that these auxiliary blocks form a verifiable chain with the rest of the proof  $\pi$  in a manner that the blocks can be arranged in topological order such that a pointer exists from each non-Genesis block to a preceding block and each preceding block is pointed to by a unique block. Such auxiliary blocks are collected as follows: For every block  $B$  of the subchain  $\mathcal{C}'$ , the immediate previous ( $E'$ ) and next ( $E$ ) blocks in the proof  $\pi$  are found. Then, a chain of blocks  $R$  which connects  $E$  back to  $B'$  is found by the algorithm `followDown`. Finally, observe that if  $E'$  is of level  $\mu$ , then there can be no other  $\mu$  level superblock between  $E'$  and  $B'$ , otherwise it would have been included in  $\pi$ . Therefore, block  $B'$  already contains a pointer to  $E'$  in its interlink, completing the chain.

---

**Algorithm 5** The followDown function which produces the necessary blocks to connect a superblock  $hi$  to a preceding regular block  $lo$ , enabling infix proofs.

---

```

1: function followDown( $hi, lo, \text{depth}$ )
2:    $aux \leftarrow [hi]$ 
3:    $B \leftarrow hi$ 
4:    $aux \leftarrow []$ 
5:    $\mu \leftarrow \text{level}(hi)$ 
6:   while  $B \neq lo$  do
7:      $B' \leftarrow \text{blockById}[B.\text{interlink}[\mu]]$ 
8:     if  $\text{depth}[B'] < \text{depth}[lo]$  then
9:        $\mu \leftarrow \mu - 1$ 
10:    else
11:       $aux \leftarrow aux \cup aux'$ 
12:       $B \leftarrow B'$ 
13:    end if
14:  end while
15:  return  $aux$ 
16: end function

```

---

The way to connect a superblock to a previous lower-level block is implemented in Algorithm 6 as the followDown algorithm. Notice that block  $B'$  cannot be of higher or equal level than block  $E$ , otherwise it would be equal to  $E$  and the followDown algorithm would return immediately. The algorithm then proceeds as follows: Starting at block  $E$  which here is marked as  $hi$ , it tries to follow a pointer to as far as possible. If following the pointer surpasses block  $B'$ , which here is marked as  $lo$ , then the following is aborted and a lower level is tried, which will cause a smaller step within the skiplist. If a pointer was followed without surpassing  $B'$ , the operation continues from the new block, until eventually  $B'$  will be reached, which concludes the algorithm.

---

**Algorithm 6** The Prove algorithm for infix proofs

---

```


1: function ProveInfix $_{m,k}(\mathcal{C}, \mathcal{C}', \text{depth})$ 
2:    $(\pi, \chi) \leftarrow \text{Prove}_{m,k}(\mathcal{C})$ 
3:   for  $B' \in \mathcal{C}'$  do
4:     for  $E \in \pi$  do
5:       if  $\text{depth}[E] \geq \text{depth}[B']$  then
6:          $R \leftarrow \text{followDown}(E, B', \text{depth})$ 
7:          $aux \leftarrow aux \cup R$ 
8:         break
9:       end if
10:       $E' \leftarrow E$ 
11:    end for
12:  end for
13:  return  $(aux \cup \pi, \chi)$ 
14: end function

```

---

An example of the kind of auxiliary sequences that followDown produces is seen in Figure 4. This is only a portion of the proof shown at the point where the superblock levels are at level 4, but a descend to level 0 was necessary so that a regular block would be included in the chain. Note that the level 0 block can jump immediately back up to level 4 because it has a high-level pointer.

**Fig. 4.** An infix proof descending down to level 0 and back up to level 4. Only blue blocks are included in the proof. Blue blocks of level 4 are part of  $\pi$ , while the blue blocks of level 1 and 3 are produced by followDown to get to the block of level 0 which is part of  $\mathcal{C}'$ . Notice how in this proof the blocks happened to be distributed slightly different from expectation and there are two consecutive 0-level blocks.



## 7.2 Security

The security of infix proofs follows immediately from Theorem 4. To see this, notice that the verifier will still verify the suffix proof. Upon determining which of the blockchains is the longest securely, it will subsequently verify that the subchain blocks have been properly included within the referenced chain. However, if the chain is honest, this proof of inclusion cannot be produced by the adversary.

## 7.3 Succinctness

As long as the number of blocks on which the predicate depends is polylogarithmic with respect to the chain length, our proofs remain succinct. Specifically, the proof size for the suffix has exactly the same size. Then the part of the proof that is of interest is the output of the followDown algorithm. However, notice that this algorithm will on average produce as many blocks as the difference of levels between  $B'$  and  $E$ , which is at most logarithmic in the chain size. Hence the proof sizes will be in expectation  $(m + |\mathcal{C}'|) \log(|\mathcal{C}|)$ , which remains succinct if  $|\mathcal{C}'| \in O(\text{polylog}(|\mathcal{C}|))$ .

# 8 Gradual Deployment Paths

Our construction requires an upgrade to the cryptocurrency consensus layer. We envision that new cryptocurrencies will adopt our construction in order to support efficient light clients.

However, existing cryptocurrencies could also benefit by adopting our construction as an upgrade. In this section we outline several possible upgrade paths. We also contribute a novel upgrade approach, called a “velvet fork,” which allows for gradual deployment without harming unupgraded nodes or miners at all.

## 8.1 Hard Forks and Soft Forks

First of all, keep in mind that the interlink data structure will be useful to be included in the form of a Merkle tree of interlink pointers, as discussed in the next section on implementation optimizations.

The most obvious way to upgrade a cryptocurrency to support our protocol is a hard fork: the block header could be modified to include our additional interlink structure, and the validation rules would be modified to require that new blocks (after a “flag day”) contain a correctly-formed interlink hash.

The safety of the hard fork approach is often debated: it is not “forward compatible,” in the sense that nodes failing to upgrade in time would not be able to process newly created blocks. Hard forks are also associated with “schisms”: in one notable case, Ethereum applied a hard fork in order to reverse the effects of an exploited vulnerability; this decision was controversial, and dissenters splintered off into a new network, Ethereum Classic [26]. However, both Ethereum and Ethereum Classic have subsequently applied hard fork upgrades without incident. The Monero cryptocurrency has applied hard forks on a predictable schedule [27].

A soft fork construction in practice would require including the interlink data structure not in the block header, but in the data of the coinbase transaction. Clearly it is enough to store a single hash of the whole interlink structure in the coinbase data. The only requirement for the PoPoWs to work is that the PoW commits to all the pointers within interlink so that the adversary cannot cause a chain reorganization except with negligible probability. That way, each superchain is really a chain which goes all the way back to genesis. If we take that route, then each PoPoW will be required to present not only the block header, but also a proof-of-inclusion path within the Merkle tree of transactions proving that the coinbase transaction is indeed part of the block. Once that is established, the coinbase data can be presented, and the verifier will thereby know that the hash of the interlink data structure is correct. Given the fact that in the current Bitcoin implementation there is a fixed limit for block sizes, we observe that including such proofs-of-inclusion will only increase the PoPoW sizes by a constant factor per block, allowing for the communication complexity to remain at  $\Theta(\text{polylog}(|\mathcal{C}|))$ .

## 8.2 Velvet Forks

We now describe a novel upgrade path that avoids the need for a fork at all. The key idea is that clients can make use of our scheme, even if only some blocks in the blockchain include the interlink structure. Given that intuitively these changes constitute rule modifications to the consensus layer, we call this technique a *velvet fork*.

The way to achieve this is by requiring upgraded miners to include the interlink data structure in the form of a new Merkle tree root hash in their coinbase data, similar to a soft fork. An unupgraded miner will as usual ignore this data as comments. However, we further require the upgraded miners to accept all previously accepted blocks, regardless of whether they have included the interlink data structure or not. In fact, even if the interlink data structure is included and contains invalid data, we require the upgraded miners to accept their containing blocks. Malformed interlink data could be simply of the wrong format, for example pointers of particular level appearing in the wrong position in the Merkle tree, or the pointers could be pointing to superblocks of incorrect levels. Furthermore, the pointers could be pointing to superblocks of the correct level, but not to the most recent block. By requiring upgraded miners to accept all such blocks, we do not modify the set of accepted blocks: Upgraded and unupgraded miners accept the exact same set of blocks. Therefore, the upgrade is simply a “recommendation” for miners and not an actual change in the consensus rules. Hence, while a hard fork makes new upgraded blocks invalid to unupgraded clients and a soft fork makes new unupgraded blocks invalid to upgraded clients, the velvet fork has the effect that blocks produced by either upgraded or unupgraded clients are valid for either. Hence, the blockchain is never forked. Only the codebase is upgraded, and the existing data on the blockchain is interpreted differently.

The reason this can work is because provers and verifiers of our protocol can check the validity of the claims of miners who make false interlink chain claims. An upgraded prover can check whether a block contains correct interlink data and use it. If a block does not

contain correct interlink data, the prover can opt not to use those pointers at all in their proofs. As the Verifier verifies all claims of the prover, adversarial miners cannot cause harm by including invalid interlink data. The only thing that the Verifier cannot verify in terms of interlink claims is whether the claimed superblock of a given level is in fact the most recent previous superblock of that level. However, an adversarial prover cannot make use of that to construct winning proofs, as they are only able to present shorter chains in that case. The honest prover can simply ignore such pointers as if they were not included at all.

The velvet prover works as usual, but additionally maintains a *realLink* data structure as part of its full node, which stores the correct interlink for each block. Whenever a new winning chain is received from the network, the prover checks it for blocks that it hasn't seen before. For those blocks, it maintains its own *realLink* data structure which it updates accordingly to make sure it is correct regardless of what the interlink data structure of the received block claims.

The velvet `constructInnerChain` is implemented identically as before, except that instead of following the interlink pointer blindly it now calls the helper function called *followUp*. The `followUp` algorithm is shown in Algorithm 7. It accepts a block  $B$  and a level  $\mu$  and creates a connection from  $B$  back to the most recent preceding level- $\mu$  block. It does this by following the interlink pointer if it is available and correct. Otherwise, it follows the `previd` link which is available in unupgraded blocks, and tries to follow the interlink pointer again from there. Finally, the velvet prover simply calls the `constructInnerChain` function and includes the auxiliary connecting nodes within the final proof. No changes in the verifier are needed.

---

**Algorithm 7** The `followUp` function which produces the necessary blocks to connect two superblocks in velvet forks.

---

```

1: function followUp( $B, \mu, \text{realLink}, \text{blockById}$ )
2:    $\text{aux} \leftarrow [B]$ 
3:   while  $B \neq \text{Gen}$  do
4:     if  $B.\text{interlink}[\mu] = \text{realLink}[\text{id}(B)][\mu]$  then
5:        $\text{id} \leftarrow B.\text{interlink}[\mu]$ 
6:     else ▷ Invalid interlink
7:        $\text{id} \leftarrow B.\text{interlink}[0]$ 
8:     end if
9:      $B \leftarrow \text{blockById}[\text{id}]$ 
10:     $\text{aux} \leftarrow \text{aux} \cup B$ 
11:    if  $\text{level}(B) = \mu$  then
12:      return  $B, \text{aux}$ 
13:    end if
14:  end while
15:  return  $B, \text{aux}$ 
16: end function

```

---

It is straight-forward to see that velvet NIPoPoWs preserve security. But observe that if a constant minority of miners has upgraded their nodes to add NIPoPoW support, then succinctness is also preserved, as there is only a constant factor penalty as proven in the following theorem.

**Theorem 8.** *Velvet non-interactive proofs-of-proof-of-work on honest chains by honest provers remain succinct as long as a constant percentage  $g$  of miners has upgraded, with overwhelming probability.*

*Proof.* From Theorem 7 we know that the proofs  $\pi$  contain only a  $O(\text{polylog}(m))$  amount of blocks. For each of these blocks, the velvet client needs to include a followUp tail of blocks. Assume a percentage  $0 < g \leq 1$  of miners have upgraded with NIPoPoW support. Then the question of whether each block in the honest chain is upgraded follows a Bernoulli distribution. If the velvet proof were to be larger than  $\Delta$  times the soft fork proof in the number of blocks included, then this would require at least one of the followUp tails to include at least  $\Delta$  sequential unupgraded blocks. But since the upgrade status of each block is independent, the probability of this occurring is  $g^\Delta$ , which is negligible in  $\Delta$ .  $\square$

The reason we were able to upgrade using a velvet fork was because the changes we made were helpful but verifiable by those looking at the chain. Note that we would not have been able to pull off this upgrade without modifications to the consensus layer in the sense that the interlink data structure could not have been maintained somewhere independently of the blockchain: It is critical that the proof-of-work commits to the interlink data structure. However, interestingly, the interlink data structure does not even need to be part of the coinbase transaction and could even be produced and included in regular transactions by users (such as OP\_RETURN transactions). In that manner, the miners can be completely oblivious to it, while users and provers make use of the feature. The way to realize this would be to ask users to regularly create transactions containing the most recent interlink pointers so that they are included in the next block. If the transaction makes it to the next block, it can be used by the prover who keeps track of these. Otherwise, if it becomes part of a subsequent block, in which case some of the pointers it contains are invalid, it can be simply be ignored or only partially used.

*Supporting clients with different beliefs.* We note that the interlink format does not depend at all on the parameters  $m$  and  $k$ . Therefore, regardless of the upgrade procedure, it is not necessary to reach consensus agreement on a particular choice of these parameters. Instead, the choice of  $m$  and  $k$  can be exposed as a user-configurable parameter to clients.

## 9 Implementation & Parameters

**Table 1.** Estimated size of NIPoPoWs, applied to the Bitcoin blockchain today ( $\approx 450\text{k}$  blocks) for various values of the security parameter  $m$  setting  $k = 6$ .

$m$	NIPoPoW size (KB)	# of blocks	# of hashes
6	70	108	1440
15	146	231	2925
30	270	426	5400
50	412	656	8250
100	750	1206	15000
127	952	1530	19050

*Interlink Optimizations* We now discuss the exact manner in which the interlink data structure should be hashed to provide a single hash within the coinbase transaction. A Merkle tree should be used to hash the interlink data structure into a single hash [11]. Observe that NIPoPoWs form a single chain of various levels which can skip certain blocks. However, it does not form a tree. Therefore, for each block included in the proof, only a single pointer needs to be presented to convince the Verifier. Near genesis, the pointers that need to be included correspond to high levels; while near the most recently generated

block, the pointers correspond to low levels. One important observation is this: In the formal construction of the proof, the highest level superchain with at least  $m$  blocks is included, and assume it is of level  $\mu$ . The next level superchain,  $\mu - 1$  is completely included and has an expected number of  $2m$  blocks. But note that since all  $\mu$ -level superblocks are also  $(\mu - 1)$ -level superblocks, only the  $\mu - 1$  level needs to be in fact included. Among the expected  $2m$  blocks of level  $\mu - 1$ , the last  $m$  will need to be supported by level  $\mu - 2$ . Using the same argument, since  $(\mu - 1)$ -level superblocks are also  $(\mu - 2)$  level superblocks, in expectation only  $m$  blocks need to be included in level  $\mu - 1$ . The argument continues inductively, until an expected number of  $2m$  blocks needs to be included at level 0 immediately before the  $\chi$  suffix. This gives us an exact estimation on the size of the proof: The proof will contain exactly  $m(\log(|\mathcal{C}|) - \log(m))$  blocks in expectation,  $m$  at each of the  $\mu - 1$  levels and  $2m$  at level 0.

By organizing the interlink pointers into a Merkle tree of  $\log(|\mathcal{C}|)$  items, a proof-of-inclusion in the pointers Merkle tree can be provided in  $\log \log(|\mathcal{C}|)$  space, noting that the 0-level pointer need not be included in this Merkle tree, but the genesis block needs to be. Then the root of the pointers Merkle tree can be proved to have been included in the block header in  $\log(|\bar{x}|)$  using the standard Merkle tree of transactions, where  $\bar{x}$  is the vector of all transactions in the given block. This makes the proof size require  $\log(|\bar{x}|) + \log \log(|\mathcal{C}|)$  hashes per block for a total of  $m(\log(|\mathcal{C}|) - \log(m))(\log(|\bar{x}|) + \log \log(|\mathcal{C}|))$  hashes. In addition,  $m(\log(|\mathcal{C}|) - \log(m))$  block headers and coinbase transactions are needed. More concretely, given that currently  $|\mathcal{C}| = 464,185$ , we have that  $\log(|\mathcal{C}|) = 18$  and  $\log \log(|\mathcal{C}|) = 5$ . Typically,  $|\bar{x}| = 2000$  in the current bitcoin setting, which makes  $\log(|\bar{x}|) = 11$ . For the  $k$ -suffix, only  $k$  block headers need to be included. We set  $k = 6$  as is traditional in the bitcoin setting and see that block headers are 80 bytes and hash outputs are 32 bytes. However, we note that block headers do not need to include the previous block header hash for the  $k$ -suffix, as each previous block header hash can be calculated by the Verifier, limiting header sizes to 48 bytes for the  $k$  last blocks. Furthermore, the  $k$ -suffix does not require the presentation of the coinbase transaction data. These are also true for the  $2m$  blocks in  $\overline{II}[0]$ . Note also that for the proof prefix  $\overline{II}$ , the coinbase transaction hash need not be included: It can be evaluated by the Verifier when building the Merkle tree during the proof-of-inclusion. Similarly, the root of the pointers Merkle tree can be omitted from the coinbase transaction data when transmitting the proof. Additionally, the leaf of the pointers Merkle tree can also be omitted, as it can be obtained by the Verifier by simply hashing the previous block in the proof. In fact, no block ids need to be transmitted at all, as they can all be built from the Verifier starting from the knowledge of Genesis.

From these observations, we estimate our scheme's proof sizes for various parameterizations of  $m$  in Table 1.

## 9.1 Choosing Secure Parameters


In order to determine concrete values for the security parameter  $m$ , we now turn our attention to a particular adversarial strategy of interest and analyze its probability of success under specific conditions. It is important to keep in mind that this adversarial strategy is only one possible strategy and an advanced adversary can perform more sophisticated attacks. However, the fact that this attack is reasonable and possible to simulate allows us to extract specific values for  $m$ . The attack is an extension of the stochastic processes originally described in [1] and further explored in [28].

The experiment works as follows: Initially, the value  $m$  is fixed and some adversarial computational power percentage  $q$  over the total network computational power is chosen.  $k$  is chosen based on  $q$  according to Nakamoto [1]. The number of blocks  $y$  during which

parallel mining will occur is also fixed. Then the experiment begins with the adversary and the honest parties sharing a common blockchain which ends in block  $B$ . As soon as  $B$  is mined, the adversary starts mining in secret and in parallel with the honest parties on his own private fork on top of  $B$ . She ignores the honest chain and does not publish her private fork, so that the two chains remain disjoint after  $B$ . As soon as  $y$  blocks have been mined in total, the adversary attempts a double spend by creating two conflicting transactions which are committed to an honest block and an adversarial block respectively on top of each current chain. Finally, the adversary mines  $k$  blocks on top of the double spending transaction within her private chain. As soon as these  $k$  blocks have been mined, she publishes her private chain in an attempt to overcome the honest chain.

Based on the above experiment, we measure the probability of success of the adversary. We experiment with various values of  $m$  for  $y = 100$ , indicating 100 blocks of secret parallel mining. We make the simplifying assumption that honest party communication is perfect and immediate, meaning that all honest party rounds that are successful are also uniquely successful. We ran 1,000,000 Monte Carlo executions<sup>7</sup> of the experiment for each value of  $m$  from 1 to 30. We ran the simulation for values of computational power percentage  $q = 0.1$ ,  $q = 0.2$  and  $q = 0.3$ . The results are plotted in Figure 5.

**Fig. 5.** Simulation results for a private mining attacker with  $q \in \{0.1, 0.2, 0.3\}$  computational power, suffix security  $k$  according to Nakamoto and parallel mining parameter  $y = 100$ . Probabilities are plotted on a logarithmic scale. The threshold probability of [1] is indicated by the horizontal line.



Based on this data, we conclude that  $m = 5$  is sufficient to achieve a 0.001 probability of failure against an adversary with 10% mining power. To secure against a powerful adversary who controls more than 30% of the network’s mining power, a choice of  $m = 15$  is needed.

## 10 Applications and Evaluation

### 10.1 Multi-blockchain wallets.

One potential application of our technique is an efficient multi-cryptocurrency client. Consider for example a merchant that wishes to accept payments in any cryptocurrency, not just the most popular ones. A naïve approach would be for the merchant to simply install an SPV client for each known cryptocurrency. However, this approach would entail

<sup>7</sup> The URL to the GitHub repository of this MIT-licensed experiment has been redacted for anonymity and will be provided in the proceedings version of this paper.



downloading the header chain for each cryptocurrency, and periodically (e.g., each day) syncing up by fetching any newly generated block headers. Another alternative would be to use an online service supporting multiple currencies, such as an exchange, although this introduces reliance on a third party. Similarly, existing multi-cryptocurrency apps, such as Jaxx and Coinomi rely on third party networks rather than peer-to-peer networks.

*A multi-blockchain wallet based on NIPoPoWs.* The NIPoPoW-based client maintains a most recent  $k$ -stable block hash for each of its supported cryptocurrencies, initially the genesis block for each. Each time a payment is received, the client connects to peers on the corresponding network and asks for a NIPoPoW proof relative to the most recently stored block hash. For cryptocurrencies where payments are received very frequently, the NIPoPoW-based client might download nearly every block header, just like an ordinary SPV client; however, for cryptocurrencies used infrequently, the NIPoPoW-based client would be able to skip over many of the blocks.

*Simulation.* We developed a simulation to evaluate the potential resources savings resulting from the use of a NIPoPoW-based client. We model the arrival of payments in each cryptocurrency as a Poisson process, and assume that the market capitalization of a cryptocurrency is a proxy for usage, i.e. in our model most payments received are Bitcoin transactions, with transactions from the remaining constituting a long tail.

At the current time of writing, a total of 731 cryptocurrencies are listed on coin market directories<sup>8</sup>. This list includes many defunct cryptocurrencies that are no longer functioning at all. We narrow our focus to the 80 cryptocurrencies that have their own proof-of-work blockchains (i.e., no proof of stake) and that have a market cap of over USD \$100,000.

In Table 2 we show aggregate statistics about these 80 cryptocurrencies, grouped according to their proof-of-work puzzle. Bitcoin’s 10 minute block time is relatively slow; other cryptocurrencies typically use faster block times, e.g. 12 seconds for Ethereum and 2 minutes for Litecoin. Thus while the entire chain of proof-of-work puzzles in Bitcoin today only amounts to about 40 MB, taken together, the 80 cryptocurrencies comprise nearly 10 GB of proofs-of-work, and generate nearly 10 MB more each day.

In our simulation we choose NIPoPoW parameters based as in the previous section, with  $m = 24$  and  $k = 6$ . In Table 3 we show the resulting bandwidth costs from simulating a period of 60 days, with varying rates of payments received.

For the naïve SPV client, the total bandwidth cost is dominated by the need to fetch the entire chain of block headers, which the NIPoPoW client does not need to do. The marginal cost for naïve SPV depends on the number of blocks generated each day, as well as the transaction inclusion proofs associated with each payment. The NIPoPoW based client provides the most savings when the number of transactions per day is smallest, reducing the necessary bandwidth per day (not including the initial sync up) by 90%.

## 10.2 Certificate Transparency

In Catena [7] the Bitcoin blockchain is used as an equivocation-resistant public log in which to publish commitments to SSL certificates. The Catena client prototype is based on the BitcoinJ SPV library, and therefore requires downloading the entire Bitcoin proof-of-work chain. Catena could therefore immediately be improved by using NIPoPoWs. First, we note that the Catena authors anticipate needing to launch a dedicated Header Relay Network [7] to accommodate the extra bandwidth demands from new Catena clients. A variant based on NIPoPoWs could obviate this, since it eliminates the need to bootstrap

---

<sup>8</sup> <https://coinmarketcap.com/>

**Table 2.** Cost of header chains for all active Proof-of-Work-based cryptocurrencies (collected from coinwarz.com)

Hash	# Instances	Size today	Growth rate
Scrypt	44	4.3 GB	5.5 MB / day
SHA-256	15	1.4 GB	937.0 kB / day
X11	5	581.1 MB	556.3 kB / day
Quark	3	647.9 MB	518.4 kB / day
CryptoNight	2	199.0 MB	115.2 kB / day
EtHash	2	588.6 MB	921.6 kB / day
Groestl	2	300.3 MB	184.2 kB / day
NeoScrypt	2	310.6 MB	153.6 kB / day
Equihash	2	17.7 MB	92.2 kB / day
Keccak	1	161.1 MB	115.2 kB / day
X13	1	30.0 MB	57.6 kB / day
Lyra2REv2	1	57.4 MB	46.1 kB / day
Total	80	8.5 GB	9.2 MB / day

**Table 3.** Simulated bandwidth of multi-blockchain clients after two months (Averaged over 10 trials each)

Tx/day	Naive SPV		NIPoPoW		Savings
	Total	(Per day)	Total	(Per day)	
100	5.5 GB	(5.5 MB)	31.7 MB	(507.4 kB)	0.99 (0.91)
500	5.5 GB	(5.7 MB)	68.2 MB	(1.1 MB)	0.99 (0.81)
1000	5.5 GB	(6.0 MB)	99.1 MB	(1.6 MB)	0.98 (0.73)
3000	5.6 GB	(7.0 MB)	192.5 MB	(3.1 MB)	0.97 (0.56)

new clients by transmitting the entire 40MB Bitcoin header chain. Second, the steady state cost of operating a Catena client depends on how frequently certificate digests are published. For example, one usage scenario cited by Catena [7] is Keybase, a service which publishes certificate digests every 6 hours. During a 6 hour period, Bitcoin would generate 6 kilobytes of headers, whereas a NIPoPoW proof covering the same range would require less than half this size. The savings would increase further if Catena were implemented using any proof-of-work blockchain with more frequent blocks, such as Ethereum or Litecoin.

### 10.3 Sidechains

At the time of writing, it is widely known that Bitcoin faces significant scaling hurdles [29], but upgrading Bitcoin is notoriously difficult. A widely anticipated solution, first proposed by Back et al. [9], is to treat the Bitcoin blockchain as a host for “sidechains,” which are proof-of-work blockchains separate from Bitcoin, but that can be backed by Bitcoin deposits. Our NIPoPoW construction solves an open problem posed by Back et al., and thus enables this vision.

In the following, we explain sidechains in an incremental fashion, by first providing some context on other protocols involving cross-blockchain interactions. The most well known is the “atomic swap” [8], due to Tier Nolan, by which a quantity of one cryptocurrency can be atomically traded for a quantity of another (e.g., Bitcoin for Litecoin). This protocol relies only on configuring the two blockchains to both rely on a common event, namely the publishing of a certain hash preimage, which is already supported by Bitcoin script.

As a more complicated example, BTCRelay is an Ethereum [3] smart that acts as an SPV client for the Bitcoin blockchain. Users submit Ethereum transactions containing Bitcoin block headers and transaction inclusion Merkle tree proofs, which the BTCRelay contracts validates and stores. This allows Ethereum smart contracts to condition their

behavior based on events in the Bitcoin blockchain. If Bitcoin were upgraded to support NIPoPoWs as described in Section 8, then the BTCRelay functionality could be provided at less cost.

Building on this idea, the sidechains protocol can be seen as a pair of blockchains, the sidechain and the parentchain, that each act as an SPV client of each other. It is then straightforward to cross-chain applications, such as a virtual asset on the sidechain that is backed by a quantity of currency on the parent chain (i.e., a two-way peg). This is accomplished by defining two new transaction types: “deposit” transactions, which have the effect of locking coins on the parent chain and creating new assets on the side chain, as well as “withdrawal” transactions that have the opposite effect, unlocking coins on the parent chain and destroying coins on the sidechain. Deposit transactions are committed first on the parent chain, and then delivered to the sidechain using an SPV proof; vice versa for withdrawal transactions. Note that the parent chain would not need to know about most transactions on the sidechain except for the withdrawal transactions — hence this provides a potential scalability solution if significant transaction volume can be delegated to sidechains.

In the simplest case, the naïve SPV approach would work; each blockchain header of the parent chain would be submitted in a transaction payload to the sidechain (c.f. Back et al. [9], Section 3.2). However, similar to the scenario in Section 10.1, a NIPoPoW-based approach could offer significant performance gains when a single parent chain acts as host to a large number of sidechains.

## 11 Future work

Our security proof was constructed in the Backbone [10] setting in which the block difficulty is treated as a constant. It is easy to extend our construction to the variable difficulty setting. The only difference is that chains (and proofs) would have to be compared according to the total difficulty they possess, i.e., blocks at superlevel  $\mu$  will be  $\mu$  times more difficult than 0-level blocks. We conjecture that this construction can be proven secure using a similar argumentation as the one presented. The model in [30] can be used as a basis and we leave the full analysis for future work.

We worked with a construction that allows a certain class of chain predicates to be proved: In particular, stable predicates which are functions only of the suffix of the chain or depend on a small segment within the chain. However, slightly more general predicates can also be proved succinctly. We leave the complete characterization of succinctly provable predicates for future work.

## References

1. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008)
2. Nakamoto, S.: Bitcoin open source implementation of p2p currency. P2P Foundation **18** (2009)
3. Wood, G.: Ethereum: A secure decentralised generalised transaction ledger. Ethereum Project Yellow Paper **151** (2014)
4. Dwork, C., Naor, M.: Pricing via processing or combatting junk mail. In: Annual International Cryptology Conference, Springer (1992) 139–147
5. Back, A., et al.: Hashcash-a denial of service counter-measure (2002)
6. Shin, L.: For first time bitcoin accounts for less than half of market cap of all cryptocurrencies. <https://www.forbes.com/sites/laurashin/2017/05/16/for-first-time-bitcoin-accounts-for-less-than-half-of-market-cap-of-all-cryptocurrencies/> (May 2017)
7. Tomescu, A., Devadas, S.: Catena: Efficient non-equivocation via bitcoin. In: IEEE Symp. on Security and Privacy. (2017)
8. Nolan, T.: Alt chains and atomic transfers. [bitcointalk.org](http://bitcointalk.org) (May 2013)

9. Back, A., Corallo, M., Dashjr, L., Friedenbach, M., Maxwell, G., Miller, A., Poelstra, A., Timón, J., Wuille, P.: Enabling blockchain innovations with pegged sidechains. URL: <http://www.opensciencereview.com/papers/123/enablingblockchain-innovations-with-pegged-sidechains> (2014)
10. Garay, J., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol: Analysis and applications. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer (2015) 281–310
11. Kiayias, A., Lamprou, N., Stouka, A.P.: Proofs of proofs of work with sublinear complexity. In: International Conference on Financial Cryptography and Data Security, Springer (2016) 61–78
12. Miller, A.: The high-value-hash highway, bitcoin forum post (2012)
13. Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J.A., Felten, E.W.: Sok: Research perspectives and challenges for bitcoin and cryptocurrencies. In: Security and Privacy (SP), 2015 IEEE Symposium on, IEEE (2015) 104–121
14. Poelstra, A.: On stake and consensus. (2015)
15. Friedenbach, M.: Compact spv proofs via block header commitments. <https://www.mail-archive.com/bitcoin-development@lists.sourceforge.net/msg04318.html> (2014)
16. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications. In: Proceedings of the twentieth annual ACM symposium on Theory of computing, ACM (1988) 103–112
17. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Proceedings of the 1st ACM conference on Computer and communications security, ACM (1993) 62–73
18. Canetti, R.: Security and composition of multiparty cryptographic protocols. *Journal of CRYPTOLOGY* **13**(1) (2000) 143–202
19. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on, IEEE (2001) 136–145
20. Canetti, R.: Security and composition of multiparty cryptographic protocols. *Journal of CRYPTOLOGY* **13**(1) (2000) 143–202
21. Narayanan, A., Bonneau, J., Felten, E., Miller, A., Goldfeder, S.: Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction. Princeton University Press (2016)
22. Antonopoulos, A.M.: Mastering Bitcoin: unlocking digital cryptocurrencies. " O'Reilly Media, Inc." (2014)
23. Eyal, I., Sirer, E.G.: Majority is not enough: Bitcoin mining is vulnerable. In: International conference on financial cryptography and data security, Springer (2014) 436–454
24. Buterin, V., et al.: A next-generation smart contract and decentralized application platform. white paper (2014)
25. Bahack, L.: Theoretical bitcoin attacks with less than half of the computational power (draft). arXiv preprint arXiv:1312.7013 (2013)
26. DuPont, Q.: Experiments in algorithmic governance: A history and ethnography of "the dao," a failed decentralized autonomous organization. Bitcoin and Beyond: Cryptocurrencies, Blockchains and Global Governance. Routledge (2017)
27. Spagni, R.: A formal approach towards better hard fork management. <https://forum.getmonero.org/4/academic-and-technical/303/a-formal-approach-towards-better-hard-fork-management> (May 2015)
28. Rosenfeld, M.: Analysis of hashrate-based double spending. arXiv preprint arXiv:1402.2009 (2014)
29. Croman, K., Decker, C., Eyal, I., Gencer, A.E., Juels, A., Kosba, A., Miller, A., Saxena, P., Shi, E., Sirer, E.G., et al.: On scaling decentralized blockchains. In: International Conference on Financial Cryptography and Data Security, Springer (2016) 106–125
30. Garay, J.A., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol with chains of variable difficulty. (2016)