Blink: An Optimal Proof of Proof-of-Work

Lukas Aumayr*
TU Wien

Zeta Avarikioti
TU Wien,
Common Prefix

Matteo Maffei
TU Wien,
CDL-BOT

Giulia Scaffino*
TU Wien,
CDL-BOT ,
Common Prefix

Dionysis Zindros
Stanford University,
Common Prefix

ABSTRACT

Designing light clients for Proof-of-Work blockchains has been a foundational problem since Nakamoto's SPV construction in the Bitcoin paper. Over the years, communication was reduced from $O(C)$ down to $O(\text{polylog}(C))$ in the system's lifetime C. We present Blink, the first *provably secure* $O(1)$ *light client that does not require a trusted setup.*

1 INTRODUCTION

It is impractical for a blockchain user, such as a wallet, to download and verify the whole chain due to communication, computation, and storage constraints. In the seminal Bitcoin white paper [28], Satoshi Nakamoto predicted this need for efficiency and designed a *light client* called the Simplified Payment Verification (SPV) protocol, which decouples the download of the execution layer data (transactions) from the consensus layer data (block headers). An SPV client retrieves all block headers and verifies them according to the longest chain consensus rule. This process requires communication that grows linearly with the systems' lifetime as the header chain grows at a roughly linear rate.

Several subsequent works optimized this concept, introducing *superlight clients* whose communication complexity is only polylogarithmic (*succinct*) in the lifetime of the system [12, 20, 21, 24]. Nevertheless, these protocols are not out-of-the-box compatible with Bitcoin but instead require a consensus fork.

Designing a client with constant communication complexity has remained an elusive goal over the past dozen years. This paper fills this gap.

Contributions. In this work, we present Blink, a novel *interactive PoW light client with constant communication complexity*. In a nutshell, the Blink client connects to a set of full nodes, one of which is honest. The client locally samples a random value η, includes it in a transaction Tx_η, and sends it to the full nodes. For instance, Tx_η can simply be a payment to a vendor's fresh address, which was sampled with high entropy. Then, Blink waits for Tx_η to be included in a (high-entropy) block and confirmed. The full nodes respond to the client with a proof π consisting of $2k+1$ consecutive blocks, with the high-entropy block in the middle and k blocks before and after it (see Figure 1); k is the security parameter [17], e.g., the conventional 6 confirmation blocks in Bitcoin. Importantly, full nodes do not send the full header chain to the client. The constant-sized proof π ensures that *the first block in the proof is stable in the chain* and, therefore, it can be considered as a checkpoint or, in other words, as a new genesis block $\mathcal{G}'$.

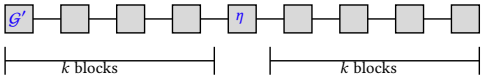

Figure 1: Structure of the Blink's proof π. The proof π consists of $2k+1$ consecutive blocks, with the block including the entropy η in the middle, and k blocks before and after it. The first block $\mathcal{G}'$ in π is stable in the chain and acts as a new genesis block.

We highlight that *Blink does not require any trusted setup*, and we prove it secure under an honest majority of computational power, i.e., against less than 1/2 adversaries. We analyze security in the static PoW model introduced by the Bitcoin Backbone [17], and we adopt the light client state security definitions introduced in [31]. In this model, we refine the problem of Proofs of Proofs-of-Work [24] and prove that Blink has optimal communication cost, building the *first provably secure Optimal Proof of Proof-of-Work (OPoPoW) without trusted setup.*

Furthermore, Blink is a powerful tool that can be leveraged to develop a plethora of applications with enhanced efficiency compared to state-of-the-art protocols. Specifically, we present the following applications, all with constant communication costs:

(1) Bootstrapping PoW blockchain clients, full nodes, and miners
(2) Active payment verification
(3) Past payment and ledger state verification
(4) Bridging PoW blockchains

In further detail, Blink naturally provides a *bootstrapping method*: an SPV client or a light miner[1] broadcasts Tx_η, and upon receiving π, they can efficiently select the tip of the current longest chain $\mathcal{G}'$, and start running their protocol on top of it. As a result, the bandwidth cost of bootstrapping is reduced from linear to constant with respect to the lifetime of the system.

Second, Blink allows *trustless and efficient verification of on-chain payments*. In particular, upon identifying the new genesis $\mathcal{G}'$, an SPV protocol is executed on top of $\mathcal{G}'$, finalizing the payment as soon as k confirmation blocks appear on top of Tx_η in the longest chain. Hence, the payment protocol has the same latency as a

*These two authors contributed equally to the work.

[1]Light miners do not validate transactions included in the chain before they booted up thus, to be sure they start mining on the correct tip of the chain, they need to run an efficient protocol to identify the current longest chain [22]. They start fully verifying transactions after bootstrapping.

standard SPV client, but only constant communication complexity as at most $3k+1$ blocks are relayed in total for finalizing a payment.

Third, assuming block headers include a commitment to the state of the ledger (e.g., Ethereum PoW and ZCash) or include an efficient way of verifying the history of transactions (i.e., ancestry proofs such as block interlinking in the form of Merkle Mountain Ranges [3, 4] or vector commitments [13]), the client enables the extraction of any historical state of the ledger from $\mathcal{G}'$ (including the current one). This means that users can use the Blink-based payment protocol to *read any past transaction as well as any historical state of a smart contract.* This approach reduces the communication overhead from polylogarithmic, as seen in state-of-the-art protocols [12, 24], to constant (in the system's lifetime).

Finally, Blink can serve as a building block for optimistic bridges, where π is used as a fraud proof. This way, Blink enables the *first trustless, secure PoW bridge with constant communication* for relaying a transaction from a source to a destination chain. We also prove that a recent work claiming a trustless constant-size bridge construction [30] is, in fact, insecure.

We provide a *proof-of-concept implementation* of Blink, and evaluate its communication cost for the conventional confirmation block value $k = 6$ and the block height at the time of writing. We underscore that Blink improves on all previous light client solutions in terms of bandwidth: SPV requires 67.3MB, NIPoPoWs requires 10KB, FlyClient requires 5KB, ZK ZeroSync requires 197KB, whereas Blink requires only 1.6KB. All of the solutions have the same latency as they all have to wait for k confirmation blocks.

Related Work. The description of Nakamoto's SPV client appears already in the paper that introduced Bitcoin [28]. A series of optimizations followed. The first succinct construction was the interactive *Proofs of Proof-of-Work protocol* [21] with polylogarithmic communication costs. Later work removed this interactivity and achieved security against $1/2$ adversaries but succinctness only in the optimistic setting (against no adversaries) [24]. This construction was subsequently optimized [20], made practical [15], and redesigned with backwards compatibility in mind [25]. The optimistic setting limitation was alleviated in a follow-up work, achieving succinctness against all adversaries up to a $1/3$ threshold [23]. An alternative construction was also proposed, enabling security and succinctness against a $1/2$ adversary, and adding support for variable difficulty [12]. All these solutions require polylogarithmic communication, whereas Blink requires only constant.

Recently, generic (recursive) zero-knowledge (ZK) techniques were utilized to build constant communication light clients [5, 11, 32]. However, these approaches incur prohibitively high computational costs (or necessitate specialized blockchain deployments [5, 32] utilizing ZK-friendly cryptographic primitives [19]) and additionally require a trusted setup to generate and prove verification keys (which can only be removed if polylogarithmic communication is acceptable). Contrarily, Blink does not impose high communication costs nor a trusted setup.

To develop a constant communication light client without a trusted setup, the idea of using only a small segment of the chain near the tip was proposed [2]. However, the proposed construction was shown to be susceptible to pre-mining attacks and thus insecure [30]. Recently, another construction was introduced called Glimpse [30], combining the idea of [2] with the injection of a *high-entropy* transaction (which was originally introduced in [36, Chapter 5] but for a different purpose) to prove the provided segment of the chain is "fresh" and not pre-mined. Nevertheless, Glimpse remains insecure as we show in this work. We also leverage these ideas to design Blink, the first provably secure light client with constant communication that does not require a trusted setup.

Finally, a similar quest for proof of stake light clients has achieved polylogarithmic complexity in an interactive setting [10]. For a review of the long-standing light client problem, see [14]. Light clients are also a cornerstone for building trustless bridges between chains, a question that has been explored in a multitude of works [18, 27, 33, 34]. In this work, we demonstrate how Blink can be utilized to construct a trustless and efficient optimistic bridge.

Comparison. In Table 1, we compare the characteristics of existing light client protocols, including Blink. We denote by C the lifetime of the system (informally, the length of the blockchain) and by k the security parameter. According to the Bitcoin Backbone model, k is the *common prefix* parameter, which is constant for a protocol execution, albeit with the trade-off of logarithmically increasing the probability of failure in the lifetime of the system.

We first observe that Glimpse [30] achieves $O(k)$ communication but it is not secure in the honest majority assumption (as shown in this work); its exact resilience, if any, remains unknown. ZK clients, on the other hand, achieve $O(1)$ communication but necessitate a trusted setup; unlike Blink in which such assumption is not necessary. We further expose a trade-off between communication overhead and interactivity: prior state-of-the-art PoPoWs are non-interactive but require $O(k \text{ polylog}(C))$ communication [12, 21, 23, 24]. Contrarily, Blink only requires $O(k)$ communication but needs one round of interaction.

2 PROTOCOL DESIGN

In this section, we introduce Blink, the first provably secure, optimal PoW light client that does not require a trusted setup. We begin with a high-level overview of our protocol's objectives and introduce a protocol abstraction that embodies these goals. Next, we present Blink, describing in simple terms the rationale behind its design and security. Throughout this work, we will use the term block to mean a block header.

2.1 Optimal Proof of Proof-of-Work Client

A client protocol is an interactive protocol between a set of provers $P \in \mathcal{P}$ maintaining a ledger, and a verifier V, i.e., the client. If the provers convince V about the current state of the ledger without asking V to download the whole ledger or execute all the transactions, then the client is a *light client.* In particular, if the verifier only receives a *constant* amount of data independently of the ledger's lifetime, then the light client protocol has *optimal* communication.

A client is convinced about the state of a ledger or, simply, of a blockchain, if it receives a block B fulfilling the following properties: (a) B is safe, i.e., it will never be reverted in the view of an honest node; (b) B is live, i.e., B was created recently and therefore the client has an up-to-date view of the state of the blockchain.

Our goal is to design a client protocol that, with only constant communication complexity, satisfies the security notions defined

	SPV[28]	KLS[21],NIPoPoW [24] FlyClient[12], Mining LogSpace[23]	ZK Clients[5, 11, 32]	Glimpse [30]	**Blink**
Communication	$O(C)$	$O(k \text{ polylog}(C))$	$O(1)$	$O(k)$	$O(k)$
No Trusted Setup	✓	✓	✗	✓	✓
Adv. Resilience	1/2	1/2	1/2	✗(?)	1/2
Interactivity	No	No	No	Yes	Yes

Table 1: Comparison of light client solutions

in (a) and (b). Figure 2 illustrates a client protocol abstraction that realizes our objectives. It showcases the interaction between the set of provers ($\mathcal{P}$) and the verifier (V) highlighting the pivotal components of our construction which ensure security: the *constant-sized proof* π that P sends to V and the *extraction of block* B from π, allowing V to read the current state of the ledger.

Initiation
(1) V connects to a set of nodes in $\mathcal{P}$
Proof Construction
(2) Nodes in $\mathcal{P}$ may interact with V
(3) Nodes in $\mathcal{P}$ send to V a (set of) *constant-size* proof(s) π that includes a block B that is safe and live
Block Extraction
(4) V verifies π
(5) V extracts B from π and terminates

Figure 2: Abstraction of an OPoPoW client protocol

Throughout the remainder of this work, we will omit discussing the initiation step, as it remains the same. For simplicity, we treat ledgers extracted from blockchains only, i.e., we assume the ledger is generated as the output of a blockchain protocol (and not, e.g., a DAG protocol). We generalize the discussion in later sections.

2.2 Blink Client

The ultimate goal of an OPoPoW client is to identify a recent, correct block of the ledger, by only receiving a constant-sized amount of data from the set of provers. Towards this, we start with a naive client construction; we identify security threats and propose solutions until we converge to a secure client protocol.

A Naive Construction. Let us start analyzing one of the simplest constructions one might think of. The provers give the last $k+1$ consecutive blocks in their longest chain to the client, who in turn verifies the validity of these blocks and accepts the first (the oldest) block in the proof as safe and live. We recall that in PoW blockchains, blocks are considered final after they have k confirmation blocks, where k is the safety parameter, e.g., in Bitcoin folklore blocks are considered final after 6 confirmations. According to [17], k is a constant for a protocol execution, which, however, implies that the blockchain's security bounds degrade logarithmically with its lifetime. Since the client checks their validity, all blocks in the proof fulfill the PoW difficulty requirements. Trivially, this construction is broken: adversarial provers can have pre-mined $k+1$ fake blocks stored somewhere, and when the client boots up, they provide the client with a block that is either not part of the longest chain or is outdated. Upon receiving different $k+1$ blocks from honest and adversarial provers, the client cannot identify the correct chain with higher probability than random guessing.

Preventing Upfront Mining Attacks. To prevent this upfront mining attack, the client V can locally sample a random string η and give it to the provers along with a time window T, within which V accepts a proof π [30]. Then, provers can then broadcast an *entropy transaction* Tx_η which embeds η to the blockchain network, and wait for it to be included in a block. We will call this block $\dot{\mathsf{B}}$ and since it contains η, this is a high-entropy block. Before the timeout T expires, if k blocks are built on top of $\dot{\mathsf{B}}$, P sends to V a proof π consisting of $\dot{\mathsf{B}}$ followed by its k confirmation blocks. Finally, V accepts $\dot{\mathsf{B}}$. Figure 3 illustrates the protocol presented in [30]; λ is a security parameter. While randomizing the proof π solves the upfront mining attack, it does not lead us to a secure client protocol.

Proof Construction
(5) V samples $\eta \xleftarrow{\$} \{0,1\}^\lambda$
(6) V selects a time T in the future that corresponds to the expected creation time of $k+1$ blocks
(7) V sends η and T to every $P \in \mathcal{P}$ along with a request to return a light client proof π of length $k+1$ conditioned to η, within time T
(8) $\mathcal{P}$ construct an entropy transaction Tx_η containing η and broadcast it to the blockchain network
(9) As soon as a party $P \in \mathcal{P}$ has a proof π consisting of a block $\dot{\mathsf{B}}$ containing Tx_η with k confirmations blocks, P sends π to V
Block Extraction
(10) V accepts π if it was received within time T, $\dot{\mathsf{B}}$ contains Tx_η and has k confirmation blocks on top of it
(11) V extracts $\dot{\mathsf{B}}$ from π and terminates

Figure 3: Naive client protocol [30].

Indeed, we observe that an adversary $\mathcal{A}$ has a probability $\frac{t}{n} < \frac{1}{2}$ to be elected as PoW block proposer, with n the total number of participants in the PoW game, out of which t are controlled by the adversary. This means that $\mathcal{A}$ has a non-negligible probability to censor Tx_η in the first $k-1$ blocks after Tx_η was broadcast. If T is such that fewer than $2k$ consecutive blocks are produced in T with overwhelming probability, the adversary can violate the liveness of the client with probability $\frac{t}{n}$, because honest parties cannot produce a valid proof of $k+1$ blocks within time T. Figure 4 demonstrates this attack.

Preventing Liveness Attacks. To protect the client from this liveness attack, one could take different directions: (A) remove the time window T (alternatively, increase it such that at least $2k$ blocks are produced in T with overwhelming probability), or (B) accept proofs of length less than k. In (A), V accepts the first proof π it receives, with π consisting of $\dot{\mathsf{B}}$ and at least k confirmation blocks on top of it. In (B), V accepts the proof π that, by T, has the most

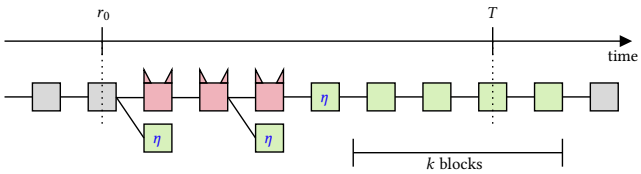

Figure 4: Consider $k = 4$. The light client boots at round r_0 and broadcast the entropy η. With significant probability, the adversary can censor Tx_η in the first $k - 1$ blocks after Tx_η was broadcast. It results that honest parties might not find a proof π of sufficient length by the timeout T.

confirmation blocks on top of $\dot{\mathsf{B}}$. In Figure 5, we present the client protocol for the (A) and (B) variants. While both these attempts safeguard the liveness of the client, the client's safety is broken: V might accept a block that is not part of any honest party's chain.

<u>Proof Construction</u>

(5) V samples $\eta \xleftarrow{\$} \{0,1\}^\lambda$
(6) V sends η to every $P \in \mathcal{P}$ along with a request to return a light client proof π of length $k + 1$ conditioned to η
(7) $\mathcal{P}$ construct an entropy transaction Tx_η containing η and broadcast it to the blockchain network
(8) As soon as a party $P \in \mathcal{P}$ has a proof π consisting of a block $\dot{\mathsf{B}}$ containing Tx_η with k confirmations blocks, P sends π to V

<u>Block Extraction</u>

(9) V accepts the first π it receives where $\dot{\mathsf{B}}$ contains Tx_η and has k confirmation blocks on top of it
(10) V extracts $\dot{\mathsf{B}}$ from π and terminates

<u>Proof Construction</u>

(5) V samples $\eta \xleftarrow{\$} \{0,1\}^\lambda$
(6) V selects a time T in the future
(7) V sends η to every $P \in \mathcal{P}$ along with a request to return a light client proof π conditioned to Tx_η within time T
(8) $\mathcal{P}$ construct an entropy transaction Tx_η containing η and broadcast it to the blockchain network
(9) At time T, each party $P \in \mathcal{P}$ sends to V its π, consisting of $\dot{\mathsf{B}}$ containing Tx_η along with all the subsequent confirmation blocks that P is aware of

<u>Block Extraction</u>

(10) V accepts the π that has the most confirmation blocks on top of $\dot{\mathsf{B}}$ containing Tx_η and was received within time T
(11) V extracts $\dot{\mathsf{B}}$ from π and terminates

Figure 5: Insecure attempts (A) and (B), top and bottom respectively.

We now describe the safety attack for (A), but a similar logic applies to (B) as well. After V broadcasts Tx_η, honest parties immediately include it on-chain, while the adversary $\mathcal{A}$ starts mining on a private chain that censors Tx_η. $\mathcal{A}$ can mine $k - l$ blocks in its private chain, with $0 < l < k - 1$, while honest parties only mine $\dot{\mathsf{B}}$ with at most $k - l - 2$ confirmations. This can happen with non-negligible probability, as we are considering subchains with fewer than k blocks [17], with k being the safety security parameter. Then, $\mathcal{A}$ broadcasts their private chain, causing all honest parties to switch to the adversarial chain due to the longest chain rule. Honest parties subsequently include Tx_η in their new longer chain and keep mining on top of it. In the meantime, $\mathcal{A}$ starts privately mining on top of the abandoned chain that included Tx_η early on. Now, to create a valid proof, $\mathcal{A}$ only needs to privately mine $l + 2 < k + 1$ blocks, while honest parties need to mine $k + 1$ blocks. As a result, $\mathcal{A}$ can generate a valid proof faster than honest parties, and trick the client to accept a proof consisting of blocks that will not be part of the honest chain, thereby breaking security. Figure 6 illustrates this attack.

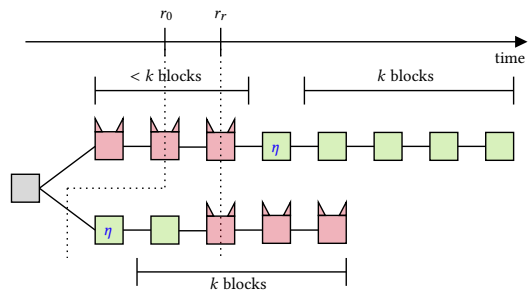

Figure 6: Consider $k = 4$ and $l = 1$. The client broadcasts η at r_0. $\mathcal{A}$ privately mines a subchain of 3 blocks censoring Tx_η, while honest parties include Tx_η and only mine 2 blocks overall. At r_r, $\mathcal{A}$ releases the private chain, which is adopted by honest parties as per the longest chain selection rule. Honest parties now need to mine 5 blocks to find a valid π. Contrarily, $\mathcal{A}$ needs to only mine 3 blocks. Hence, $\mathcal{A}$ finds π first.

The Blink Proof. Before detailing Blink, we observe that the safety attack in Figure 6 relies on $\mathcal{A}$ privately mining in order to delay the inclusion of Tx_η in the main chain. However, such censoring can only succeed for a limited time, specifically less than k consecutive blocks, as extending a private chain beyond this would lead to a safety violation [17]. In other words, $\mathcal{A}$ may create up to $k - 1$ blocks faster than honest parties (with non-negligible probability) but not more than that: any honest majority will create k blocks faster than any minority adversary, rendering the attack in Figure 6 infeasible.

The client can securely accept a block of the blockchain, if they can identify it as a *safe block*, i.e., a block that is already k deep in at least one honest party's chain [17]. Furthermore, the safe block also needs to be *live*, i.e., recent enough to be sufficiently close to the tip of the chain.

We know that the adversary can only censor Tx_η for $k - 1$ blocks and it takes k additional blocks for Tx_η to become safe (Figure 6). Therefore, we *modify π to be of length $2k + 1$ and to specifically contain Tx_η in the middle block $\dot{\mathsf{B}}$, i.e., at position $k + 1$*, as depicted in Figure 1. However, if the client accepts $\dot{\mathsf{B}}$ of the first valid proof received, safety is again violated by the same attack described before: $\mathcal{A}$ can create π before honest parties by using the $k+(k-l-1)$ blocks from the abandoned honest chain and mining $l + 1$ new blocks; meanwhile, honest parties must mine $k > l + 1$ new blocks. Nonetheless, π now being of length $2k + 1$, it necessarily contains a safe block, i.e., a block that is at least k deep in the chain to be stable for all honest parties; this is true even if the π the client receives comes from $\mathcal{A}$. In particular, π contains at least a block that was safe even before Tx_η was broadcast: the honest subchain starting from the block $\dot{\mathsf{B}}$ included early on is at most of length $k - 2$ and at

least of length 1, thus the first block in π is was already part of the honest parties' stable chain (cf. Figure 7). Naturally, the first block in π is attached to genesis, otherwise honest parties would not have extended it. This holds true regardless of the strategy $\mathcal{A}$ follows: Honest parties only abandon their chain if they see a longer one.

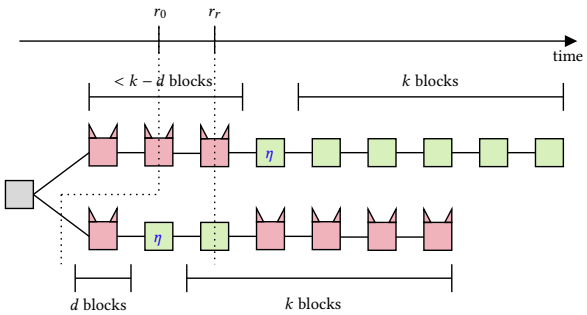

Figure 7: Consider $k = 5$. As in Figure 6, except that $\mathcal{A}$ censors Tx_η by d blocks also on the lower branch, such that $d \leq k - 1$ and s.t. the overall number of adversarial blocks before Tx_η on all branches is smaller than k. This shows why it is not sufficient to take less than k blocks before Tx_η.

We note that for any π coming from an honest party, any block before the entropy block $\dot{\mathsf{B}}$ is safe, as there are at least $k + 1$ confirmations. Thereby, *the first block* B *of any* $2k + 1$ *proof* π *is always safe, i.e., it has at least* k *confirmations in the view of an honest party.* As a result, the client can safely *accept the first block in the first valid* π *it receives.*

Blink Protocol. In Figure 8 we showcase the pseudocode of the Blink protocol, while in Algorithm 1 we put forth the algorithm run by the Blink client, employing Algorithm 2; similarly, in Algorithm 3 we present the code run by provers. We use $m \dashrightarrow A$ to indicate that message m is sent to party A and $m \dashleftarrow A$ to indicate that message m is received from party A.

We observe that the client reads a proof of length $2k + 1$, which is constant in the system's lifetime, and accepts a block that is $2k$ blocks old, incurring a waiting time of k blocks, similarly to an SPV. *Blink is the first PoW light client protocol that achieves optimal proof size with only at most one round of communication between provers and verifier.*

<u>Proof Construction</u>

(5) V samples $\eta \xleftarrow{\$} \{0,1\}^\lambda$
(6) V sends η to every $P \in \mathcal{P}$
(7) $\mathcal{P}$ construct an entropy transaction Tx_η containing η and broadcast it to the blockchain network
(8) As soon as a party $P \in \mathcal{P}$ has k confirmation blocks on top of the block $\dot{\mathsf{B}}$ containing Tx_η, P sends to V π consisting of $\dot{\mathsf{B}}$ with k blocks before and k blocks after it

<u>Block Extraction</u>

(9) V accepts the first π it receives consisting of $2k + 1$ consecutive well-formed blocks where the middle block contains η, i.e., $\pi[k] = \dot{\mathsf{B}}$
(10) V extracts the first block of proof π, i.e., $\mathsf{B} := \pi[0]$, and terminates

Figure 8: Pseudocode of the Blink protocol

Algorithm 1 The algorithm ran by the verifier V, i.e., the Blink client. We split the proof π into (π_0, π_1), with π_0 allowing to identify a stable and recent block of the blockchain, i.e., the new genesis $\mathcal{G}'$, and π_1 being the Merkle proof that verifies inclusion of η into the middle block of π_0.

```
function Verifier_G()
    η ←-- {0,1}^λ
    for P ∈ 𝒫 do
        η --→ P
        while True do
            π ←-- P                ▷ Only constant amount of data downloaded
            (π_0, π_1) = π
            if Valid_G(π, η) then
                return π_0[0]
            end if
        end while
    end for
end function
```

Algorithm 2 The algorithm ran by V to check the validity of the blocks in the proof. Let x be the root of the transaction Merkle tree in a block, and s be its parent hash.

```
function Valid_G(π, η)
    (π_0, π_1) ← π
    if |π_0| < k + 1 then
        return False
    end if
    if ¬MerkleVerify(π_1, η) ∨ π_1.root ≠ π_0[k + 1].x then
        return False
    end if
    h = π_0[0].s
    for B ∈ π_0 do
        if B.s ≠ h then                    ▷ Ancestry failure
            return False
        end if
        h = H(B)
        if h ≥ T then                      ▷ Hardcoded target T, static setting
            return False                   ▷ PoW failure
        end if
        return G = π_0[0] ∨ |π_0| = 2k + 1
    end for
end function
```

Algorithm 3 The algorithm ran by the provers $P \in \mathcal{P}$.

```
function Prover()
    η ←-- V
    Tx_η ←-- MakeTx(η)
    Tx_η --→ Network                       ▷ Wait for Tx_η to be k-confirmed
    π_0 ←-- C[−(2k + 1):]                  ▷ By Common Prefix, Tx ∈ C[−(2k + 1):]
    π_1 ←-- MerkleProve(C[k + 1], η)
    π ← (π_0, π_1)
    π --→ V
end function
```

3 APPLICATIONS

In this section, we showcase how Blink can be used for different applications, ranging from verification of payments and state verification to bootstrapping and bridging.

Proof Construction

(5) V samples $\eta \xleftarrow{\$} \{0,1\}^\lambda$
(6) V sends η to every $P \in \mathcal{P}$
(7) $\mathcal{P}$ construct an entropy transaction Tx_η containing η and broadcast it to the blockchain network
(8) As soon as a party $P \in \mathcal{P}$ has k confirmation blocks on top of the block $\dot{\mathsf{B}}$ containing Tx_η, P sends to V π consisting of $\dot{\mathsf{B}}$ with k blocks before and k blocks after it
(9) V accepts the first π it receives consisting of $2k+1$ consecutive well-formed blocks where the middle block contains η, i.e., $\pi[k] = \dot{\mathsf{B}}$
(10) Upon accepting π, V extracts the new genesis $\mathcal{G}' := \pi[0]$ and sends $\mathcal{G}'$ to all $P \in \mathcal{P}$
(11) Each $P \in \mathcal{P}$ keeps sending all the blocks descending from $\mathcal{G}'$ in their chain

State Extraction

(12) V maintains the longest chain C descending from $\mathcal{G}'$
(13) When Tx_η is k deep in C, V extracts the state from the block including Tx_η and terminates

Figure 9: Pseudocode of the payment verification with Blink

3.1 Payment Verification

Consider a vendor that wants to check whether a particular buyer has made a payment for the purchase of a good. The vendor will only ship the goods after the client's payment has been verified. The Blink protocol, as described in Figure 8, only gives security for the first block in the proof and not, in particular, for the block containing Tx_η: indeed, the proof π accepted by the client might come from the adversary and, thus, the entropy block might not belong to the stable chain. In the payment setting, however, it is desirable to define *security of the stable entropy block* B_η: Tx_η is the transaction of the payment to the vendor, with η now being an address freshly sampled at random by the vendor. Assume the buyer has paid the correct amount to the vendor's new address. To argue about the finality of the payment, i.e., the finality of Tx_η, we recall the strong security guarantee that Blink achieves: Blink allows us to define a recent, trustlessly identified, stable block. This block behaves as a secure checkpoint or, in other words, as a new genesis $\mathcal{G}'$: it is in the stable chain of honest parties, i.e., it will never be reverted, and the consensus rules applied to $\mathcal{G}'$ are consistent to the consensus rules applied to the genesis block $\mathcal{G}$. We now show how to extend the Blink protocol to verify payments. Upon accepting a proof π and identifying $\mathcal{G}'$, the client can send $\mathcal{G}'$ to all provers, and provers start sending to the client all the blocks descending from $\mathcal{G}'$. The client now maintains the longest chain descending from $\mathcal{G}'$, essentially running an SPV algorithm with $\mathcal{G}'$ as a starting point. When Tx_η is in a block that is k-deep in the longest chain (this will happen, at most, $3k$ consecutive blocks on top of $\mathcal{G}'$), the client considers the payment final and terminates.

With one additional round of communication, Blink can now verify payments with a constant-sized proof. We observe that the client latency is the same one of a standard SPV client, i.e., k blocks when there is no adversarial attack, and $2k$ when under attack. In Figure 9 we show the pseudocode for the Blink-based protocol for payment verification. We note that the Blink construction can be used out-of-the-box to verify payments in the Bitcoin Backbone protocol in the static difficulty setting. We refer the reader to Section 6 for variable difficulty and practical deployment.

3.2 Bootstrapping via Blink

In blockchains, there is an interplay between different types of parties: *consensus nodes*, *full nodes*, and *clients*. *Consensus nodes*, also called miners, receive transactions from the network (environment) and execute a distributed protocol that outputs a ledger, i.e., a finite, ordered sequence of transactions identical for all nodes. *Full nodes* do not participate in the distributed ledger protocol; instead, they receive the ledger from consensus nodes, execute transactions to verify their validity, and maintain the ledger. Finally, *clients* connect to full nodes to retrieve a specific state element from the ledger, e.g., an account balance. Bootstrapping these nodes usually requires a lot of time (from several hours to several days) and resources because, starting from genesis, they need to download and execute all the transactions in the ledger (full and consensus nodes) or verify all the blocks in the ledger (SPV-based clients).

In Section 3.1 we used Blink to identify a recent stable block that behaves as a new genesis $\mathcal{G}'$ and, commencing from this block, our client started running an SPV protocol, i.e., the one often run by (light) clients. Blink can thus serve as an efficient bootstrapping protocol that allows the identification of a new stable block $\mathcal{G}'$ and, from that block (e.g., using the state commitment therein), runs the protocol of a consensus, full, or light node. In this way, nodes do not have to execute the entire transaction history or download past blocks but start executing only from transactions $2k$ blocks in the past.

3.3 State Verification

In this work, we demonstrated how to convince a light client about the state of a ledger, incurring only constant communication overhead. As specified in Section 2.1, Blink operates on the premise that each block embeds a constant-sized commitment to the current state of the ledger. Commitments come in different flavors (Merkle tree-based commitments, accumulators, vector commitments), and they are used to download and verify the UTXO set or account balances after the block, including it, has been successfully executed.

Having a chain with state commitments enables Blink to be used to verify more than just payments: *Blink allows to verify account balances and read the current state of on-chain contracts.* For a discussion on chains that have state commitments and how to introduce them to systems like Bitcoin, see Section 6.

3.4 Historical Transaction Verification

While it is uncommon to verify very old transactions, it might be necessary for some applications to verify, e.g., a few weeks old transactions. In these cases, once Blink identifies the new genesis $\mathcal{G}'$, one could travel back the chain block by block until hitting the block containing the transaction to be verified. While Blink has constant communication, this is a naive approach for checking past transactions that comes with a linear overhead: the older the transaction, the more blocks the client has to download. More advanced techniques called *proof of ancestry*, achieve better performances in proving that a block is an ancestor of another block: these include using Merkle Mountain Ranges (MMRs), i.e., extensions of Merkle trees that allow for efficient appends in logarithmic openings, or vector commitments with constant opening. It follows that

Blink allows to succinctly synchronize with the current state of the ledger and, from there, using a proof of ancestry, to travel back the transaction history until verifying the desired old transaction. When verifying historical transactions, the communication of Blink remains constant but can be combined with a linear, logarithmic, or constant proof of ancestry.

3.5 Bridging with Blink

After more than 15 years of research and work from academia and industry alike, the blockchain space has grown in a variety of 100+ chains, each presenting different and unique features in terms of consensus, privacy, throughput, applications, and programmability. To leverage these diverse opportunities and to enhance users' flexibility in the crypto world, light clients have recently become a pivotal component for bridges as well, allowing them to efficiently and securely read the state of a chain within new resource-constrained environments: blockchain themselves.

Successful bridges move a high volume of transactions: ideally, at least one transaction per block. In this case, every block that includes a cross-chain transaction must be relayed by the bridge from the source to the destination chain, in an SPV-like fashion. However, contrarily to an SPV client, the on-chain costs of the bridge can be minimized by avoiding verifying blocks by default. Instead, blocks can be optimistically accepted and only verified on-demand, i.e., in case a dispute is raised. This is what an *optimistic bridge* does. We demonstrate how to *use Blink for creating succinct fraud proofs to resolve disputes.*

Consider a PoW source blockchain C_S including state commitments in its blocks and allowing for efficient ancestry proofs. *Relayers* of the bridge can optimistically relay a stable block B from C_S to the destination blockchain C_D, by submitting B along with a random string η_R they sampled to the smart contract, where the bridge is deployed. Should a *challenger* notice misbehavior, they have a time window to start a challenge in which they pinpoint the contested block B and they reveal a random string η_C to the bridge contract. The challenger proceeds to publish a transaction Tx_η on C_S, which includes $\eta := \eta_R \oplus \eta_C$ (where $\oplus$ is bit-wise xor). Both parties need to contribute with a random string to prevent each of them from cheating, i.e., pre-mining a fake proof. The bridge contract will accept, from anyone, the first valid proof π containing Tx_η, and via ancestry proof it can verify whether or not B is an ancestor of the first block in π by checking the block height. If it is not, B is removed from the bridge contract. Honest behavior can be incentivized through collateral that is slashed or redistributed in case of misbehavior.

4 MODEL

4.1 Notation

The bracket notation $[n]$ refers to the set $\{1, \ldots, n\}$ for a natural number n. $A[i]$ denotes the i-th element (starting from 0) of a sequence A, while negative indices like $A[-i]$ refer to the i-th element from the end. $A[i : j]$ represents the subsequence of A from index i (inclusive) to j (exclusive), while $A[i :]$ and $A[: j]$ represent the subsequences from i onwards and up to j, respectively. The notation $|A|$ denotes the size of the sequence A. The symbols $A \preceq B$ and $A \prec Y$ indicate that A is a prefix or a strict prefix of B or Y, respectively.

We denote with $C_r^{\bigcap} := \bigcap_{P \in \mathcal{H}} C_r^P$ the intersection of the view of all honest parties' chains at round r. Similarly, we denote with $C_r^{\bigcup} := \bigcup_{P \in \mathcal{H}} C_r^P$ the union of the chains of all honest parties, that yields a blocktree. For simplicity, we extend our slicing notation that chops off the last k elements of a sequence, i.e., $[: -k]$, to trees as well. For trees, it works as follows. For every leaf in a tree, select that leaf and the $k-1$ preceding nodes. Then, for every leaf, remove all selected nodes. The slicing notation for trees will be helpful later on, when distinguishing between a stable chain in the view of all honest parties and a stable chain in the view of at least one honest party. It follows that $C_r^{\bigcap}[: -k]$ is the intersection of the view of the blockchain of all honest parties at round r, pruned of the last k blocks; likewise, $C_r^{\bigcup}[: -k]$ is the union of the view of the blockchain of all honest parties at round r, pruned of the last k blocks. In Lemma A.15 (Appendix A.1), we prove that $C_r^{\bigcup}[:-k] = \bigcup_{P \in \mathcal{H}} C_r^P[:-k]$.

We say a block *extends* another block, if the former has the latter as ancestor and has a higher block height. We say a block *descends* from another block, if the former extends the latter or they are the same block. Finally, two blocks are *parallel* when they have the same height.

4.2 Ledger Model

We assume a synchronous network and we consider the protocol execution to proceed in discrete rounds.

Definition 4.1 (Ledger). A *ledger* is a sequence of transactions.

Definition 4.2 (Distributed Ledger Protocol). A *distributed ledger protocol* is an Interactive Turing Machine which exposes the following methods:

- execute: Executes a single round of the protocol, during which the machine can communicate with the network.
- write *(*Tx*)*: Takes transaction Tx as input.
- read *()*: Outputs a ledger.

A distributed protocol that returns a total order of the input transactions for all consensus nodes, satisfies two key properties: safety and liveness. The notation $\mathcal{L}_r^P$ denotes the output of read () invoked on party P at the end of round r.

Definition 4.3 (Safety). A distributed ledger protocol is *safe* if:

- (Self-consistency) For any honest party P and any rounds $r_1 \leq r_2$, it holds that $\mathcal{L}_{r_1}^P \preceq \mathcal{L}_{r_2}^P$.
- (View-consistency) For any honest parties P_1, P_2 and any round r, it holds that either $\mathcal{L}_r^{P_1} \preceq \mathcal{L}_r^{P_2}$ or $\mathcal{L}_r^{P_2} \preceq \mathcal{L}_r^{P_1}$.

Definition 4.4 (Liveness). A distributed ledger protocol is *u-live* if all transactions written to any honest party at round r, appear in the ledgers of all honest parties by round $r + u$.

The ledger uniquely defines the system's current state. An empty ledger is equivalent to a constant genesis state, denoted as st_0. To ascertain the state of a non-empty ledger, transactions from the ledger are sequentially applied to the state, starting from the genesis state. This transaction application to the existing state is encapsulated by

a transition function δ. For a given ledger $\mathcal{L} = \{tx_1, \ldots, tx_n\}$, the state of the system is $\delta(\ldots\delta(\delta(st_0, tx_1), tx_2)\ldots, tx_n)$

We use the shorthand notation δ^* to apply a sequence of transactions $tx = \{tx_1, \ldots, tx_n\}$ to a state. Specifically, $\delta^*(st_0, tx) = \delta(\ldots\delta(\delta(st_0, tx_1), tx_2)\ldots, tx_n)$.

Prover-Verifier Model. A client protocol is an interactive protocol between the client, acting as verifier V, and a set of full nodes, acting as provers $P \in \mathcal{P}$. We focus on a client V that bootstraps on the network for the first time and it is only aware of the genesis state.

We assume that the client is honest and connects to at least one honest prover, in accordance with the standard non-eclipsing assumption. While honest parties adhere to the correct protocol execution, the adversary can execute any probabilistic polynomial-time algorithm.

We can now define state security for client protocols, as originally introduced in [31]. Assuming safety, we use $\mathcal{L}_r^{\bigcup}$ to denote the longest among all the ledgers kept by honest parties at the end of round r, and $\mathcal{L}_r^{\bigcap}$ to denote the shortest among them.

Definition 4.5 (Ledger Client State Security [31]). An interactive Prover-Verifier protocol (P, V) is *state secure* with safety parameter v, if the state commitment $\langle \mathrm{st} \rangle$ output by V at the end of the protocol execution at round r satisfies safety and liveness as defined below.

There exists a ledger $\mathcal{L}$ such that $\langle \delta^*(\mathrm{st}_0, \mathcal{L}) \rangle = \langle \mathrm{st} \rangle$, and $\forall r' \geq r + v$:

Safety: $\mathcal{L}$ is a prefix of $\mathcal{L}_{r'}^{\bigcup}$.

Liveness: $\mathcal{L}$ is an extension of $\mathcal{L}_r^{\bigcap}$.

When a client gets knowledge of the state of the ledger without downloading the entire ledger or executing all transactions, it is a *light client*. Ideally, a light client learns the desired state element by downloading asymptotically less data than a full node. We measure the performance of a client protocol by defining the *communication cost* for the verifier. In other words, for a specific client protocol we measure the data received by the verifier in the proof construction (π) phase.

Definition 4.6 (Client Communication Cost). We define $\mathrm{cost}(\mathcal{E}, V)$ to be the communication cost (in bits) of an execution $\mathcal{E}$ of a protocol $\Pi(\mathcal{P}, V)$ for party V.

We say that a client protocol has *optimal communication cost* if $\mathrm{cost}(\mathcal{E}, V) = O(1)$, i.e., the verifier receives only a constant amount of data. In particular, we will show later that Blink is a light client with optimal communication cost $\mathrm{cost}(Blink) = O(k) = O(1)$, where k is the safety security parameter that is constant for a protocol execution [17].

Definition 4.7 (Optimal Proof-of Proof-of-Work Protocol (OPoPoW)). A light client protocol is an *Optimal Proof-of Proof-of-Work* protocol when it is secure (Definition 4.5) and has optimal communication cost (Definition 4.6).

4.3 PoW Blockchain Model

A blockchain protocol is a distributed ledger protocol that operates typically as follows: Consensus nodes receive and broadcast chains composed of blocks. Each node P maintains a view of the blockchain, denoted by C^P, which invariably starts with the genesis block G. Nodes verify these chains by ensuring they comply with the validity and consensus rules. These chains include fixed-size transactions arranged in a specific order. Every node interprets its chain to produce a transaction sequence, i.e., to output its ledger. Moreover, a consensus node receives new, unconfirmed transactions from the network, and attempts to add them to its ledger by proposing a new block that includes them. The nodes' local views the ledger can vary from node to node because of the network latency. Honest nodes adhere to the consensus protocol, while adversarial nodes may diverge from it. Nevertheless, under specific assumptions, a blockchain protocol may guarantee that the local chains of different parties satisfy the two key properties of ledgers, namely safety and liveness, albeit typically in a probabilistic manner.

To model the proof-of-work setting, the *q-bounded synchronous setting* defined in [17] can be leveraged. The protocol is analyzed in the static model, where the number of consensus nodes n remains fixed throughout the protocol execution, albeit not known to the nodes themselves. Furthermore, each of them is assumed to have an equal computational power (flat model). The protocol proceeds in synchronous communication rounds. We highlight that the static model implies *static difficulty*, i.e., the PoW difficulty remains the same throughout the protocol execution. The limited capability of the nodes to generate PoW solutions is captured by their restricted access to the hash function $H(\cdot)$ modeled as a Random Oracle; each node is allowed q queries per round. The adversary controls up to $t < \frac{n}{2}$ nodes, meaning they are allowed $t \cdot q$ queries per round. The adversary can insert messages, manipulate their order, and launch Sybil attacks, creating seemingly honest messages. However, the adversary cannot censor honest parties' messages, ensuring that all honest parties receive honestly broadcast messages.

The Bitcoin Backbone model [17] identifies three security properties of a blockchain: *common prefix*, *chain quality*, and *chain growth*. Informally, common prefix dictates that at any point in time, any two honest parties' chains after pruning the last k blocks are either the same or one is a prefix of the other. Chain growth expresses that the blockchain makes progress at least at the pace at which the honest parties produce blocks. Finally, chain quality captures the ratio of honestly produced blocks in the system in any long enough chunk of the chain. The formal definitions can be found in Appendix A.1. A blockchain protocol satisfying common prefix, chain quality, and chain growth also maintains a secure ledger, as per Definition 4.3 and Definition 4.4, under the so-called *k-deep confirmation rule*. This rule states that all nodes consider a block safe when it is part of their local chain pruned by the last k blocks. As expected, both safety and liveness hold probabilistically.

Chain Client Security. As a blockchain defines a specific distributed ledger protocol, full nodes, and clients function as described above. Inheriting the same interactive model, we now define the client security for blockchain protocols. To do so, we first define the notion of admissible blocks as a stepping stone.

Definition 4.8 ((u, k)-Admissible Block at r). Parameterized by $u \in \mathbb{N}$ and $k \in \mathbb{N}$, we call admissible block at r any block B observed at round r fulfilling the following properties:

- **Safety**: $\mathrm{B} \in C_{r+u}^{\bigcup}[: -k]$
- **Liveness**: $\mathrm{B} \notin C_r^{\bigcap}[: -k]$

In our definition of (u, k)-admissible blocks at r, the parameters u and k are free parameters. In our proofs, it turns out that this admissibility holds if u is the "wait time" parameter of liveness, and k is the "depth" parameter of safety/persistence of [17]. Thus, for readability we omit (u, k) and mean admissibility in the round in which the client terminates, if not stated otherwise.

Definition 4.9 (Chain Client Security). An interactive Prover-Verifier protocol (P, V) for clients is *secure* if any block B output by the Verifier at the end of the protocol execution at round r^* is *admissible* for some round $r \leq r^*$.

In other words, the client accepts a block B at round r^*, if for some round $r \leq r^*$ the following holds: B is seen as stable by at least one honest party at round $r + u$ (safety), and B is not yet seen by all parties at round r (liveness).

State Commitments. We consider PoW blockchains in which block headers include state commitments, denoted by $\langle st \rangle$. State commitments are a succinct representation of the state of the ledger, and they are assumed to be of constant size. In the account model of, e.g., Ethereum, an example of state commitment is the Merkle root of account balances; in the UTXO model of, e.g., Bitcoin, an example is the Merkle root of the Sparse Merkle Tree where the value of each leaf corresponds to a UTXO of the UTXO set [28, 31]. Equipped with this functionality, client protocols satisfying Definition 4.9 also satisfy Definition 4.5. We stress however that state commitments are necessary in Blink only for the extraction of the ledger's state but not for the secure proof creation.

5 ANALYSIS

In this section, we present the main theorems and formal analysis of our paper. We start by giving a high-level overview on the proof strategy, followed by the formal proofs. Due to space constraints, some preliminary definitions and lemmas used in the proofs are deferred to Appendix A.

Analysis Overview. The main theorem we prove in this paper is as follows.

Theorem 5.1. *Blink achieves ledger client state security (Definition 4.5).*

Towards proving Theorem 5.1, we start proving the admissibility of $\pi[0]$. We identify a special type of block, which we call *convergence event at a round* r (Definition A.18). A convergence event is an honestly produced block that has (by round r) no parallel block that is acceptable. We call a block an *acceptable block* (Definition A.17) if it is valid and there is at least one honest party who might potentially switch to a chain including it. These convergence event blocks have some interesting properties. In particular, (i) all acceptable blocks at some round r need to descend from all convergence events at round r with smaller block height (Lemma A.20); (ii) a block that is a convergence event $\tilde{\mathsf{B}}$ in a round in which there exists a valid block $\hat{\mathsf{B}}$ with a height of at least k more than $\tilde{\mathsf{B}}$ (even if $\hat{\mathsf{B}}$ is only known to the adversary), $\tilde{\mathsf{B}}$ is destined to become stable for all honest parties (Lemma A.21); (iii) the nearest ancestral convergence event to any block is always fewer than k blocks away (Theorem A.22). Note that these desirable properties hold regardless of our construction, and might be of independent interest.

Towards proving the safety of $\pi[0]$, we show that $\pi[k:]$ always extends a so-called anchor block $\tilde{\mathsf{B}}$, which is the nearest convergence event at the time that π is found and sent to the client (Theorem 5.7). Since $\pi[k]$ is fewer than k blocks away from its nearest ancestral convergence event (Theorem A.22), we know that $\tilde{\mathsf{B}} \in \pi$. Also, $\tilde{\mathsf{B}}$ will become stable (Lemma A.21), and thus $\pi[0]$ is safe. Intuitively, liveness holds since $\pi[k]$ is fresh as it contains the newly sampled η and $\pi[0]$ is exactly k blocks away and thus also new; we formally prove this in Theorem 5.9.

Towards chain client safety, we start arguing about the first proof π of length $2k + 1$ the client accepts at round $r^* + 1$, with $\pi[k]$ containing η. As a first step, we say that π must extend an *anchor block* $\tilde{\mathsf{B}}$ (Anchor, Theorem 5.7). In turn, $\tilde{\mathsf{B}}$ extends a block B' which is stable for all honest parties already at round r_0. Intuitively, this holds because when η is broadcast, honest parties will only produce blocks extending $\tilde{\mathsf{B}}$. As a result of honest majority, a proof extending the anchor is found first.

Liveness and safety yield that $\pi[0]$ is an admissible block at round r^* (Theorem 5.9) and we prove that the longest chain rule applied to the genesis block is consistent with the longest chain rule applied to $\pi[0]$ (New Genesis, Lemma 5.10). Finally, we show that, eventually, all honest parties will have an admissible block including Tx_η and that such a block is close to $\pi[0]$. It follows that running a (succinct) SPV algorithm on top of $\mathcal{G}' = \pi[0]$ will guarantee to Blink admissibility of a block B_η including Tx_η, when B_η is buried k blocks deep in the longest chain. We prove that Tx_η becomes stable for all honest parties (Lemma 5.14) after u rounds and that its distance to $\mathcal{G}'$ is upper-bounded by $3k$ blocks (Lemma 5.15).

To conclude, we prove by reduction that any client construction that fulfilling the chain client security definition, having state commitments, also fulfills the ledger client state security definition (Corollary 5.12).

Theorem 5.2. *Blink has optimal* communication cost*, i.e.,* $O(k)$.

The *communication cost* (Definition 4.6) measures the bits sent/received by V during an execution $\mathcal{E}$ of a protocol $\Pi(\mathcal{P}, V)$. For each P, to which V is connected, there is the following overhead. To identify the new genesis block, V sends η which has a size of $O(1)$ and receives (at most) one proof consisting of $2k + 1$ blocks for each $P \in \mathcal{P}$. This makes for a total size of $O(k)$. To predicate security of the block including η, the client sends the new genesis $\mathcal{G}'$ to all full nodes and it keeps receiving blocks descending from $\mathcal{G}'$, until the entropy block is k deep in the longest chain - this will happen after, at most, $3k$ blocks from $\mathcal{G}'$. This makes for a total size of $O(k)$.

Note that light client constructions connect to a subset of all full nodes. Depending on how many nodes the light client connects to, the overhead increases. This is true for other light client constructions as well. Regardless, the *communication cost* of *Blink* is $O(1)$, i.e., constant in the chain length C. From Theorems 5.1 and 5.2 it follows, that *Blink* is an Optimal Proof-of Proof-of-Work Protocol (OPoPoW, Definition 4.7).

5.1 Safety and Liveness of Blink

We model time to proceed in discrete rounds. Our network model stipulates that messages sent in a round r reach the recipient in

round $r+1$. Like other nodes, the client can send and receive messages.

Consider a client booting up at round $r_0 - 1$ and broadcasting the entropy η. η is received by the blockchain nodes at round r_0. We say the proof π is generated at round r^* and received by the client at round $r^* + 1$. Upon receiving the proof, the client sends $\pi[0]$ to full nodes and waits for Tx_η to become stable. Finally, the client terminates when Tx_η is stable in the chain of honest parties, i.e., at round $r^{**} \geq r^* + 3$.

Should the blockchain have fewer than k blocks at round r_0, a proof with fewer than k blocks before η is valid if its first block is the genesis block. However, if the chain is shorter than k blocks, the chain itself is already succinct and a light client is not needed.

Consider the blocktree of the execution at round r_0. We define $\mathsf{B}' \in C_{r_0}^{\cap}$ as the block with the greatest height which is a convergence event at r_0.

Lemma 5.3. B' *exists.*

Proof. The genesis block satisfies the definition of B'. □

We denote the round in which B' was produced as r', with $r' < r_0$. From Lemma A.21, we know that all honest blocks produced after r' extend B'.

Now, consider the blocktree of the execution at round r^*. We define $\tilde{\mathsf{B}}$ as the block with the greatest height that descends from B', was mined before r_0, and it is a convergence event at r^*. Because this block is similar to the blocks named $\tilde{\mathsf{B}}$ in Theorems A.22 and A.23, we re-use the name $\tilde{\mathsf{B}}$. We say $\tilde{\mathsf{B}}$ is produced at round $\tilde{r}$, with $r' \leq \tilde{r} < r_0$. We define $\tilde{S} := \{\tilde{r}, \ldots, r^*\}$.

Lemma 5.4. $\tilde{\mathsf{B}}$ *exists.*

Proof. B' satisfies the definition of $\tilde{\mathsf{B}}$. □

Lemma 5.5. *Acceptable blocks produced in $\tilde{S}$ descend from $\tilde{\mathsf{B}}$.*

Proof. This follows directly from Lemmas A.19 and A.20 (Appendix A.2).

□

As a consequence of Lemma 5.5 and Observation 2, all honest blocks produced in $\tilde{S}$ descend from $\tilde{\mathsf{B}}$.

Lemma 5.6. *All honest blocks produced in uniquely successful rounds within $\{\tilde{r}+1, \ldots, r_0\}$ have a parallel acceptable (by r^*) adversarial block.*

Proof. Because of the maximality (in terms of height) of $\tilde{\mathsf{B}}$, all blocks extending $\tilde{\mathsf{B}}$ and mined in uniquely successful rounds before r_0 have a parallel, acceptable adversarial block. □

For a set of consecutive rounds S, let $X(S)$ be honest queries, i.e., rounds in which at least one honest node found a block, $Y(S)$ be uniquely successful honest queries, i.e., rounds in which exactly one honest node found a block, and $Z(S)$ be adversarial queries, i.e., rounds in which the adversary found a block. We denote with $|X(S)|$, $|Y(S)|$, and $|Z(S)|$ the number of successful queries in $X(S)$, $Y(S)$, and $Z(S)$. These sets are defined in [17] or Appendix A.1.

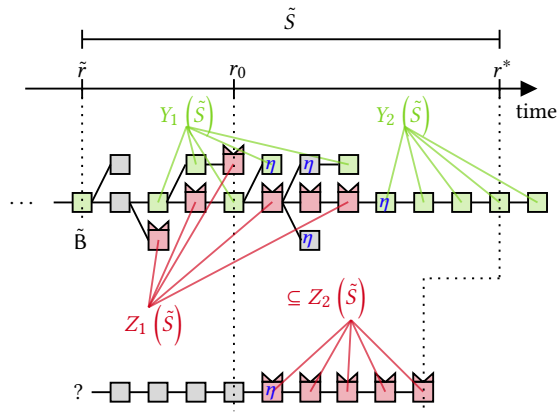

Figure 10: This figure illustrates the proof of Theorem 5.7.

Theorem 5.7 (Anchor). *In a typical execution, the block with η and its k subsequent blocks of the proof π that the client accepts, i.e., $\pi[k:]$, always extend $\tilde{\mathsf{B}}$.*

Proof. Let $Y(\tilde{S})$ be the set of honest uniquely successful queries within $\tilde{S}$, and $Z(\tilde{S})$ be the set of successful adversarial queries within $\tilde{S}$. Consider Figure 10 and let us define the following disjoint sets, Y_1, Y_2 and Z_1, Z_2, where $Y_1 \cup Y_2 = Y(\tilde{S})$ and $Z_1 \cup Z_2 = Z(\tilde{S})$.

(1) The queries of Z_1 produce blocks that extend $\tilde{\mathsf{B}}$.
(2) The queries of Z_2 produce blocks that do not extend $\tilde{\mathsf{B}}$.
(3) The queries of Y_1 produce blocks parallel to (at least) one of the blocks in Z_1 acceptable at r^*.
(4) The queries of Y_2 produce blocks not parallel to any of the blocks in Z_1 acceptable at r^*.

Claim 1. *If $|Y_2| = k+1$, at round r^*+1 the client has received a proof π with the blocks $\pi[k:]$ extending $\tilde{\mathsf{B}}$.*

This is true because in Y_2 there are no successful adversarial queries in $\tilde{S}$ producing blocks extending $\tilde{\mathsf{B}}$ and having parallel blocks. Furthermore, by definition of $\tilde{\mathsf{B}}$ and by causality, there cannot be successful adversarial queries outside of $\tilde{S}$ producing blocks extending $\tilde{\mathsf{B}}$. Yet, there can exist successful adversarial queries in Z_2 which produce blocks not extending $\tilde{\mathsf{B}}$.

Claim 2. *After r_0, honest parties do not extend blocks in Z_2.*

Blocks in Z_2 do not extend $\tilde{\mathsf{B}}$, and thus are not acceptable by r^*. Therefore honest parties do not extend them within $\tilde{S}$.

After r_0, honest nodes will include η in a block, if η was not included before. It follows that the block in Y_2 with the smallest height descends from a block including η. The k blocks produced by the remaining queries in Y_2 extend the block with η by one block each, as they are uniquely successful and there are no parallel, adversarial acceptable by r^* blocks.

Claim 3. *Independently of k, the block in Y_2 with the greatest height has all other blocks of Y_2 as ancestors.*

Towards a contradiction of Theorem 5.7, suppose that at round r^*+1 the client accepts a proof π which is generated at round r^* and where $\pi[k:]$ does not extend $\tilde{\mathsf{B}}$. For the client to receive such a proof, the number of blocks produced between r_0 and r^*

not extending $\tilde{\mathsf{B}}$, thus in Z_2, has to be larger than or equal to $k+1$. Therefore, also $|Z_2| \geq |Y_2|$. $|Y_2|$ can grow at most of 1 per round: if $|Y_2|$ was of $k+1$ in a previous round $r_p < r^*$, the light client would have received the proof in $r_p + 1$, contradicting the minimality of r^*. Now we count these sets. We have that $|Z| = |Z_1| + |Z_2|$ and $|Y| = |Y_1| + |Y_2|$. By definition of Y_1, we know that $|Y_1| \leq |Z_1|$. It follows, that $|Z| = |Z_1| + |Z_2| \geq |Y_1| + |Y_2| = |Y|$. However, from Lemma A.7 we know that $|\tilde{S}| \geq \lambda$ and thus, typicality bounds apply to this set of rounds. This means that $|Z| < |Y|$, which is a contradiction. This concludes the proof of Theorem 5.7. □

Lemma 5.8. $\tilde{\mathsf{B}} \in \pi$.

Proof. Because the block containing η, $\pi[k]$ or $\dot{\mathsf{B}}$, which was produced in round $\dot{r}$, is acceptable and has a height larger than any block that was honestly produced before it, we know from Theorem A.23 that the nearest convergence event at $\dot{r}$ has a height difference smaller than k blocks. □

Theorem 5.9. *In a typical execution, the first element $\pi[0]$ in the proof π accepted by Blink client at round r^* is an admissible block (cf. Definition 4.8).*

Proof. (Safety) From Lemma 5.8 we know that π includes $\tilde{\mathsf{B}}$. From Theorem A.22, we know that $\tilde{\mathsf{B}}$ is safe (i.e., $\tilde{\mathsf{B}} \in C^{\bigcup}_{r_0+u}[:-k]$). Since $\pi[0]$ is either $\tilde{\mathsf{B}}$ or an ancestor of $\tilde{\mathsf{B}}$, $\pi[0]$ is safe as well, i.e., $\pi[0] \in C^{\bigcup}_{r_0+u}[:-k]$.

(Liveness) Let l' be the height of B'. Define $\mathsf{B}'' := C^{\bigcap}_{r_0}[-k-1]$, and denote its height with height l''. Since B' is by definition either B'' (if the latter is uniquely successful and has no adversarial blocks at the same height by round r_0) or else an earlier block, it follows that $l'' \geq l'$.

At round r_0, honest users each have a local chain with height of at least $l'' + k$, because B'' is stable for all honest parties at round r_0. Since $\pi[k]$ includes η it has to be mined after r_0, which is the round in which η was released. This means, for the height l_k of $\pi[k]$, it holds that $l_k > l'' + k$.

As $\pi[0]$, with height l_0, is k blocks before $\pi[k]$, it holds that $l_0 = l_k - k$. Therefore $l_0 + k > l'' + k$, which means that $l_0 > l''$. However, since B'' was the last block in the stable intersection at round r_0, this implies $\pi[0] \notin C^{\bigcap}_{r_0}[:-k]$.

Therefore, at round r^* when the client accepts the a proof π, $\pi[0]$ is an admissible block. □

We observe that, after r_0, every honest chain tip descends from $\pi[0]$. We refer to $\pi[0]$ as new genesis block $\mathcal{G}'$.

Lemma 5.10 (New Genesis). *The longest chain rule applied to the genesis block $\mathcal{G}$ is consistent with the longest chain rule applied to $\mathcal{G}'$, with $\mathcal{G}'$ being an admissible block.*

Proof. Suppose there exists a longest chain that contains $\mathcal{G}$ but does not contain $\mathcal{G}'$. From admissible safety, we know that $\mathcal{G}'$ is stable for at least one honest user U, i.e., $\mathcal{G}' \in C^{\bigcap}_{r}[:-k]$. Since the longest chain does not contain $\mathcal{G}'$, honest users will adopt it in the next round, including the user U who has reported $\mathcal{G}'$ as stable. This violates common prefix. □

Corollary 5.11 (Chain Client Security for Blink). *Blink is* chain client secure *according to Definition 4.9.*

Given a client protocol Π which outputs a block B, one can build another protocol Π' that runs Π and reports the state commitment in B.[2]

Corollary 5.12. *For any client protocol Π that is chain client secure, the corresponding protocol Π' constructed in the above manner is* ledger client state secure *(Definition 4.5).*

This follows from a simple reduction since Π' merely reports the state commitment of B. If the state commitment was such that Π' is not ledger client state secure, the corresponding B cannot have been admissible. This concludes the proof of the main theorem Theorem 5.1, which is stated again here:

Theorem 5.13 (Ledger Client State Security for Blink). *Blink is* ledger client state secure *with the safety parameter v (Definition 4.5) being the wait time parameter u of liveness (Definition 4.8).*

5.2 Safety and Liveness of $\mathsf{B}_\eta := \pi[k]$

We now consider the case where Blink is used to verify a payment (or anything else that is in B_η), and we show that the corresponding proof size remains constant. We recall that in this use-case, after adopting $\mathcal{G}'$ and sending it to the provers, the Blink client maintains the longest chain descending from $\mathcal{G}'$. We now show that the entropy block will be eventually stable for all honest parties at most $3k$ consecutive blocks away from $\mathcal{G}'$.

Lemma 5.14 (Stability of Tx_η). *In a typical execution, a block B_η including Tx_η becomes stable for all honest parties at most at round $r_0 + u$, i.e., $\mathsf{B}_\eta \in C^{\bigcap}_{r_0+u}[:-k]$.*

Proof. It follows from the ledger liveness in Definition 4.4. □

Lemma 5.15 (Vicinity of Tx_η). *In a typical execution, a block B_η including Tx_η becomes stable for all honest parties at most $3k$ consecutive blocks away from $\mathcal{G}'$.*

Proof. Let r_g be the round at which the new genesis block is produced. By construction, k consecutive blocks are produced between r_g and r_0. By Lemma 5.14, the entropy transaction Tx_η becomes stable for all honest parties at most at $r_s = r_0 + u$. By the liveness of the chain (chain quality and chain growth), at r_s, at most $2k-1$ consecutive blocks are produced between $\mathcal{G}'$ and B_η (Corollary A.11), and Tx_η is at least k blocks deep in every honest party's chain. It follows that after at most $3k$ consecutive blocks are produced, B_η is stable for all honest parties. □

6 PRACTICALITY, LIMITATIONS, AND EXTENSIONS OF BLINK

State Commitments. In Section 3, we presented an application of Blink to build a light client that can be convinced about the current state of a ledger with optimal communication cost. This way, we enable the confirmation of historical transactions in the ledger, tracing back to its genesis. However, this application operates on the premise that each block embeds a state commitment to the

[2]For instance, this can easily be achieved for any blockchain protocol that has state commitments.

current ledger state. While several blockchains like ZCash, Nimiq, and Ethereum PoW uphold this premise, the most notable PoW blockchain, Bitcoin, does not incorporate state commitments in its block headers. NIPoPoWs, i.e., the polylogarithmic clients described in [12, 24], have the potential to be added retroactively via a velvet fork [26, 35]. The idea of introducing state commitments for Blink via velvet fork is appealing, however, its practical application is still undetermined.

Multiple Clients. Blink addresses the problem of one light client connecting to multiple full nodes and asking for the current state of the chain. In case we have multiple such requests, it is possible to compress the different entropy transactions using standard techniques. For example, multiple random strings can be ordered in a Merkle tree, and only the Merkle root is published on-chain within the entropy transaction. For this to be safe, each light client instance needs to have a Merkle proof of inclusion of its randomness in the tree.

Entropy Transaction Fees. Blink incurs on-chain fees which can be paid by light clients within entropy transactions. These fees can be paid in the form, for instance, of an anyone-can-spend output. The way the light client pays the on-chain cost for the entropy transaction can also be addressed in other ways on the application level: For instance, dedicated contracts or untrusted services can be designed such that clients' costs are mitigated.

Interactivity. Blink demands one round of interactivity between the client and the full nodes, unlike its predecessors that operate non-interactively [12, 21, 24]. This is the trade-off we incur for achieving a constant-sized proof instead of a polylogarithmic one as in [12, 21, 24]. We could remove the interactivity by introducing additional assumptions, for example: (i) a trusted committee service operates the client, similarly to the service provided by Chainlink for oracles, (ii) a random beacon acts as global entropy source and provides a service for Blink clients. However, both solutions come with drawbacks, i.e., centralization or a strong non-practical cryptographic primitive, respectively. It remains an open question whether designing a non-interactive light client with constant communication is possible without extra assumptions.

Variable Difficulty. Blink is analyzed in the static setting [17], i.e., the PoW difficulty remains the same throughout the protocol execution. In practice, Bitcoin uses a variable difficulty recalculation. Blink can still be used safely if we assume that parties agree on a difficulty beforehand, look it up on a trusted service (e.g., some blockchain explorer), or make some assumptions on the computational power of a potential adversary. Ideally, however, we can design a construction that is secure in the variable difficulty setting [16]. This challenge can be overcome by utilizing difficulty balloons to measure the current difficulty in a succinct fashion [36]. This approach, which is not unlike ours, utilizes entropy proofs to estimate (within some error) the current PoW difficulty of the network, by which point we can apply Blink as is. However, we anticipate that such an approach would only be secure under a weaker adversary that controls up to 1/3 of the computational power of the system. To provide an intuition behind this threshold, consider an adversary $t < 1/2$ that acts as follows: while measuring the difficulty, the adversary can abstain, thus creating a false sense of how many blocks she can produce in any given set of rounds. Thereby, she can take advantage of this false estimation to mine privately the required proof thereby violating the safety of Blink. We estimate that this adversarial advantage may be mitigated if honest nodes can produce double as many PoWs as the adversary.

Another approach would be to modify our light client construction by changing the selection rule for the proof: now the client would choose the proof with the most work after the intersection of all proofs within a given time window. We conjecture such an approach may alleviate the possible attack vectors of a minority adversary ($t < 1/2$), and we plan to explore it in future work.

7 EVALUATION

We evaluate the feasibility of Blink by measuring its proof size and its waiting time for Bitcoin. A Proof-of-Concept implementation of Blink can be found at [8] and all entropy transactions broadcast during this evaluation can be inspected at this Bitcoin address [1].

Our client uses the python bitcoin-utils library [7] to create the entropy transactions, and the python request HTTP library [9] to communicate via RPC APIs [6].

Experimental Setup. We deployed two mainnet Bitcoin full nodes running Bitcoin Core 25.0 and acting as untrusted provers: one was operated in-house on our own hardware (Central Europe) and the other one on a Vultr virtual machine (UK). We use two different deployments to emulate more realistic network conditions. The nodes maintain the entire history of all transactions of the ledger and they allow us to broadcast transactions to the Bitcoin network as well as to retrieve blocks, transactions, and Merkle proofs.

We ran our custom implementation client on commodity hardware. The client begins by sampling uniformly at random a 160-bit string η and creating the entropy transaction Tx_η by placing η in an `OP_RETURN` output. The size of Tx_η is 222 bytes. Then, the client connects to the two Bitcoin full nodes, broadcasts Tx_η, and waits for it to be k-confirmed (we set $k = 6$ according to Bitcoin folklore). When one of the two full nodes reports Tx_η k-deep, the client downloads and verifies the Blink proof π of size $2k + 1$ block headers, i.e., it checks blocks' parent-child relation and the PoW inequality.

Proof Size. We measure all the data received by the client from the full node that first reports Tx_η with k confirmations. This data amounts to 7912 bytes (7544 for π_0, and 368 for π_1, Algorithm 1).

The 7912 bytes of network data transmission required is due to the use of the inefficient JSON format and to the available standard RPC endpoints of the bitcoind full node. Using an optimized data transmission that avoids superfluous data, the total amount of data transmitted over the network can be brought down to 1646 bytes per prover connection (1040 bytes for the 13 block headers of 80 bytes each, 384 bytes for the Merkle inclusion proof consisting of 12 sibling SHA256 hashes of 256 bits each, and 222 bytes for the transaction Tx_η). In Table 2, for height 841368, we compare this to a full node that requires 684GB, an SPV client that requires 67.3MB, NIPoPoW and FlyClient clients that require 10.0KB and ~5KB, respectively, and to a PoW ZK-STARK-based client (ZeroSync[29]) that requires 197KB. We note that *the differences between these clients will be more pronounced as the blockchain grows.* We further note that proving Bitcoin's state with ZeroSync costs 4k USD (one-time cost), whereas Blink only incurs the cost of running a full node, e.g., ~ 15 USD a day.

Full node	SPV[28]	KLS[21], NIPoPoW [24], Mining LogSpace[23]	FlyClient[12]	ZK ZeroSync Client [29]	**Blink**
684GB	67.3MB	10KB	~5KB	197KB	1.6KB

Table 2: Comparison of light client solutions for Bitcoin mainnet at height 841368, using the parameter $k = 6$

Waiting Time. We measure the time it takes the client algorithm to run, averaging it over 10 runs. We broadcast the entropy transaction with a high-priority fee, which allows Tx_η to be included in the next 1 or 2 blocks. The average waiting time of the client to accept a proof is 59 minutes, with a standard deviation of 17 minutes. This is in accordance with the Bitcoin folklore belief of 6 blocks per hour. Any node that waits for 6 confirmations incurs the same waiting time, regardless of whether it is a full node or a light client. However, full nodes and SPV clients need to download a linear amount of data in the system's lifetime, while Blink requires only constant data in the chain's length to be downloaded.

8 CONCLUSION

This work presents Blink, the first Optimal Proof of Proof-of-Work client with constant communication complexity and without trusted setup. Blink allows to securely identify a state of the ledger which is safe and live by solely downloading a proof of $2k + 1$ consecutive blocks. We showcase how Blink can be leveraged in several different applications, ranging from verification of payments and state verification to bootstrapping and bridging. We prove Blink secure in the Bitcoin Backbone model against an adversary with minority computational power. Finally, we implemented Blink to verify its feasibility and we measured its proof size (experimental 7.9KB, 1.6 KB theoretical) and waiting time (59 ± 17 minutes).

ACKNOWLEDGMENTS

The authors thank Joachim Neu and Kostis Karantias for the helpful discussions in the early phase of the work. The work was partially supported by CoBloX Labs, by the European Research Council (ERC) under the European Union's Horizon 2020 research (grant agreement 771527-BROWSEC), by the Austrian Science Fund (FWF) through the SFB SpyCode project F8510-N and F8512-N, and the project CoRaF (grant agreement ESP 68-N), by the Austrian Federal Ministry for Digital and Economic Affairs, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association through the Christian Doppler Laboratory Blockchain Technologies for the Internet of Things (CDL-BOT), and by the WWTF through the project 10.47379/ICT22045.

REFERENCES

[1] [n. d.]. Bitcoin 137WhkasQG5zE1zpHZZijkGkSiEW2mo4qy Address . https://blockstream.info/address/137WhkasQG5zE1zpHZZijkGkSiEW2mo4qy.
[2] 2018. How to Validate Bitcoin Payments in Ethereum (for only 700k gas!). https://medium.com/summa-technology/cross-chain-auction-technical-f16710bfe69f.
[3] 2018. Merkle Mountain Ranges. https://github.com/opentimestamps/opentimestamps-server/blob/master/doc/merkle-mountain-range.md.
[4] 2018. Merkle Mountain Ranges (MMR). https://docs.grin.mw/wiki/chain-state/merkle-mountain-range/.
[5] 2023. Mina Docs. https://docs.minaprotocol.com/about-mina.
[6] 2024. Bitcoin RPC APIs, Chain Query. https://chainquery.com/bitcoin-cli.
[7] 2024. Bitcoin utils. https://pypi.org/project/bitcoin-utils/.
[8] 2024. OPoPoW Blink Client Implementation. https://anonymous.4open.science/r/OPoPoW-Blink-Client/README.md.
[9] 2024. Python Request Library. https://pypi.org/project/requests/.
[10] Shresth Agrawal, Joachim Neu, Ertem Nusret Tas, and Dionysis Zindros. 2023. Proofs of Proof-of-Stake with Sublinear Complexity. arXiv:2209.08673 [cs.CR]
[11] Joseph Bonneau, Izaak Meckler, Vanishree Rao, and Evan Shapiro. 2020. Coda: Decentralized Cryptocurrency at Scale. https://eprint.iacr.org/2020/352.pdf.
[12] Benedikt Bünz, Lucianna Kiffer, Loi Luu, and Mahdi Zamani. 2020. FlyClient: Super-Light Clients for Cryptocurrencies. In *2020 IEEE Symposium on Security and Privacy (SP)*. 928–946. https://doi.org/10.1109/SP40000.2020.00049
[13] Dario Catalano and Dario Fiore. 2013. Vector Commitments and Their Applications. In *Public-Key Cryptography – PKC 2013*. Springer Berlin Heidelberg.
[14] Panagiotis Chatzigiannis, Foteini Baldimtsi, and Konstantinos Chalkias. 2021. SoK: Blockchain Light Clients. In *IACR Cryptology ePrint Archive*. https://api.semanticscholar.org/CorpusID:245908572
[15] Stelios Daveas, Kostis Karantias, Aggelos Kiayias, and Dionysis Zindros. 2020. A Gas-Efficient Superlight Bitcoin Client in Solidity. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. 132–144.
[16] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. 2017. The Bitcoin Backbone Protocol with Chains of Variable Difficulty. In *Advances in Cryptology – CRYPTO 2017*, Jonathan Katz and Hovav Shacham (Eds.). Springer International Publishing.
[17] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. 2024. The Bitcoin Backbone Protocol: Analysis and Applications. In *Journal of the ACM (to appear)*.
[18] Peter Gaži, Aggelos Kiayias, and Dionysis Zindros. 2018. Proof-of-Stake Sidechains. Cryptology ePrint Archive, Paper 2018/1239. https://eprint.iacr.org/2018/1239 https://eprint.iacr.org/2018/1239.
[19] Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. 2021. Poseidon: A New Hash Function for Zero-Knowledge Proof Systems. In *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, Michael D. Bailey and Rachel Greenstadt (Eds.). USENIX Association, 519–535. https://www.usenix.org/conference/usenixsecurity21/presentation/grassi
[20] Kostis Karantias, Aggelos Kiayias, and Dionysis Zindros. 2020. Compact Storage of Superblocks for NIPoPoW Applications. In *Mathematical Research for Blockchain Economy*, Panos Pardalos, Ilias Kotsireas, Yike Guo, and William Knottenbelt (Eds.). Springer International Publishing.
[21] Aggelos Kiayias, Nikolaos Lamprou, and Aikaterini-Panagiota Stouka. 2016. Proofs of Proofs of Work with Sublinear Complexity. In *Financial Cryptography and Data Security*, Jeremy Clark, Sarah Meiklejohn, Peter Y.A. Ryan, Dan Wallach, Michael Brenner, and Kurt Rohloff (Eds.). Springer Berlin Heidelberg.
[22] Aggelos Kiayias, Nikos Leonardos, and Dionysis Zindros. 2021. Mining in Logarithmic Space. In *CCS*. ACM, 3487–3501.
[23] Aggelos Kiayias, Nikos Leonardos, and Dionysis Zindros. 2021. Mining in Logarithmic Space *(CCS '21)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3460120.3484784
[24] Aggelos Kiayias, Andrew Miller, and Dionysis Zindros. 2020. Non-interactive Proofs of Proof-of-Work. In *Financial Cryptography and Data Security*, Joseph Bonneau and Nadia Heninger (Eds.). Springer International Publishing.
[25] Aggelos Kiayias, Andrianna Polydouri, and Dionysis Zindros. 2021. The velvet path to superlight blockchain clients. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*. 205–218.
[26] Aggelos Kiayias, Andrianna Polydouri, and Dionysis Zindros. 2021. The Velvet Path to Superlight Blockchain Clients. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies (AFT '21)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3479722.3480999
[27] Aggelos Kiayias and Dionysis Zindros. 2019. Proof-of-Work Sidechains. In *IACR Cryptology ePrint Archive*. https://api.semanticscholar.org/CorpusID:53243925
[28] Satoshi Nakamoto. 2009. Bitcoin: A Peer-to-Peer Electronic Cash System. http://bitcoin.org/bitcoin.pdf.
[29] Linus Robin, George Lukas, Milson Andrew, and Steffens Tino. 2024. ZeroSync - STARK proofs for Bitcoin. https://github.com/ZeroSync/header_chain.
[30] Giulia Scaffino, Lukas Aumayr, Zeta Avarikioti, and Matteo Maffei. 2023. Glimpse: On-Demand PoW Light Client with Constant-Size Storage for DeFi. In *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9-11, 2023*, Joseph A. Calandrino and Carmela Troncoso (Eds.). USENIX Association. https://www.usenix.org/conference/usenixsecurity23/presentation/scaffino
[31] Ertem Nusret Tas, Dionysis Zindros, Lei Yang, and David Tse. 2022. Light Clients for Lazy Blockchains. Cryptology ePrint Archive, Paper 2022/384. https://eprint.iacr.org/2022/384
[32] Psi Vesely, Kobi Gurkan, Michael Straka, Ariel Gabizon, Philipp Jovanovic, Georgios Konstantopoulos, Asa Oines, Marek Olszewski, and Eran Tromer. 2023. Plumo: An Ultralight Blockchain Client. https://celo.org/papers/plumo.
[33] Tiancheng Xie, Jiaheng Zhang, Zerui Cheng, Fan Zhang, Yupeng Zhang, Yongzheng Jia, Dan Boneh, and Dawn Song. 2022. zkBridge: Trustless Cross-chain Bridges Made Practical. In *ACM SIGSAC Conference on Computer and Communications Security, CCS*. https://doi.org/10.1145/3548606.3560652
[34] Alexei Zamyatin, Dominik Harz, Joshua Lind, Panayiotis Panayiotou, Arthur Gervais, and William Knottenbelt. 2019. XCLAIM: Trustless, Interoperable, Cryptocurrency-Backed Assets. In *2019 IEEE Symposium on Security and Privacy*

(SP).
[35] A. Zamyatin, N. Stifter, A. Judmayer, P. Schindler, E. Weippl, and W. J. Knottenbelt. 2018. A Wild Velvet Fork Appears! Inclusive Blockchain Protocol Changes in Practice. Springer-Verlag, Berlin, Heidelberg. https://doi.org/10.1007/978-3-662-58820-8_3
[36] Dionysis Zindros. 2020. *Decentralized Blockchain Interoperability.* Ph. D. Dissertation. University of Athens.

A ANALYSIS

A.1 Background from the Bitcoin Backbone [17]

We now introduce notation, definitions, theorems, and lemmas stated in [17] which will be necessary for our analysis.

The properties of blockchain protocols defined in the backbone model are presented below. Such properties are defined as predicates over the random variable $\text{view}^{t,n}_{\Pi,A,Z}$ by quantifying over all possible adversaries A and environments Z that are polynomially bounded. Note that blockchain protocols typically satisfy properties with a small probability of error in a security parameter κ (or others). The probability space is determined by random queries to the random oracle functionality and by the private coins of all interactive Turing machine instances.

Definition A.1 (Common Prefix Property [17]). The common prefix property Q_{cp} with parameter $k \in \mathbb{N}$ states that for any pair of honest players P_1, P_2 adopting the chains C_1, C_2 at rounds $r_1 \leq r_2$ in $\text{view}^{t,n}_{\Pi,A,Z}$ respectively, it holds that $C_1^{|k|} \preceq C_2$.

Definition A.2 (Chain Quality Property [17]). The chain quality property Q_{cq} with parameters $\mu \in \mathbb{R}$ and $\ell \in \mathbb{N}$ states that for any honest party P with chain C in $\text{view}^{t,n}_{\Pi,A,Z}$, it holds that for any ℓ consecutive blocks of C, the ratio of honest blocks is at least μ.

Definition A.3 (Chain Growth Property [17]). The chain growth property Q_{cg} with parameters $\tau \in \mathbb{R}$ and $s \in \mathbb{N}$ states that for any honest party P that has a chain C in $\text{view}^{t,n}_{\Pi,A,Z}$, it holds that after any s consecutive rounds, it adopts a chain that is at least $\tau \cdot s$ blocks longer than C.

Closely following[17], we will call a query $q \in \mathbb{N}$ of a party successful if it returns a valid solution to the PoW. For each round $i, j \in [q]$, and $k \in [t]$, we define Boolean random variables X_i, Y_i, and Z_{ijk} as follows. If at round i an honest party obtains a PoW, then $X_i = 1$, otherwise $X_i = 0$. If at round i exactly one honest party obtains a PoW, then $Y_i = 1$, otherwise $Y_i = 0$. Regarding the adversary, if at round i, the j-th query of the k-th corrupted party is successful, then $Z_{ijk} = 1$, otherwise $Z_{ijk} = 0$. Define also $Z_i = \sum_{k=1}^{t} \sum_{j=1}^{q} Z_{ijk}$. For a set of rounds S, let $X(S) = \sum_{r \in S} X_r$ and similarly define $Y(S)$ and $Z(S)$. Further, if $X_i = 1$, we call i a successful round and if $Y_i = 1$, a uniquely successful round. We denote with f the probability that at least one honest party succeeds in finding a PoW in a round.

Definition A.4 (Typical Execution [17]). An execution is (ϵ, λ)-typical (or just typical), for $\epsilon \in (0, 1)$ and integer $\lambda \geq 2/f$, if, for any set S of at least λ consecutive rounds, the following hold.

(a) $(1-\epsilon)\mathbb{E}[X(S)] < X(S) < (1+\epsilon)\mathbb{E}[X(S)]$ and $(1-\epsilon)\mathbb{E}[Y(S)] < Y(S)$.
(b) $Z(S) < \mathbb{E}[Z(S)] + \epsilon\mathbb{E}[X(S)]$.
(c) No insertions, no copies, and no predictions occurred.

Let n be the number of consensus nodes, out of which t are controlled by the adversary. Let Q be an upper bound on the number of computation or verification queries to the random oracle. Let L be the total number of rounds in the execution, and λ, κ security parameters. Finally, we denote with ν the min-entropy of the value that the miner attempts to insert in the chain.

Theorem A.5 (Theorem 4.5 in [17]). *An execution is not typical with probability less than*

$$\epsilon_{\text{typ}} = 4L^2 e^{-\Omega(\epsilon^2 \lambda f)} + 3Q^2 2^{-\kappa} + [(n-t)L]^2 2^{-\nu}.$$

Lemma A.6 (Lemma 4.6 in [17]). *The following hold for any set S of at least λ consecutive rounds in a typical execution. For $S = \{i : r < i < s\}$ and $S' = \{i : r \leq i \leq s\}$, $Z(S') < Y(S)$.*

Lemma A.7 (Lemma 4.8 in [17], (aka Patience Lemma)). *In a typical execution, any $k \geq 2\lambda f$ consecutive blocks of a chain have been computed in more than $\frac{k}{2f}$ consecutive rounds.*

Lemma A.8 (Lemma 4.1 in [17], (aka Pairing Lemma)). *Suppose the k-th block* B *of a chain C was computed by an honest party in a uniquely successful round. Then the k-th block a chain C' either is* B *or has been computed by the adversary.*

Lemma A.9 (Lemma 4.2 in [17], (aka Chain Growth)). *Suppose that at round r an honest party has a chain of length l. Then, by round $s \geq r$, every honest party has adopted a chain of length at least $l + \sum_{i=r}^{s-1} X_i$.*

Theorem A.10 (Theorem 4.11 in [17], (aka Chain Quality)). *In a typical execution the chain quality property holds with parameters $\ell \geq 2\lambda f$ and*

$$\mu = 1 - \frac{1+f}{(1-f)(1-\epsilon)} \cdot \frac{t}{n-t} - \frac{(1+f)\epsilon}{1-\epsilon}$$
$$> 1 - \frac{1}{1-2\delta/3} \cdot \frac{t}{n-t} - \frac{\delta/3}{1-\delta/3} \xrightarrow{\delta \to 0} \frac{n-2t}{n-t}$$

Corollary A.11 (Corollary 4.12 in [17]). *In a typical execution the following hold.*

- *Any $\lceil 2\lambda f \rceil$ consecutive blocks in the chain of an honest party contain at least one honest block.*
- *For any λ consecutive rounds, the chain of an honest party contains an honest block computed in one of these rounds.*

In our analysis, we assume a typical execution in all proofs. We note that from Theorem A.5 typical execution fails with negligible probability, resulting in our proofs holding with overwhelming probability.

A.2 Preliminaries

In this section, we introduce some definitions, observations, and lemmas that will be used as building blocks in the formal analysis of Blink security (Section 5.1).

Let H be a hash function modeled as a Random Oracle, and let T be the target hash value used by parties for solving the PoW. Given a chain C and a block b to be inserted in the chain, consider the hash $h = H(C[-1], b)$ of these values, and let ctr be a counter.

Definition A.12 (PoW Inequality). The PoW inequality holds if $\text{H}(ctr, h) < \text{T}$.

If a ctr fulfilling the PoW inequality is found, the chain C is extended by the block b (which includes ctr). If no suitable ctr is found, the chain remains unaltered.

Definition A.13 (Valid Chain). A chain C is (syntactically) *valid* if:

- $C = \emptyset$, or
- $C[: -1]$ is valid and the PoW inequality holds for $h = H(C[-2], C[-1])$.

Definition A.14 (Valid Block). A block is *valid* if it belongs to a valid chain.

Lemma A.15. *The following equality holds:*

$$C_r^{\bigcup}[:-k] = \bigcup_{P \in \mathcal{H}} C_r^P[:-k] \tag{1}$$

Proof. We observe that $C_r^{\bigcup}$ is a tree where each leaf C_r^P corresponds to the view of the chain of (at least) one honest party P at some round r. $C_r^{\bigcup}[:-k]$ is the result of taking $C_r^{\bigcup}$ and removing the last k blocks from each of the leaves of the tree. $\bigcup_{P \in \mathcal{H}} C_r^P[:-k]$ is the result of taking all the chains of honest parties at round r, chopping off the last k blocks and taking the union of these chains. By common prefix, honest parties' chains can only diverge by less than k blocks; therefore, $C_r^{\bigcup}[:-k]$ is a chain such that $C_r^{\bigcup}[:-k] = \bigcup_{P \in \mathcal{H}} C_r^P[:-k]$, with some honest parties being aware of all the blocks in it, and some others lagging behind. □

Definition A.16 (Acceptable Chain at r). A valid chain C is *acceptable at round r*, if

- $C = \emptyset$, or
- $C[: -1]$ is acceptable at r, and either $C \preceq C_r^{\bigcap}[: -k]$ or $C_r^{\bigcap}[: -k] \preceq C$.

An important notion we use is an *acceptable block*. Intuitively, an acceptable block is a block to which honest parties can switch to without violating common prefix. Honest nodes will never switch to chains containing non-acceptable blocks.

Definition A.17 (Acceptable Block at r). A block is *acceptable at r* if it belongs to an acceptable chain at r.

Observation 1. *If a block is stable in an honest party's view, it is also acceptable.*

Observation 2. *All honestly produced blocks are acceptable in the round in which they are produced.*

Observation 3. *All honestly produced blocks only descend from blocks that are acceptable in the round in which the former are produced.*

Observation 4. *Any block* B *produced in round r_B and acceptable in round $r \geq r_B$, is also acceptable in all rounds in the set of consecutive rounds $\{r_B, \ldots, r\}$.*

Definition A.18 (Convergence Event at r). A block B is a *convergence event at round r* if it is produced in a uniquely successful round r_B and, by round $r \geq r_B$, it does not have a parallel acceptable block in any round in the set of consecutive rounds $\{r_B, \ldots, r\}$.

Observation 5. *A convergence event is always honestly produced.*

Lemma A.19. *If a block* B *produced in round r_B is a convergence event at round r, it is a convergence event in all rounds in the set of consecutive rounds $\{r_B, \ldots, r\}$.*

Proof. By definition, there are no acceptable blocks at $\{r_B, \ldots, r\}$ parallel to B. Therefore, B fulfills the definition of convergence event at all rounds $\{r_B, \ldots, r\}$. □

Lemma A.20. *An acceptable block* B *at r must descend from all convergence events at r with a height smaller or equal to* B*'s height.*

Proof. Towards a contradiction, suppose there exists a convergence event $\hat{\mathsf{B}}$ at r, such that B does not descend from $\hat{\mathsf{B}}$. There must be a block B' parallel to $\hat{\mathsf{B}}$ from which B descends. Because B is acceptable at r, by definition, B' needs to be acceptable at r. However, both B' acceptable and $\hat{\mathsf{B}}$ being a convergence event, imply $\hat{\mathsf{B}} = \mathsf{B}'$. Thus, B' descends from $\hat{\mathsf{B}}$, reaching a contradiction. □

Lemma A.21. *Let r be the round in which a block* B *was produced. For any block* B *and any round r', for which* B *is a convergence event at r' and* $\mathsf{B} \in C_{r'}^{\bigcap}[: -k]$*, blocks acceptable at any round after r always extend* B.

Proof. From Observation 2 and Lemma A.20, honest parties will extend B in the rounds between r and r' (included). B is stable for all honest parties at round r'. Therefore, after r' all honest parties only extend B, otherwise common prefix is violated. □

We denote with $|X(S)|$, $|Y(S)|$, and $|Z(S)|$ the number of successful queries X, Y, and Z in a set of consecutive rounds S.

Theorem A.22 (General Eventual Stability). *Consider a convergence event $\tilde{\mathsf{B}}$ at r^*, which was produced in round $\tilde{r}$ and it has height $\tilde{l}$. If there exists a block with height $l \geq \tilde{l} + k$ in $C_{r^*}^{\bigcup}$ then, in a typical execution, $\tilde{\mathsf{B}}$ becomes stable for at least one honest party at most at round $\tilde{r} + u$, i.e., $\tilde{\mathsf{B}} \in C_{\tilde{r}+u}^{\bigcup}[: -k]$.*

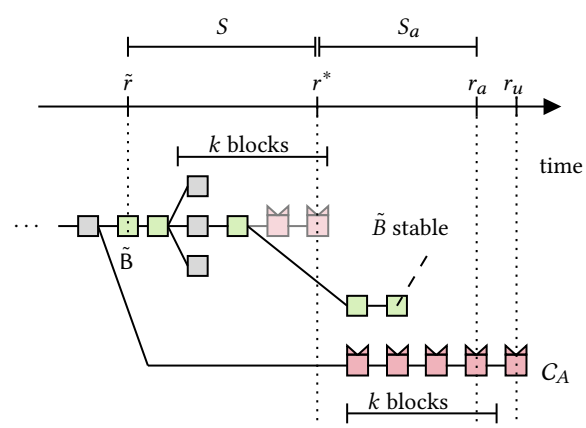

Figure 11: This figure illustrates the proof of Theorem A.22.

Proof. Consider Figure 11. Let $\tilde{l}$ be the height of $\tilde{\mathsf{B}}$ and $\tilde{r}$ the round at which $\tilde{\mathsf{B}}$ was produced. Since $\tilde{\mathsf{B}}$ is a convergence event,

we know that it was honestly produced. Thus, at any round $r > \tilde{r}$, honest parties have adopted a chain of length at least $\tilde{l}$.

By round r^*, since $\tilde{\mathsf{B}}$ is a convergence event and due to causality, the acceptable blocks with height larger than $\tilde{l}$ have been mined at or after $\tilde{r}$. Since the blocktree at round r^* contains a block with a height $l \geq \tilde{l} + k$, at least k consecutive blocks were mined in the set of consecutive rounds $S' := \{\tilde{r}, \ldots, r^*\}$. Let $S := \{\tilde{r} - 1, \ldots, r^* + 1\}$. We can thus apply the patience lemma (Lemma A.7) to this set of rounds, which means that $|S| > \lambda$ and typicality bounds apply. In particular, $|X(S)| > |Z(S')|$, which implies $|X(S)| > \frac{k}{2}$. From chain growth (Lemma A.9), we know that in every round r in which there is at least one honest block found, i.e. $X_r = 1$, honest parties increase the length of their chains by (at least) 1. It follows that in any round $r > r^*$, honest parties have adopted a chain longer than $\tilde{l} + \frac{k}{2}$.

Towards contradiction, suppose that $\tilde{\mathsf{B}} \notin C^{\cup}_{\tilde{r}+u}[: -k]$. This means that there exists a round r_u in which all honest parties have adopted a stable chain C_A of length $l_A \geq \tilde{l} + k$ which excludes $\tilde{\mathsf{B}}$. We note that *all* honest parties must have adopted C_A, otherwise common prefix would be violated. It follows that: (i) $r_u < r + u$ because otherwise, by chain quality and chain growth, at round $r + u$, $\tilde{\mathsf{B}}$ would be stable; (ii) $r^* < r_u$ because, by definition of convergence event at r^*, $\tilde{\mathsf{B}}$ does not have any parallel acceptable adversarial block at round r^*. The blocks $C_A[-(k+1) :]$ have a height of at least $\tilde{l}$ and are produced after r^*. We now proceed with a counting argument for the set of rounds $S_a := \{r^*, \ldots, r_a\}$, where $r_a \leq r_u$ is the first round in which C_A contains at least k blocks with a height higher or equal to $\tilde{l}$. Again, since (at least) k consecutive blocks were mined in S_a and we can apply Lemma A.7 to this set of rounds, which means that $|S_a| > \lambda$ and typicality bounds apply. In particular, $|X(S_a)| > |Z(S'_a)|$, which implies $|X(S_a)| > \frac{k}{2}$.

From Lemma A.9, we know that there are at least $|X(S_a)|$ consecutive blocks extending $\tilde{\mathsf{B}}$. From Lemma A.7, we know that $|S_a| \geq \lambda$, which means that typicality bounds apply, i.e., $|X(S_a)| > |Z(S_a)|$, hence $|X(S_a)| > \frac{k}{2}$ and $Z(S_a) < \frac{k}{2}$. The chain C_A which extends B' but not $\tilde{\mathsf{B}}$, has a length of at most $\tilde{l} - 1 + \frac{k}{2}$, as honest parties do not extend shorter chains. Therefore, at round r_a, all honest parties cannot have adopted C_A, because they have a chain of length at least $\tilde{l} + k$, which includes $\tilde{\mathsf{B}}$. This concludes the contradiction.

□

Theorem A.23 (General Vicinity). *Consider any acceptable block $\dot{\mathsf{B}}$ at round r, produced in $\dot{r}$ and having a height larger than any honestly produced block in any round before $\dot{r}$. Let $\tilde{\mathsf{B}}$ be a convergence event at $\dot{r}$, such that $\tilde{\mathsf{B}}$ is the closest convergence event to $\dot{\mathsf{B}}$ in terms of height, and such that the height $\tilde{l}$ of $\tilde{\mathsf{B}}$ is smaller or equal to the height $\dot{l}$ of $\dot{\mathsf{B}}$, i.e., $\tilde{l} \leq \dot{l}$. In a typical execution, $\dot{l} - \tilde{l} < k$.*

Proof. Consider Figure 12. Let $\tilde{r} \leq \dot{r}$ be the round in which $\tilde{\mathsf{B}}$ was produced. We now look at the blocks $\{\mathsf{B}\}_{Y(S)}$ that were honestly produced in the uniquely successful rounds in $S := \{\tilde{r} + 1, \ldots, \dot{r} - 1\}$, i.e., $Y(S)$. By definition, every block $\mathsf{B} \in \{\mathsf{B}\}_{Y(S)}$ has a height smaller than $\dot{\mathsf{B}}$. However, due to Lemma A.20, every block $\mathsf{B} \in \{\mathsf{B}\}_{Y(S)}$ also extends $\tilde{\mathsf{B}}$ and thus has a height larger than $\tilde{\mathsf{B}}$.

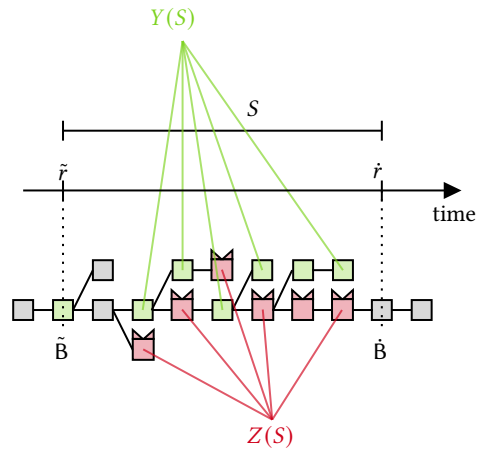

Figure 12: This figure illustrates the proof of Theorem A.23.

Because $\tilde{\mathsf{B}}$ is the nearest convergence event at $\dot{r}$, any block $\mathsf{B} \in \{\mathsf{B}\}_{Y(S)}$ needs to have a parallel, acceptable at $\dot{r}$ (and thus mined at or before $\dot{r}$) block. Otherwise, B would be the nearest convergence event to $\dot{\mathsf{B}}$. Because these parallel blocks are acceptable, they need to extend $\tilde{\mathsf{B}}$ (Lemma A.20) and thus by causality, need to have been produced at or after $\tilde{r}$ and at or before $\dot{r}$, which means they are produced in $S' := \{\tilde{r}, \ldots, \dot{r}\}$. Additionally, from Lemma A.8, we know that these parallel blocks need to be adversarially produced. Thus, there needs to be at least one successful adversarial query within S' for each uniquely successful round in S, i.e., $|Z(S')| \geq |Y(S)|$.

Due to causality, the blocks between $\tilde{\mathsf{B}}$ and $\dot{\mathsf{B}}$ need to have been produced in $S'' := \{\tilde{r}, \ldots, \dot{r} - 1\}$. Suppose towards a contradiction, the difference in height between $\dot{\mathsf{B}}$ and $\tilde{\mathsf{B}}$ is k or more. From Lemma A.7 we know that $|S''| > \lambda$, thus $|S| \geq \lambda$, and thus, typicality bounds apply to this set of rounds. Thus, by Lemma A.6 it holds that $|Z(S')| < |Y(S)|$, which contradicts the above. □