

Bootstrapping the Blockchain — Directly

Juan A. Garay
Yahoo Research
garay@yahoo-inc.com

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

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

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

October 31, 2017

Abstract

The Bitcoin backbone protocol [Eurocrypt 2015] extracts basic properties of Bitcoin’s underlying *blockchain* data structure, such as “common prefix” and “chain quality,” and shows how fundamental applications including consensus and a robust public transaction ledger can be built on top of them. The underlying assumptions are “proofs of work” (POWs), adversarial hashing power strictly less than $1/2$ *and* no adversarial pre-computation—or, alternatively, the existence of an unpredictable “genesis” block.

In this paper we first show how to remove the latter assumption, presenting a “bootstrapped” Bitcoin-like blockchain protocol relying on POWs that builds genesis blocks “from scratch” in the presence of adversarial pre-computation. Importantly, the round complexity of the genesis block generation process is independent of the number of participants.

Next, we consider applications of our construction, including a PKI generation protocol and a consensus protocol without trusted setup assuming an honest majority (in terms of computational power). Previous results in the same setting (unauthenticated parties, no trusted setup, POWs) required a round complexity linear in the number of participants.

1 Introduction

As the first decentralized cryptocurrency, Bitcoin [33] has ignited much excitement, not only for its novel realization of a central bank-free financial instrument, but also as an alternative approach to classical distributed computing problems, such as reaching agreement distributedly in the presence of misbehaving parties. Formally capturing such reach has been the intent of several recent works, notably [21], where the core of the Bitcoin protocol, called the Bitcoin *backbone*, is extracted and analyzed. The analysis includes the formulation of fundamental properties of its underlying *blockchain* data structure, which parties (“miners”) maintain and try to extend by generating “proofs of work” (POW, aka “cryptographic puzzle” [16, 37, 3, 24])¹, called *common prefix* and *chain quality*. It is then shown in [21] how applications such as consensus (aka Byzantine agreement) [36, 31] and a robust public transaction ledger (i.e., Bitcoin) can be built “on top” of such properties, assuming that the hashing power of an adversary controlling a fraction of the parties is strictly less than $1/2$.

*Research supported by ERC project CODAMODA #25915.

¹In Bitcoin, solving a proof of work essentially amounts to brute-forcing a hash inequality based on SHA-256.

Importantly, those properties hold assuming that all parties—honest and adversarial—“wake up” and start computing at the same time, or, alternatively, that they compute on a common random string only made available at the exact time when the protocol execution is to begin (see further discussion under related work below). Indeed, the coinbase parameter in Bitcoin’s “genesis” block, hardcoded into the software, contains text from *The Times* 03/Jan/2009 issue [5], arguably unpredictable.

While satisfactory in some cases, such a trusted setup/behavioral assumption might be unrealistic in other POW-based systems where details may have been released a lot earlier than the actual time when the system starts to run. A case in point is Ethereum², which was discussed for over a year before the system officially kicked off. That’s from a practical point of view. At a foundational level, one would in addition like to understand what kind of cryptographic primitives can be realized without any trusted setup assumption and based on POWs, and whether that is in particular the case for the Bitcoin backbone functionality and its enabling properties mentioned above.

The former question was recently considered by Andrychowicz and Dziembowski [1], who, building on previous suggestions by Aspnes *et al.* [2] of using POWs as an identity-assignment tool and constructions by Fitzi *et al.* [19, 12] showing how to morph “graded” consistency into global consistency, showed how to create a consistent PKI using POWs and no other trusted setup, which can then be used to run secure computation protocols (e.g., [38, 23]) and realize any cryptographic functionality assuming an honest majority among parties. While this in principle addresses the foundational concerns, it leaves open the questions of doing it in *scalable* way—i.e., with round complexity independent of the number of parties, and in the context of blockchain protocols in particular, designing one that is *provably* secure without a trusted setup.

Our contributions. In this paper we answer the above questions. First, we present a Bitcoin-like protocol that neither assumes a simultaneous start nor the existence of an unpredictable genesis block, and has round complexity essentially independent of the number of participants³. Effectively, the protocol, starting “from scratch,” enables the coexistence of multiple genesis blocks with blockchains stemming from them, eventually enabling the players to converge to a single blockchain. This takes place despite the adversary being allowed (polynomial in the security parameter) pre-computation time. We work in the same model as [21] and we assume a 1/2 bound on adversarial hashing power. We call this protocol the *bootstrapped* (Bitcoin) backbone protocol. A pictorial overview of the protocol’s phases, preceded by a period of potential precomputation by the corrupt players, is given in Figure 1.




Figure 1: *Timeline and phases of the bootstrapped Bitcoin backbone protocol.*

Second, we present applications of our bootstrapped construction, starting with its primary/original one: a distributed ledger, i.e., a public and permanent summary of all transactions that honest par-

²The Ethereum Project, <https://www.ethereum.org/>.

³“Essentially” because even though there will be a dependency of the round complexity of the setup phase on the probability of computing POWs, which in turn depends on the number of parties, this dependency can be made small enough so as to be considered a constant. See Remark 3.

ties can agree on as well as add their own, despite the potentially disruptive behavior of parties harnessing less than 1/2 of the hashing power. This entails proving that the ledger’s required security properties (Persistence and Liveness — cf. [21]) hold in a genesis block-less setting.

Next, we consider the problem of setting up a PKI in our unauthenticated network setting *from scratch*. As mentioned above, the idea of using POWs as an identity-assignment tool was put forth by Aspnes *et al.* [2]. Here we build on this idea as well as on the “2-for-1 POWs” technique from [21] to use our bootstrapped protocol to assign identities to parties. The assignment relation will possibly assign more than one identity to the same party, while guaranteeing that the majority of them is assigned to honest parties. Such an identity infrastructure/“pseudonymous PKI” has numerous applications, including the bootstrapping of a proof-of-stake protocol [28, 30], and the election of honest-majority “subcommittees,” which would enable the application of traditional Byzantine fault-tolerant techniques for ledger creation and maintenance (cf. [7]) to permissionless (as opposed to permissioned) networks.

Finally, applying the 2-for-1 POWs technique we can also solve the consensus (aka Byzantine agreement) problem [36, 31] probabilistically and from scratch, even if the adversary has almost the same hashing power as the honest parties⁴, and with round complexity independent of the number of parties. Indeed, all our protocols have round complexity linear in the security parameter, and enjoy simultaneous termination.

Related work. Nakamoto [32] proposed Bitcoin, the first decentralized currency system based on POWs while relaxing the anonymity property of a digital currency to mere pseudonymity. This work was followed by a multitude of other related proposals including Litecoin⁵, Primecoin [29], and Zerocash [4], and further analysis improvements (e.g., [18, 17]), to mention a few.

As mentioned above, we work in a model that generalizes the model put forth by Garay *et al.* [21], who abstracted out and formalized the core of the Bitcoin protocol—the Bitcoin *backbone*. As presented in [21], however, the protocol considers as valid any chain that extends the empty chain, which is not going to work in our model. Indeed, if the adversary is allowed polynomial-time pre-computation, he can prepare a very long, private chain; then, by revealing blocks of this chain at the rate that honest players compute new blocks, he can break security. As also mentioned above, to overcome this problem one can assume that at the time honest parties start the computation, they have access to a fresh common random string (a “genesis” block). Then, if we consider as valid only the chains that extend this block, all results proved in [21] follow, since the probability that the adversary can use blocks mined before honest players “woke up” is negligible in the security parameter. In this paper we show how to establish such genesis block directly, and in a number of rounds essentially independent of the number of participants.

To our knowledge, the idea of using POWs to distributedly agree on something (specifically, a PKI) in an unauthenticated setting with no trusted setup was first put forth by Aspnes *et al.* [2], who suggested to use them as an identity-assignment tool as a way to combat *Sybil* attacks [14], and in such a way that the number of identities assigned to the honest and adversarial parties can be made proportional to their aggregate computational power, respectively. For example, by assuming that the adversary’s computational power is less than 50%, one of the algorithms in [2] results in a number of adversarial identities less than half of that obtained by the honest parties. By running this procedure in a pre-processing stage, it is then suggested in [2] that a standard authenticated broadcast protocol (specifically, the one by Dolev and Strong [13]) could be run. Such protocols, however, would require that the PKI be *consistent*, details of which are not laid out in [2].

⁴Thus marking a contrast with the $\frac{2}{3}$ lower bound for consensus on the number of honest parties in the traditional network setting with no setup [6].

⁵<http://www.litecoin.com>.

They are in [1], where Andrychowicz and Dziembowski address the more general goal of secure computation in this setting based on POWs, as mentioned earlier; the POWs are used to build a “graded” PKI, where keys have “ranks.” The graded PKI is an instance of a “graded agreement,” or “partial consistency” problem [19, 12, 20], where honest parties do not disagree “by much,” according to some metric. In [19], Fitzi calls this the *b-set-neighboring* problem (“proxcast” in [12]), with b the number of possible “grades,” and shows how to achieve global consistency by running the *b-set-neighboring* protocol multiple times. In [1], the fact is used that an unreliable broadcast is available among honest parties to achieve the same—global consistency on a PKI, where the number of identities each party gets is proportional to his hashing power, as suggested in [2].

The protocol in [1], however, suffers from a total running time that depends on the number of parties, because of two factors: (1) the way in which it uses POWs, and (2) the use of the Dolev-Strong authenticated broadcast protocol (run multiple times in parallel based on the graded PKI), which takes a linear number of rounds. Regarding (1), and in more detail, in order to assign exactly one key per party, a low variance POW scheme is used. This implies that the time needed by an honest party to mine a POW is going to be proportional to the ratio of the adversarial hashing power to the hashing power of the weakest honest party. Otherwise, the “rushing” adversary would be able to compute more identities in the additional time she has due to the latency of the communication infrastructure.⁶ Regarding (2), we note that potentially an expected-constant-round protocol could be used instead of Dolev-Strong, although the parallel composition of n instances would require more involved techniques [11].

Furthermore, having a PKI allows parties to generate an unpredictable beacon (in the random oracle model), which is then suggested in [1] as a genesis block-generation method for a new cryptocurrency. Yet, no formal treatment of the security of the resulting blockchain protocol is presented, and—as already mentioned—the round complexity of the suggested genesis block generation procedure is linear in the number of participants, both in contrast to our work.

As in [1], Katz *et al.* [26] also consider achieving pseudonymous broadcast and secure computation from POWs (“cryptographic puzzles”) and the existence of digital signatures without prior PKI setup, but under the assumption of an existing unpredictable beacon.

Finally, Pass *et al.* [35] consider a partially synchronous model of communication where parties are not guaranteed to receive messages at the end of each round but rather after a specified delay Δ (cf. [15]), and show that the backbone protocol can be proven secure in this setting. In principle, our results about the bootstrapped backbone protocol can be extended to their setting as shown in [22].

Organization of the paper. The rest of the paper is organized as follows. In Section 2 we describe the network and adversarial model, introduce some basic blockchain notation, and enumerate the various security properties. In Section 3 we present the bootstrapped Bitcoin backbone protocol and its analysis. Applications are presented in Section 4: a robust public transaction ledger, and PKI generation and consensus without trusted setup and with round complexity independent of the number of parties. Due to space limitations, some of the proofs are presented in the supplementary material.

⁶On the flip side, the benefit of the approach in [1] is that when all honest parties have the *same* hashing power, a PKI that maps each party to exactly one identity and preserves an honest majority on the keys can be achieved. However, in today’s environments where even small devices (e.g., mobile phones, smart watches) have powerful CPUs with possibly different clock frequencies, this assumption is arguably weak.

2 Model and Definitions

We describe our protocols in a model that extends the synchronous communication network model presented in [21] for the analysis of the Bitcoin backbone protocol (which in turn is based on Canetti’s formulation of “real world” execution for multi-party cryptographic protocols [8, 9]). As in [21], the protocol execution proceeds in rounds with inputs provided by an environment program denoted by \mathcal{Z} to parties that execute the protocol.

Next we provide a high level overview of the model, focusing on the differences that are intrinsic to our setting where the adversary has a precomputation advantage. The adversarial model in the network is actively malicious following the standard cryptographic approach. The adversary is *rushing*, meaning that in any given round it gets to see all honest players’s messages before deciding its strategy. Message delivery is provided by a “diffusion” mechanism that is guaranteed to deliver all messages, without however preserving their order and allowing the adversary to arbitrarily inject its own messages. Importantly, the honest parties are not guaranteed to have the same view of the messages delivered in each round, except for the fact that all honest messages from the previous round are delivered. Furthermore, the adversary is allowed to change the source information on every message (i.e., communication is not authenticated). In the protocol description, we will use `DIFFUSE` as the message transmission command to capture the “send-to-all” functionality that is available in our setting.⁷ Note that, as in [21], an adversarial sender may abuse `DIFFUSE` and attempt to confuse honest parties by sending and delivering inconsistent messages to them.

In contrast to [21], where all parties (the honest ones and the ones controlled by the adversary), are activated for the first time in the execution of the protocol in the same round⁸, in our model the environment will choose the round at which all the honest parties will become active; the corrupted parties, on the other hand, are activated in the first round. Once honest parties become active they will remain active until the end of the execution. In each round, after the honest parties become active, the environment activates each one by providing input to the party and receives the party’s output when it terminates. When activated, parties are able to read their input tape `INPUT()` and communication tape `RECEIVE()`, perform some computation that will be suitably restricted (see below) and issue a `DIFFUSE` message that is guaranteed to be delivered to all parties at the beginning of the next round.

In more detail, we model the execution in the following manner. We employ the parameterized system of ITM’s from [9] (2013 version) that is comprised of an initial ITM \mathcal{Z} , called the environment, and C , a control function that is specified below. We remark that our control function C is suitably restricted compared to that of [9, 10] to take into account restrictions in the order of execution that are relevant to our setting.

The execution is defined with respect to a protocol Π , a set of parties P_1, \dots, P_n and an adversary \mathcal{A} . The adversary is allowed to corrupt parties adaptively up to a number of $t < n$ parties. The protocol Π has access to two resources or “ideal functionalities,” the random oracle, and the diffusion channel. Initially, the environment may pass input to either the adversary \mathcal{A} or spawn an instance running the protocol Π which will be restricted to be assigned to the lexicographically smallest honest party (such restrictions are imposed by the control function [9]). After a party P_i is activated, the environment is restricted to activate the lexicographically next honest party, except in the case when no such party is left, in which case the next program to be activated is the adversary \mathcal{A} ; subsequently, the round-robin execution order between the honest parties will be repeated.

Whenever a party is activated the control function allows for q queries to be made to the random oracle while in the case of an activation of \mathcal{A} a number of $t \cdot q$ queries are allowed where t is the

⁷In [21] the command name `BROADCAST` is used for this functionality, which we sometimes also will use informally.

⁸After their first-time activation, the environment keeps activating parties in every round (cf. [8]).

number of corrupted parties. Honest parties are also allowed to annotate their queries to the random oracle for verification purposes, in which case an unlimited amount of queries is permitted (that still counts towards the overall running time of the system execution). Note that the adversary is not permitted to take advantage of this feature of the execution. With foresight, this asymmetry will be necessary, since otherwise it would be trivial for the adversary to break the properties of our protocols by simply “jamming” the incoming communication tape of the honest parties with messages whose verification would deplete their access quota to the random oracle per activation. Furthermore, for each party a single invocation to the diffusion channel is permitted. The diffusion channel maintains the list of messages diffused by each party, and permits the adversary \mathcal{A} to perform a “fetch” operation so that it obtains the messages that were sent. When the adversary \mathcal{A} is activated, the adversary will interact with the diffusion channel, preparing the messages to be delivered to the parties and performing a fetch operation. This write and fetch mode of operation with the communication channel enables the channel to enforce synchrony among the parties running the protocol (cf. [25]).

The term $\{\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^P(\kappa, z)\}_{\kappa \in \mathbb{N}, z \in \{0,1\}^*}$ denotes the random variable ensemble describing the view of party P after the completion of an execution with environment \mathcal{Z} , running protocol Π , and adversary \mathcal{A} , on auxiliary input $z \in \{0,1\}^*$. We often drop the parameters κ and z and simply refer to the ensemble by $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^P$ if the meaning is clear from the context. Following the resource-bounded computation model of [9], it holds that the total length of the execution is bounded by a polynomial in the security parameter κ and the length of the auxiliary string $|z|$, provided that the environment is *locally bounded* by a polynomial (cf. Proposition 3 in [9]). Note that the above execution model captures adversarial precomputation since it permits the environment to activate the adversary an arbitrary number of times (bounded by a polynomial in the security parameter κ of course) before the round-robin execution of the honest parties commences.

We note that the above modeling obviates the need for a strict upper bound on the number of messages that may be transmitted by the adversary in each activation (as imposed by [1]). In our setting, honest parties, at the discretion of the environment, will be given sufficient time to process all the messages delivered via the diffusion channel including all messages that are injected by the adversary.

The concatenation of the view of all parties ever activated in the execution, say, P_1, \dots, P_n , is denoted by $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}$. As in [21], we are interested in protocols Π that do not make explicit use of the number of parties n or their identities. Further, note that because of the unauthenticated nature of the communication model the parties may never be certain about the number of participants in a protocol execution.

In our correctness and security statements we will be concerned with *properties* of protocols Π running in the above setting (as opposed to simulation-based notions of security). Such properties will be defined as predicates over the random variable $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}(\kappa, q, z)$ by quantifying over all locally polynomial-bounded adversaries \mathcal{A} and environments \mathcal{Z} (in the sense of [9]). Note that all our protocols will only satisfy properties with a small probability of error in κ as well as in a parameter k that is selected from $\{1, \dots, \kappa\}$. (Note that, in practice, one may choose k to be much smaller than κ , e.g., $k = 6$.)

2.1 Blockchain notation

Next, we introduce some basic blockchain notation, following [21]. A *block* is any triple of the form $B = \langle s, x, ctr \rangle$ where $s \in \{0,1\}^\kappa$, $x \in \{0,1\}^*$, $ctr \in \mathbb{N}$ are such that satisfy predicate $\text{validblock}_q^D(B)$ defined as

$$(H(ctr, G(s, x)) < D) \wedge (ctr \leq q),$$

where H, G are cryptographic hash functions (e.g., SHA-256) modelled as random oracles. The parameter $D \in \mathbb{N}$ is also called the block’s *difficulty level*. The parameter $q \in \mathbb{N}$ is a bound that in the Bitcoin implementation determines the size of the register ctr ; in our treatment we allow this to be arbitrary, and use it to denote the maximum allowed number of hash queries in a round. We do this for convenience and our analysis applies in a straightforward manner to the case that ctr is restricted to the range $0 \leq ctr < 2^{32}$ and q is independent of ctr .

A *blockchain*, or simply a *chain* is a sequence of *blocks*. The rightmost block is the *head* of the chain, denoted $\text{head}(\mathcal{C})$. Note that the empty string ε is also a chain; by convention we set $\text{head}(\varepsilon) = \varepsilon$. A chain \mathcal{C} with $\text{head}(\mathcal{C}) = \langle s', x', ctr' \rangle$ can be extended to a longer chain by appending a valid block $B = \langle s, x, ctr \rangle$ that satisfies $s = H(ctr', G(s', x'))$. In case $\mathcal{C} = \varepsilon$, by convention any valid block of the form $\langle s, x, ctr \rangle$ may extend it. In either case we have an extended chain $\mathcal{C}_{\text{new}} = \mathcal{C}B$ that satisfies $\text{head}(\mathcal{C}_{\text{new}}) = B$.

Consider a chain \mathcal{C} of length m and any nonnegative integer k . We denote by $\mathcal{C}^{\lceil k}$ the chain resulting from the “pruning” of the k rightmost blocks. Note that for $k \geq \text{len}(\mathcal{C})$, $\mathcal{C}^{\lceil k} = \varepsilon$. If \mathcal{C}_1 is a prefix of \mathcal{C}_2 we write $\mathcal{C}_1 \preceq \mathcal{C}_2$.

2.2 Basic security properties of the blockchain

We are going to show that the blockchain data structure built by our protocol satisfies a number of basic properties, as formulated in [21, 27]. At a high level, the first property, called *common prefix*, has to do with the existence, as well as persistence in time, of a common prefix of blocks among the chains of honest players [21]. Here we will consider a stronger variant of the property, presented in [27, 34], which allows for the black-box proof of application-level properties (such as the *persistence* of transactions entered in a public transaction ledger built on top of the Bitcoin backbone—cf. Section 4).

Definition 1 ((Strong) Common Prefix Property). The *strong common prefix property* Q_{cp} with parameter $k \in \mathbb{N}$ states that the chains $\mathcal{C}_1, \mathcal{C}_2$ reported by two, not necessarily distinct honest parties P_1, P_2 , at rounds r_1, r_2 , with $r_1 \leq r_2$, satisfy $\mathcal{C}_1^{\lceil k} \preceq \mathcal{C}_2$.

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

Definition 2 (Chain Quality Property). The *chain quality property* Q_{cq} with parameters $\mu \in \mathbb{R}$ and $k, k_0 \in \mathbb{N}$ states that for any honest party P with chain \mathcal{C} in $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}(\kappa, z)$, it holds that for any k consecutive blocks of \mathcal{C} , excluding the first k_0 blocks, the ratio of adversarial blocks is at most μ .

Further, in the derivations in [21] an important lemma was established relating to the rate at which the chains of honest players were increasing as the Bitcoin backbone protocol was run. This was explicitly considered in [27] as a property under the name *chain growth*. Similarly to the variant of the common prefix property above, this property along with chain quality were shown sufficient for the black-box proof of application-level properties (in this case, transaction ledger *liveness*; see Section 4).

Definition 3 (Chain Growth Property). The chain growth property Q_{cg} with parameters $\tau \in \mathcal{R}$ (the “chain speed” coefficient) and $s, r_0 \in \mathbb{N}$ states that for any round $r > r_0$, where honest party P has chain \mathcal{C}_1 at round r and chain \mathcal{C}_2 at round $r + s$ in $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}(\kappa, z)$, it holds that $|\mathcal{C}_2| - |\mathcal{C}_1| \geq \tau \cdot s$.

3 The Bootstrapped Backbone Protocol

We begin this section by presenting the “bootstrapped” Bitcoin backbone protocol, followed by its security analysis. In a nutshell, the protocol is a generalization of the protocol in [21], which is enhanced in two ways: (1) an initial challenge-exchange phase, in which parties contribute random values, towards the establishment of an unpredictable genesis block, despite the precomputation efforts of corrupt players, and (2) a ranking process and chain-validation predicate that, in addition to its basic function (checking the validity of a chain’s content), enables the identification of “fresh” candidate genesis blocks. The ranking process yields a graded list of genesis blocks and is inspired by the “key ranking” protocol in [1], where it is used to produce a “graded” PKI, as mentioned in Section 1.

Before describing the bootstrapped backbone protocol in detail, we highlight its unique features.

- *No trusted setup and individual genesis block mining.* Parties start without any prior coordination and enter an initial challenge-exchange phase, where they will exchange random values that will be used to construct “freshness” proofs for candidate genesis blocks. The parties will run the initial challenge-exchange phase for a small number of rounds, and subsequently will try to mine their own genesis blocks individually. Once they mine or accept a genesis block from the network they will engage in mining further blocks and exchanging blockchains as in Bitcoin’s blockchain protocol. On occasion they might switch to a chain with a different genesis block. Nevertheless, as we will show, quite soon they will stabilize in a common prefix and a single genesis block.
- *Freshness of genesis block impacts chains’ total weight.* Chains rooted at a genesis block will incorporate its weight in their total valuation. Genesis blocks can be quite “heavy” compared to regular blocks and their total valuation will depend on how fresh they are. Their weight in general might be as much as a linear number of regular blocks in the security parameter. Furthermore, each regular block in a chain accounts for 3 units in terms of the total weight of the chain, something that, as we show, will be crucial to account for differences in terms of weight that are assigned to the same genesis block by different parties running the protocol (cf. Remark 1).
- *Personalized chain selection rule.* Given the co-existence of multiple genesis blocks, a ranking process is incorporated into the chain selection rule that, in addition to its basic function (checking the validity of a chain’s content) and picking the longest chain, it now also takes into account the freshness degree of a genesis block from the perspective of each player running the protocol. The ranking process effectively yields a graded list of genesis blocks and is inspired by the “key ranking” protocol in [1], where it is used to produce a “graded” PKI (see further discussion below). The weight value for each genesis block will be thus proportional to its *perceived* “freshness” by each party running the protocol (the fresher the block the higher its weight). It follows that honest players use different chain selection procedures since each predicate is “keyed” with the random coins that were contributed by each player in the challenge-exchange phase (and thus guaranteed to be fresh from the player’s perspective). This has the side effect that the same genesis block might be weighed differently by different parties. Despite these differences, we show that eventually all parties accept the same chains as valid and hence will unify their chain selection rule in the course of the protocol.
- *Robustness is achieved after an initial period of protocol stabilization.* All our modifications integrate seamlessly with the Bitcoin backbone protocol [21], and we are able to show that our blockchain protocol is a robust transaction ledger, in the sense of satisfying the properties of persistence and liveness. Nevertheless, contrary to [21], the properties are satisfied only after an

initial period of rounds where persistence is uncertain and liveness might be slower; this is the period where the parties still stabilize the genesis block and they might be more susceptible to attacks. Despite this, a ledger built on top of our blockchain will be available immediately after the challenges exchange phase. Furthermore, once the stabilization period is over the robust transaction ledger behavior is guaranteed with overwhelming probability (in the length of the challenges-exchange phase).

3.1 Protocol Description

The bootstrapped Bitcoin backbone protocol is executed by an arbitrary number of parties over an unauthenticated network (cf. Section 2). For concreteness, we assume that the number of parties running the protocol is n ; however, parties need not be aware of this number when they execute the protocol. Communication over the network is achieved by utilizing a send-to-all `DIFFUSE` functionality that is available to all parties (and may be abused by the adversary in the sense of delivering different messages to different parties). After an initial (“challenge”) phase, each party is to maintain a data structure called a “blockchain,” as defined above. Each party’s chain may be different, but, as we will prove, under certain well-defined conditions, the chains of honest parties will share a large common prefix.

The protocol description intentionally avoids specifying the type of values that parties try to insert in the chain, the type of chain validation they perform (beyond checking for its structural properties with respect to the hash functions $G(\cdot), H(\cdot)$), and the way they interpret the chain. In the protocol description, these actions are abstracted by the external functions $V(\cdot), I(\cdot), R(\cdot)$ which are specified by the application that runs “on top” of the backbone protocol.

The protocol is specified as Algorithm 1. At a high level, the protocol first executes a challenge-exchange phase for $l+1$ rounds (l will be determined later), followed by the basic backbone functions, i.e., mining and broadcasting blocks; a crucial difference here with respect to the original backbone protocol is that the chain validation process must also verify candidate genesis blocks, which in turn requires updating the validation function as the protocol proceeds. (This, however, only happens in the next l rounds after the challenge phase.) The protocol’s supporting algorithms are specified next.

The challenge-exchange phase. In order to generate an unpredictable genesis block, players first execute a “challenge-exchange” phase, where they broadcast, for a given number of rounds ($l + 1$), randomly generated challenges that depend on the challenges received in the previous rounds. The property that is assured is that an honest player’s k -round challenge, $1 \leq k \leq l$, depends on the $(k - 1)$ -round challenges of all honest players. This dependence is made explicit through the random oracle. The code of the challenge-exchange phase is shown in Algorithm 2.

Validation predicate update. In the original backbone protocol [21], the chain validation function (called `validate`—see below) performs a validation of the structural properties of a given chain \mathcal{C} , and remains unchanged throughout the protocol. In our case, however, where there is no initial fresh common random string, the function plays the additional role of checking for valid genesis blocks, and players have to update their validation predicate as the protocol advances (for the first l rounds after the challenge phase).

Indeed, using the challenges distributed in the challenge-exchange phase of the protocol, players are able to identify fresh candidate genesis blocks that have been shared during that phase and are accompanied by a valid proof. In addition, the valid genesis blocks are ranked with a negative dependence on the round they were received. In order to help other players to also identify the same genesis blocks, players broadcast the valid genesis blocks they have accepted together with

Algorithm 1 The *bootstrapped backbone* protocol, parameterized by the *input contribution function* $I(\cdot)$, the *chain reading function* $R(\cdot)$, and parameter l .

```

1:  $\mathcal{C} \leftarrow \epsilon$ 
2:  $st \leftarrow \epsilon$ 
3:  $round \leftarrow 1$  ▷ Global variable  $round$ 
4:  $Gen \leftarrow \emptyset$  ▷ Set of candidate genesis blocks
5:  $Rank \leftarrow \langle \epsilon \rangle$ 
6:  $(\mathbf{c}, \mathbf{A}, c) \leftarrow \text{exchangeChallenges}(1^\kappa)$ 
7: while TRUE do
8:    $k \leftarrow round - l - 2$ 
9:    $M_{Gen} \leftarrow \{(\langle s', x', ctr' \rangle, \langle A'_{l+1}, \dots, A'_{l+1-k} \rangle)\}$  from RECEIVE()
10:   $M_{Chain} \leftarrow \text{chains } \mathcal{C}' \text{ found in RECEIVE}()$ 
11:   $(Gen, Rank) \leftarrow \text{updateValidate}(\mathbf{c}, \mathbf{A}, M_{Gen}, Gen, Rank)$ 
12:   $\tilde{\mathcal{C}} \leftarrow \text{maxvalid}(\mathcal{C}, M_{Chain}, Gen, Rank)$ 
13:   $\langle st, x \rangle \leftarrow I(st, \tilde{\mathcal{C}}, round, \text{INPUT}(), \text{RECEIVE}())$ 
14:   $\mathcal{C}_{\text{new}} \leftarrow \text{pow}(x, \tilde{\mathcal{C}}, c)$ 
15:  if  $\mathcal{C} \neq \mathcal{C}_{\text{new}}$  then
16:    if  $\mathcal{C} = \epsilon$  then ▷ New genesis block has been produced
17:      DIFFUSE(  $(\mathcal{C}_{\text{new}}, \langle A_{l+1}, \dots, A_{l+1-(k+1)} \rangle)$  )
18:       $\mathcal{C} \leftarrow \mathcal{C}_{\text{new}}$ 
19:      DIFFUSE( $\mathcal{C}$ )
20:   $round \leftarrow round + 1$ 
21:  if INPUT() contains READ then
22:    write  $R(\mathbf{x}_{\mathcal{C}})$  to OUTPUT()

```

the additional information needed by the other players for verification. The validation predicate update function is shown in Algorithm 3. Recall that Gen is the set of candidate genesis blocks.

Chain validation. A chain is considered valid if in addition to the checks performed by the basic backbone protocol regarding the chain’s structural properties, its genesis block is in the Gen list, which is updated by the `updateValidate` function (Algorithm 3). The chain validation function is shown in Algorithm 4.

Chain selection. The objective of the next algorithm in Algorithm 1, called `maxvalid`, is to find the “best possible” chain when given a set of chains. The accepted genesis blocks have different *weights* depending on when a player received them. In addition, it is possible that the same genesis block is received by honest players in two different rounds (as we show later, those rounds have to be consecutive). In order to take into account the “slack” introduced by the different views honest players may have regarding the same block, as well as the different weights different blocks may have, we let the *weight of a chain* \mathcal{C} be equal to the *weight of its genesis block plus three times its length minus one*. The chain selection function is shown in Algorithm 5.

Algorithm 2 The *challenge-exchange* function. Note that variable *round* is global, and originally set to 1.

```

1: function exchangeChallenges( $1^\kappa$ )
2:    $c_1 \xleftarrow{R} \{0, 1\}^\kappa$ 
3:   DIFFUSE( $c_1$ )
4:    $round \leftarrow round + 1$ 
5:   while  $round \leq l + 1$  do
6:      $A_{round} \leftarrow \kappa$ -bit messages found in RECEIVE()
7:      $r_{round} \xleftarrow{R} \{0, 1\}^\kappa$ 
8:      $A_{round} \leftarrow A_{round} || r_{round}$ 
9:      $c_{round} \leftarrow H(A_{round})$  ▷ Compute challenge
10:    DIFFUSE( $c_{round}$ )
11:     $round \leftarrow round + 1$ 
12:  return ( $\langle c_1, \dots, c_l \rangle, \langle A_2, \dots, A_{l+1} \rangle, c_{l+1}$ )

```

Algorithm 3 The *validation predicate update* function.

```

1: function updateValidate( $\mathbf{c}, \mathbf{A}, M_{Gen}, Gen, Rank$ )
2:    $k \leftarrow round - l - 2$ 
3:   if  $k \geq l$  then
4:     return  $Gen, Rank$  ▷ No updates after round  $2l + 2$ 
5:   for each ( $\langle s', x', ctr' \rangle, \langle A'_{l+1}, \dots, A'_{l+1-k} \rangle$ ) in  $M_{Gen}$  do
6:     if  $\text{validblock}_q^D(\langle s, x, ctr \rangle) \wedge \langle s, x, ctr \rangle \notin Gen$  then
7:        $flag \leftarrow (H(A'_{l+1}) = s) \wedge (c_{l-k} \in A'_{l+1-k})$ 
8:       for  $i = l + 1 - k$  to  $l$  do
9:         if  $H(A'_i) \notin A'_{i+1}$  then
10:           $flag \leftarrow \text{FALSE}$ 
11:       if  $flag = \text{TRUE}$  then
12:          $Gen \leftarrow Gen \cup \langle s, x, ctr \rangle$ 
13:          $Rank[\langle s, x, ctr \rangle] \leftarrow l - k$ 
14:         DIFFUSE( $\langle s, x, ctr \rangle, \langle A'_{l+1}, \dots, A'_{l+1-k}, A_{l-k} \rangle$ ) ▷ Augment  $A'$  sequence with own
15:          $A$  value.
16:       return  $Gen, Rank$ 

```

Algorithm 5 The function that finds the “best” chain. The input is a set of chains and the list of genesis blocks.

```

1: function maxvalid( $\mathcal{C}_1, \dots, \mathcal{C}_k, Gen$ )
2:    $temp \leftarrow \varepsilon$ 
3:    $maxweight \leftarrow 0$ 
4:   for  $i = 1$  to  $k$  do
5:     if validate( $\mathcal{C}_i, Gen$ ) then
6:        $w_{Gen_i} \leftarrow Rank(\text{tail}(\mathcal{C}_i))$ 
7:        $weight \leftarrow w_{Gen_i} + 3(|\mathcal{C}_i| - 1)$ 
8:       if  $maxweight < weight$  then 11
9:          $maxweight \leftarrow weight$ 
10:         $temp \leftarrow \mathcal{C}_i$ 
11:  return  $temp$ 

```

Algorithm 4 The *chain validation predicate*, parameterized by q, D , the hash functions $G(\cdot), H(\cdot)$, and the *content validation predicate* $V(\cdot)$. The input is \mathcal{C} .

```

1: function validate( $\mathcal{C}, Gen$ )
2:    $b \leftarrow V(\mathbf{x}_{\mathcal{C}}) \wedge (\mathcal{C} \neq \varepsilon) \wedge (\text{tail}(\mathcal{C}) \in Gen)$ 
3:   if  $b = \text{True}$  then
4:      $\langle s, x, ctr \rangle \leftarrow \text{head}(\mathcal{C})$ 
5:      $s' \leftarrow H(ctr, G(s, x))$ 
6:     repeat
7:        $\langle s, x, ctr \rangle \leftarrow \text{head}(\mathcal{C})$ 
8:       if  $\text{validblock}_q^D(\langle s, x, ctr \rangle) \wedge (H(ctr, G(s, x)) = s')$  then
9:          $s' \leftarrow s$  ▷ Retain hash value
10:         $\mathcal{C} \leftarrow \mathcal{C}^{\uparrow 1}$  ▷ Remove the head from  $\mathcal{C}$ 
11:       else
12:          $b \leftarrow \text{False}$ 
13:     until  $(\mathcal{C} = \varepsilon) \vee (b = \text{False})$ 
14:   return  $b$ 

```

The proof-of-work function. Finally, we need to modify the proof-of-work function in [21], so that when a genesis block is mined, the challenge computed in the last round of the challenge-exchange phase will be included in the block. This, in addition to the proof of genesis information sent in the backbone protocol, is required so that other honest players accept this block as valid and rank it accordingly. The code is presented in Algorithm 6.

Algorithm 6 The *proof of work* function, parameterized by q, D and hash functions $H(\cdot), G(\cdot)$. The input is (x, \mathcal{C}, c) .

```

1: function pow( $x, \mathcal{C}, c$ )
2:   if  $\mathcal{C} = \varepsilon$  then
3:      $s \leftarrow c$  ▷  $c$  is required to prove freshness
4:   else
5:      $\langle s', x', ctr' \rangle \leftarrow \text{head}(\mathcal{C})$ 
6:      $s \leftarrow H(ctr', G(s', x'))$ 
7:      $ctr \leftarrow 1$ 
8:      $B \leftarrow \varepsilon$ 
9:      $h \leftarrow G(s, x)$ 
10:    while  $(ctr \leq q)$  do
11:      if  $(H(ctr, h) < D)$  then ▷ Proof of work found
12:         $B \leftarrow \langle s, x, ctr \rangle$ 
13:      break
14:       $ctr \leftarrow ctr + 1$ 
15:     $\mathcal{C} \leftarrow \mathcal{C}B$  ▷ Extend chain
16:  return  $\mathcal{C}$ 

```

Figure 2 presents the overall structure (phases and corresponding rounds) of the bootstrapped

backbone protocol. Next, we turn to its analysis.

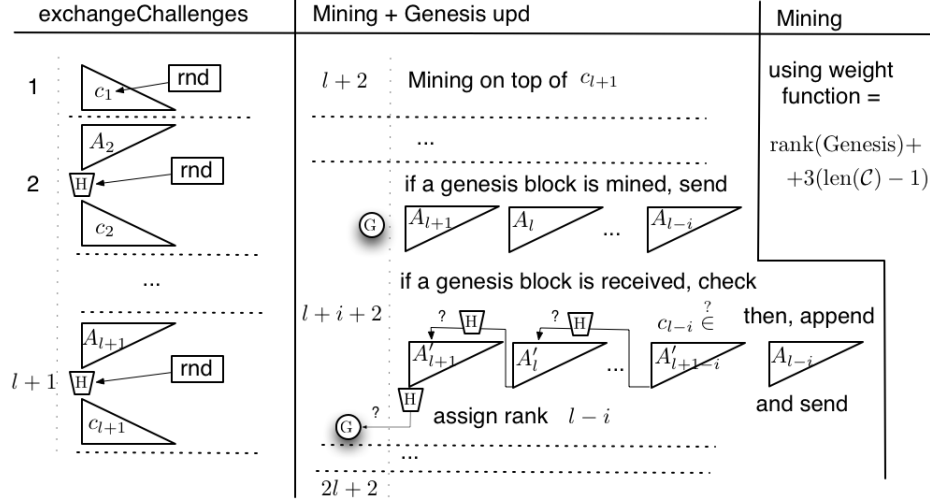


Figure 2: The different phases and corresponding rounds of the bootstrapped backbone protocol.

Remark 1. To understand some of our design choices we briefly give some examples of simpler protocols that don't work. For the first example, assume that we only have one round of challenge exchange i.e. l equal to 1. With some non-negligible probability, the adversary can send one block to half of the honest players and another block to the other half. By splitting the honest players in two groups such that no one in the first group will choose the chain of the second and vice versa, agreement becomes impossible. Moreover, l must be large enough so that at least one honest party computes a genesis block with overwhelming probability. Otherwise the adversary can choose to remain silent and no genesis block will be mined with non-negligible probability.

For the second example assume that blocks weigh less than 3 units, as in the original protocol. Also, assume that somehow the problem of the first example was avoided and honest parties only adopted chains with genesis blocks that everyone had in their genesis block list. In this case, uniquely successful rounds would not imply agreement on a single chain (see Figure 3), as the adversary would have been able to take advantage of the different views that honest players have regarding the weight of genesis blocks. However, if we set the block weight to 3, this event becomes impossible and makes the analysis a lot easier.

3.2 Analysis of the Bootstrapped Backbone Protocol

First, some additional definitions that will become handy in the analysis. We saw in the previous section that genesis blocks are assigned weights, and, further, that a single genesis block may have different weights for different parties depending on when they received it. We extend this notion to chains of blocks.

Definition 4. Let $w_P(B)$ be the weight that P assigned to genesis block B . We define the weight of a chain \mathcal{C} with genesis block B (with respect to party P) to be:

$$w_P(\mathcal{C}) = w_P(B) + 3(|\mathcal{C}| - 1).$$

If block B was not received by P until round $2l + 1$, or if $\mathcal{C} = \epsilon$, then $w_P(\mathcal{C}) = -1$.




Figure 3: An example where blocks weigh 2 units. In the table the weights of the respective chains are depicted. Initially player P_1 has adopted chain C_1 and player P_2 chain C_2 . Then a uniquely successful round happens and C_2 is extended to C_2' . Notice that, P_1 will not adopt C_2' since it has the same weight as C_1 . If the new block weighted 3 units, all players would have adopted chain C_2' .

In [21], all parties assign the same weight to the same chain, i.e., the length of the chain; thus, for all parties P_i, P_j we have that $w_{P_i}(\mathcal{C}) = w_{P_j}(\mathcal{C})$. In contrast, in our case the genesis block of each chain may have different weight for different parties, akin to some bounded amount of “noise” that is party-dependent being added to the chain weights. We are going to show that if the amount of noise is at most 1, then by letting each new block weigh 3 units our protocol satisfies the chain growth, common prefix and chain quality properties.

Definition 5. Regarding chains and their weight:

- Define $h_{\mathcal{C}} = \max_P \{w_P(\mathcal{C})\}$ and $\ell_{\mathcal{C}} = \min_P \{w_P(\mathcal{C})\}$.
- Let $\mathcal{C}(B)$ denote the truncation of chain \mathcal{C} after its block B .
- For a block B of a chain \mathcal{C} , define $h_{\mathcal{C}}(B) = h_{\mathcal{C}(B)}$ and similarly for $\ell_{\mathcal{C}}(B)$. (Sometimes we will abuse notation and write $\ell(B)$ instead of $\ell_{\mathcal{C}}(B)$. As long as no collision happens $\ell(B)$ is well defined. The same holds for $h(B)$.)
- For chains \mathcal{C}_1 and \mathcal{C}_2 , define $\mathcal{C}_1 \cap \mathcal{C}_2$ to be the chain formed by their common prefix.

The following are important concepts introduced in [21], which we are also going to use in our analysis:

Definition 6. A round is called:

- *successful* if at least one honest party computes a solution in this round;
- *uniquely successful* if exactly one honest party computes a solution in this round.

Whether a block was mined by an honest party or by the adversary will be critical for our analysis.

Definition 7. In an execution blocks are called:

- *honest*, if mined by an honest party,
- *adversarial*, if mined by the adversary, and
- *u.s. blocks*, if mined in a uniquely successful round by an honest player.

Recall that our model is “flat” in terms of computational power in the sense that all honest parties are assumed to have the same computational power while the adversary has computational power proportional to the number of players that it controls. The total number of parties is n and the adversary is assumed to control up to t of them (honest parties do not know any of these parameters). Obtaining a new block is achieved by finding a hash value that is smaller than the difficulty parameter D . Thus, the success probability that a single hash query produces a solution is $p = \frac{D}{2^\kappa}$, where κ is the length of the hash. The total hashing power of the honest players is $\alpha = pq(n - t)$, the hashing power of the adversary is $\beta = pqt$, and the total hashing power is $f = \alpha + \beta$. Moreover, in [21], a lower bound on the probability that a round is uniquely successful was established; denoted by γ and equal to $\alpha - \alpha^2$. Notice that γ is also a bound for the probability of a round being just successful.

For each round j , we define the Boolean random variables X_j and Y_j as follows. Let $X_j = 1$ iff j was a *successful round*, i.e., at least one honest party computed a POW at round j , and let $Y_j = 1$ iff j was a *uniquely successful round*, i.e., exactly one honest party computed a POW at round j . With respect to a set of rounds S , let $Z(S)$ denote the number of POWs obtained by the adversary during the rounds in S (i.e., in $qt|S|$ queries). Also, let $X(S) = \sum_{j \in S} X_j$ and define $Y(S)$ similarly. Note that $\gamma|S| \leq \mathbb{E}[Y(S)] \leq \mathbb{E}[X(S)] \leq \alpha|S|$ and $\mathbb{E}[Z(S)] = \beta|S|$.

Lemma 8. *If $|S| = k$ and $\gamma \geq (1 + \delta)\beta$ for some $\delta \in (0, 1)$, then*

$$\Pr[Y(S) > (1 + \frac{5\delta}{9})Z(S)] > 1 - e^{-\Omega(\delta^2 k)}.$$

Proof. By the Chernoff bound we have that:

$$\Pr[Y(S) \leq (1 - \frac{\delta}{8})\mathbb{E}[Y(S)]] \leq e^{-\frac{\delta^2 \gamma k}{128}} \text{ and } \Pr[Z(S) \geq (1 + \frac{\delta}{9})\mathbb{E}[Z(S)]] \leq e^{-\frac{\delta^2 \beta k}{243}}.$$

Suppose none of the above events happens. Then, from the union bound, we get that with probability $1 - e^{-(2 \min(\frac{\beta}{243}, \frac{\gamma}{128})\delta^2 k - \ln(2))}$ it holds that

$$Y(S) > (1 - \frac{\delta}{8})\gamma k \geq (1 - \frac{\delta}{8})(1 + \delta)\beta k \geq (1 + \frac{5\delta}{9})(1 + \frac{\delta}{9})\beta k > (1 + \frac{5\delta}{9})Z(S).$$

□

Remark 2. For ease of exposition, in our analysis we will assume that there are no collisions; that is, for any two different queries to the random oracle, always a different response is returned. This would generally be a problem since for example it would break independence of X_i, X_j , for $i \neq j$, and we would not be able to apply the Chernoff bound in the previous lemma. However, since the probability of a collision happening, as well as all other events we consider, is at most $e^{-\Omega(\kappa)}$, we can always use the union bound to include the event of no collision occurring to our other assumptions. In addition, we assume that no two queries to the oracle are the same, as formalized by the Input Entropy condition in [21].

Properties of the genesis block generation process. We now establish a number of properties of the genesis block generation process.

Lemma 9 (Graded Consistency). *If any honest party P_i accepts genesis block B with rank $w_{P_i}(B) > 1$, then all honest parties accept B with rank at least $w_{P_i}(B) - 1$.*

Proof. Let $w_{P_i}(B) = k > 1$. Since P_i accepted B with rank k at some round r , he must have received a message of the form $(B, E_{l+1}, \dots, E_{k+1})$, where

- B is a valid block that contains $H(E_{l+1})$;
- E_{k+1} contains c_k and for $k+2 < j \leq l+1$, E_j contains $H(E_{j-1})$; and
- c_k is the challenge computed by P_i at round k .

Since $k > 0$, according to Algorithm 3, P_i is going to broadcast $(B, E_{l+1}, \dots, E_{k+1}, A_k)$, where $H(A_k) = c_k$ is contained in E_{k+1} and A_k contains all the messages received by P_i at round k . All honest-party challenges of round $k-1$ were received in this round; therefore, all honest parties have accepted or will accept block B by the next round and the lemma follows. \square

Lemma 10 (Validity). *Genesis blocks computed by honest parties before round $2l+2$, will be accepted by all honest parties in the next round.*

Proof. Suppose honest party P_i mined genesis block B at round m . According to Algorithm 1, B contains the challenge he has computed at the last round of the challenge-exchange phase. In addition, when the party broadcasts it, it includes the message sets A_{l+1}, \dots, A_r , where A_j contains the messages received by P_i at round j and $r = 2l + 2 - m$. Since P_i is honest, the following hold:

- B is a valid block that contains $H(A_{l+1})$;
- for $r+1 < j \leq l+1$, A_j contains $H(A_{j-1})$;
- if c_r is the challenge sent by some honest party at round r , then c_r is contained in A_{r+1} ; and
- all honest parties are going to receive the message.

Thus, all honest parties are going to accept B at round $m+1$ and the lemma follows. \square

Lemma 11 (Freshness). *Let $r \leq l+2$. Every block computed before round r cannot be part of some chain with genesis block B , where $w_P(B) \geq r-1$ for some honest party P , with overwhelming probability in the security parameter κ .*

Weak chain growth. We now turn our attention to the weight of chains and prove a *weak* chain-growth property. In the original Bitcoin backbone protocol [21], it was proved that chains grow at least at the rate of successful rounds, independently of the adversary’s behavior. Here, at least initially, the chains of honest parties grow in a “weak” manner, in the sense that the adversary is able to slow down this growth by using his own blocks. Later on, we will show that after some specific round our protocol also achieves optimal chain growth.

Lemma 12. *Let round r such that $l+2 \leq r < 2l+2$, and suppose that at round r an honest party, say, P_1 has a chain \mathcal{C} such that $w_{P_1}(\mathcal{C}) = d$. Then, by round s , where $r \leq s < 2l+2$, every honest party P will have received a chain \mathcal{C}' of weight at least $w_P(\mathcal{C}') = d - 2 + 3 \sum_{i=r}^{s-1} Y_i - \sum_{i=r}^{s-1} Z_i$.*

Proof. (Sketch) Note that every time a uniquely successful round happens, the minimum weight over all parties’ chains will increase by 2. Moreover, if the adversary has not diffused any block in the same round the minimum weight increases by 3. By applying this result iteratively, the lemma follows. We refer the reader to the supplementary material for the full proof. \square

Universal chain validity. A novelty of our construction is that the same genesis block may have different weight for different parties. Unfortunately, it could be the case that due to the adversary’s influence, a genesis block is valid for one party but invalid for another. This could lead to disagreement, in the sense that some honest parties may adopt a chain that others don’t because it is not valid for them. We will show that with overwhelming probability such an event cannot occur for our protocol; as such, chain validity is a “universal” property; if some honest party accepts a chain \mathcal{C} as valid, then \mathcal{C} will also be valid for all other parties.

Notice, that in order to prove the following lemma we need l to be greater than a value that depends on $1/\gamma$, i.e. the expected time it takes for honest parties to mine a block, and the security parameter κ (see also Remark 1). Intuitively l should be large enough so that (i) honest parties mine at least one block at this time interval, and (ii) any adversarial chain that is based on a genesis block broadcast at the end of the bootstrapping phase will never be adopted by honest parties (because such genesis block will have too small weight in comparison).

Lemma 13. *Suppose that for some $\delta \in (0, 1)$, $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, and $\gamma \geq (1 + \delta)\beta$, and that at round r an honest party P has chain \mathcal{C} . Then \mathcal{C} will also be valid for all other parties from this round on with probability $1 - e^{-\Omega(\delta^2 k)}$.*

The complete version of the weak chain growth lemma follows from the argument we've made above.

Corollary 14. *Suppose that for some $\delta \in (0, 1)$, $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, and $\gamma \geq (1 + \delta)\beta$. Let round r such that $r \geq l + 2$, and suppose that at round r an honest party, say, P_1 has a chain \mathcal{C} such that $w_{P_1}(\mathcal{C}) = d$. Then, by round s , where $r \leq s$, every honest party P will have received a chain \mathcal{C}' of weight at least $w_P(\mathcal{C}') = d - 2 + 3 \sum_{i=r}^{s-1} Y_i - \sum_{i=r}^{s-1} Z_i$ with probability at most $1 - e^{-\Omega(\delta^2 k)}$.*

Remark 3. Note further that the dependency of γ on n does not undermine the scalability of the round complexity of our protocol. This claim is argued on the basis that the difficulty level D can be set proportional to $1/n$, so that γ can be treated as a constant and then l is in essence independent of n (note that both parameters would be polynomials in κ).

A bound on adversarially precomputed blocks. The honest parties begin mining right after the challenge-exchange phase. Note that it does not help the adversary to precompute blocks before the challenge-exchange phase, except for the small probability of the event that some of his blocks happen to extend future blocks. We have shown that the adversary cannot create a private chain that honest parties will adopt if he starts mining at the first round of the challenge-exchange phase. It is though possible to start mining *after* the first round in order to gain some advantage over the honest parties. The following lemma provides a bound on the number of blocks mined during the challenge-exchange phase with sufficient weight so that they can be later used by the adversary.

Lemma 15 (Precomputed blocks). *Assume $3(1 + \delta)f < 1$ and $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, for some $\delta \in (0, 1)$. Let R be the set that contains any adversarial block B mined before round $l + 2$, where $h(B) > l - 1 - (1 - \delta)\delta^2 k$. Then $\Pr[|R| > \frac{5\delta}{9} k\beta] \leq e^{-\Omega(\delta^4 k)}$.*

Proof. (Sketch) We first show that the adversary cannot take advantage of blocks which belong to chains whose genesis block was computed early on in the challenge exchange phase. Hence, with overwhelming probability she can only use blocks computed near the end of the challenge exchange phase; remember that the weight of a genesis block is small if it is mined early in the challenge exchange phase. By applying appropriate Chernoff bounds the result follows. We refer the reader to the supplementary material for the full proof. \square

We are now ready to prove the security properties listed in Section 2.2.

Common Prefix. Every time a uniquely successful round happens all honest players converge to one chain, unless the adversary broadcasts some new block. This turns out to be a very important fact and a consequence of it is described in the next lemma.

Lemma 16. *Suppose block B in chain \mathcal{C} is a u.s. block and consider a chain \mathcal{C}' such that $B \notin \mathcal{C}'$. If $\ell_{\mathcal{C}'} \geq \ell_{\mathcal{C}}(B) - 1$ then there exists a unique adversarial block B' such that $\ell_{\mathcal{C}'}(B') \in [\ell_{\mathcal{C}}(B) - 1, \ell_{\mathcal{C}}(B) + 1]$. Moreover, if B is not a genesis block, then B' will also not be a genesis block.*

Proof. Assume block B was mined at some round r . If B is not a genesis block, then for any honest block B'' mined before round r it should hold that $\ell(B'') \leq \ell(B) - 2$. Otherwise, at round r no honest party would choose the parent of B to mine new blocks. If B is a genesis block, then no other honest party has mined a block in some previous round. On the other hand, for any honest block B'' mined after round r it must hold that $\ell(B'') \geq \ell(B) - 1 + 3 = \ell(B) + 2$, since honest parties will only extend chains of length at least $\ell(B) - 1$ after this round. Thus, if a block with weight in the given interval exists, it must be adversarial.

For the sake of contradiction, suppose B is not a genesis block while B' is a genesis block and let B'' be the parent of B . Then $h_C(B'') < \ell_{C'}(B')$ since $h_C(B'') \leq \ell_C(B) - 2$. This implies that every honest party received B' before block B'' . But then, no honest party would mine on the parent of B , because he would have lower weight than B' , which leads to a contradiction. Hence, the lemma follows. \square

Next, we use Lemma 16 in order to show that the existence of a fork implies that the adversary must have mined blocks proportional in number to the time the fork started. The full proof is in the supplementary material.

Theorem 17. *Assume $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, $\gamma \geq (1 + \delta)\beta$, for some real $\delta \in (0, 1)$. Let S be the set of the chains of the honest parties from round $2l + 2$ and onwards of the bootstrapped backbone protocol. Then the probability that S does not satisfy the strong common-prefix property with parameter k is at most $e^{-\Omega(\delta^4 k)}$.*

Proof. (sketch.) We are going to use Lemma 16 to match u.s. blocks to adversarial ones. Function ℓ will help us show that the matched blocks are distinct; every pair of matched blocks is very close with respect to the ℓ function, while pairs of u.s. blocks can be very far under specific conditions. Initially, we construct such a matching whenever we have a fork between two chains $\mathcal{C}_1, \mathcal{C}_2$, either by matching adversarial blocks alternatively to each of the chains or by matching consecutive u.s. blocks on the same chain to consecutive adversarial blocks in the other chain. For a more in depth look of this procedure the reader is deferred to the claim in Lemma 28 in the supplementary material.

Using this result, we prove that if a “deep” enough fork exists, the adversary must have mined more blocks than his hashing power allows, which leads to a contradiction. In more detail, the initial fork implies the existence of some honestly mined block B that is part of the common prefix of the two chains. Starting from B we construct a matching of all u.s. blocks mined after it, by picking “deeper” and “deeper” forks and repeatedly applying the matching procedure. Using the upper bound on precomputed blocks established in Lemma 15, we can show that the adversary is able to mine a sufficient number of blocks only with negligible probability. Hence, the theorem follows. \square

Chain Growth. We proved that after round $2l + 1$ the strong common-prefix property is satisfied. This implies that all players share a common genesis block after this round. The next lemma shows that this is sufficient in order to get chain growth at the same level as in the original Backbone protocol.

Lemma 18. *Suppose that at round r an honest party P_1 has a chain \mathcal{C} of weight $w_{P_1}(\mathcal{C}) = d$ and all honest parties after round $r - 1$ adopt chains that share the same genesis block B . Then, by round $s \geq r$, every honest party P will have received a chain \mathcal{C}' of weight at least $w_P(\mathcal{C}') = d - 1 + 3 \sum_{i=r}^{s-1} X_i$.*

Proof. Since all parties adopt chains with the same genesis block after round $r - 1$, and P_1 has adopted a chain \mathcal{C} of weight d , there are two cases: either (1) $\ell_C = d - 1$ and any chain that honest parties adopt after round $r - 1$ has a weight according to their view that is congruent to d or $d - 1$

modulo 3, or (2) $\ell_C = d$ and the weight is congruent to d or $d + 1$ modulo 3. This observation is implied from the fact that each extra block adds 3 units of weight to the chain and the B can only have two different weights under the views of honest parties.

It is sufficient to study only one of the two cases so w.l.o.g. suppose that the weight of the chains is congruent to d or $d - 1$ modulo 3. The proof is by induction on $s - r \geq 0$. For the basis ($s = r$), observe that if at round r P_1 has a chain \mathcal{C} of weight $w_{P_1}(\mathcal{C}) = d$, then he broadcast \mathcal{C} at an earlier round (than r). It follows that every honest party P will receive \mathcal{C} by round r and $w_P(\mathcal{C}) \geq d - 1$.

For the inductive step, note that by the inductive hypothesis every honest party P has received a chain \mathcal{C}' of weight at least $w_P(\mathcal{C}') = d' = d - 1 + 3 \sum_{i=r}^{s-2} X_i$ by round $s - 1$. When $X_{s-1} = 0$ the statement follows directly, so assume $X_{s-1} = 1$. Observe that every honest party queried the oracle with a chain of weight at least d' at round $s - 1$. It follows that every honest party P successful at round $s - 1$ broadcast a chain \mathcal{C}' of weight at least $w_P(\mathcal{C}') = d' + 3$. For every other party P' it holds that $w_{P'}(\mathcal{C}') \geq d' + 2 \geq d - 1 + 3 \sum_{i=r}^{s-1} X_i - 1$. However, no chain that an honest party adopts can have length $d' + 2$, because $d' + 2$ is congruent to $d - 2$ modulo 3. Thus all honest parties adopt chains that have length at least $d' + 3$ and the lemma follows. \square

It can be easily shown that Lemma 18 implies the chain growth property after round $2l + 1$.

Theorem 19. *Assume $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, $\gamma \geq (1 + \delta)\beta$, for some real $\delta \in (0, 1)$. The bootstrapped Bitcoin protocol satisfies the chain growth property for $r_0 = 2l + 2$ with speed coefficient $(1 - \delta)\gamma$ and probability at least $1 - e^{-\Omega(\delta^4 s)}$.*

Chain Quality. We first observe a consequence of Theorem 17.

Lemma 20. *Assume $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, $\gamma \geq (1 + \delta)\beta$, for some real $\delta \in (0, 1)$. From round $2l + 2$ and onwards of the bootstrapped backbone protocol, the probability that the adversary has a chain which is more than k blocks longer than the chain of some honest party is at most $e^{-\Omega(\delta^4 k)}$.*

Proof. Given any execution and an adversary that at a round r has a chain \mathcal{C} which is k blocks longer than the chain \mathcal{C}' of an honest party P , we can define an adversary such that at round $r + 1$ the common-prefix property does not hold for parameter k . The adversary simply sends \mathcal{C} to $P' \neq P$ at round r . \square

Theorem 21. *Assume $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, $\gamma \geq (1 + \delta)\beta$, for some real $\delta \in (0, 1/2)$. Suppose \mathcal{C} belongs to an honest party and consider any k consecutive blocks of \mathcal{C} computed after round $2l + 2$ of the bootstrapped backbone protocol. The probability that the adversary has contributed more than $(1 + \frac{\delta}{2})\frac{\beta}{\gamma} \cdot k \leq (1 - \frac{\delta}{3})k$ of these blocks is less than $e^{-\Omega(\delta^5 k)}$.*

Corollary 22. *Assume $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, $\gamma \geq (1 + \delta)\beta$, for some real $\delta \in (0, 1/2)$. The bootstrapped Bitcoin protocol satisfies the chain-quality property with parameters $\mu = (1 + \frac{\delta}{2})\frac{\beta}{\gamma}$, $k_0 = 2f(1 + \delta)(l + 1)$, and k , with probability at least $1 - e^{-\Omega(\delta^5 k)}$.*

Proof. Note that the next two events occur with probability at least $1 - e^{-\Omega(\delta^{2l})}$, for any $\delta \in (0, 1)$. The honest parties in the first $l + 1$ rounds have computed at most $\alpha(1 + \delta)(l + 1)$ blocks. The adversary, who might have been mining also during the challenges phase, has computed at most $2\beta(1 + \delta)(l + 1)$. The statement then follows from Theorem 21, since $\alpha(1 + \delta)(l + 1) + 2\beta(1 + \delta)(l + 1) < 2f(1 + \delta)(l + 1)$. \square

4 Applications of the Bootstrapped Backbone Protocol

In this section we present applications of our construction, starting with its primary/original one: a distributed ledger, i.e., a public and permanent summary of all transactions that honest parties can agree on as well as add their own, despite the potentially disruptive behavior of parties harnessing less than $1/2$ of the hashing power. This entails proving that the ledger’s required security properties (Persistence and Liveness — cf. [21]) hold in a genesis block-less setting.

Next, we consider the problem of setting up a PKI in our unauthenticated network setting *from scratch*, i.e., without any trusted setup. As mentioned in Section 1, the idea of using POWs as an identity-assignment tool was put forth by Aspnes *et al.* [2]. Here we build on this idea as well as on the “2-for-1 POWs” technique from [21] to use our bootstrapped protocol to assign identities to parties. The assignment relation will possibly assign more than one identity to the same party, while guaranteeing that the majority of them is assigned to honest parties.

Finally, applying the 2-for-1 POWs technique we can also solve the consensus (aka Byzantine agreement) problem [36, 31] without any trusted setup, even if the adversary has almost the same hashing power as the honest parties, and in a number of rounds independent of the number of parties. Indeed, all our protocols run have round complexity linear in the security parameter, and enjoy simultaneous termination.

Compared to other works, most notably [1], our approach is different in the order in which it sets up a “bulletin board” and assigns identities to parties. We choose to first establish the former—i.e., the ledger—and then assign the identities; in contrast, in [1] identities are established first in a graded manner, and then using that infrastructure the parties can implement a broadcast channel.

We now turn to the applications in detail.

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

In order to turn the backbone protocol into a protocol realizing a public transaction ledger suitable definitions were given for functions $V(\cdot), R(\cdot), I(\cdot)$ in [21]. Namely, $V(\langle x_1, \dots, x_m \rangle)$ is true if its input is a valid ledger. Function $R(\mathcal{C})$ returns the contents of the chain if they constitute a valid ledger, otherwise it is undefined. Finally, $I(st, \mathcal{C}, round, \text{INPUT}(), \text{RECEIVE}())$ returns the largest subsequence of transactions in the input and receive tapes that constitute a valid ledger, with respect to the contents of the chain the party already has, together with a randomly generated neutral transaction. We denote the instantiation of our protocol with these functions by $\Pi_{\text{PL}}^{\text{Boot}}$. For more details we refer to [21].

Definition 23. A protocol Π implements a *robust public transaction ledger* in the q -bounded synchronous setting without trusted setup if there is a round r_0 so that the following two properties are satisfied:

- *Persistence:* Parameterized by $k \in \mathbb{N}$ (the “depth” parameter), if in a certain round after r_0 an honest player reports a ledger that contains a transaction tx in a block more than k blocks away

from the end of the ledger, then tx will always be reported in the same position in the ledger by any honest player from this round on.

- *Liveness*: Parameterized by $u, k \in \mathbb{N}$ (the “wait time” and “depth” parameters, resp.), provided that a transaction either (i) issued by Txgen , or (ii) is neutral, is given as input to all honest players continuously for u consecutive rounds after round r_0 , then there exists an honest party who will report this transaction at a block more than k blocks from the end of the ledger.

Chain quality, chain growth and the strong common prefix property were shown in [27] to be sufficient to implement such a ledger⁹ in a black-box manner. Our protocol satisfies all these properties after a specific condition is met. Chain quality holds after the $2f(1 + \delta)(l + 1)$ block in the chain of any player, as Corollary 22 dictates, and common prefix and chain growth hold after round $2l + 2$, according to Theorem 17. Finally, due to chain growth, after at most $(2(1 + \delta)(1 - \delta)f/\gamma + 2)(l + 1) \leq 14(l + 1)$ rounds all necessary conditions will have been met with overwhelming probability.

Lemma 24 (Persistence). *Assume $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, $\gamma \geq (1 + \delta)\beta$, for some real $\delta \in (0, 1/2)$. Then for all $k \in \mathbb{N}$ protocol $\Pi_{\text{PL}}^{\text{boot}}$ satisfies Persistence after round $2l + 2$ with probability $1 - e^{-\Omega(\delta^5 k)}$, where k is the depth parameter.*

Lemma 25 (Liveness). *Assume $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, $\gamma \geq (1 + \delta)\beta$, for some real $\delta \in (0, 1/2)$. Further, assume oracle Txgen is unambiguous. Then for all $k \in \mathbb{N}$ protocol $\Pi_{\text{PL}}^{\text{boot}}$ satisfies Liveness after round $14(l + 1)$ with wait time $u = \frac{3}{(1-\delta)\gamma} \cdot \max(k, \frac{1}{1-(1+\frac{\delta}{2})\frac{\beta}{\gamma}})$ rounds and depth parameter k with probability at least $1 - e^{-\Omega(\delta^5 k)}$.*

Corollary 26. *Assume $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, $\gamma \geq (1 + \delta)\beta$, for some real $\delta \in (0, 1/2)$. Then protocol $\Pi_{\text{PL}}^{\text{boot}}$ implements a robust transaction ledger with parameter $r_0 = 14(l + 1)$.*

Fast PKI setup. Next, we use the ledger to generate an honest majority PKI from scratch in a number of rounds that is linear in the security parameter. The first idea that we are going to use is that of a *2-for-1 POW* described in [21]. At a high level, the technique allows to do combined mining for two POW schemes in the price of one. In more detail, we can add additional information in the queries to the random oracle, and if the response to the query is less than some value T_1 , then we consider it a valid POW of type 1; if it is greater than some value T_2 we consider it as a valid POW of type 2. T_1 and T_2 should be appropriately chosen so that the events of success in either of these POWs are independent. The second POW is used to “mine” transactions, in the same way blocks are mined. This guarantees that the *number of transactions* is proportional to the hashing power of each player. By having parties broadcast their transactions on one hand, and making sure that at least one honest block that contains these transactions is in the chain of all honest parties due to liveness on the other hand, the protocol in [21] manages to achieve consensus assuming an honest-majority hashing power.

In our case, transactions will contain the public keys, and in this way we will obtain an honest-majority PKI. However, in contrast with [21], we cannot let parties start mining transactions from the beginning of the execution, since the adversary would have some additional precomputation time. Instead, we are going to wait for the public ledger to be established, and then use some of the blocks added by honest parties to guarantee that all transactions were mined recently enough.

⁹A similar definitional approach was pursued in [34].

In more detail, any POW will be represented by a triple $\langle w, ctr, label \rangle$. The verification procedure for “block level” POWs (“block POWs” for short) will be of the form

$$H(ctr, \langle G(w), label \rangle) < T_1,$$

while the verification procedure for the “transaction level” POWs will be of the form

$$[H(ctr, \langle label, G(w) \rangle)]^R < T_2,$$

where $[a]^R$ denotes the reverse of the bitstring a . In w we are going to encode the information needed for each application. For example, in block POWs, w will contain the transactions related to this block as well as the hash of the previous block. Note that by making one hash query of the form $H(ctr, \langle G(w_0), G(w_1) \rangle)$ and only two comparisons, we will be mining POWs of both types at the same time. Moreover, if $\lceil \log(T_1) \rceil + \lceil \log(T_2) \rceil$ is less than κ , where κ is the size of the hash’s output, then the events of succeeding in any of the two POWs are independent, since they depend on different bits of the hash which are sampled independently and uniformly at random by the random oracle.

Next, we describe our protocol $\Pi_{\text{PL}}^{\text{PKI}}$ for an honest party P . L_1, L_2 are constants such that $L_1 < L_2$.

- *Initialization.* P runs $\Pi_{\text{PL}}^{\text{boot}}$, as described so far, until she receives a chain of length at least L_1 . We choose L_1 so that it is guaranteed that all security properties hold, and about k new blocks have been inserted in the common-prefix of the chains of all honest players.
- *2-for-1 mining.* Let \mathcal{C} be P ’s chain at the end of the initialization phase. From now on, she is going to do 2-for-1 POW mining, and include in her transaction POWs (i) the hash of the $(L_1 - k)$ -th block of \mathcal{C} , and (ii) a randomly generated public key for which she has stored the corresponding secret key. Obviously, a new key must be generated every time she starts mining a new transaction. Whenever P mines a new transaction, she diffuses it to the network, and whenever she receives one, she includes it in the transactions of the block she is mining.

The first time P receives a chain of length greater or equal to L_2 , she runs the Key extraction procedure (below). The phase ends at round $\frac{L_2}{(1-\delta)\gamma}$, where P runs the Termination procedure.

- *Key extraction.* P extracts and stores a set of keys from her current chain according to the following rules: If chain \mathcal{C}' is her chain at this round, she stores any public key which belongs to a transaction that (i) is in the first $L - k$ blocks of \mathcal{C}' , and (ii) the hash of the block contained in the transaction matches the hash of the $(L_1 - k)$ -th block in her chain.
- *Termination.* P outputs the keys from the key extraction phase and terminates.

Next, we prove that a consistent PKI with an honest majority is generated at the end of the execution protocol $\Pi_{\text{PL}}^{\text{PKI}}$. Two properties are guaranteed: (1) honest parties output the same set of keys and (2) more than half of these keys have been generated by them. For the rest of this section let α_2, β_2, f_2 be the corresponding values of α, β, f for the difficulty level T_2 , e.g. $f_2 = nq\frac{T_2}{2^k}$. The full proof of the following is provided in the supplementary material.

Theorem 27. *Assume $3(1 + \delta)f < 1$, $l > \frac{(1-\delta)k/\gamma+3}{1-3(1+\delta)f}$, $\gamma \geq (1 + \delta)\beta$, for some real $\delta \in (0, 1/2)$ and $\lceil \log(T_1) \rceil + \lceil \log(T_2) \rceil \leq \kappa$. Then, for parameters $L_1 = 14(l + 1)(1 + \delta)f + 2k$ and $L_2 = L_1 + 2k \cdot (1 + \frac{10}{\delta})\frac{(1+\delta)f}{(1-\delta)\gamma}$ the following hold for protocol $\Pi_{\text{PL}}^{\text{PKI}}$ with probability $1 - e^{-\Omega(\delta^5 k)}$:*

- *All honest players output the same set of public keys, the size of which is*

$$k \frac{\alpha_2}{\gamma} \frac{20}{\delta} \leq N \leq 60k \frac{f_2}{\gamma} (1 + \frac{10}{\delta});$$

- the majority of the keys are generated by honest parties; and
- $\Pi_{\text{PL}}^{\text{PKI}}$ has round complexity linear in κ .

Proof. (Sketch) First, note that the adversary can start precomputing transactions about $2k/\gamma$ rounds before the honest parties. Otherwise, she will be unable to predict the hash of their chain as dictated by our protocol, since by the chain quality property the chain of each honest player will contain an honest block near the tail of the chain. Moreover, again by the chain quality and common prefix properties, the adversary will stop mining transactions at most $2k/\gamma$ rounds after the honest parties. After this round, she will be unable to insert her transactions deep enough in the chain for the honest parties to take them into account. Finally, by choosing an appropriate value for ϵ , we are sure that the number of keys mined by the honest parties is greater than the number of keys mined by the adversary. \square

Remark 4. To better understand $\Pi_{\text{PL}}^{\text{PKI}}$ we compute different parameters of the system for the Bitcoin network parameters. Assume that $f = 2\%$, $\alpha = 1.33\%$, $\gamma = 1.31\%$, $\beta = 0.6\%$, $k = 10$, and $\delta = 0.25$. The choice of f approximately corresponds to a rate of one block per 10 minutes with a round duration of about 12 seconds; the adversary’s hashing power is half of that of the honest parties. Then, $l \approx 623$, which corresponds in terms of rounds to about 2 hours. Moreover, if we set f_2 to be equal to f/k we have¹⁰ that $80 < N < 600$. We note that the parameters of Bitcoin are quite conservative and that’s why our runtime suffers. In principle, by carefully analyzing and re-engineering our protocol we can get tighter bounds; many of the design decisions we got here, were made to aid the readability of our work.

Remark 5. If we want to have better control on the number of keys obtained by each party or on the total number of keys in our PKI, we can change the difficulty of the transaction-level POW. The probability that some honest party succeeds in mining at least one transaction is:

$$\Pr[\geq 1 \text{ key}] = 1 - \Pr[0 \text{ keys}] = 1 - \left(1 - \frac{T_2}{2^\kappa}\right)^{q \frac{20k}{(1-\delta)\gamma\delta}} \geq 1 - e^{-\frac{T_2}{2^\kappa} \cdot q \frac{20k}{(1-\delta)\gamma\delta}}.$$

Hence, by setting $T_2 > \frac{\ln(\frac{1}{a})2^\kappa(1-\delta)\gamma\delta}{q \cdot 20k}$, each party will obtain at least one key with probability at least $1 - a$, for any $a \in (0, 1)$. Note here that T_2 and κ must be carefully chosen to retain the independence of the 2 POWs.

Consensus and other applications. Next, we describe how $\Pi_{\text{PL}}^{\text{PKI}}$ can be used in other contexts. First, a direct application of our protocol is in the context of *proof of stake* protocols. In this type of protocols, blocks are mined by randomly selecting stake holders with probability proportional to their stake. A typical requirement for bootstrapping such protocols (e.g. [28, 30]), is that in the initial state of the economy the majority of the coins is controlled by honest parties. By assigning one coin to each public key produced by our protocol, we can efficiently and securely bootstrap a proof of stake protocol.

A more general application of $\Pi_{\text{PL}}^{\text{PKI}}$ is in solving consensus (aka Byzantine agreement) [36, 31], with no trusted setup, and in a number of rounds independent of the number of parties. If parties submit transactions containing their input instead of public keys, it follows that by taking the majority of their output they are going to achieve Byzantine agreement. That is, everyone will agree on the same value (the Agreement property), and if all honest parties have the same input v , they are all going to output v (Validity).

¹⁰The upper bound is derived from a tighter formula from our analyses found in the supplementary material.

Finally, our protocol for the establishment of an honest-majority PKI enables the application of traditional Byzantine fault-tolerant techniques for ledger creation and maintenance based on “sub-committees” as opposed to mining (cf. [7]) to permissionless networks. Instead of having arbitrary membership authorities, these committees can be elected using our protocol with the guarantee of an honest majority. Note that by changing the difficulty of the transaction-level POW we can force the number of parties in the committee to be in a specific predefined interval.

References

- [1] M. Andrychowicz and S. Dziembowski. Pow-based distributed cryptography with no trusted setup. In *CRYPTO 2015*, pages 379–399, 2015.
- [2] J. Aspnes, C. Jackson, and A. Krishnamurthy. Exposing computationally-challenged Byzantine impostors. Technical Report YALEU/DCS/TR-1332, Yale University Department of Computer Science, July 2005.
- [3] A. Back. Hashcash. <http://www.cypherspace.org/hashcash>, 1997.
- [4] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from bitcoin. *IACR Cryptology ePrint Archive*, 2014:349, 2014.
- [5] Bitcoinwiki. Genesis block. https://en.bitcoin.it/wiki/Genesis_block.
- [6] M. Borderding. Levels of authentication in distributed agreement. In *Distributed Algorithms, 10th International Workshop, WDAG '96*, 1996.
- [7] C. Cachin. Architecture of the hyperledger blockchain fabric. In *Workshop on Distributed Cryptocurrencies and Consensus Ledgers*, 2016.
- [8] R. Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
- [9] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. *IACR Cryptology ePrint Archive*, 2000:67, 2000.
- [10] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS 2001*, pages 136–145. IEEE Computer Society, 2001.
- [11] R. Cohen, S. Coretti, J. Garay, and V. Zikas. Probabilistic termination and composability of cryptographic protocols. In *Annual Cryptology Conference*, pages 240–269. Springer, 2016.
- [12] J. Consideine, M. Fitzi, M. K. Franklin, L. A. Levin, U. M. Maurer, and D. Metcalf. Byzantine agreement given partial broadcast. *J. Cryptology*, 18(3):191–217, 2005.
- [13] D. Dolev and H. R. Strong. Authenticated algorithms for byzantine agreement. *SIAM J. Comput.*, 12(4):656–666, 1983.
- [14] J. R. Douceur. The sybil attack. IPTPS '01. Springer-Verlag, 2002.
- [15] C. Dwork, N. A. Lynch, and L. J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, 1988.
- [16] C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In *CRYPTO 1992*, pages 139–147, 1992.
- [17] I. Eyal, A. E. Gencer, E. G. Sirer, and R. van Renesse. Bitcoin-ng: A scalable blockchain protocol. *CoRR*, abs/1510.02037, 2015.
- [18] I. Eyal and E. G. Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography*, 2014.

- [19] M. Fitzi. *Generalized communication and security models in Byzantine agreement*. PhD thesis, ETH Zurich, Zürich, Switzerland, 2003.
- [20] J. A. Garay, J. Katz, C. Koo, and R. Ostrovsky. Round complexity of authenticated broadcast with a dishonest majority. In *(FOCS 2007)*,, pages 658–668, 2007.
- [21] J. A. Garay, A. Kiayias, and N. Leonardos. The bitcoin backbone protocol: Analysis and applications. In *EUROCRYPT 2015, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, pages 281–310, 2015.
- [22] J. A. Garay, A. Kiayias, N. Leonardos, and G. Panagiotakos. Bootstrapping the blockchain-directly. *IACR Cryptology ePrint Archive*, 2016:991, 2016.
- [23] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229. ACM, 1987.
- [24] A. Juels and J. G. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *NDSS*. The Internet Society, 1999.
- [25] J. Katz, U. Maurer, B. Tackmann, and V. Zikas. Universally composable synchronous computation. *IACR Cryptology ePrint Archive*, 2011:310, 2011.
- [26] J. Katz, A. Miller, and E. Shi. Pseudonymous secure computation from time-lock puzzles. *IACR Cryptology ePrint Archive*, 2014:857, 2014.
- [27] A. Kiayias and G. Panagiotakos. Speed-security tradeoffs in blockchain protocols. Technical report, IACR: Cryptology ePrint Archive, 2015.
- [28] A. Kiayias, A. Russell, B. David, and R. Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. *Cryptology ePrint Archive*, Report 2016/889, 2016. <http://eprint.iacr.org/2016/889>.
- [29] S. King. Primecoin: Cryptocurrency with prime number proof-of-work. <http://primecoin.io/bin/primecoin-paper.pdf>, July 2013.
- [30] S. King and S. Nadal. Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. *self-published paper*, August, 19, 2012.
- [31] L. Lamport, R. E. Shostak, and M. C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- [32] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <http://bitcoin.org/bitcoin.pdf>, 2008.
- [33] S. Nakamoto. Bitcoin open source implementation of P2P currency. <http://p2pfoundation.ning.com/forum/topics/bitcoin-open-source>, February 2009.
- [34] R. Pass, L. Seeman, and abhi shelat. Analysis of the blockchain protocol in asynchronous networks. *Cryptology ePrint Archive*, Report 2016/454, 2016. <http://eprint.iacr.org/2016/454>.
- [35] R. Pass, L. Seeman, and A. Shelat. Analysis of the blockchain protocol in asynchronous networks. In *EUROCRYPT 2017*, pages 643–673, 2017.
- [36] M. C. Pease, R. E. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.
- [37] R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, Cambridge, MA, USA, 1996.
- [38] A. C.-C. Yao. Protocols for secure computations (extended abstract). In *FOCS*, pages 160–164. IEEE, 1982.

Supplementary Material

A Proofs

A.1 Proof of Lemma 11

Proof. We first look into the case where the block in the statement is a genesis block. So for the sake of contradiction, suppose the adversary mines some genesis block B before round r , and this block is accepted by some honest party P with a value greater or equal to $r - 1$. In the worst case, that means that the adversary also created sets A_{l+1}, \dots, A_r such that:

- B is a valid block that contains $H(A_{l+1})$;
- for $r + 1 < j \leq l + 1$, A_j contains $H(A_{j-1})$; and
- if c_{r-1} is the $r - 1$ round challenge of P , c_{r-1} is in A_r .

Due to the random nonce honest parties add to their challenges at every round, the probability that the adversary can guess c_{r-1} before round r is negligible in κ . Hence, to put it more simply, the adversary must compute sets A_{l+1}, \dots, A_r such that $H(A_{l+1} || H(\dots H(A_r || c_r)))$ is equal to c , where c is the value in some of the blocks he has precomputed. Assuming his precomputation time is bounded by a polynomial in κ , say $p_1(\kappa)$, and that he also has $p_2(\kappa)$ steps in order to compute these sets, by the union bound the probability of success is

$$\Pr[\text{success}] \leq \frac{p_1(\kappa) + p_2(\kappa)}{2^\kappa} (p_1(\kappa) + p_2(\kappa)) \leq \text{negl}(\kappa)$$

where $\frac{p_1(\kappa) + p_2(\kappa)}{2^\kappa}$ is an upper bound on the probability that a query made to the oracle by the adversary satisfies the conditions described before, and $p_1(\kappa) + p_2(\kappa)$ is an upper bound on the number of queries he makes. Hence, the adversary cannot create such a genesis block with overwhelming probability in κ .

Otherwise, suppose that there exists some non-genesis adversarial block B' , that has been mined before round r and is part of a chain with genesis block B , where for some honest party P , $w_p(B) \geq r - 1$. If no collision has occurred, B must have been mined before B' , and thus as we proved for the first case, the probability that a genesis block with these properties exists is negligible in κ . Hence, the lemma follows. \square

A.2 Proof of Lemma 13

Proof. For the sake of contradiction, suppose there exists some honest party P' such that P' has received chain \mathcal{C} at round r and it is not valid for him. The only reason this may happen is that P' has not accepted \mathcal{C} 's genesis block B . By Lemma 10 all honest parties know the genesis blocks mined by other honest parties, thus B must have been computed by the adversary. We have two cases. In the first case, round r is before round $2l + 2$. Recall that $2l + 1$ is the last round where the validation predicate is updated. Then, since P has adopted \mathcal{C} in the previous round, he must have also broadcast B in the previous round. Thus, all honest parties will accept B as a valid genesis block at round r and will also accept \mathcal{C} as valid, which is a contradiction.

Otherwise, suppose $r = 2l + 2$. Again, if B was received before round $2l + 1$ by some honest party, \mathcal{C} will look valid to all parties. So B must have been received for the first time at round $2l + 1$ by P ; no honest party accepts new genesis blocks after this round. We will show that with overwhelming probability in k , no honest party will ever accept a chain based on B .

Without loss of generality, suppose that P is the first honest party that accepts a chain based on B at round r . Let E_1 be the event where the honest parties mine a genesis block after round $l+2+((1-\delta)k/\gamma-1)$ for the first time. It holds that the probability of E_1 is at most $(1-\gamma)^{(1-\delta)k/\gamma} < e^{-(1-\delta)k}$. So suppose that E_1 does not happen and at round $l+2+((1-\delta)k/\gamma-1)$ ($< r$) the honest parties have computed at least one genesis block that has weight at least $l - (1-\delta)k/\gamma$. Hence, in this case, it follows from Lemma 12 that every honest party at round r will have a chain of weight at least $l - (1-\delta)k/\gamma - 2 + 3(Y(S') - 1) - Z_1(S')$, where $S' = \{l+2, \dots, r\}$ and $Z_1(S')$ is the set of blocks the adversary has broadcast to slow down chain growth during rounds in S' .

On the other hand, since block B is adversarial, and chain \mathcal{C} is accepted for the first time by an honest party at round r , all of its blocks must be adversarial; possibly \mathcal{C} contains just B . By definition block B weighs 1 for P . Thus, by Lemma 11 the adversary can start extending \mathcal{C} at round 2. However, the blocks that the adversary uses to slow down chain growth cannot also be used for \mathcal{C} , because they belong to chains whose genesis block has been announced earlier. So let $Z_2(S)$ denote the blocks mined by the adversary in $S = \{2, \dots, r\}$ that are not in $Z_1(S')$. In order for \mathcal{C} to be accepted by some honest party, it must hold that

$$1 + 3(Z_2(S) - 1) \geq l - (1-\delta)k/\gamma - 2 + 3(Y(S') - 1) - Z_1(S').$$

Since $Z_1(S)$ and $Z_2(S')$ are disjoint and $S' \subseteq S$, the above implies:

$$3Z(S) \geq l - (1-\delta)k/\gamma + 3Y(S') - 3. \quad (1)$$

Let E_2 be the event that

$$Y(S') \leq Z(S'). \quad (2)$$

From Lemma 8, if $|S'| > k$, which is the case here, $\Pr(E_2) \leq e^{-\Omega(\delta^2 k)}$. Also, let E_3 be the event that $Z(S \setminus S') \geq (1+\delta)l\beta$. Again, by an application of the Chernoff bound we have that $\Pr(E_3) \leq e^{-\Omega(\delta^2 k\beta)}$, since $|S \setminus S'| = l \geq k$. Suppose now that none of E_2 or E_3 holds. Then, since by our assumptions we have that

$$3(1+\delta)l\beta < l - (1-\delta)k/\gamma - 3, \quad (3)$$

it follows that inequality 1 cannot hold and thus no honest party will ever accept a chain based on B . By an application of the union bound the event $E_1 \vee E_2 \vee E_3$ has probability at most $e^{-\Omega(\delta^2 k)}$ and the lemma follows.

A subtle point here is that since the lemma holds for $r = 2l+2$, it follows that Lemma 12 should hold for $r = 2l+2$. The same proof then can be applied for $r = 2l+3$ in this lemma, and inductively it follows that both lemmas hold for any round of the execution. However, since we have to apply repeatedly the universal validity proof for different r , we must argue about the probability that the statement holds for any r . Let $E_2(r)$ be the parameterized version of E_2 , where $E_2(r)$ is the event where $Y(S') \leq Z(S')$ for $S' = \{l+2, \dots, r\}$. Then, for some $\epsilon > 0$ it holds that

$$\bigvee_{i \geq 2l+2} E_2(i) \leq \sum_{i \geq 2l+2} e^{-(\epsilon\delta^2 k - \ln(2))} \leq e^{-(\epsilon\delta^2 k - \ln(2) + \ln(1 - e^{-\epsilon\delta^2}))} \leq e^{-\Omega(\delta^2 k)}$$

Thus, again by the union bound, the event $E_1 \vee \bigvee_{i \geq 2l+2} E_2(i) \vee E_3$ has probability at most $e^{-\Omega(\delta^2 k)}$. If this event does not occur, as we have argued universal validity holds for any r . Hence, the lemma follows. \square

A.3 Proof of Lemma 12

Proof. Since $r < 2l + 2$, the genesis blocks of the chains that honest players have at this or any previous round, must have weight at least 1. Hence, by Lemma 9, for any chain \mathcal{C}' of these chains, it should hold that $h_{\mathcal{C}'} \leq \ell_{\mathcal{C}'} + 1$. Let $\ell(r) = d$ iff d is the minimum value of the set $\{\ell_{\mathcal{C}} \mid P \text{ is honest and at round } r \text{ has chain } \mathcal{C}\}$. Then we can show the following:

Claim 1. Suppose round r is uniquely successful and $\ell(r) = d$. Then for any round $s > r$ it holds that $\ell(s) \geq d + 2$. Moreover, if the adversary has not broadcast by round s any chain \mathcal{C} that contains an adversarial block $B = \text{head}(\mathcal{C})$ such that $\ell_{\mathcal{C}}(B) = d + 2$, it holds that $\ell(s) \geq d + 3$.

Proof of Claim. The proof is quite straightforward. For the first part, since $\ell(r) = d$ and r is uniquely successful, an honest party will broadcast a chain \mathcal{C} at round r where $\ell_{\mathcal{C}} \geq d + 3$. Thus, at round $r + 1$ all parties will receive a chain that has weight at least $d + 3$ according to their view. This implies that, at worst, they may adopt a chain of the same weight, hence in any case it holds that $\ell(s) \geq d + 2$.

Suppose that by round s , the adversary has not broadcast any block $B' = \text{head}(\mathcal{C}')$ such that $\ell_{\mathcal{C}'}(B') = d + 2$ and \mathcal{C}' is valid. For the sake of contradiction, suppose that there exists some round $s > r$ such that $\ell(s) < d + 3$. Since at round $r + 1$ honest parties receive \mathcal{C} , they will all adopt a chain that weighs in their view at least $d + 3$. Otherwise, they would adopt \mathcal{C} . Moreover, they will never adopt a chain with smaller weight. Hence, the only way $\ell(r + 1) = d + 2$ is if a chain that has weight $d + 2$ for some honest party was broadcast at some round. By our assumption, an honest party has mined the head of this chain. Since $\ell(r) = d$, he must have done that before round r , otherwise the chain would weigh at least $d + 3$ for any honest party. However, if he mined this chain before round r , at round r all honest parties would have received this chain and $\ell(r) = d + 2$, which is a contradiction. Hence, the claim follows. \dashv

Observe that if at round r P_1 has a chain \mathcal{C} of weight $w_{P_1}(\mathcal{C}) = d$, then he broadcast \mathcal{C} at an earlier round. It follows that every honest party P will receive \mathcal{C} by round r and $w_P(\mathcal{C}) \geq d - 1$. It is easy to see that if each honest party P at some round r' has received a chain \mathcal{C} where $w_P(\mathcal{C}) \geq k$, then for every round $s' \geq r'$ it holds that $\ell(s') \geq k - 1$. Thus for every round $s' \geq r$ it holds that $\ell(s') \geq d - 2$.

We now have two cases. In the first case, $\sum_{i=r}^{s-1} Y_i \leq \sum_{i=r}^{s-1} Z_i$. The claim above guarantees that every time a uniquely successful round r' happens, $\ell(r' + 1) \geq \ell(r') + 2$. Thus, by repeatedly applying this argument we immediately get that:

$$\begin{aligned} \ell(s) &\geq d - 2 + 2 \sum_{i=r}^{s-1} Y_i \geq d - 2 + 3 \sum_{i=r}^{s-1} Y_i - \sum_{i=r}^{s-1} Y_i \\ &\geq d - 2 + 3 \sum_{i=r}^{s-1} Y_i - \sum_{i=r}^{s-1} Z_i, \end{aligned}$$

which implies that at round s all honest parties have received a chain that has sufficient weight according to the lemma.

Otherwise, $\sum_{i=r}^{s-1} Y_i > \sum_{i=r}^{s-1} Z_i$. Note that for every uniquely successful round, in order for the condition of the claim above to hold, the adversary must broadcast different blocks that have weight at least $\ell(r) + 2 = d$. Thus, for at least $\sum_{i=r}^{s-1} Y_i - \sum_{i=r}^{s-1} Z_i$ uniquely successful rounds the condition

of the claim will not hold and for any such round r' , $\ell(r' + 1) \geq \ell(r') + 3$. Thus,

$$\begin{aligned} \ell(s) &\geq d - 2 + 3\left(\sum_{i=r}^{s-1} Y_i - \sum_{i=r}^{s-1} Z_i\right) + 2\sum_{i=r}^{s-1} Z_i \\ &\geq d - 2 + 3\sum_{i=r}^{s-1} Y_i - \sum_{i=r}^{s-1} Z_i. \end{aligned}$$

□

A.4 Proof of Lemma 15

Proof. Let $B = \text{head}(\mathcal{C})$ be a block that is contained in R and B' be the genesis block of \mathcal{C} . First, we are going to show that if B' was computed before round $l + 1 - 2(1 - \delta)\delta^2k$, then B must have been computed after round $l + 1 - (1 - \delta)\delta^2k$. Suppose that block B' was computed at round $l + 1 - r^*$ of the challenges phase, for $r^* > 2(1 - \delta)\delta^2k$; thus, by Lemma 11 it holds that $h_{\mathcal{C}}(B') < l - r^*$. In order for \mathcal{C} to have the required weight, the adversary must have mined at least

$$\lceil (l - (1 - \delta)\delta^2k - 1 - (l - r^* - 1))/3 \rceil = \lceil (r^* - (1 - \delta)\delta^2k)/3 \rceil$$

blocks, starting from round $l + 1 - r^*$. Let E_1 be the event that the adversary after $r^* - (1 - \delta)\delta^2k$ rounds has computed more than $(1 + \delta)(r^* - (1 - \delta)\delta^2k)\beta$ blocks. By an application of the Chernoff bound, $\Pr(E_1) \leq e^{-\Omega(\delta^4k)}$. By our assumptions, this number of blocks is not sufficient to get a chain of the required weight, that is

$$(1 + \delta)(r^* - (1 - \delta)\delta^2k)\beta < \lceil (r^* - (1 - \delta)\delta^2k)/3 \rceil.$$

Therefore, if E_1 does not hold, then the adversary will start mining B after round $l + 1 - (1 - \delta)\delta^2k$, for any B in R .

Next, we are going to bound the number of blocks the adversary can compute in $2(1 - \delta)\delta^2k$ rounds; recall that we are interested in blocks that were mined after round $l - 2(1 - \delta)\delta^2k$ and before round $l + 2$. Let E_2 be the event that the adversary mines at least $2(1 + \delta/8)(1 - \delta)\delta^2k\beta (< \frac{5\delta}{9}k\beta)$ blocks in $2(1 - \delta)\delta^2k$ rounds. By an application of the Chernoff bound, $\Pr(E_2) \leq e^{-\Omega(\delta^4k)}$. Since by an application of the union bound $E_1 \vee E_2$ happens with probability at most $e^{-\Omega(\delta^4k)}$, the lemma follows. □

A.5 Proof of Theorem 17

We first prove a weaker lemma, called common-prefix lemma in [21]. After proving this lemma for our own model, we can apply the same ideas as in [21] to get a proof for the theorem.

Lemma 28. *Let $3(1 + \delta)f < 1$, $l > \frac{(1 - \delta)k/\gamma + 3}{1 - 3(1 + \delta)f}$, $\gamma \geq (1 + \delta)\beta$ for some $\delta \in (0, 1)$, and consider two chains \mathcal{C}_1 and \mathcal{C}_2 such that $\ell_{\mathcal{C}_2} \geq \ell_{\mathcal{C}_1}$. If \mathcal{C}_1 is adopted by an honest party at a round not before round $2l + 2$, and \mathcal{C}_2 is either adopted by an honest party or diffused at the same round, then the probability that $\mathcal{C}_1^{\lceil k} \not\preceq \mathcal{C}_2$ or $\mathcal{C}_2^{\lceil k} \not\preceq \mathcal{C}_1$ is at most $e^{-\Omega(\delta^4k)}$.*

Proof. We first show that a fork between two chains implies that the adversary must have mined a number of blocks proportional to the uniquely successful blocks associated with these chains.

Claim 2. Let $\mathcal{C}_1, \mathcal{C}_2$ be chains at some round r and B'_0, \dots, B'_k be u.s. blocks in chains $\mathcal{C}'_0, \dots, \mathcal{C}'_k$ in increasing order of round mined. Then, if $\ell_{\mathcal{C}_1 \cap \mathcal{C}_2} < \ell(B'_0) - 1$ and for all $\mathcal{C} \in \{\mathcal{C}_1, \mathcal{C}_2\} \setminus \mathcal{C}'_k$: $\ell_{\mathcal{C}} \geq \ell(B'_k) - 1$, there exist different adversarial blocks B_0, \dots, B_k such that for $i \in \{0, \dots, k\}$: $\ell(B_i) \in [\ell(B'_0) - 1, \ell(B'_k) + 1]$ and $B_i \in \{\mathcal{C}'_0, \dots, \mathcal{C}'_k\} \cup \{\mathcal{C}_1, \mathcal{C}_2\}$.

Proof of Claim. We iterate over $U = \{(B'_0, \mathcal{C}'_0), \dots, (B'_k, \mathcal{C}'_k)\}$ in the order of increasing index. Note that by the Claim in Lemma 12, if (B, \mathcal{C}) and (B', \mathcal{C}') are two consecutive elements of U , then

$$\ell_{\mathcal{C}'}(B') - \ell_{\mathcal{C}}(B) \geq 2 \quad (4)$$

Consider $(B, \mathcal{C}) \in U$ and suppose all the previous elements have been associated with a distinct adversarial block. In particular, let $(\bar{B}, \bar{\mathcal{C}})$ be the previous one associated to (B', \mathcal{C}') . To choose the adversarial block (B', \mathcal{C}') to associate with (B, \mathcal{C}) we consider the following cases (an example is presented in Figure 4 covering most of the cases). In each case B' is determined by \mathcal{C}' and Lemma 16.

- If $\mathcal{C} \notin \{\mathcal{C}_1, \mathcal{C}_2\}$ and $\bar{\mathcal{C}} \in \{\mathcal{C}_1, \mathcal{C}_2\}$, then we have two cases. If $\bar{B} \notin \mathcal{C}$ and $\bar{B}' \notin \mathcal{C}$, then let $\mathcal{C}' = \{\mathcal{C}_1, \mathcal{C}_2\} \setminus \bar{\mathcal{C}}$ and let B^* in chain $\mathcal{C}^* = \mathcal{C}$ be the block guaranteed from Lemma 16 for block \bar{B} in chain \mathcal{C} . This block will be used in a subsequent step. Otherwise, $\bar{B} \in \mathcal{C}$ or $\bar{B}' \in \mathcal{C}$. Then \mathcal{C}' should be chosen appropriately from $\{\mathcal{C}_1, \mathcal{C}_2\}$, so that in the next transition from a chain not in $\{\mathcal{C}_1, \mathcal{C}_2\}$ to a chain in $\{\mathcal{C}_1, \mathcal{C}_2\}$, the corresponding block from Lemma 16 does not intersect with the previously chosen block. This is always possible since any of the two chains can be selected.
- If $\mathcal{C} \in \{\mathcal{C}_1, \mathcal{C}_2\}$ and $\bar{\mathcal{C}} \notin \{\mathcal{C}_1, \mathcal{C}_2\}$, then $\mathcal{C}' \in \{\mathcal{C}_1, \mathcal{C}_2\} \setminus \mathcal{C}$. If \mathcal{C}^* is defined from a previous application of the first rule, we match B with B^* .
- If $\mathcal{C} \in \{\mathcal{C}_1, \mathcal{C}_2\}$ and $\bar{\mathcal{C}} \in \{\mathcal{C}_1, \mathcal{C}_2\}$, then let $\mathcal{C}' \in \{\mathcal{C}_1, \mathcal{C}_2\} \setminus \mathcal{C}$.
- If $\mathcal{C} \notin \{\mathcal{C}_1, \mathcal{C}_2\}$ and $\bar{\mathcal{C}} \notin \{\mathcal{C}_1, \mathcal{C}_2\}$, then $\{\mathcal{C}_1, \mathcal{C}_2\} \setminus \bar{\mathcal{C}}$. The adversarial block guaranteed by Lemma 16 for \mathcal{C}' is no common to $\mathcal{C}_1, \mathcal{C}_2$ due to $\ell_{\mathcal{C}_1 \cap \mathcal{C}_2} < \ell(B'_0) - 1$.
- If $B = B_0$, then if $\mathcal{C} \in \{\mathcal{C}_1, \mathcal{C}_2\}$, $\mathcal{C}' \in \{\mathcal{C}_1, \mathcal{C}_2\} \setminus \mathcal{C}$. Otherwise, as in the first case, \mathcal{C}' should be chosen appropriately from $\{\mathcal{C}_1, \mathcal{C}_2\}$.

We need to verify that the above procedure does not assign the same block to two distinct elements of U , (B_u, \mathcal{C}_u) and (B_v, \mathcal{C}_v) .

Note first that this is not possible if

$$|\ell_{\mathcal{C}_u}(B_u) - \ell_{\mathcal{C}_v}(B_v)| \geq 3.$$

For example, by Equation (4), this is true if they are not consecutive in U . To see this, observe that by Lemma 16,

$$\ell_{\mathcal{C}'_u}(B'_u) \in [\ell_{\mathcal{C}_u}(B_u) - 1, \ell_{\mathcal{C}_u}(B_u) + 1],$$

while

$$\ell_{\mathcal{C}'_v}(B'_v) \in [\ell_{\mathcal{C}_v}(B_v) - 1, \ell_{\mathcal{C}_v}(B_v) + 1].$$

Since these intervals are disjoint due to the inequality above, it follows that $B'_u \neq B'_v$.

Thus, we only need to consider the case

$$\ell_{\mathcal{C}_{u+1}}(B_{u+1}) - \ell_{\mathcal{C}_u}(B_u) = 2.$$

It is not hard to see that this situation cannot occur when B_{u+1} is a descendant of B_u . Moreover, when the blocks assigned are on different chains, it is guaranteed that they are different. These covers all different cases, except two. The first one, is when the first rule is applied and \bar{B}' is in \mathcal{C} . Then, B' will be different than \bar{B}' because, either is on a different chain $(\{\mathcal{C}_1, \mathcal{C}_2\} \setminus \mathcal{C})$ or they are on the same chain and $\ell_{\mathcal{C}}(B) - \ell_{\mathcal{C}}(\bar{B}') \geq 3$. The other case is when \mathcal{C}^* has been defined by the first rule and the second rule is applied. In this case however, the honest block is matched to B^* , so we can safely ignore the matched block and still have a complete matching. Note that, B^* is unique since no block consecutive to the one matched on B^* is matched in chain \mathcal{C}^* . Hence, again it is impossible that these two blocks are the same.

—

Now we can proceed to the core of the proof. The idea is that if a fork exists, we will use the previous claim multiple times and get a matching between a sufficiently large amount of uniquely successful blocks and adversarial blocks. Then, we will show that it is only with small probability that the adversary has mined this number of blocks. For the rest of this proof assume that \mathcal{C}_1 was adopted by an honest party at round r and it diverges from \mathcal{C}_2 at some round r' .

We start by defining a “bad” event that will only happen with negligible probability. Afterwards, we will show that if this event does not occur, then there cannot be a fork as described on the statement of the Lemma. Let BAD be the event where at least one of the following events happens:

- a successful rounds happens for the first time after round $l + 1 + (1 - \delta)\delta^2(r - r')$
- the event of Lemma 15 does happen
- for any $s \leq r - r'$, for the set of rounds $S = \{s, \dots, r\}$:

$$Y(S) \leq (1 + \frac{5\delta}{9})Z(S)$$

The probability that any of these events happens is at most $e^{-\Omega(\delta^4\beta(r-r'))}$. It follows from an application of the union bound to the disjunction of these events, that the probability that BAD happens is at most $e^{-\Omega(\delta^4\beta(r-r'))}$.

Suppose $\mathcal{C}_1, \mathcal{C}_2$ diverge at round r' and assume BAD does not occur. Let block $B'_0 \in \mathcal{C}_1 \cap \mathcal{C}_2$ be the most recently mined u.s. block where all subsequently mined u.s. blocks are descendants of B'_0 . If no such block exists, assume there exists a block B'_0 , mined in round 0, that is the parent of all genesis blocks, as in Figure 5, but does not affect the weight of the chains it belongs too. Thus, in this case, whenever we write for example \mathcal{C}_1 we mean \mathcal{C}_1 augmented with B'_0 . In any case, any u.s. block mined after B'_0 will be a descendant of B'_0 and $B'_0 \in \mathcal{C}_1 \cap \mathcal{C}_2$.

Now let B'_1 be the most recently mined u.s. block in some chain \mathcal{C}'_1 where (1) the last block in $\mathcal{C}_1 \cap \mathcal{C}'_1$ is in $\mathcal{C}_1 \cap \mathcal{C}_2$ and (2) for any u.s. block B' in chain \mathcal{C}' mined after B'_0 it holds that the last block in $\mathcal{C}_1 \cap \mathcal{C}'_1$ is the same or an ancestor of the last block in $\mathcal{C}_1 \cap \mathcal{C}'$. Note that $B'_1 \notin \mathcal{C}_1$, otherwise it would satisfy the definition of B'_0 which is a contradiction. Moreover, for the same reason, the last block of $\mathcal{C}_1 \cap \mathcal{C}'_1$ either is B'_0 or it is not a u.s. block. Hence, if B' is the first u.s. block mined after B'_0 that is in some chain \mathcal{C}' , it follows that $\ell_{\mathcal{C}_1 \cap \mathcal{C}'_1} \leq \ell_{\mathcal{C}'}(B') - 3 < \ell_{\mathcal{C}'}(B') - 1$. Additionally, since some honest party has chain \mathcal{C}_1 at round r , it holds that $\ell_{\mathcal{C}_1} \geq \ell_{\mathcal{C}'}(B'_1) - 1$. Thus, we can apply Claim 2 for the chains \mathcal{C}_1 and \mathcal{C}'_1 from the first u.s. block mined after B'_0 up to block B'_1 .

We apply this process as many times as possible. So B'_2 would be the most recently mined u.s. block in some chain \mathcal{C}'_2 where (1) the last block in $\mathcal{C}_1 \cap \mathcal{C}'_2$ is in $\mathcal{C}_1 \cap \mathcal{C}_2$ and (2) for any u.s. block B' in chain \mathcal{C}' mined after B'_1 it holds that the last block in $\mathcal{C}_1 \cap \mathcal{C}'_2$ is the same or an ancestor of the last block in $\mathcal{C}_1 \cap \mathcal{C}'$. Then we can again apply Claim 2 for the chains \mathcal{C}_1 and \mathcal{C}'_2 , from the first

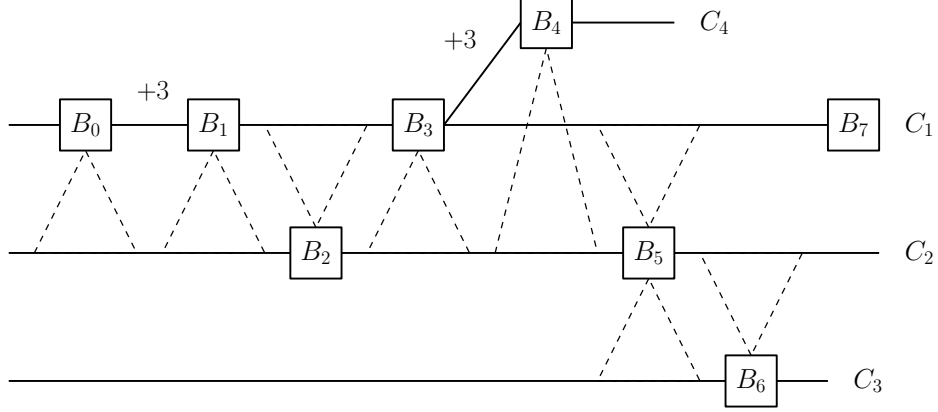


Figure 4: An example of the matching described in the claim in Lemma 28. B_0, \dots, B_7 are u.s. blocks and the dashed triangles point to the places where according to Lemma 16 there should exist an adversarial block. In this example we can use Lemma 16 two times for block B_5 . The “extra” block is matched to B_7 .

honest block mined in a u.s. round after B'_1 up to block B'_2 . We will argue that the adversarial blocks matched in the two applications of Claim 2 so far will be different. Let B' be the next u.s. block mined after B'_1 . Notice that in the worst case B'_1 has been matched to a block B'' in \mathcal{C}_1 . Also, B' will be a descendant of the (real) genesis block of \mathcal{C}_1 . Hence, $l(B') - l(B'') = 3k$ since they share the same genesis block. If $l(B') - l(B'') = 0$ then it follows that $l(B') - l(B'_1) < 2$ which is impossible. Otherwise, $l(B') - l(B'') \geq 3$, and thus the block that is going to be matched to B' by Lemma 16 cannot be B'' . This process ends when no block B'_i , for some positive i , with the desired properties exists. Notice that it may be the case that the process ends for $i = 0$, no block matching the specification of B'_1 exists.

So for any remaining u.s. block B in some chain \mathcal{C} , mined after B'_i , it holds that the last block of $\mathcal{C} \cap \mathcal{C}_1$ is not in $\mathcal{C}_1 \cap \mathcal{C}_2$. Note, that if B' is that last block mined up to round $r - 1$, and since \mathcal{C}_1 has been adopted by an honest party, it holds that $\ell_{\mathcal{C}_2} \geq \ell_{\mathcal{C}_1} \geq \ell(B') - 1$. Hence, we can apply Claim 2 for chains $\mathcal{C}_1, \mathcal{C}_2$ and for all remaining u.s. blocks. Thus, there exists a mapping between all u.s. blocks mined after B'_0 and distinct adversarial blocks that are descendants of B'_0 .

Now we have two cases. If B'_0 actually exists, all adversarial blocks of the matching must have been mined after the round B'_0 was mined. Also, since $B'_0 \in \mathcal{C}_1 \cap \mathcal{C}_2$, it must have been mined before round r' . Therefore, there exists a set of rounds $S = \{r_0, \dots, r', \dots, r\}$ such that $Z(S) \geq Y(S)$. This implies that BAD does occur, which is a contradiction.

Otherwise, the adversary may use blocks that he has precomputed. Since BAD does not hold, the honest parties have computed at least one genesis block that weighs at least $l - (1 - \delta)\delta^2(r - r')$. Then, all adversarial blocks in the matching must weigh at least $l - (1 - \delta)\delta^2(r - r') - 1$, and by Lemma 15 there are at most $5\delta/9$ such blocks. This implies that for $S = \{1, \dots, r\}$ it holds that $Y(S) \leq (1 + \frac{5\delta}{9})Z(S)$, which is a contradiction.

Thus, $\neg BAD$ implies that \mathcal{C}_1 and \mathcal{C}_2 diverge at some round greater or equal to r' . By contraposition we get that if \mathcal{C}_1 and \mathcal{C}_2 diverge at some round before r' , BAD is implied. Thus, the probability of this event is at most the probability of BAD happening, and the lemma follows.

As a final note, if P_2 has chain \mathcal{C}_2 at the beginning of round r and $head(\mathcal{C}_1)$ is a u.s. block, it may not have a corresponding adversarial block in our matching. In this case, there exist players P'_1, P'_2 that adopted the chain ending in the parent of $head(\mathcal{C}_1)$ and \mathcal{C}_2 at round $r - 1$, and these two chains diverge at round $r - r'$. Hence, again the proof holds with probability at most $e^{-\Omega(\delta^4\beta(r-r'-1))} =$




Figure 5: An example of two scenarios for the matching described by Lemma 28. Notice that in the second one an artificial block B'_0 has been introduced to aid our analysis.

$$e^{-\Omega(\delta^4 \beta(r-r'))}.$$

Finally, by an application of the Chernoff bound for $r' > r - k/(1 + \delta)f$, it follows that the probability that there exists a fork longer than k blocks is less than $1 - e^{-\Omega(\delta^4 k)}$. The lemma follows. \square

Next, using the common-prefix lemma we can prove that our protocol satisfies the strong common-prefix property.

Proof of Theorem 17. The proof is similar to that of [21]. Essentially, we proved the common-prefix lemma and now we want a proof for the common-prefix theorem in their terms. We give a brief sketch of the proof noting any differences that arise in our model. Assume that there exist chains $\mathcal{C}_1, \mathcal{C}_2$ adopted by parties P_1, P_2 at rounds r_1, r_2 such that $r_1 \leq r_2$ and $\mathcal{C}_1^{\lceil k} \not\leq \mathcal{C}_2$. Let $r \in [r_1, r_2]$ be the smallest round such that there exists an honest party P' that adopts a chain \mathcal{C}'_2 such that $\mathcal{C}'_1^{\lceil k} \not\leq \mathcal{C}_2$. If $r = r_1$, then by it holds that either $\ell_{\mathcal{C}_1} \leq \ell_{\mathcal{C}'_2}$ or $\ell_{\mathcal{C}'_2} \leq \ell_{\mathcal{C}_1}$ and by lemma 28 this event happens with negligible probability in k .

Otherwise, if $r > r_1$, let \mathcal{C}'_1 be the chain that party P' adopted at round $r - 1$. Then, it holds that

$$(\mathcal{C}'_2^{\lceil k} \preceq \mathcal{C}'_1) \wedge (\mathcal{C}'_1^{\lceil k} \preceq \mathcal{C}'_1) \wedge (\ell(\mathcal{C}'_2^{\lceil k}) \geq \ell(\mathcal{C}'_1^{\lceil k})) \implies \mathcal{C}'_1^{\lceil k} \preceq \mathcal{C}'_2^{\lceil k}$$

The first conjunct follows from the fact that \mathcal{C}'_2 was diffused at round $r - 1$ and lemma 28. The second one follows by the definition of r . The last one follows from the fact that P' selected chain \mathcal{C}'_2 despite knowing \mathcal{C}_1 , hence it must hold that $\ell_{\mathcal{C}'_2}$ is greater than $h_{\mathcal{C}_1}$ which implies the third conjunct.

The implication comes from the fact that $\mathcal{C}'_1^{[k]}$ and $\mathcal{C}'_2^{[k]}$ are part of the prefix of \mathcal{C}'_1 , and \mathcal{C}_1 is shorter of equal length to \mathcal{C}'_2 (remember that they both share the same genesis block). The theorem follows. \square

A.6 Proof of Theorem 21

Proof. Let us denote by B_i the i -th block of the chain \mathcal{C} of an honest party P at some round r and consider any k consecutive blocks B_u, \dots, B_v . Define K as the least number of consecutive blocks B_u, \dots, B_w that include the k given ones (i.e., $v \leq w$) and have the property that there exists a round at which an honest party was trying to extend the chain ending at block B_w . Observe that K is well defined since \mathcal{C} belongs to an honest party. Define also r_1 as the round that B_u was created, r_2 as the first round that an honest party attempts to extend B_w , and let $S = \{r : r_1 \leq r \leq r_2\}$.

Now let x denote the number of blocks from honest parties that are included in the k blocks and—towards a contradiction—assume that

$$x \leq \left[1 - \left(1 + \frac{\delta}{2}\right) \frac{\beta}{\gamma}\right] k \leq \left[1 - \left(1 + \frac{\delta}{2}\right) \frac{\beta}{\gamma}\right] K.$$

Let Z be the random variable that corresponds to the POWs obtained by the adversary during the rounds in S and X the successful rounds of the honest players in the same sequence of rounds.

Suppose first that all the K blocks $\{B_j : u \leq j \leq w\}$ have been computed during the rounds in the set S . Then

$$Z \geq K - x \geq \left(1 + \frac{\delta}{2}\right) \frac{\beta}{\gamma} K \geq \left(1 + \frac{\delta}{2}\right) \frac{\beta}{\gamma} \left(X - \frac{\gamma \delta k}{8f}\right)$$

The first inequality comes from the fact that the adversary computed $K - x$ of the K blocks. The second one comes from the postulated relation between x and K . To see the last inequality, assume $X - \frac{\gamma \delta k}{8f} > K$. Note first that by Lemma 20 all honest parties have chains with weight at most $\frac{3\gamma \delta k}{8f}$ less than the weight up to B_u . Next note that by Theorem 19 after round r_1 all honest parties have chains starting with the same genesis block. Thus, assuming $X - \frac{\gamma \delta k}{8f} > K$, the assumption that an honest party is on B_w at r_2 contradicts Lemma 18.

To obtain the stated bound, note that if $|S| < (1 - \delta)K/f$, then, since f is bounded away from 1 by a constant, the Chernoff bound implies that in $|S|$ rounds the total number of solutions is at least K with probability at most $e^{-\Omega(\delta^2 k)}$. Otherwise, $|S| \geq (1 - \delta)K/f \geq (1 - \delta)k/f$ and the bound follows from an application of the Chernoff bound, since $\mathbb{E}[Z] = \beta|S|$, while (using $\mathbb{E}[X] \geq \gamma|S|$, $\mathbb{E}[X] \geq (1 - \delta)\gamma k/f > \gamma k/2f$, and $(1 + \frac{\delta}{2})(1 - \frac{\delta}{4}) \geq (1 + \frac{\delta}{8})$)

$$\mathbb{E}\left[\left(1 + \frac{\delta}{2}\right) \frac{\beta}{\gamma} \left(X - \frac{\gamma \delta k}{8f}\right)\right] \geq \left(1 + \frac{\delta}{2}\right) \frac{\beta}{\gamma} \cdot \mathbb{E}\left[X - \frac{\delta}{4} \cdot \mathbb{E}[X]\right] \geq \left(1 + \frac{\delta}{8}\right) \beta |S|.$$

To finish the proof we need to consider the case in which these K blocks contain blocks that the adversary computed in rounds outside S . To manage this for a block he computed before the rounds in S implies he predicted the hash of a block in $\{B_j : u' \leq j \leq v'\}$; this occurs with probability negligible in $\log D$. If the block was computed after the rounds in S , then it was inserted between two existing blocks; this implies a collision. \square

A.7 Proof of Theorem 27

Proof. First, due to the assumption that $\lceil \log(T_1) \rceil + \lceil \log(T_2) \rceil \leq \kappa$ we prove that the events that a party succeeds in mining a POW of any of the two types in a single query are independent.

Claim 3. *The events that a party succeeds in mining a POW of any of the two types are independent.*

Proof. Let A, B be the events where a party succeeds in mining a POW of type 1 and 2 respectively. Let U be the result of the combined mining oracle query. Event A depends on the first $\lceil \log(T_1) \rceil$ bits of U , while event B depends on the last $\lceil \log(T_2) \rceil$ bits. Since each bit of U is sampled independently and uniformly at random, and $\lceil \log(T_1) \rceil + \lceil \log(T_2) \rceil \leq \kappa$ it follows that A, B do not depend on the same random bits and are thus independent. \square

Let chain $\mathcal{C}_1, \mathcal{C}_2$ be the chains that honest party P adopts at rounds r_1, r_2 when she starts the 2-for-1 POW mining phase and when she extracts the public keys respectively. We are going to define a **Bad** event that occurs with probability $1 - e^{-\Omega(\delta^5 k)}$ and do the rest of our analysis assuming this event does not occur. Let **Bad** be the disjunction of the following events:

- Less than L_1 blocks are mined in total up to round $\frac{L_1}{(1+\delta)f}$
- Less than $L_1 - 2k$ blocks are mined in total up to round $\frac{L_1 - 2k}{(1+\delta)f}$
- Less than L_2 blocks are mined in total up to round $\frac{L_2}{(1+\delta)f}$
- honest parties mine less than $X_{\text{trans}} = (r_2 - r_1 - \frac{2k}{(1-\delta)\gamma})(1 - \frac{\delta}{8})\alpha_2$ transactions from round r_1 up to round $r_2 - \frac{2k}{(1-\delta)\gamma}$
- the adversary mines more than $Z_{\text{trans}} = (r_2 - r_1 + \frac{2k}{(1-\delta)\gamma})(1 + \frac{\delta}{9})\beta_2$ transactions from round $r_1 - \frac{2k}{(1-\delta)\gamma}$ up to round r_2
- all parties have chains of length less than L_1 at round $\frac{L_1}{(1-\delta)\gamma}$
- all parties have chains of length less than L_2 at round $\frac{L_2}{(1-\delta)\gamma}$
- the common prefix property with parameter k doesn't hold for set of chains that honest parties adopt at the rounds they start the 2-for-1 mining phase
- the common prefix property with parameter k doesn't hold for set of chains that honest parties adopt at the respective extraction rounds
- the chain quality property doesn't hold for the last k blocks of $\mathcal{C}_1^{\lceil k}$
- the chain quality property doesn't hold for the last k blocks of $\mathcal{C}_2^{\lceil k}$
- the chain growth property doesn't hold for P at rounds $r_1 - \frac{2k}{(1-\delta)\gamma}$ and r_1
- the chain growth property doesn't hold for P at rounds $r_2 - \frac{2k}{(1-\delta)\gamma}$ and r_2

The first five events, happen with probability $e^{-\Omega(\delta^2 k)}$ by an application of the Chernoff bound. As we have proven in the previous sections, the rest of the events happen with probability at most $e^{-\Omega(\delta^5 k)}$. It follows that by an application of the union bound that **Bad** occurs with probability at most $e^{-\Omega(\delta^5 k)}$.

Since **Bad** does not occur we notice the following about the different phases of the protocol:

- the 2-for-1 mining phase will start after round $\frac{L_1}{(1+\delta)f}$ and before round $\frac{L_1}{(1-\delta)\gamma}$;
- all parties will receive a chain of length greater than L_2 after round $\frac{L_2}{(1+\delta)f}$ and before round $\frac{L_2}{(1-\delta)\gamma}$

We first prove that all honest parties will output the same set of public keys. Due to the common prefix property the chains of all honest parties will be the same up to block $L_2 - k$ at the round they extract the public keys. Moreover, if an honest party outputs a public key, it means that the transaction that contains this key also contains the hash of the $(L_1 - k)$ -th block of its' chain. Again, by the common-prefix property, this block will be the same for all parties. Therefore, the set of public keys output by all honest parties is exactly the same.

Next, we are going to prove that the majority of these keys have been generated by honest parties. Since $\frac{L_1 - 2k}{(1+\delta)f} \geq 14(l+1)$ it follows that the last $2k$ blocks of \mathcal{C}_1 have been mined after round

$14(l+1)$, and thus the chain quality property holds for this sequence of blocks. Hence, in the last k blocks of $\mathcal{C}_1^{\lceil k}$ there exists at least one honest block. Moreover, this honest block must have been mined after round $r_1 - \frac{2k}{(1-\delta)\gamma}$, otherwise P would have received a chain of length at least L_1 before round r_1 , which is a contradiction to the definition of r_1 . This implies that the adversary cannot have started mining transactions before this round, due to the unpredictability of the hash of the honest block.

Similarly, it holds that there exists at least one honest block in the last k blocks of $\mathcal{C}_2^{\lceil k}$ and this block must have been mined after round $r_2 - \frac{2k}{(1-\delta)\gamma}$, otherwise P would have received a chain of length L_2 earlier. Hence, there are going to be at least X_{trans} keys extracted from transactions computed by the honest parties in their output, and at most Z_{trans} by the adversary. Furthermore, it holds that:

$$r_2 - r_1 \geq \frac{L_2 - L_1}{(1+\delta)f} \geq \frac{2k}{(1-\delta)\gamma} \cdot \left(1 + \frac{10}{\delta}\right) \quad (5)$$

Hence, putting it all together:

$$\begin{aligned} X_{\text{trans}} - Z_{\text{trans}} &\geq (r_2 - r_1 - \frac{2k}{(1-\delta)\gamma})(1 - \frac{\delta}{8})\alpha_2 - (r_2 - r_1 + \frac{2k}{(1-\delta)\gamma})(1 + \frac{\delta}{9})\beta_2 \\ &\geq (r_2 - r_1)((1 - \frac{\delta}{8})\alpha_2 - (1 + \frac{\delta}{9})\beta_2) - \frac{2k}{(1-\delta)\gamma}((1 - \frac{\delta}{8})\alpha_2 + (1 + \frac{\delta}{9})\beta_2) \\ &\geq \frac{2k}{(1-\delta)\gamma} \cdot \left(1 + \frac{10}{\delta}\right)((1 - \frac{\delta}{8})\alpha_2 - (1 + \frac{\delta}{9})\beta_2) - \frac{2k}{(1-\delta)\gamma}((1 - \frac{\delta}{8})\alpha_2 + (1 + \frac{\delta}{9})\beta_2) \\ &\geq \frac{2k}{(1-\delta)\gamma} \cdot [(1 + \frac{10}{\delta})((1 - \frac{\delta}{8})\alpha_2 - (1 + \frac{\delta}{9})\beta_2) - ((1 - \frac{\delta}{8})\alpha_2 + (1 + \frac{\delta}{9})\beta_2)] \\ &\geq \frac{2k}{(1-\delta)\gamma} > 0 \end{aligned}$$

Where the forth inequality is implied by Inequality 5, and the 6th one by the fact that $\frac{10}{\delta} > \frac{2(1+\frac{\delta}{9})\beta_2}{(1-\frac{\delta}{8})\alpha_2 - (1+\frac{\delta}{9})\beta_2}$. The theorem follows.

By an application of the Chernoff bound, the number of keys is less than

$$\begin{aligned} (r_2 - r_1 + \frac{2k}{(1-\delta)\gamma}) \cdot (1+\delta)f_2 &\leq \frac{2k}{(1-\delta)\gamma} \cdot \left(\frac{(1+\delta)f}{(1-\delta)\gamma} \left(1 + \frac{10}{\delta}\right) + 1\right) \cdot (1+\delta)f_2 \\ &\leq 3 \frac{2k \cdot f_2}{\gamma} (3 \cdot 3(1 + \frac{10}{\delta}) + 1) \leq 60k \frac{f_2}{\gamma} (1 + \frac{10}{\delta}) \end{aligned}$$

and greater than

$$(r_2 - r_1 + \frac{2k}{(1-\delta)\gamma}) \cdot (1-\delta)\alpha_2 \geq \frac{2k}{(1-\delta)\gamma} \cdot \frac{10}{\delta} \cdot (1-\delta)\alpha_2 \geq k \frac{\alpha_2}{\gamma} \frac{20}{\delta}$$

□