

The Bitcoin Backbone Protocol: Analysis and Applications

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

Aggelos Kiayias*
University of Athens
aggelos@di.uoa.gr

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

October 1, 2014

Abstract

Bitcoin is the first and most popular decentralized cryptocurrency to date. In this work, we extract and analyze the core of the Bitcoin protocol, which we term the Bitcoin *backbone*, and prove two of its fundamental properties which we call *common prefix* and *chain quality*. Our proofs hinge on appropriate and novel assumptions on the “hashing power” of the adversary relative to network synchronicity; our results are shown to be tight under high synchronization.

Next, we propose and analyze applications that can be built “on top” of the backbone protocol, specifically focusing on Byzantine agreement (BA) and on the notion of a public transaction ledger. Regarding BA, we observe that Nakamoto’s suggestion falls short of solving it, and present a simple alternative which works assuming that the adversary’s hashing power is bounded by $1/3$. The public transaction ledger captures the essence of Bitcoin’s operation as a cryptocurrency, in the sense that it guarantees the “liveness” and “persistence” of committed transactions. Based on this notion we describe and analyze the Bitcoin system as well as a more elaborate BA protocol, proving them secure assuming high network synchronicity and that the adversary’s hashing power is strictly less than $1/2$, while the adversarial bound needed for security decreases as the network desynchronizes.

1 Introduction

Bitcoin, introduced in [Nak08a], is a decentralized payment system that is based on maintaining a public transaction ledger in a distributed manner. The ledger is maintained by anonymous participants (“players”) called *miners*, executing a protocol that maintains and extends a distributed data structure called the *blockchain*. The protocol requires from miners to solve a “proof of work” (POW, aka “cryptographic puzzle” — see, e.g., [DN92, RSW96, Bac97, JB99]), which essentially amounts to brute-forcing a hash inequality based on SHA-256, in order to generate new blocks for the blockchain. The blocks that comprise the blockchain contain sets of transactions that are generated at will by owners of bitcoins, who issue transactions that credit any entity of their choice who accepts payments in bitcoin. Payers broadcast transactions and miners include the transactions they receive into the blocks they generate. Miners are rewarded for maintaining the blockchain by receiving bitcoins; it is in this manner bitcoins are created and distributed among the miners who are the first recipients of newly minted bitcoins.

An important concern in Bitcoin (or any e-payment system for that matter) is the prevention of *double-spending* attacks. Specifically, in the context of Bitcoin, a double-spending attack can occur

*Research partly supported by ERC project CODAMODA.

when the attacker initially credits an account, receives service or goods by the account holder, but then manages to reorganize the transaction ledger so that the transaction that credits the account holder is reverted. In this way, the attacker keeps her bitcoin while receiving services and thus she is able to spend it again somewhere else.

In [Nak08a], Nakamoto provides an initial set of arguments of why the Bitcoin system will prevent double-spending attacks. Specifically, he argues that if a payee waits for the transaction that gives her credit to advance into the blockchain a number of k blocks, then the probability that an attacker can build an alternative blockchain that “reorganizes” the public blockchain (which contains the credit transaction) drops exponentially with k . Nakamoto argues this by modeling the attacker and the set of honest players as two competing actors performing a random walk moving toward a single direction with probabilistic steps. He demonstrates that the k blocks the payee waits are enough to ensure a negligible (in k) probability of the attacker catching up with the honest players.

Nevertheless, the above analysis can be easily seen to be oversimplified: in particular, it does not account for the fact that in Bitcoin’s decentralized setting the attacker may attempt to introduce disagreement between the honest miners, thus splitting their hashing power on different POW instances. Nakamoto himself appeared to recognize the relevance of agreement in the context of Bitcoin, arguing in a forum post [Nak08b] that actually “Bitcoin’s basic concept” of building and exchanging a blockchain is capable of solving Byzantine agreement (BA) [PSL80, LSP82] in the presence of an actively malicious adversary.¹ However a thorough analysis establishing the exact security properties of the Bitcoin system has yet to appear.

Our results. In this paper we extract, formally describe, and analyze the core of the Bitcoin protocol. We call this protocol the *Bitcoin backbone*, as we describe it in a way that is versatile and extensible and can be used to solve other problems as well — not just the problem of maintaining a public transaction ledger. The Bitcoin backbone protocol is executed by players that build a blockchain following the Bitcoin source code [Nak09] and allows a set of players to maintain a blockchain in a distributed fashion. The protocol is parameterized by three external functions $V(\cdot)$, $I(\cdot)$, $R(\cdot)$ which we call the *input validation predicate*, the *input contribution function*, and the *chain reading function*, respectively. At a high level, $V(\cdot)$ determines the proper structure of the information that is stored into the blockchain, $I(\cdot)$ specifies how the contents of the blocks are formed by the players, and $R(\cdot)$ determines how a blockchain is supposed to be interpreted in the context of the application. Note that the structure, contents, and interpretation of the blockchain are not important for the description of the backbone protocol and are left to be specified by the three external functions above, which are application-specific (we provide examples of these functions in Section 5).

We analyze the Bitcoin backbone protocol when the players operate in a synchronous communication network (more details below and in Section 2) in the presence of an adversary that controls a subset of the players. We assume that the protocol is executed by a fixed number n of players; note, however, that this number is not necessarily known to the protocol participants. The players themselves *cannot* authenticate each other and therefore there is no way to know the source of a message; we capture this by allowing the adversary to “spoof” the source address of any message that is delivered. We assume that messages are eventually delivered and all parties in the network are able to synchronize in the course of a “round.” The notion of round is not important for the

¹In [Nak08b] Nakamoto refers to the problem as “Byzantine Generals,” which is often used to refer to the single-source version of the problem, while in fact he is referring to the case where every party has an input value (Byzantine agreement). In the cryptographic setting, the problems are not equivalent in terms of the number of tolerated misbehaving parties t ($t < n$ vs. $t < n/2$, respectively).

description of the backbone protocol (which can also be executed in a loose and asynchronous fashion in the same way that Bitcoin works), however, it is important in terms of Bitcoin’s inherent computational assumption regarding the players’ ability to produce POWs.

Specifically, we assume that in a single round, all parties involved are allowed the same number of queries to a cryptographic hash function, as well as to communicate with the other participants. The hash function is modeled as a random oracle [BR93]. For simplicity we assume a “flat model,” where all parties have the same quota of hashing queries per round, say q ; the non-flat model where parties have differing hashing power capabilities can be easily captured by clustering the flat-model parties into larger virtual entities that are comprised by more than one flat-model player. In fact “mining pools” in Bitcoin can be thought of such aggregations of flat-model players. The adversary itself represents such pool as it controls $t < n$ players; for this reason, the adversary’s quota per round is $t \cdot q$ hashing queries. Note that in this setting, the fact $t < n/2$ directly corresponds to the adversary controlling strictly less than half of the system’s total “hashing power” that all players collectively harness, thus, we will use terms such as “honest majority” and “(1/2)-bounded adversary” interchangeably.

In our analysis of the Bitcoin backbone protocol we formalize and prove two fundamental properties it possesses. The properties are quantified by three parameters γ , β and f ; γ and β roughly correspond to the collective hashing power per round of the honest players and the adversary, respectively, while f represents the expected number of POWs that may be found per round by the Bitcoin network participants as a whole.

- The *common prefix property*. We prove that if $\gamma > \lambda\beta$ for some $\lambda \in [1, \infty)$ that satisfies $\lambda^2 - f\lambda + 1 \geq 0$, then the blockchains maintained by the honest players will possess a large common prefix. More specifically, if two honest parties “prune” (i.e., cut off) k blocks from the end of their local chains, the probability that the resulting pruned chains will not be mutual prefixes of each other drops exponentially in k (see Definition 2 for the precise formulation). Provided that f is very close to 0 this enables us to choose λ very close to 1 and thus establish the common prefix property as long as an honest majority of participants in the flat-model setting is guaranteed (equivalently, when the adversary controls strictly less than 50% of the hashing power). On the other hand, when the network “desynchronizes” and f gets closer to 1, achieving a common prefix requires $\lambda \rightarrow \phi$, where ϕ is the golden ratio, which in turn suggests much stricter bounds on the adversarial behavior (in fact, the upper bound on the adversary for our analysis approaches 0).
- The *chain-quality property*. We prove that if $\gamma > \lambda\beta$, for some $\lambda \in [1, \infty)$, then the ratio of blocks in the chain of any honest player that are contributed by honest players is at least $(1 - \frac{1}{\lambda})$. Again observe that if λ is close to 1, we obtain that the blockchain maintained by honest players is guaranteed to have few, but still some, blocks contributed by honest players; a higher λ would be necessary to guarantee bigger percentages of blocks contributed by honest players in the blockchain. We also show that this result is basically tight, i.e., that the adversary is capable of following a strategy (that deviates from the strategy of honest players) that enables the introduction of that many blocks in the blockchain, under a favorable assumption on the propagation of adversarial blocks in the network.

While the above two security properties may seem rather abstract since they refer to properties of the data structure that is maintained distributively by the parties, we demonstrate that they are in fact quite powerful and show that the Bitcoin backbone protocol armed with the above properties can be used as a basis for solving other problems, including the problem of distributively maintaining a “robust” public transaction ledger. In Figure 1 we show how the two properties imply the properties of the applications that are explained below.




Figure 1: An overview of the backbone protocol’s applications: Nakamoto’s BA protocol $\Pi_{\text{BA}}^{\text{nak}}$, our BA protocols $\Pi_{\text{BA}}^{1/3}$ and $\Pi_{\text{BA}}^{1/2}$, and the public ledger protocol Π_{PL} . All properties must be satisfied with overwhelming probability. In each box we state the name of the property as well as the maximum ratio of the adversarial hashing power that we can prove the protocol withstands (based on the corresponding backbone property). The value ϵ stands for a negligible quantity.

Byzantine agreement for (1/3)-bounded adversaries. As a first application, we show how a randomized BA protocol can be built on top of the Bitcoin backbone protocol more or less directly, and based solely on the POW assumption. We instantiate the $V(\cdot), I(\cdot), R(\cdot)$ functions so that parties form blockchains and act according to the following rules: each party i attempts to insert its own input $v_i \in \{0, 1\}$ into the blockchain; a blockchain is valid only if blocks contain elements in $\{0, 1\}$; the protocol terminates when the blockchain has reached a sufficient length; and, the blockchain is read by the honest parties by pruning k elements from its end and returning the majority bit appearing in the resulting blockchain’s prefix. We show how the common prefix property and the chain-quality property of the backbone protocol ensure Agreement and Validity (BA’s basic properties; see Section 2) with high probability, thus turning the Bitcoin backbone protocol into a probabilistic BA protocol.

Observe that for the above protocol to work the chain-quality property should ensure that a majority of blocks in the blockchain originate from the honest players (otherwise Validity is lost). Our chain quality property enables this with overwhelming probability assuming the adversarial power is bounded by $1/3$. This approach is different from Nakamoto’s proposal [Nak08b] for BA, which, as we also show, only guarantees Validity with overwhelming probability if the adversary has a negligible amount of hashing power. On the positive side, we stress that Nakamoto’s protocol fails gracefully when the adversarial power gets close to 50% as Validity can be shown with constant probability (but not overwhelming).

Public transaction ledgers and BA for honest majority. Next, we focus on how a “robust public transaction ledger” can be built on top of the Bitcoin backbone. We instantiate the $V(\cdot), I(\cdot), R(\cdot)$ functions so that parties form blockchains and act according to the following rules: each party (which in this context is called a “miner”) receives a set S of transactions on its input tape and attempts to insert those in its blockchain, omitting any transactions in S that are already included in it. (A Bitcoin transaction is, for example, a statement of the type “account A credits account B a z number of bitcoins,” which is signed using the secret key that corresponds to account A ’s Bitcoin address; each account has a unique Bitcoin address.) Reading a blockchain, on the other hand, amounts to returning the total sequence of transactions that is contained in the blockchain

of the miner (and note that miners may disagree about the chain they report).

We show how the common prefix property and the chain-quality property ensure two properties needed by the ledger, which we call *Persistence* and *Liveness*, assuming an honest majority and arbitrary adversarial behavior. Persistence states that once a transaction goes more than k blocks “deep” into the blockchain of one honest player, then it will be included in *every honest player’s* blockchain with overwhelming probability, and it will be assigned a permanent position in the ledger. On the other hand, Liveness says that all transactions originating from honest account holders will eventually end up at a depth more than k blocks in an honest player’s blockchain, and hence the adversary cannot perform a selective denial of service attack against honest account holders. For both properties to hold we require an honest majority (i.e., that the adversary’s hashing power is strictly less than 50%) assuming high network synchronicity (i.e., that the expected number of POW solutions per round $f \rightarrow 0$). If this is violated, Persistence requires stricter bounds on adversarial hashing power in order to be preserved following the bounds of the common prefix property. Note that for the liveness property, unforgeability of digital signatures is also a necessary assumption.

In the context of Bitcoin, our analysis implies that Bitcoin provides an operational transaction ledger under the assumptions: (i) the adversary controls less than half of the total hashing power, and (ii) the network synchronizes much faster relative to the POW solution rate, (iii) digital signatures cannot be forged. On the other hand, when the network desynchronizes our results cannot support that the ledger is maintained by assuming an honest majority. This negative result is consistent with the experimental analysis provided by Decker and Wattenhoffer [DW13], who predicted a drop below 50% in the required adversarial bound for any setting when information propagation is problematic. Our result also provides some justification for the “slow” rate of 10-minute increments used in Bitcoin block generation. Specifically, information propagation in the Bitcoin network is on the order of seconds² so the ratio (essentially f) of this time window over the average 10-minute period is reasonably close to “small” and thus transaction persistence can be shown for roughly an honest majority. On the other hand, cryptocurrencies including Litecoin, Primecoin and others, reacting to the demand to offer faster transaction processing times, opted for a much faster response rate (some as little as 1 minute), which results in more precarious situations, e.g., $f > 0.1$, which is far from being “negligible” and thus cannot support our analysis that a common prefix would be guaranteed by merely assuming an honest majority. We finally note that the Persistence and Liveness properties we put forth and prove should not be interpreted as proofs that Bitcoin’s objectives are all met. In particular, they do not guarantee that miners are properly incentivized to carry out the backbone protocol, and hence they can only offer guarantees in a setting of an *honest majority* as opposed to a setting where all parties act rationally; see Related work below as well as our summary section for further discussion.

Finally, we present a BA protocol assuming an honest majority, by suitably exploiting the properties of the robust transaction ledger above. The protocol substitutes Bitcoin’s transactions with a type of transactions that are themselves based on POWs, and hence uses POWs in two distinct ways: for the maintenance of the ledger and for the generation of the transactions. We show that the ledger’s Persistence implies Agreement, and that Liveness implies Validity, because assuming the ledger is maintained for long enough, a majority of transactions originating from the honest parties will be included (despite the fact that honest parties may control a minority of blocks in the blockchain). The protocol requires special care in the way it employs POWs since the adversary should be incapable of “shifting” work between the two POW tasks that it faces in each round. To solve this problem, we introduce a special strategy for POW-based protocol composition which we call “2-for-1 POWs.”

²See e.g., <http://bitcoinstats.com/network/propagation/>.

Related work. Realizing a digital currency with a centralized entity but while achieving strong privacy was proposed early on by Chaum in [Cha82]. A number of other works improved various aspects of this concept, however the approach remained centralized. Nakamoto [Nak08a] proposed the first decentralized currency system based on POWs while relaxing the anonymity property of the payment system to mere pseudonymity. This work was followed by a multitude of other related proposals including Litecoin³, Primecoin [Kin13], and Zerocash [BSCG⁺14], to mention a few. Our analysis of the Bitcoin backbone covers all these works as well, since they are based on exactly the same protocol.

It is interesting to juxtapose our positive results to the results of Eyal and Sirer [ES14], who introduce an attack strategy called “selfish mining” that shows how the number of blocks contributed to the blockchain by an adversary can exceed the percentage of the hashing power the adversary possesses. Their results are consistent and complementary to ours. The crux of the issue is (in our terminology) in terms of the chain-quality property, as its formulation is quite permissive: in particular we show that if the adversary controls a suitably bounded amount of hashing power, then it is also suitably bounded in terms of the number of blocks it has managed to insert in the blockchain that honest players maintain. Specifically, recall that we prove that if the hashing power of the adversary satisfies $\beta < \frac{1}{\lambda}\gamma$ (where γ roughly corresponds to the hashing power of the honest players), then the adversary may control at most a $\frac{1}{\lambda}$ percentage of the blocks in the chain. For instance, if the adversary controls up to 1/3 of the hashing power (i.e., $\lambda = 2$), then it will provably control less than 50% of the blocks in the honest players’ blockchain. As it can be easily seen, this does not guarantee that the rate of a party’s hashing power translates to an equal rate of rewards (recall that in Bitcoin the rewards are linearly proportional to the number of blocks that a party contributes in the chain). We define as *ideal chain quality* the property that for any coalition of parties (following any mining strategy) the percentage of blocks in the blockchain is exactly proportional to their collective hashing power. The chain quality property that we prove is not ideal and the results of [ES14] show that in fact there is a strategy that magnifies the percentage of a malicious coalition. Still, their mining attack does much worse than our bound. To close the gap, we sketch (cf. Remark 3) a simple selfish mining strategy that matches our upper bound and hence our chain quality result is tight in our model⁴ assuming the number of honest parties is large.

Byzantine agreement (BA, aka distributed consensus) [PSL80, LSP82] considers a set of n parties connected by reliable and authenticated pair-wise communication links and with possible conflicting initial inputs that wish to agree on a common output in the presence of the disruptive (even malicious) behavior of some of them. The problem has received a considerable amount of attention under various models. In this paper we are interested in randomized solutions to the problem (e.g., [BO83, Rab83, BG93, FM97, FG03, KK09])⁵ as in the particular setting we are in, deterministic BA algorithms are not possible. In more detail, we consider BA in the *anonymous synchronous setting*, i.e., when processors do not have identifiers and cannot correlate messages to their sources, even across rounds. This model for BA was considered by Okun, who classified it as “anonymous model without port awareness,” and proved the aforementioned impossibility result, that deterministic algorithms are impossible for even a single failure [Oku05b, Oku05a]. In addition, Okun showed that probabilistic BA is feasible by suitably adapting Ben-Or’s protocol [BO83] for the standard,

³<http://www.litecoin.com>.

⁴Our model allows the unfavorable event of adversarial messages winning all head-to-head races in terms of delivery with honestly generated messages in any given round.

⁵We remark that, in contrast to the approach used in typical randomized solutions to the problem, where achieving BA is reduced to (the construction of) a shared random coin, the probabilistic aspect here stems from the parties’ likelihood of being able to provide proofs of work. In addition, as our analysis relies on the random oracle model [BR93], we are interested in computational/cryptographic solutions to the problem.

non-anonymous setting (cf. [Oku05b])⁶; the protocol, however, takes exponentially many rounds. It turns out that by additionally assuming that the parties are “port-aware” (i.e., they can correlate messages to sources across rounds), deterministic protocols are possible and some more efficient solutions were proposed in [OB08].

The anonymous synchronous setting was also considered by Aspnes *et al.* [AJK05] who pointed to the potential usefulness of proofs of work (e.g., [DN92, RSW96, Bac97, JB99]) as an identity assignment tool, 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 [AJK05] 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 that a standard authenticated BA protocol could be run. Such protocols, however, would require the establishment of a consistent PKI (as well as of digital signatures), details of which are not laid out in [AJK05].

In contrast, and as mentioned above, building on our analysis of the Bitcoin backbone protocol, we propose two BA protocols solely based on POWs that operate in $O(k)$ rounds with error probability $e^{-\Omega(k)}$. The protocols solve BA with overwhelming probability under the assumption that the adversary controls less than $1/3$ and $1/2$ of the computational power, respectively.

The connection between Bitcoin and probabilistic BA was also considered by Miller and LaViola in [ML14] where they take a different approach compared to ours, by not formalizing how Bitcoin works, but rather only focusing on Nakamoto’s suggestion for BA [Nak08b] as a standalone protocol. As we observe here, and also recognized in [ML14], Nakamoto’s protocol does not quite solve BA since it does not satisfy Validity with overwhelming probability. The exact repercussions of this fact are left open in [ML14], while with our analysis, we provide explicit answers regarding the transaction ledger’s actual properties and the level of security that the backbone realization can offer.

Finally, related to the anonymous setting, the feasibility of secure computation without authenticated links was considered by Barak *et al.* in [BCL⁺11] in a more extreme model where all messages sent by the parties are controlled by the adversary and can be tampered with and modified (i.e., not only source addresses can be “spoofed,” but also messages’ contents can be altered and messages may not be delivered). It is shown in [BCL⁺11] that it is possible to limit the adversary so that all he can do is to partition the network into disjoint sets, where in each set the computation is secure, and also independent of the computation in the other sets. Evidently, in such model one cannot hope to build a global transaction ledger.

Organization of the paper. The rest of the paper is organized as follows. In Section 2 we present our model within which we formally express the Bitcoin backbone protocol and prove its basic properties. The backbone protocol builds “blockchains” based on a cryptographic hash function; we introduce notation for this data structure as well as the backbone protocol itself in Section 3, followed by its analysis in Section 4. Sections 5 and 6 are dedicated to the applications built on top of the backbone protocol — (simple) BA protocols and robust transaction ledger, respectively. Specifically, Section 5 covers Nakamoto’s (insufficient) suggestion for BA as well as our solution for $1/3$ adversarial power, while in Section 6 we present our treatment of a robust public ledger formalizing the properties of Persistence and Liveness and how they apply to Bitcoin. Finally, we also include in this section our BA protocol for $1/2$ adversarial power. Some directions for future research are offered in Section 7.

⁶Hence, BA in this setting shares a similar profile with BA in the asynchronous setting [FLP85].

2 Model and Definitions

We describe our protocols in a standard multiparty synchronous communication setting (e.g., Canetti’s formulation of “real world” execution [Can00]) with the relaxation that the underlying communication graph is not fully connected and messages are delivered through a “diffusion” mechanism that reflects Bitcoin’s peer-to-peer structure. Our adversarial model in the network is “adaptive”, meaning that the adversary is allowed to take control of parties on the fly, and “rushing”, meaning that in any given round the adversary gets to see all honest players’ messages before deciding his strategy, and, furthermore, also allows the adversary to change the source information on every message. Note that the adversary cannot change the contents of the messages nor prevent them from being delivered. Effectively, this parallels communication over TCP/IP in the Internet where messages between parties are delivered reliably, but nevertheless malicious parties may “spoof” the source of a message they transmit and make it appear as originating from an arbitrary party (including another honest party) in the view of the receiver. This aspect of the communication model, where processors cannot correlate messages to their sources, even across rounds, was considered by Okun [Oku05a], who classified it as “anonymous model without port awareness.” In this setting we use BROADCAST as the message transmission command that captures the “send-to-all” functionality allowed by our communication model. Note that an adversarial sender may abuse BROADCAST and attempt to confuse honest parties by sending and delivering inconsistent messages to them.

The parties’ inputs are provided by the environment \mathcal{Z} which also receives the parties’ outputs. Parties that receive no input from the environment remain inactive, in the sense that they will not act when their turn comes in each round. The environment may provide input to a party at any round and may also modify that input from round to round. We denote by $\text{INPUT}()$ the input tape of each party.

In each round, parties are able to read their input tape $\text{INPUT}()$ and communication tape $\text{RECEIVE}()$, perform some computation that will be suitably restricted (see below) and issue⁷ a BROADCAST message that is guaranteed to be delivered to all parties in the beginning of the next round. As stated above the adversary can do multiple broadcasts per round and in fact deliver to each honest party a different message or even multiple messages.

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. If n parties P_1, \dots, P_n execute Π , the concatenation of the view of all parties $\langle \text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{P_i} \rangle_{i=1, \dots, n}$ is denoted by $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}$. With foresight, we note that, in contrast to the standard setting where parties are aware of the number of parties executing the protocol, 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 order to capture the parties’ limited ability to produce POWs, we assume that all parties may have access to an oracle $H(\cdot)$ and allowed to perform a number of queries q per round, where q is a function of the security parameter κ ; we refer to such parties as q -bounded. Note that this is a “flat-model” interpretation of the parties’ computation power, where all parties are assumed equal. In the real world, different honest parties may have different “hashing power;” nevertheless, our flat-model does not sacrifice generality since one can imagine that real honest parties are simply clusters

⁷For simplicity, we assume that the broadcast operation is atomic and hence the corruption of a party may not happen while the operation is taking place (cf. [HZ10, GKKZ11]).

of some arbitrary number of honest flat-model parties. The adversary \mathcal{A} is allowed to perform $t \cdot q$ queries per round, where $t \leq n$ is the number of corrupted parties. The environment \mathcal{Z} , on the other hand, is not permitted any queries to $H(\cdot)$. The rationale for this is that we would like to bound the “CPU power” [Nak08a] of the adversary to be proportionate to the number of parties it controls while making it infeasible for them to be aided by external sources or by transferring the hashing power potentially invested in concurrent or previous protocol executions. It follows that in our analysis we will focus on the “standalone” setting, where a single protocol instance is executed in isolation.

We refer to the above restrictions on the environment, the parties and the adversary as the *q-bounded synchronous setting*. The view of the parties participating in the protocol will be denoted by $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{P, H(\cdot)}(\kappa, q, z)$ and the concatenation of all parties’ views by $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{H(\cdot)}(\kappa, q, z)$.

In our theorems we will be concerned with *properties* of protocols Π in the *q-bounded synchronous setting*. Such properties will be defined as predicates over the random variable $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{H(\cdot)}(\kappa, q, z)$ by quantifying over all possible adversaries \mathcal{A} and environments \mathcal{Z} . 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 can be freely selected in $\{1, \dots, \kappa\}$. The probability space is determined by the oracle $H(\cdot)$ as well as any random choices made by the protocol itself (if any). Further details about the model are given in Appendix A.

Byzantine agreement. As a simple illustration of the formulation above we define the properties of a Byzantine agreement (BA) protocol.

Definition 1. A protocol Π solves BA in the *q-bounded synchronous setting* provided it satisfies the following two properties:

- *Agreement.* There is a round after which all honest parties return the same output if queried by the environment.
- *Validity.* The output returned by an honest party P equals the input of some party P' that is honest at the round P ’s output is produced.

We note that in our protocols, the participants are capable of detecting agreement and furthermore they can also detect whether other parties detect agreement, thus *termination* can be easily achieved by all honest parties. The formulation of Validity above is intended to capture security/correctness against adaptive adversaries. The notion (specifically, the requirement that the output value be one of the honest parties’ inputs) has also been called “Strong Validity” [Nei94], but the distinction is only important in the case of non-binary inputs. In either case, it is known that in the synchronous cryptographic setting the problem has a solution if and only if $n > |V|t$, where V is the input/decision domain [FG03]. Our POW-based protocols work for both versions of the problem.

3 The Bitcoin Backbone Protocol

We start by introducing blockchain notation. Let $G(\cdot), H(\cdot)$ be cryptographic hash functions with output in $\{0, 1\}^\kappa$. 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

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

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 $\text{head}(\mathcal{C}) = \langle s', x', r' \rangle$ can be extended to a longer chain by appending a valid block $B = \langle s, x, r \rangle$ that satisfies $s = H(r', G(s', x'))$. In case $\mathcal{C} = \varepsilon$, by convention any valid block of the form $\langle s, x, r \rangle$ may extend it. In either case we have an extended chain $\mathcal{C}_{\text{new}} = \mathcal{C}B$ that satisfies $\text{head}(\mathcal{C}_{\text{new}}) = B$.

The *length* of a chain $\text{len}(\mathcal{C})$ is its number of blocks. 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” 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 $\mathcal{C}_1 \preceq \mathcal{C}_2$.

We note that Bitcoin uses chains of variable difficulty, i.e., the value D may change across different blocks within the same chain according to some rule that uses the x values stored in the chain. This does not affect our analysis, and hence we will only consider chains of fixed difficulty D (i.e., we do not consider the dynamic setting). We are now ready to describe the protocol.

3.1 The backbone protocol

The Bitcoin backbone protocol is executed by an arbitrary number of parties over an unauthenticated network. 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. As mentioned in Section 2, communication over the network is achieved by utilizing a send-to-all BROADCAST functionality that is available to all parties (and maybe abused by the adversary in the sense of delivering different messages to different parties). Each party maintains 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. (Figure 2 depicts the local view of each party as well as the shared portion of their chains.)

In the protocol description we intentionally avoid 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. These functions are handled 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 Bitcoin backbone protocol is specified as Algorithm 4. Before describing it in detail we first introduce the protocol’s three supporting algorithms.

Chain validation. The first algorithm, called `validate` performs a validation of the structural properties of a given chain \mathcal{C} . It is given as input the values q and D , as well as a hash function $H(\cdot)$. It is parameterized by a predicate $V(\cdot)$, called the *input validation predicate*. For each block of the chain, the algorithm checks that the proof of work is properly solved and that the counter ctr does not exceed q . It further collects all the inputs from the chain’s blocks and assembles them into a vector $\mathbf{x}_{\mathcal{C}}$. If all blocks verify and $V(\mathbf{x}_{\mathcal{C}})$ is true then the chain is deemed valid; otherwise it is rejected. Note that we purposely leave the predicate $V(\cdot)$ undetermined.




Figure 2: Overview of the basic operation of the Bitcoin backbone protocol. Miner M_1 receives from the environment a READ instruction that results in the application of the $R(\cdot)$ function on the contents of its chain which are equal to the vector $\langle x_1, x_2, x_3, x_4, x_5 \rangle$. Miner M_2 receives from the environment an INSERT instruction and uses the function $I(\cdot)$ to determine the value y_5 that it subsequently successfully inserts in its local block chain by solving a proof of work; this results in a broadcast of the newly extended chain. Finally miner M_3 receives the newly extended chain and validates it both structurally as well as using the input validation predicate $V(\cdot)$. M_3 will adopt this chain if M_3 deems it better than its local chain as specified by the backbone protocol. Note that the joint view of M_1, M_2, M_3 is inconsistent but there is agreement on the prefix $\langle x_1, x_2, x_3 \rangle$.

Algorithm 1 The chain validation predicate, parameterized by q, D , the hash function $H(\cdot)$, and the input validation predicate $V(\cdot)$. The input is chain \mathcal{C} .

```

1: function validate( $\mathcal{C}$ )
2:    $b \leftarrow \text{True}$ 
3:    $\mathbf{x}_{\mathcal{C}} \leftarrow \langle \varepsilon \rangle$ 
4:   while  $\mathcal{C} \neq \varepsilon$  do
5:      $\langle s, x, ctr \rangle \leftarrow \text{head}(\mathcal{C})$ 
6:     if  $\text{validblock}_q^D(\langle s, x, ctr \rangle)$  then
7:        $\mathcal{C} \leftarrow \mathcal{C}^{\uparrow 1}$ 
8:        $\mathbf{x}_{\mathcal{C}} \leftarrow \text{append}(\mathbf{x}_{\mathcal{C}}, x)$ 
9:     else
10:       $b \leftarrow \text{False}$ 
11:    end if
12:  end while
13:  return  $(b \wedge V(\mathbf{x}_{\mathcal{C}}))$ 
14: end function

```

\triangleright Remove the head from \mathcal{C}
 \triangleright extend $\mathbf{x}_{\mathcal{C}}$ with x

Chain comparison. The objective of the second algorithm, called `maxvalid`, is to find the “best possible” chain when given a set of chains. The algorithm is straightforward and is parameterized by a `max(·)` function that applies some ordering in the space of chains. The most important aspect is the chains’ length, in which case `max(C1, C2)` will return the longest of the two. In case `len(C1) = len(C2)`, some other characteristic can be used to break the tie. In our case, `max(·, ·)` will always return the first operand⁸; alternatively, other options exist, such as lexicographic order or picking a chain at random. The analysis we will perform will essentially be independent of the tie-breaking rule⁹.

Algorithm 2 The function that finds the “best” chain, parameterized by function `max(·)`. The input is $\{\mathcal{C}_1, \dots, \mathcal{C}_k\}$.

```

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

```

Proof of work. The third algorithm, called `pow`, is the main “workhorse” of the backbone protocol. It takes as input a chain and attempts to extend it via solving a proof of work. This algorithm is parameterized by two hash functions $H(\cdot), G(\cdot)$ (which in our analysis will be modeled as random oracles)¹⁰, as well as two positive integers q, D ; q represents the number of times the algorithm is going to attempt to brute-force the hash function inequality that determines the POW instance, and D determines the “difficulty” of the POW. The algorithm works as follows. Given a chain \mathcal{C} and a value x to be inserted in the chain, it hashes these values to obtain h and initializes a counter ctr . Subsequently, it increments ctr and checks to see whether $H(ctr, h) \leq D$; if a suitable ctr is found then the algorithm succeeds in solving the POW and extends chain \mathcal{C} by one block inserting x as well as ctr (which serves as the POW). If no suitable ctr is found, the algorithm simply returns the chain unaltered.

⁸Note that the way we deploy `maxvalid`, amounts to parties always giving preference to their local chain as opposed to any incoming chain. This is consistent with current Bitcoin operation; however, some debate about alternate tie-breaking rules has ensued in Bitcoin forums, e.g., see [Cun13].

⁹It is worth to point out that the behavior of `maxvalid(·)` is associated with some stability aspects of the backbone protocol and currently there are proposals to modify it (e.g., by randomizing it — cf. [ES14]). It is an interesting question whether any improvement in our results can be achieved by randomizing the `maxvalid` operation.

¹⁰In reality the same hash function (SHA-256) instantiates both G and H ; however, it is notationally more convenient to consider them as distinct.

Algorithm 3 The *proof of work* function, parameterized by q , D and hash functions $H(\cdot), G(\cdot)$. The input is (x, \mathcal{C}) .

```

1: function pow( $x, \mathcal{C}$ )
2:    $ctr \leftarrow 1$ 
3:   if  $\mathcal{C} = \varepsilon$  then                                     ▷ Determine proof of work instance.
4:      $\tilde{s} \leftarrow 0$ 
5:   else
6:      $\langle \tilde{s}, \tilde{x}, \tilde{r} \rangle \leftarrow \text{head}(\mathcal{C})$ 
7:   end if
8:    $h \leftarrow G(\tilde{s}, x)$ 
9:   while ( $ctr \leq q$ ) do
10:    if ( $H(ctr, h) \leq D$ ) then                             ▷ Proof of work succeeded.
11:       $B \leftarrow \langle \tilde{s}, x, ctr \rangle$ 
12:       $\mathcal{C} \leftarrow \mathcal{C}B$ 
13:      break
14:    end if
15:     $ctr \leftarrow ctr + 1$ 
16:  end while
17:  return  $\mathcal{C}$ 
18: end function

```

The backbone protocol. Given the three algorithms above, we are now ready to describe the Bitcoin backbone protocol. This is the protocol that is executed by the miners and which is assumed to run “indefinitely” (in our security analysis will apply when the total running time is polynomial in κ). It is parameterized by two functions, the input contribution function $I(\cdot)$ and the chain reading function $R(\cdot)$, which is applied to the values stored in the chain.

Each miner maintains a local chain \mathcal{C} , attempting to extend it by invoking the POW algorithm **pow** described above. Prior to updating the chain, the miner checks its communication tape **RECEIVE()** to see whether a “better” chain has been received. This is done using the **maxvalid** function, depending on which the local chain is substituted.

The value that the miner attempts to insert in the chain is determined by function $I(\cdot)$. The input to $I(\cdot)$ is the state st , the current chain \mathcal{C} , the contents of the miner’s input tape **INPUT()** (recall that they can be written by the environment \mathcal{Z} at the beginning of any round) and communication tape **RECEIVE()**, as well as the current round number *round*. The protocol expects two types of entries in the input tape, **READ** and (**INSERT**, *value*); other inputs are ignored.

We purposely leave the functions $I(\cdot), R(\cdot)$ undetermined in the description of the backbone protocol, as their specifics will vary according to the application. One may choose, for example, $I(\cdot)$ to be as simple as copying the contents of the **INSERT** input symbols from **input()** into x and keeping $st = \epsilon$, or performing a complex operation parsing \mathcal{C} and maintaining old inputs in st . We provide explicit examples of $I(\cdot)$ and $R(\cdot)$ in Section 5. When the input x is determined, the protocol attempts to insert it into the chain \mathcal{C} by invoking **pow**. In case the local chain \mathcal{C} is modified during the above steps, the protocol transmits (“broadcasts”) the new chain to the other parties. Finally, in case a **READ** symbol is present in the communication tape, the protocol applies function $R(\cdot)$ to its current chain and writes the result onto the output tape **OUTPUT()**. This way, the round ends and a new round begins, continuing indefinitely.

Algorithm 4 The Bitcoin backbone protocol, parameterized by the *input contribution function* $I(\cdot)$ and the *chain reading function* $R(\cdot)$.

```

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

```

3.2 (Desired) Properties of the backbone protocol

We next define the two main properties of the backbone protocol that we will prove. The first property is called the *common prefix property* and is parameterized by a value $k \in \mathbb{N}$. It considers an arbitrary environment and adversary in the q -bounded setting, and it holds as long as any two honest parties' chains are different only in its most recent k blocks.

Definition 2 (Common Prefix Property). The common prefix property Q_{cp} with parameter $k \in \mathbb{N}$ states that for any pair of honest players P_1, P_2 maintaining the chains $\mathcal{C}_1, \mathcal{C}_2$ in $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{H(\cdot)}(\kappa, q, z)$, it holds that

$$\mathcal{C}_1^{\lceil k} \preceq \mathcal{C}_2 \text{ and } \mathcal{C}_2^{\lceil k} \preceq \mathcal{C}_1.$$

The second property, which we call the *chain quality property*, aims at expressing the number of honest-player contributions that are contained in a sufficiently long and continuous part of an honest player's chain. Specifically, for parameters $k \in \mathbb{N}$ and $\mu \in (0, 1)$, the rate of adversarial input contributions in a continuous part of an honest party's chain is bounded by μ . This is intended to capture that at any moment that an honest player looks at a sufficiently long part of its blockchain, that part will be of sufficient "quality," i.e., the number of adversarial blocks present in that portion of the chain will be suitably bounded.

Definition 3 (Chain Quality Property). The chain quality property Q_{cq} with parameters $\mu \in \mathbb{R}$ and $\ell \in \mathbb{N}$ states that for any honest party P with chain \mathcal{C} in $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{H(\cdot)}(\kappa, q, z)$, it holds that for any ℓ consecutive blocks of \mathcal{C} the ratio of adversarial blocks is at most μ .

It is easy to see that any set of, say, h honest parties, obtain as many blocks as their proportion of the total hashing power, i.e., h/n . We say that a protocol Π satisfies *ideal chain quality* if this is the case for adversarial parties as well, i.e., $\mu = t/n$ with respect to those parties. The ideal chain quality is not achieved by the Bitcoin backbone protocol, cf. Remark 3.

4 Analysis of the Bitcoin Backbone

We now proceed to the analysis of the protocol presented in the previous section. Let $\{0, 1\}^\kappa$ be the range of $H(\cdot)$. Each party tries to provide a POW by issuing queries to $H(\cdot)$, which succeed with probability $p = D/2^\kappa$, where D is the difficulty level. By the properties of the random oracle $H(\cdot)$, any collection of queries will be treated as a collection of independent Bernoulli trials with success probability p . In order to support this we will assume that the function $I(\cdot)$ (which determines the input of the players that is to be inserted in the blockchain) ensures (at least with overwhelming probability) that the inputs are unique. There are two simple ways to enforce this: either have $I(\cdot)$ add a sufficiently long random nonce to x , or, in case parties have unique identities, it may be parameterized by it and introduce it as part of x . In either case, this value will be ignored by the other functions $V(\cdot), R(\cdot)$ as it need not be useful in the application. It is easy to see that if a κ -long nonce is used the output will be unique except for probability at most $\bar{q}^2 \cdot 2^{-\kappa}$ where \bar{q} is the total number of queries submitted to the random oracle; we will ignore this small term in our analysis.

4.1 Definitions and preliminary lemmas

Recall that n is the number of parties, t of which can be corrupted by the adversary. We introduce the following parameters for notational convenience:

$$\alpha = pq(n-t), \quad \beta = pqt, \quad \gamma = \alpha - \alpha^2, \quad f = \alpha + \beta.$$

The first parameter, α , reflects the hashing power of the honest parties. It is an upper bound on the expected number of solutions that the honest parties compute in one round. Similarly, β , is the expected number of solutions that the corrupted parties compute in one round. Notice the asymmetry that while the honest parties will not compute more than one solution per round, a corrupted party may use all its q queries and potentially compute more than one solution. The parameter γ will serve as a lower bound on the following two probabilities. The first one is that at least one honest party computes a solution in a round:

$$1 - (1-p)^{q(n-t)} \geq 1 - e^{-\alpha} \geq \gamma;$$

we will call such round a *successful round*. The second one is the probability that exactly one honest party does so; we will call such round a *uniquely successful round*. We lower bound the probability of such a round by the probability that out of $q(n-t)$ coin tosses exactly one comes up heads. Thus, the probability is at least:

$$(n-t)qp(1-p)^{q(n-t)-1} \geq \alpha(1-\alpha+p) \geq \gamma.$$

The ratio $\alpha/\beta = (n-t)/t$ will be of interest for the analysis. When α is small (as it will be when f is small), then $\gamma \approx \alpha$ and we will be justified to concentrate on the ratio γ/β . To understand how well γ estimates the probability of a uniquely successful round, call it γ' , we observe the following upper bound:

$$\begin{aligned} \gamma' &= (n-t)(1-(1-p)^q)(1-p)^{q(n-t)-1} \leq (n-t)pqe^{-\alpha+pq} \\ &\leq \alpha(1-\alpha+pq + (\alpha-pq)^2/2) = \alpha - \alpha^2(1 - \frac{1}{n-t}) + \frac{\alpha^3}{2}(1 - \frac{1}{n-t})^2, \end{aligned}$$

where we use Facts 1 and 2 (see Appendix B). From this it follows that $\gamma' \leq \alpha - \alpha^2 + \alpha^3/2 + O(1/(n-t))$.

The following definition will be crucial in the analysis of the common-prefix property.

Definition 4 (Uniform rounds). We call a round *uniform* if, at that round, every honest party invokes the $\text{pow}(\cdot)$ algorithm with a chain of the same length (i.e., $\text{len}(\tilde{C})$ at line 7 of Algorithm 4 is the same for all honest parties).

We will call a query of a party *successful* if it submits a pair (ctr, h) such that $H(ctr, h) \leq D$. Without loss of generality, let P_1, \dots, P_t be the set of corrupted parties (knowledge of this set will not be used in any argument). For each round i , $j \in [q]$, and $k \in [t]$, we define Boolean random variables X_i and $Z_{ijk} \in \{0, 1\}$ as follows. If at round i an honest party obtains a POW, then $X_i = 1$, otherwise $X_i = 0$. Regarding the adversary, if at round i , the j -th query of the k -th corrupted party is successful, then $Z_{ijk} = 1$, otherwise $Z_{ijk} = 0$. Further, if $X_i = 1$, we call i a *successful round*. If a round is uniform (Def. 4) and uniquely successful, we say it is a *uniquely successful uniform round*.

Next, we will prove two preliminary lemmas that will be helpful in our analysis. The first one states that, at any round, the length of any honest party's chain will be at least as large as the number of successful rounds. As a consequence, the chain of honest parties will grow at least at the rate of successful rounds. The second lemma is a simple application of Chernoff bounds and states that, with high probability, the honest parties will have, at any round, at least λ as many successful rounds as the adversary has. The usefulness of this lemma will be in showing that honest parties will be building a blockchain at a rate the adversary will find it hard to overcome.

Lemma 5. *Suppose that at round r the chain of an honest party is of length ℓ . Then, after round $s \geq r$, the chain of any honest party will have length at least $\ell + \sum_{i=r}^s X_i$.*

Proof. By induction on $s - r \geq 0$. For the basis, observe that at the round at which the honest party's chain reaches length ℓ (round r or earlier), he broadcasts it. Therefore, after round r every honest party would have a chain of length at least ℓ . For the inductive step, if it is the case $X_s = 0$, then the inductive hypothesis suffices. Otherwise ($X_s = 1$), observe that by hypothesis every honest party has a chain of length at least $\ell' = \ell + \sum_{i=1}^{s-1} X_i$. Therefore, at least one honest party broadcasts at round s a chain of length at least $\ell' + 1$. Since $\ell' + 1 = \ell + \sum_{i=1}^s X_i$, this completes the proof. \square

Lemma 6. *Assume $\gamma \geq (1 + \delta)\lambda\beta$ for some $\delta \in (0, 1)$ and $\lambda \geq 1$. The probability that during s rounds the number of successful rounds exceeds by a factor $(1 + \frac{\delta}{2})\lambda$ the number of solutions computed by the adversary is at least $1 - e^{-\Omega(\delta^2 s)}$.*

Proof. Without loss of generality we assume the s rounds start at round 1. Let $X = \sum_{i=1}^s X_i$ and $Z = \sum_{i=1}^s \sum_{j \in [q]} \sum_{k \in [t]} Z_{ijk}$. By an application of Chernoff bounds (Appendix B) we obtain

$$\Pr[X \leq (1 - \frac{\delta}{4})\gamma s] \leq e^{-\Omega(\delta^2 s)} \quad \text{and} \quad \Pr[Z \geq (1 + \frac{\delta}{5})\beta s] \leq e^{-\Omega(\delta^2 s)}.$$

It follows that the union of these events has a measure exponentially small in s . However, if none of them hold, then

$$X > (1 - \frac{\delta}{4})\gamma s \geq (1 - \frac{\delta}{4})(1 + \delta)\lambda\beta s > (1 + \frac{\delta}{2})(1 + \frac{\delta}{5})\lambda\beta s > (1 + \frac{\delta}{2})\lambda Z.$$

\square

We are now ready for the treatment of the protocol's properties outlined in Section 3.2.

4.2 The common-prefix property

This property is established in Theorem 9, whose main argument is in turn given in Lemma 8. We start with a lemma leading to that argument. The lemma will be used to argue that uniform rounds favor the honest parties. Informally, the idea is that a uniquely successful uniform round forces an adversary trying to make honest parties' chains "diverge" to produce POWs. In the second lemma we take advantage of this, to show that if the adversary has appropriately bounded computational power, then there will be enough uniquely successful uniform rounds to prevent him from mounting a successful attack on the common-prefix property.

Lemma 7. *Consider a uniquely successful uniform round where the honest parties have chains of length $\ell - 1$. Then, in any subsequent round, there can be at most one chain \mathcal{C} where the ℓ -th block was contributed by an honest party.*

Proof. Let r be a uniquely-successful uniform round and \mathcal{C} , with $\text{len}(\mathcal{C}) = \ell$, be the chain computed by the party that solves the proof of work and extends its local chain of length $\ell - 1$ to ℓ . At round $r + 1$ every honest party will receive \mathcal{C} and will either adopt it or adopt another chain sent by the adversary. In any case, every honest party will have a chain of length at least ℓ , and will never query the $\text{pow}(\cdot)$ function with a chain of length $\ell - 1$ again. The statement of the lemma thus follows. \square

Note that in order for the common-prefix property to be violated at round r , at least two honest parties should have chains \mathcal{C}_1 and \mathcal{C}_2 such that $\mathcal{C}_1^{\lceil k} \not\preceq \mathcal{C}_2$ or $\mathcal{C}_2^{\lceil k} \not\preceq \mathcal{C}_1$. Therefore, the existence of many blocks computed at uniform rounds forces the adversary to provide as many blocks of its own. We need to show that, with high probability the adversary will fail to collect as many solutions by round r .

We say that two chains *diverge* at a given round, if the last block of their common prefix was computed before that round.

Our main lemma below asserts the following. Suppose the protocol is halted at round r and two honest parties have distinct chains \mathcal{C}_1 and \mathcal{C}_2 . Then, for s large enough, the probability that \mathcal{C}_1 and \mathcal{C}_2 diverge at round $r - s$ is negligible. The idea of the proof is to upper bound the number of (valid) broadcasts that the adversary can perform during these last s rounds. Note that they are in the order of βs in expectation. The crucial observation here is that if at a given round the adversary is silent, then a uniform round follows. Therefore we expect about $(1 - \beta)s$ uniform rounds, and consequently $\gamma(1 - \beta)s$ uniquely-successful uniform rounds. Recalling Lemma 7, the adversary needs to collect $\gamma(1 - \beta)s$ POWs. Thus, in the lemma's condition we choose the relation between β and γ suitably so that the adversary is incapable of accomplishing this task, except with probability exponentially decreasing in s .

Lemma 8. *Assume $f < 1$ and $\gamma \geq (1 + \delta)\lambda\beta$, for some real $\delta \in (0, 1)$ and $\lambda \geq 1$ such that $\lambda^2 - f\lambda - 1 \geq 0$. Suppose \mathcal{C}_1 and \mathcal{C}_2 are the chains of two honest parties at round r . Then, for any $s \leq r$, the probability that \mathcal{C}_1 and \mathcal{C}_2 diverge at round $r - s$ is at most $e^{-\Omega(\delta^3 s)}$.*

Proof. We define three bad events, A , B and C , which we show to hold with probability exponentially small in s . We conclude the proof by showing that if none of these bad events happens, then there cannot exist \mathcal{C}_1 and \mathcal{C}_2 diverging at round $r - s$.

The bad event A occurs if, at some round $r' \geq r - s$, the adversary broadcasts a chain \mathcal{C} with the following properties. (1) \mathcal{C} is returned by the function maxvalid of an honest party; (2) the block $\text{head}(\mathcal{C})$ was computed by the adversary before round $r - (1 + \frac{\delta}{8})s$.

We now give an upper bound on the probability that event A occurs. Let $r^* \leq r - (1 + \frac{\delta}{8})s$ be the latest round at which a block of \mathcal{C} was computed by an honest party (if none exists, then $r^* = 0$), and let ℓ denote the length of the chain up to that block. If any other block computed by an honest party exists among the blocks from length ℓ up to $\text{len}(\mathcal{C})$, then such block was computed in rounds $r - (1 + \frac{\delta}{8})s$ up to r' , and it follows that the probability that the adversary's block can extend it at round r' is negligible in $(\kappa - \log D)$. Therefore, we infer that with overwhelming probability the adversary has computed all the blocks from length ℓ to $\text{len}(\mathcal{C})$, and done so during the rounds r^* to r' . Let Z denote the total number of solutions the adversary obtained in $r' - r^*$ rounds. Let also X denote the total number of successful rounds for the honest parties in $r' - r^*$ rounds. We have

$$Z \geq \text{len}(\mathcal{C}) - \ell \geq X.$$

The first inequality was argued above and the second one follows from Lemma 5. Finally, note that, by Lemma 6, the event $Z \geq X$ has measure exponentially small in the number of rounds $r' - r^*$. Since that number satisfies $r' - r^* \geq \delta s/8$, we conclude that $\Pr[A] \leq e^{-\Omega(\delta^3 s)}$.

The second bad event occurs if the adversary has obtained a large number of solutions during $(1 + \frac{\delta}{8})s$ rounds. Specifically, let Z denote the number of successful calls to the oracle by the adversary, for a total of $(1 + \frac{\delta}{8})s$ rounds. Define B to be the event $Z \geq (1 + \frac{\delta}{9})(1 + \frac{\delta}{8})\beta s$. An application of Chernoff bounds gives

$$\Pr[Z \geq (1 + \frac{\delta}{9})(1 + \frac{\delta}{8})\beta s] \leq e^{-\Omega(\beta \delta^2 s)}.$$

The third bad event occurs when the honest parties do not obtain enough solutions from the oracle during uniform rounds. Consider any number, say, s' of rounds (not necessarily consecutive), and denote by X the number of them that were uniquely successful. We have

$$\Pr[X \leq (1 - \frac{\delta}{4})\gamma s'] \leq e^{-\Omega(\gamma \delta^2 s')}.$$

From now on we assume that none of the events A , B and C occurs. It is easy to see that if at any round the adversary does not broadcast a (new) POW, then the next round will be uniform. Using this observation for a given s consecutive rounds, we will calculate a lower bound on the number of rounds that will be uniform. The adversary may prevent a round among the s consecutive rounds from being uniform by broadcasting a solution that was found during the s consecutive rounds as well as in the past for an extended period of $(1 + \frac{\delta}{8})s$ rounds. Note that, since A does not occur, he may not use even older solutions with probability at least $1 - e^{-\Omega(\delta^3 s)}$.

The negation of the second bad event bounds the number of solutions the adversary can obtain. This implies that at least

$$s' = s - (1 + \frac{\delta}{9})(1 + \frac{\delta}{8})\beta s \geq s - (1 + \frac{\delta}{4})\beta s = (1 - \beta)s - \frac{\delta}{4}\beta s$$

rounds among the s rounds will be uniform.

Given the negation of the third bad event, there were $X > (1 - \frac{\delta}{4})\gamma s'$ uniquely successful uniform rounds during the s rounds of the protocol. By Lemma 7, it is necessary for the adversary, in order to maintain the concurrent existence of \mathcal{C}_1 and \mathcal{C}_2 , to obtain at least X solutions. Thus, for the adversary to succeed, it should hold that $Z \geq X$. Substituting in this inequality the bounds on $Z \leq (1 + \frac{\delta}{4})\beta s$ and $X > (1 - \frac{\delta}{4})\gamma s'$ given by $\neg B$ and $\neg C$, respectively, and rearranging we obtain

$$(1 + \frac{\delta}{2})\beta \geq (1 - \frac{\delta}{4})\gamma(1 - \beta). \tag{1}$$

Recall that $\beta + \gamma < f$. Since $\gamma \geq (1 + \delta)\lambda\beta$, this implies $1 - \beta > (1 + \lambda - f)/(1 + \lambda)$. Using the last two inequalities to simplify (1), this implies

$$\lambda^2 - f\lambda - 1 < 0,$$

contradicting the choice of λ in the statement of the lemma. We conclude that if $A \cup B \cup C$ does not occur, then \mathcal{C}_1 and \mathcal{C}_2 cannot diverge at round $r - s$. Finally, an application of the union bound on $A \cup B \cup C$ implies that the adversary can successfully maintain such \mathcal{C}_1 and \mathcal{C}_2 with probability at most exponentially small in s and the statement of the lemma follows. \square

The above lemma is almost what we need, except that it refers to number of rounds instead of number of blocks. In order to obtain the common-prefix property we should use the properties of the blockchains of the parties themselves as the sole measure of divergence. The next theorem establishes the connection.

Theorem 9. *Assume $f < 1$ and $\gamma \geq (1 + \delta)\lambda\beta$, for some real $\delta \in (0, 1)$ and $\lambda \geq 1$ such that $\lambda^2 - f\lambda - 1 \geq 0$. Let \mathcal{S} be the set of the chains of the honest parties at a given round of the backbone protocol. Then the probability that \mathcal{S} does not satisfy the common-prefix property with parameter k is at most $e^{-\Omega(\delta^3 k)}$.*

Remark 1. Observe that as $f \rightarrow 0$, $\lambda \rightarrow 1$. On the other hand, if $f \rightarrow 1$ then $\lambda \rightarrow \phi$, where ϕ is the golden ratio $(\frac{1+\sqrt{5}}{2})$.

Proof. If there is only one chain in \mathcal{S} then the property is satisfied trivially. Consider two chains \mathcal{C}_1 and \mathcal{C}_2 in \mathcal{S} and the least integer k^* such that

$$\mathcal{C}_1^{\lceil k^* \rceil} \preceq \mathcal{C}_2 \quad \text{and} \quad \mathcal{C}_2^{\lceil k^* \rceil} \preceq \mathcal{C}_1. \quad (2)$$

We need to show that the event $k^* \geq k$ happens with probability exponentially small in k .

Let r be the current round and let $r - s$ be the round at which the last common block of \mathcal{C}_1 and \mathcal{C}_2 was computed. The length of the chains cannot be greater than the number of solutions Y obtained from the oracle in s rounds. By the Chernoff bound,

$$\Pr[Y \geq (1 + \delta)fs] \leq e^{-\delta^2 fs/3}.$$

It follows that, with probability $1 - e^{-\delta^2 fs/3}$, $s > k^*/((1 + \delta)f)$. Thus, if $k^* \geq k$, we have a sequence of $s = \Omega(k)$ consecutive rounds with chains \mathcal{C}_1 and \mathcal{C}_2 diverging, and the theorem follows from Lemma 8. \square

Remark 2. Recall that in our analysis we are interested in the relationship between α and β . In particular, the ratio α/β reflects the power the honest parties have against the power of the adversary. $\alpha \leq \beta$ implies that the adversary can, with constant probability, preclude the honest parties that follow the protocol from doing anything useful. This is simply because such an adversary has enough power to build a chain that will often be longer than the chain the honest parties are building. Therefore, it is to be expected that the statements are meaningful only when α/β is bounded away from 1 by a constant δ . In case the network is synchronized ($f \rightarrow 0$), the value of α gets very close to the value of $\gamma = \alpha - \alpha^2$, and hence our result is tight. In case of a larger f , our analysis shows that the upper bound on the adversarial hashing power devolves and in fact approaches 0 as $f \rightarrow 1$ — in other words, in a network where a POW becomes relatively easy compared to network synchronization time, Theorem 9 provides no security guarantee whatsoever.




Figure 3: The degradation of the adversarial bound of Theorem 9 as f ranges in $(0, 1)$ in the x -axis (lower curve). When ties are broken following lexicographic order the analysis can be improved (upper curve).

In practice, this underscores the importance of calibrating the difficulty of the proof of work to maintain a small value of f (such calibration takes place in the Bitcoin system every 2016 blocks). It is an interesting question to further explore the behavior of the backbone protocol in desynchronized networks. We remark that with our analysis we can prove a much better behavior for $f \rightarrow 1$ for a modified backbone protocol that has a deterministic tie-breaking rule (e.g., chooses a chain that is the lexicographically smallest from those received¹¹). In this case we can prove, for example, that our analysis enables the common prefix property to hold when $f = 1$ assuming the adversary controls less than about 29% of the hashing power. In Figure 3 we show how the bound of Theorem 9 degenerates when the parameter f ranges in the $(0, 1)$ range as well as the improvement in the analysis that can be achieved by lexicographic tie-breaking (we omit the details of this analysis).

4.3 The chain-quality property

We now turn to the chain-quality property (Definition 3), which the theorem below establishes for a suitable bound on the number of blocks introduced by the adversary.

Theorem 10. *Assume $f < 1$ and $\gamma \geq (1 + \delta)\lambda\beta$ for some $\delta \in (0, 1)$. Suppose \mathcal{C} belongs to an honest party and consider any ℓ consecutive blocks of \mathcal{C} . The probability that the adversary has contributed more than $(1 - \frac{\delta}{3})\frac{1}{\lambda}\ell$ of these blocks is less than $e^{-\Omega(\delta^2\ell)}$.*

From the above theorem, it follows immediately that the chain quality is satisfied with parameter $\mu = \frac{1}{\lambda}$ for any segment length ℓ and probability that drops exponentially in ℓ .

Proof. Let us denote the i -th block of \mathcal{C} by B_i so that $\mathcal{C} = B_1 \dots B_{\text{len}(\mathcal{C})}$ and consider the ℓ consecutive blocks B_u, \dots, B_v that correspond to some suffix of $\mathcal{C}^{\lceil k}$ that has length at least k . We let L denote the least number of consecutive blocks $B_{u'}, \dots, B_{v'}$ that include the ℓ given ones (i.e., $u' \leq u$ and $v \leq v'$) and have the properties (1) that the block $B_{u'}$ was computed by an honest party or is B_1 , in case such block does not exist, and (2) that there exists a round at which an honest party was trying to extend the chain ending at block $B_{v'}$. Observe that number L is well defined since $B_{\text{len}(\mathcal{C})}$ is at the head of a chain that an honest party is trying to extend.

¹¹This has in fact been debated in an number of occasions; see, e.g., [Cun13].

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

$$x \leq \left(1 - \left(1 - \frac{\delta}{3}\right)\frac{1}{\lambda}\right)\ell \leq \left(1 - \left(1 - \frac{\delta}{3}\right)\frac{1}{\lambda}\right)L.$$

Let s be the number of rounds during which the L blocks were incorporated to the chain. Given that f is bounded away from 1 by a constant, the Chernoff bound implies that in s rounds the total number of solutions is greater than s with probability at most $e^{-\Omega(s)}$. It follows that $s \geq L$ with high probability. Observe that the adversary computed $L - x$ of these blocks.

Let Z be the random variable that corresponds to the POWs obtained by the adversary during these s rounds and X the successful rounds of the honest players in the same sequence of rounds. The probability that the adversary guesses the hash of block $B_{u'}$ and is able to extend it using solutions from previous rounds is negligible in $(\kappa - \log D)$ and we ignore it. We now have:

$$Z \geq L - x \geq \left(1 - \frac{\delta}{3}\right)\frac{1}{\lambda}L \geq \left(1 - \frac{\delta}{3}\right)\frac{1}{\lambda}X \geq \frac{1}{\left(1 + \frac{\delta}{2}\right)\lambda}X$$

where the penultimate inequality follows from Lemma 5. By Lemma 6 this has measure $e^{-\Omega(\delta^2\ell)}$. \square

Remark 3. We are able to argue that Theorem 10 is tight under the simplification that ties between blockchains of equal length always favor the adversary. In particular, we assume that the function `maxvalid` at line 5 of Algorithm 4, in case of chains of equal length, will always return the suggestion of the adversary if there is one. This simplification is made without loss of generality in our model since the adversary is rushing and hence in case two chains are transmitted in a single round the adversary can always arrange it so that its own solution arrives first¹². Furthermore, if the number of honest parties is large, when an honest party discovers a solution in a round, all other honest parties will prefer the one transmitted by the adversary and thus the effect of a single honest party opting for its own block will be negligible.

The attack below is a type of “selfish mining” attack (variation of the one in [ES14]) that accomplishes the stated bound. The attack is as follows. Initially, the adversary works on the same chain as every honest party. However, whenever it finds a solution it keeps it private and keeps on extending a private chain. Whenever an honest party finds a solution, the rushing adversary releases one block from the private chain; if the private chain is depleted the adversary returns to the public chain. We now argue that this strategy exploits the conditions stated above and maximizes the adversarial blocks in the blockchain up to the upper bound of Theorem 10.

Consider s rounds of the protocol. With high probability, the adversary will obtain more than $(1 - \epsilon)\beta s$ solutions for some small $\epsilon > 0$. With each one of them it will try to block the blocks that are broadcast by honest parties. At the end of the s rounds, there may be a few “unused” blocks but these will be, with high probability, at most $\epsilon\beta s$. This is because during the rounds that the adversary acquired the blocks that it did not broadcast, none of the honest players obtained a solution; this is a low probability event. Now, the honest parties will have—with high probability—at most $(1 + \epsilon)\gamma s$ successful rounds. It follows that, for a small constant δ , the quality of the chain is $1 - \frac{1-\delta}{\lambda}$. Note that the Chernoff bound can be used to make the argument more formal and replace the expression “with high probability” with $1 - e^{-\Omega(s)}$. From this it follows that in order to obtain better chain quality one should consider mechanisms that result in more favorable (for the honest parties) behavior in the function `maxvalid`.

¹²In fact, this rushing capability was argued to be realistic in [ES14] through the dispersion of sybil nodes in the Bitcoin peer-to-peer network that echo the adversary’s messages.

5 Simple POW-based Byzantine Agreement Protocols

We now turn to applications of the Bitcoin backbone protocol, showing how it can be used as a basis to solve other problems. We start in this section by analyzing Nakamoto’s suggestion for solving BA, observing that it falls short of satisfying Definition 1; we then present our simple instantiation which solves BA. This protocol, however, only tolerates an adversarial hashing power less than $1/3$, which takes us to the next section, where we present Bitcoin’s essential task, namely, distributively maintaining a public transaction ledger, as well as a more elaborate BA protocol tolerating an adversarial power strictly less than $1/2$. An overview of our applications and the way their properties depend on those of the backbone protocol was already presented in Figure 1.

5.1 Nakamoto’s suggestion for Byzantine agreement

As our first illustration of how the Bitcoin backbone can be used we present Nakamoto’s suggestion for solving BA, as presented in a forum post [Nak08b]¹³. We describe his solution (call it $\Pi_{\text{BA}}^{\text{nak}}$) via the backbone protocol by specifying the functions $V(\cdot), I(\cdot), R(\cdot)$ in a suitable way (see Figure 4). The input validation predicate $V(\cdot)$ will be defined to require that all valid chains contain the same input value. The chain reading function $R(\cdot)$ simply returns this value in case the chain has length at least k (which is the security parameter); otherwise it is undefined. The input contribution function $I(\cdot)$ examines the contents of the current chain \mathcal{C} and the contents of the input tape $\text{INPUT}()$. In case $\mathcal{C} = \varepsilon$ the input contribution for the next block is taken *verbatim* from the input tape; otherwise, the input contribution is determined as the (unique) value that is already present in the \mathcal{C} (and in this case the input is ignored). Note that we will only consider environments \mathcal{Z} that provide an input symbol to all parties.

It follows that initially the protocol builds various chains all containing the same value. The intuition is that Agreement will follow from the fact that the honest players will eventually agree on a single chain, as long as the majority of the hashing power lies with the honest parties. While this is true, as we will demonstrate, the second necessary property does not hold: this protocol cannot provide Validity (with high probability).

| | |
|---|---|
| Input validation predicate $V(\cdot)$ | $V(\langle x_1, \dots, x_n \rangle)$ is true if and only if $x_1 = \dots = x_n$. |
| Chain reading function $R(\cdot)$ (parameterized by k) | If $V(x_{\mathcal{C}}) = \text{True}$ and $\text{len}(\mathcal{C}) \geq k$, the value of $R(x_{\mathcal{C}})$ is the (unique) value that is present in each block of \mathcal{C} , while it is undefined if $V(x_{\mathcal{C}}) = \text{False}$ or $\text{len}(\mathcal{C}) < k$. |
| Input contribution function $I(\cdot)$ | If $\mathcal{C} = \emptyset$, $I(st, \mathcal{C}, \text{round}, \text{INPUT}())$ is equal to v if the input tape contains (INSERT, v) ; otherwise ($\mathcal{C} \neq \emptyset$), it is equal to the x value present in \mathcal{C} . The state st always remains ε . |

Figure 4: Expressing Nakamoto’s BA protocol $\Pi_{\text{BA}}^{\text{nak}}$ over the Bitcoin backbone protocol via the specification of $V(\cdot), R(\cdot), I(\cdot)$.

As we now show, Agreement follows easily from the common-prefix property. Indeed, as long as there is a common prefix (irrespective of its length), it is ensured that when $R(\cdot)$ becomes defined all honest parties will produce the same output. Note however that a small tweak would be needed to ensure “work independence”: the parties need to introduce a fresh random κ -bit nonce at each block (cf. the beginning of Sec. 4). With this tweak we can easily derive the following.

¹³Note that Nakamoto’s description is quite informal. We make the most plausible interpretation of it in our formal framework.

Lemma 11 (Agreement). *Suppose $f < 1$ and $\gamma \geq (1 + \delta)\lambda\beta$, for some real $\delta \in (0, 1)$ and $\lambda \geq 1$ such that $\lambda^2 - f\lambda - 1 \geq 0$. Protocol $\Pi_{\text{BA}}^{\text{nak}}$ from Fig. 4 satisfies Agreement (cf. Definition 1) with probability at least $1 - e^{-\Omega(\delta^3 k)}$.*

Proof. Observe that chains contain unique values, therefore a disagreement between honest parties implies that two parties have completely disjoint chains (essentially, this is equivalent to a fork that happens at the onset). It follows from the common prefix property (Theorem 9) that the event of any two chains of length at least k that are completely disjoint happens with probability at most $e^{-\Omega(\delta^3 k)}$. \square

On the other hand, it is easy to see that Validity cannot be guaranteed with overwhelming probability unless the hashing power of the adversary is negligible compared to the honest players, i.e., t/n is negligible. This is because in case the adversary finds a solution first, then every honest player will extend the adversary’s solution and switch to the adversarial input hence abandoning the original input. While one can still show that Validity can be ensured with non-zero probability (and thus the protocol fails gracefully assuming honest majority), $\Pi_{\text{BA}}^{\text{nak}}$ falls short from providing a solution to BA. Interestingly, by appropriately modifying the way the backbone protocol is used, we show in the next section how a solution can be derived.

5.2 A Byzantine agreement protocol for (1/3)-bounded adversaries

We now show that the Bitcoin backbone can actually be used to satisfy BA’s properties with an error that decreases exponentially in the length of the chain, assuming however that the adversary’s hashing power is less than 1/3. There are two important differences with respect to the approach in the previous section: (i) parties never abandon their original input but instead they do insist in inserting it into the blockchain, and (ii) when the chain becomes of length $2k$, they output the majority of a length- k prefix (note that here we consider binary BA). The protocol (i.e., the specification of the functions) is presented in Figure 5.

| | |
|---|--|
| Input validation predicate $V(\cdot)$ | $V(\langle x_1, \dots, x_n \rangle)$ is true if and only if $v_1, \dots, v_n \in \{0, 1\}, \rho_1, \dots, \rho_n \in \{0, 1\}^\kappa$ where v_i, ρ_i are the values from the pair $\langle v_i, \rho_i \rangle = x_i$. |
| Chain reading function $R(\cdot)$ (parameterized by k) | If $V(\langle x_1, \dots, x_n \rangle) = \text{True}$ and $n \geq 2k$, the value $R(x_{\mathcal{C}})$ is the majority bit of v_1, \dots, v_k where $x_i = \langle v_i, \rho_i \rangle$; otherwise ($V(\langle x_1, \dots, x_n \rangle) = \text{False}$ or $n < 2k$) the output value is undefined. |
| Input contribution function $I(\cdot)$ | $I(st, \mathcal{C}, \text{round}, \text{INPUT}())$ is equal to $\langle v, \rho \rangle$ if the input tape contains (INSERT, v) ; ρ is a random κ -bit string. The state st remains always ϵ . |

Figure 5: Protocol $\Pi_{\text{BA}}^{1/3}$ over the Bitcoin backbone via the specification of $V(\cdot), R(\cdot), I(\cdot)$.

Lemma 12 (Agreement). *Suppose $f < 1$ and $\gamma \geq 2(1 + \delta)\beta$, for some real $\delta \in (0, 1)$. Protocol $\Pi_{\text{BA}}^{1/3}$ of Fig. 5 satisfies Agreement in $O(k)$ rounds with probability at least $1 - e^{-\Omega(\delta^3 k)}$.*

Proof. In order for agreement to be violated, at least two honest players should have upon termination chains \mathcal{C}_1 and \mathcal{C}_2 such that $\mathcal{C}_1^{[k]} \neq \mathcal{C}_2^{[k]}$. In particular, the set $\{\mathcal{C}_1, \mathcal{C}_2\}$ should be a set of chains that belong to honest parties and does not satisfy the common-prefix property. Thus, the statement of the lemma follows directly from Theorem 8. \square

We now turn to the Validity property. In order to prove it we need to show that, upon termination of the protocol, the chain of any honest party will contain among the first k inputs more inputs from honest players than provided by the adversary. As we will see, this is a consequence of the chain-quality property.

Lemma 13 (Validity). *Suppose $f < 1$ and $\gamma \geq 2(1 + \delta)\beta$, for some real $\delta \in (0, 1)$. Protocol $\Pi_{\text{BA}}^{1/3}$ satisfies Validity in $O(k)$ rounds with probability at least $1 - e^{-\Omega(\delta^2 k)}$.*

Proof. For the property to be satisfied we only need to ensure that in $\mathcal{C} \upharpoonright_k$ the majority of the inputs was computed by the honest parties. As in protocol $\Pi_{\text{BA}}^{1/3}$ we have $\text{len}(\mathcal{C} \upharpoonright_k) = k$, Theorem 10 with $\lambda = 2$ provides exactly what we want. \square

Note that $\Pi_{\text{BA}}^{1/3}$ solves BA only in case the adversary’s hashing power is bounded by $1/3$. In case adversarial blocks win all head-to-head races within a round, the result is tight, as argued in Remark 3. In the next section we show a more elaborate construction based on a transaction ledger which tolerates less than $1/2$.

Remark 4. As mentioned in Section 2, “Strong Validity” refers to the requirement that the output value be one of the honest parties’ inputs, and the distinction is relevant in the case of non-binary inputs, i.e., coming from an arbitrary set V , $|V| > 2$. It is easy to modify the above algorithm to also satisfy this property by making the chain reading function the element with highest plurality in the chain (ties broken favoring the lexicographically smallest element in V), as opposed to majority, and by imposing a more stringent bound on the adversary, namely, that $\gamma \geq |V|(1 + \delta)\beta$. This ensures that the expected number of blocks in the blockchain that are controlled by the adversary is less than $\frac{1}{|V|}$, and maintains validity even in the worst case that the honest parties’ inputs are equally split among all possible values but one (i.e., there are $|V| - 1$ inputs equally proportioned among the honest parties). Agreement is ensured in the same way as before via the common-prefix property. The bound is in-line with the known bounds for the standard (computational) setting, $n > |V|t$, cf. [FG03].

6 Public Transaction Ledgers

We now come to the application which the Bitcoin backbone was designed to solve: maintaining a public transaction ledger. We first formally introduce this object — a “book” where transactions are recorded — and its properties, and then we show how it can be used to implement the Bitcoin ledger and BA in the honest majority setting by properly instantiating the notion of a transaction.

6.1 Robust public transaction ledgers

A *public transaction ledger* is defined with respect to a set of valid ledgers \mathcal{L} and a set of valid transactions \mathcal{T} , each one possessing an efficient membership test. A ledger $\mathbf{x} \in \mathcal{L}$ is a vector of sequences of transactions $\text{tx} \in \mathcal{T}$. Each transaction tx may be associated with one or more *accounts*, denoted a_1, a_2, \dots

The backbone protocol parties, called *miners* in the context of this section, process sequences of transactions of the form $x = \text{tx}_1 \dots \text{tx}_e$ that are supposed to be incorporated into their local chain \mathcal{C} . The input inserted at each block of the chain \mathcal{C} is the sequence x of transactions. Thus, a ledger is a vector of transaction sequences $\langle x_1, \dots, x_m \rangle$, and a chain \mathcal{C} of length m contains the ledger $\mathbf{x}_{\mathcal{C}} = \langle x_1, \dots, x_m \rangle$ if the input of the j -th block in \mathcal{C} is x_j .

The description and properties of the ledger protocol will be expressed relative to an oracle Txgen which will control a set of accounts by creating them and issuing transactions on their behalf. In an execution of the backbone protocol, the environment \mathcal{Z} as well as the miners will have access to Txgen . Specifically, Txgen is a stateful oracle that responds to two types of queries (which we purposely only describe at a high level):

- $\text{GenAccount}(1^\kappa)$: It generates an account a .
- $\text{IssueTrans}(1^\kappa, \tilde{\text{tx}})$: It returns a transaction tx provided that $\tilde{\text{tx}}$ is some suitably formed string, or \perp .

We also consider a symmetric relation on \mathcal{T} , denoted by $C(\cdot, \cdot)$, which indicates when two transactions tx_1, tx_2 are conflicting. Valid ledgers $\mathbf{x} \in \mathcal{L}$ can never contain two conflicting transactions. We call oracle Txgen *unambiguous* if it holds that for all PPT \mathcal{A} , the probability that $\mathcal{A}^{\text{Txgen}}$ produces a transaction tx' such that $C(\text{tx}', \text{tx}) = 1$, for tx issued by Txgen , is negligible in κ .

Finally, a transaction tx is called *neutral* if $C(\text{tx}, \text{tx}') = 0$ for any other transaction tx' . The presence of neutral transactions in the ledger can be helpful for a variety of purposes, as we will see next and in the BA protocol that we build on top of the ledger. For convenience we will assume that a single random nonce $\rho \in \{0, 1\}^\kappa$ is also a valid transaction. Nonces will be neutral transactions and may be included in the ledger for the sole purpose of ensuring independence between the POW instances solved by the honest parties.

Next, we determine the three functions $V(\cdot), I(\cdot), R(\cdot)$ that will turn the backbone protocol into Π_{PL} , a protocol realizing a public transaction ledger.

| | |
|--|--|
| Input validation predicate $V(\cdot)$ | $V(\langle x_1, \dots, x_m \rangle)$ is true if and only if the vector $\langle x_1, \dots, x_m \rangle$ is a valid ledger, i.e., $\langle x_1, \dots, x_m \rangle \in \mathcal{L}$. |
| Chain reading function $R(\cdot)$ | If $V(\langle x_1, \dots, x_m \rangle) = \text{True}$, the value $R(\mathbf{x}_{\mathcal{C}})$ is equal to $\langle x_1, \dots, x_m \rangle$; undefined otherwise. |
| Input contribution function $I(\cdot)$ | $I(st, \mathcal{C}, \text{round}, \text{INPUT}())$ operates as follows: if the input tape contains (INSERT, v) , it parses v as a sequence of transactions and retains the largest subsequence $x' \preceq v$ that is valid with respect to $\mathbf{x}_{\mathcal{C}}$ (and whose transactions are not already included in $\mathbf{x}_{\mathcal{C}}$). Finally, $x = \text{tx}_0 x'$ where tx_0 is a neutral random nonce transaction. |

Figure 6: *The public transaction ledger protocol Π_{PL} , built on the Bitcoin backbone.*

We now introduce two essential properties for a protocol maintaining a public transaction ledger: (i) *Persistence* and (ii) *Liveness*. In a nutshell, Persistence states that once an honest player reports a transaction “deep enough” in the ledger, then all other honest players will report it indefinitely whenever they are asked, and at exactly the same position in the ledger (essentially, this means that all honest players agree on all the transactions that took place and in what order). In a more concrete Bitcoin-like setting (cf. Appendix 6.2), Persistence is essential to ensure that credits are final and that they happened at a certain “time” in the system’s timeline (which is implicitly defined by the ledger itself).

Note that Persistence is useful but not enough to ensure that the ledger makes progress, i.e., that transactions are eventually inserted in a chain. This is captured by the Liveness property, which states that as long as a transaction comes from an honest account holder and is provided by the environment to all honest players, then it will be inserted into the honest players’ ledgers, assuming the environment keeps providing it as an input for a sufficient number of rounds¹⁴.

¹⁴Observe that here we take the view that new transactions are available to all honest players and the way they

Definition 14. A protocol Π implements a *robust public transaction ledger* in the q -bounded synchronous setting if it satisfies the following two properties:

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

We prove the two properties separately, starting with Persistence. The proof is based on the common prefix property of the backbone protocol (recall Definition 2 and Theorem 9).

Lemma 15 (Persistence). *Suppose $f < 1$ and $\gamma \geq (1 + \delta)\lambda\beta$, for some real $\delta \in (0, 1)$ and $\lambda \geq 1$ such that $\lambda^2 - f\lambda - 1 \geq 0$. Protocol Π_{PL} satisfies Persistence with probability $1 - e^{-\Omega(\delta^3 k)}$, where k is the depth parameter.*

Proof. Let \mathcal{C} be the chain of some honest player. We show that if a transaction tx is included in $\mathcal{C}^{\lceil k}$ then this transaction will be always included in every honest player’s chain with probability at least $1 - e^{-\Omega(\delta^2 k)}$.

Theorem 9 implies that, with high probability, $\mathcal{C}^{\lceil k}$ is a common prefix of every honest party’s chain. Thus, tx is included in $\mathcal{C}'^{\lceil k}$, for any \mathcal{C}' that belongs to an honest party. Now, let r be the current round and consider the first round $r' > r$ at which an honest party with chain \mathcal{C}_1 accepts an alternative chain \mathcal{C}_2 , such that $\mathcal{C}_2^{\lceil k}$ does not include tx . Suppose tx was inserted in \mathcal{C}_1 at round r^* . Then, \mathcal{C}_1 and \mathcal{C}_2 diverge at round r^* . By Lemma 8, this occurs with probability $e^{-\Omega(\delta^3(r' - r^*))}$. The claim then follows by a union bound over all $r^* \leq r$. Letting $s = r - r^*$ and ϵ be an appropriate constant, the probability that Persistence is violated is at most $\sum_{s' \geq s} e^{-\epsilon \delta^3 s'} = e^{-\Omega(\delta^3 s)}$. Finally, as in the proof Theorem 9 we can argue that $s = \Omega(k)$. \square

We next prove Liveness, which is based on the chain-quality property (recall Definition 3 and Theorem 10).

Lemma 16 (Liveness). *Assume $f < 1$ and $\gamma \geq (1 + \delta)\lambda\beta$, for some $\delta \in (0, 1)$, $\lambda \in [1, \infty)$ and let $k \in \mathbb{N}$. Further, assume oracle Txgen is unambiguous. Then protocol Π_{PL} satisfies Liveness with wait time $u = 2k / ((1 - \delta)\gamma)$ and depth parameter k with probability at least $1 - e^{-\Omega(\delta^2 k)}$.*

Proof. We prove that assuming all honest players receive as input the transaction tx for at least $u = 2k / ((1 - \delta)\gamma)$ rounds, there exists an honest party with chain \mathcal{C} such that tx is included in $\mathcal{C}^{\lceil k}$.

Indeed, after u rounds the Chernoff bound implies that the honest parties had at least $2k$ successful rounds with probability at least $1 - e^{-\Omega(\delta^2 k)}$. Invoking Lemma 5, we infer that the chain’s length of any honest party has increased by at least $2k$ blocks. Finally, the chain-quality property (Theorem 10) implies that at least one of the blocks in the length- k suffix of $\mathcal{C}^{\lceil k}$ was computed by an honest party. Such a block would include tx since it is infeasible for adversarial \mathcal{Z}, \mathcal{A} to produce a conflicting transaction tx' (which would be the only event making an honest player drop tx from the sequence of transactions x that it attempts to insert in the blockchain). Thus, the lemma follows. \square

are propagated is handled by the environment that feeds the backbone protocol. While this makes sense in the honest/malicious cryptographic model, it has been challenged in a model where all players are rational [BDOZ12]. Analysis of the backbone protocol in a setting where transaction propagation is governed by rational players is beyond the scope of our paper.

6.2 Bitcoin-like transactions and ledger

Next, we show how to instantiate the public transaction ledger for Bitcoin, by defining the sets of transactions and valid ledgers.

Transactions and accounts are defined with respect to a digital signature scheme that is comprised of three algorithms $\langle \text{KeyGen}, \text{Sign}, \text{Verify} \rangle$. An *account* will be a pair $a = (vk, G(vk))$ where $G(\cdot)$ is a hash function and $G(vk)$ is the “address” corresponding to the account.

A transaction tx is of the form “ $\{a_1, a_2, \dots, a_i\} \rightarrow (\sigma, \{(a'_1, b'_1), \dots, (a'_o, b'_o)\})$,” where a_1, \dots, a_i are the accounts to be debited, a'_1, \dots, a'_o are the addresses of the accounts¹⁵ to be credited with funds b'_1, \dots, b'_o , respectively, and σ is a vector of verification keys and digital signatures issued under them $\langle (vk_1, \sigma_1), \dots, (vk_i, \sigma_i) \rangle$ on the same message $\{(a'_1, b'_1), \dots, (a'_o, b'_o)\}$. (We note that Bitcoin transactions can be more expressive but the above description is sufficient for the purpose of our analysis).

Next, we specify the Txgen oracle:

- $\text{GenAccount}(1^\kappa)$: It generates an account a by running KeyGen and computing the hash $G(\cdot)$ on the verification key. The account is the pair $(vk, G(vk))$, where $G(vk)$ is the account’s address. The corresponding secret key, sk , is kept in the state of Txgen .
- $\text{IssueTrans}(1^\kappa, \tilde{\text{tx}})$: It returns a transaction tx provided that $\tilde{\text{tx}}$ is a transaction that is only missing the signatures by accounts that are maintained by Txgen . (Recall the format of transactions above.) Each account is only allowed a single transaction.

Note that the above restriction on IssueTrans is without loss of generality, as in Bitcoin, entities typically maintain a number of accounts and are allowed (although not forced) to move their balances forward to a new account as they make new transactions. The conflict relation $C(\cdot, \cdot)$ over \mathcal{T} satisfies that $C(\text{tx}_1, \text{tx}_2) = 1$ if and only if $\text{tx}_1 \neq \text{tx}_2$ and tx_1, tx_2 have an overlapping input account¹⁶. Thus, we can easily prove the unambiguity of the Txgen oracle based on the unforgeability of the underlying digital signature.

Lemma 17. *Assume that $\langle \text{KeyGen}, \text{Sign}, \text{Verify} \rangle$ is an existentially unforgeable signature scheme. Then oracle Txgen is unambiguous.*

In order to define the set of valid Bitcoin ledgers we first need to determine in what sense a transaction may be valid with respect to a ledger. Then we will define the set of valid ledgers recursively as the maximal set of vectors of sequences of transactions that satisfy this condition. So here it goes.

A transaction tx is valid with respect to a Bitcoin ledger $\mathbf{x} = \langle x_1, \dots, x_m \rangle$ provided that all digital signatures verify and $\sum_{j=1}^i b_j \geq \sum_{j=1}^o b'_j$, where b_j is the balance that was credited to account a_j in the latest transaction involving a_j in \mathbf{x} . In case $e = \sum_{j=1}^i b_j - \sum_{j=1}^o b'_j > 0$, then e is a transaction fee that may be claimed separately in a special transaction of the form “ $\emptyset \rightarrow \dots$,” called a *coinbase* transaction. In more detail, a coinbase transaction has no inputs and its purpose is to enable miners to be rewarded for maintaining the ledger. The transaction is of the form “ $\emptyset \rightarrow \{(a_1, b_1), \dots, (a_o, b_o)\}$,” and $\sum_{j=1}^o b_j$ is determined based on the other transactions that are “bundled” in the block as well as a flat reward fee, as explain below.

A sequence of transactions $x = \langle \emptyset \rightarrow \{(a_1, b_1), \dots, (a_o, b_o)\}, \text{tx}_1, \dots, \text{tx}_l \rangle$ is said to be valid with respect to a ledger $\mathbf{x} = \langle x_1, \dots, x_m \rangle$, if each transaction tx_j is valid with respect to the ledger \mathbf{x}

¹⁵In bitcoin terminology every account has an address that is used to uniquely identify it. Payments directed to an account require only this “bitcoin address.” The actual verification key corresponding to the account will be revealed only when the account makes a payment.

¹⁶The conflict relation is more permissive in the actual Bitcoin ledger. We adopt the more simplified version given above as it does not change the gist of the analysis.

extended by the transactions $\text{tx}_1, \dots, \text{tx}_{j-1}$. I.e., for all $j = 1, \dots, l$ the transaction tx_j should be valid with respect to ledger

$$\langle x_1, \dots, x_m, \text{tx}_1 \dots \text{tx}_{j-1} \rangle,$$

and furthermore, the total fee $e = \sum_{j=1}^o b_j$ collected in the transaction $\emptyset \rightarrow \{(a_1, b_1), \dots, (a_o, b_o)\}$ does not exceed $r_m + \sum_{j=1}^m e_j$, which includes all the individual fees corresponding to transactions $\text{tx}_1, \dots, \text{tx}_e$, plus a value r_m that is the flat reward given for extending a ledger of length m to a ledger of length $m + 1$.¹⁷

The set of valid ledgers \mathcal{L} with respect to a reward progression $\{r_j\}_{j \in \mathbb{N}}$ contains ε (the empty ledger), and any ledger \mathbf{x} which extends a ledger in \mathcal{L} by a valid sequence of transactions. Note that the first transaction sequence of any ledger $\mathbf{x} \in \mathcal{L}$ contains a single transaction of the form $\emptyset \rightarrow \{(a_1, b_1), \dots, (a_o, b_o)\}$ that satisfies $\sum_{j=1}^o b_j = r_0$, where r_0 is the initial flat reward. This first transaction “distributes an initial amount of money” to the ledger’s initiator(s)¹⁸. It is easy to see that \mathcal{L} has an efficient membership test.

Given the existence of coinbase transactions in this application we can do away with random nonces as standalone transactions and the description of the input contribution function I in Fig. 6, is modified to include their generation each time an input sequence of transactions is determined to be inserted in the ledger. Specifically, $I(\cdot)$ will form a *coinbase transaction* $\emptyset \rightarrow \{(a, b)\}$, where $b = r_{\text{len}(c)} + \sum_{j=1}^m e_j$ and e_j is the fee corresponding to x ’s j -th transaction. Account a is a freshly created account that is obtained via running `KeyGen`. $I(\cdot)$ will append account a and the corresponding (vk, sk) to its private state st .

We will refer to the modified Π_{PL} protocol by the moniker Π_{BTC} . Π_{BTC} inherits from Π_{PL} the properties of Persistence and Liveness which will ensure the following with overwhelming probability in k .

- Apart from its latest k blocks, the transaction ledger is fixed and immutable for all honest miners.
- If a majority of miners¹⁹ receive an honest transaction and attempt to insert it following the protocol for a sufficient number of rounds (equal to parameter u , the “wait time”), it will become a permanent entry in the ledger (no matter the adversarial strategy of the remaining miners).

6.3 Byzantine agreement for honest majority

We now use the public transaction ledger formulation to achieve POW-based BA for an honest majority by properly instantiating the notion of a transaction, thus improving on the simple BA protocol tolerating a $(1/3)$ -bounded adversary presented in Section 5.

Here we consider a set of valid ledgers \mathcal{L} that contain sequences of transactions of the form $(ctr, nonce, v)$, and satisfy the predicate:

$$(H_1(ctr, G(nonce, v)) \leq D) \wedge (ctr \leq q), \tag{3}$$

where $H_1(\cdot), G(\cdot)$ are two hash functions as in the definition of the backbone protocol, and $v \in \{0, 1\}$ is a party’s input. (Recall that D is the difficulty level and q determines how many calls to $H_1(\cdot)$ a party is allowed to make per round.) To distinguish the oracles, in this section we will use $H_0(\cdot)$ to refer to the oracle used in the backbone protocol.

¹⁷Currently, the flat reward for extending the Bitcoin chain is 25BTC. The sequence r_0, r_1, \dots for Bitcoin follows a geometric progression with large constant intervals.

¹⁸In the case of Bitcoin, it was supposedly Nakamoto himself who collected this first reward of 50BTC.

¹⁹Recall that we assume a flat model w.r.t. hashing power; a majority of miners corresponds to a set of parties controlling the majority of the hashing power.

Algorithm 5 The POW-based *transaction production* protocol Π_{tx} , parameterized by q , D and hash functions $H_1(\cdot), G(\cdot)$.

```

1:  $v \leftarrow \text{INPUT}()$ 
2:  $ctr \leftarrow 1$ 
3:  $\text{tx} \leftarrow \varepsilon$ 
4:  $h \leftarrow G(\text{nonce}, v)$  ▷ nonce is a random  $\kappa$ -bit string
5: while ( $ctr \leq q$ ) do
6:   if ( $H_1(ctr, h) \leq D$ ) then ▷ Proof of work succeeded
7:      $\text{tx} \leftarrow \langle ctr, \text{nonce}, v \rangle$ 
8:     break
9:   end if
10:   $ctr \leftarrow ctr + 1$ 
11: end while
12:  $\text{BROADCAST}(\text{tx})$ 

```

Figure 7: The transaction production protocol Π_{tx} .

For the ledger we consider in this section, there will be no accounts and all transactions will be neutral — i.e., the conflict predicate $C(\cdot, \cdot)$ will be false for all pairs of transactions.

We first provide a high level description of the BA protocol assuming parties have q queries per round to each oracle $H_0(\cdot), H_1(\cdot)$. We then show how to use a single oracle $H(\cdot)$ to achieve the combined functionality of both of them while only using q queries per round.

At a high level, the protocol, $\Pi_{\text{BA}}^{1/2}$, works as follows:

- *Operation*: In each round, parties run two protocols in parallel. The first protocol is protocol Π_{PL} (Fig. 6), which maintains the transaction ledger and requires q queries to the oracle $H_0(\cdot)$. The second process is the “transaction production” protocol Π_{tx} , (Fig. 7), which continuously generates transactions satisfying predicate (3). The protocol makes q queries to the $H_1(\cdot)$ oracle.
- *Termination*: When the ledger reaches $2k$ blocks, a party prunes the last k blocks, collects all the unique POW transactions that are present in the ledger and returns the majority bit from the bits occurring in these transactions.

As described, protocol $\Pi_{\text{BA}}^{1/2}$ does not conform to the q -bounded setting since parties require q queries to oracle $H_0(\cdot)$ and q queries to oracle $H_1(\cdot)$ to perform the computation of a single round (the setting imposes a bound of q queries to a single oracle for all parties). Note that a naïve simulation of $H_0(\cdot), H_1(\cdot)$ by a single oracle $H(\cdot)$ in the $(2q)$ -bounded setting (e.g., by setting $H_b(x) = H(b, x)$) would violate the restriction imposed on each oracle individually, since nothing would prevent the adversary, for example, from querying $H_0(\cdot)$ $2q$ times. Next, we show how we can combine the two protocols into a single protocol that utilizes at most q queries to a single random oracle in a way that the adversary will remain q -bounded for each oracle. This transformation, explained below, completes the description of $\Pi_{\text{BA}}^{1/2}$.

2-for-1 POWs. We now tackle the problem of how to turn a protocol operation that uses two separate POW subprocedures involving two distinct and independent oracles $H_0(\cdot), H_1(\cdot)$ into a protocol that utilizes a single oracle $H(\cdot)$ for a total number of q queries per round. Our transformation is general and works for any pair of protocols that utilize $H_0(\cdot), H_1(\cdot)$, provided that certain

Algorithm 6 POW-based protocol fragment of Π_b , $b \in \{0, 1\}$ parameterized by q , D and hash functions $H_b(\cdot), G(\cdot)$, $b \in \{0, 1\}$. The value w_b is determined from the protocol's context.

```

1: ...                                ▷ Value  $w_b$  is determined
2:  $ctr \leftarrow 1$ 
3:  $B \leftarrow \varepsilon$ 
4:  $h_b \leftarrow G(w_b)$ 
5: while ( $ctr \leq q$ ) do
6:   if ( $H(ctr, h_b) < D$ ) then
7:      $B_b \leftarrow \langle ctr, w_b \rangle$ 
8:     break
9:   end if
10:   $ctr \leftarrow ctr + 1$ 
11: end while
12: ...                                ▷ The POW  $B$  is exploited here

```

Algorithm 7 The *double proof of work* function, parameterized by q , D and hash functions $H(\cdot), G(\cdot)$ that substitutes steps 2-11 of two POW-based protocols.

```

1: function double-pow( $w_0, w_1$ )
2:    $B_0, B_1 \leftarrow \varepsilon$ 
3:    $ctr \leftarrow 1$ 
4:   while ( $ctr \leq q$ ) do
5:      $h \leftarrow H(ctr, G(w_0), G(w_1))$ 
6:     if ( $h < D$ ) then
7:        $B_0 \leftarrow \langle ctr, w_0, G(w_1) \rangle$ 
8:       break
9:     end if
10:    if ( $[h]^R < D$ ) then
11:       $B_1 \leftarrow \langle ctr, w_1, G(w_0) \rangle$ 
12:      break
13:    end if
14:     $ctr \leftarrow ctr + 1$ 
15:  end while
16:  return  $\langle B_0, B_1 \rangle$ 
17: end function

```

Figure 8: The 2-for-1 POW transformation.

conditions are met (which are satisfied by protocol $\Pi_{\text{BA}}^{1/2}$ above). In more detail, we consider two protocols Π_0, Π_1 that utilize a POW step as shown in Algorithm 6 in Figure 8.

In order to achieve composition of the two protocols Π_0, Π_1 in the q -bounded setting with access to a single oracle $H(\cdot)$, we will substitute steps 2-11 in both protocols with a call to a new function, `double-pow`, defined below. First, observe that in Π_b , $b \in \{0, 1\}$, the POW steps 4-11 operate with input w_b and produce output in B_b if the POW succeeds. The probability of obtaining a solution is $D \cdot 2^{-\kappa}$.

The modification consists in changing the structure of the POWs from pairs of the form (ctr, w) to triples of the form $(ctr, w, label)$, where *label* is a κ -bit string. This will further require the modification of the verification step for POWs in both protocols Π_0, Π_1 in the following manner.

- Any verification step in Π_0 of a POW (ctr, w_0) which is of the form $H(ctr, G(w_0)) < D$, will now operate with a POW of the form $(ctr, w_0, label)$ and will verify the relation

$$H(ctr, G(w_0), label) < D.$$

- Any verification step in Π_1 of a POW (ctr, w_1) which is of the form $H(ctr, G(w_1)) < D$, will now operate with a POW of the form $(ctr, w_1, label)$ and will verify the relation

$$[H(ctr, label, G(w_1))]^R < D,$$

where $[a]^R$ denotes the reverse of the bitstring a .

This parallel composition strategy in the form of function `double-pow` is shown in Algorithm 7. Either or both the solutions it returns, B_0, B_1 , may be empty if no solution is found.

Protocol $\Pi_{\text{BA}}^{1/2}$ will employ `double-pow`, which will substitute the individual POW operation of the two underlying protocols Π_0, Π_1 as defined in lines 2-10 of Algorithm 6. The correctness of the above composition strategy follows from the following simple observation.

Lemma 18. *Consider a uniform random variable U over the integers in $[0, 2^\kappa)$ and an integer D such that $D = 2^t$ for some positive integer $t < \kappa/2$. Then, the events $(U < D)$ and $([U]^R < D)$ are independent and they both occur with probability $D \cdot 2^{-\kappa}$.*

Proof. It is easy to see that each event happens with probability $D \cdot 2^{-\kappa}$. The conjunction of the two events involves the choice of an integer U which satisfies $U < D$ and $[U]^R < D$. Observe that because $D = 2^t$, it follows that the conditioning on $U < D$ leaves the t least significant bits of U uniformly random while fixing the remaining $\kappa - t$ bits. It follows that the t most significant bits of $[U]^R$ are uniformly random in the conditional space $U < D$. The event $[U]^R < D$ has probability $(D/2^{\kappa-t})/D = D2^{-\kappa}$ and thus the two events are independent. \square

Theorem 19. *Assume $f < 1$ and $\gamma \geq (1 + \delta)\lambda\beta$, for some real $\delta \in (0, 1)$ and $\lambda \geq 1$ such that $\lambda^2 - f\lambda - 1 \geq 0$. Protocol $\Pi_{\text{BA}}^{1/2}$ solves BA in $O(k)$ rounds with probability at least $1 - e^{-\Omega(\delta^3 k)}$.*

Proof. First observe that due to Lemma 18, the success probability for all parties to solve a proof of work of either kind in each round is $q \cdot D2^{-\kappa}$ and the events are independent with each other.

Regarding Agreement, observe that it follows directly from Theorem 9 (common prefix) that all parties will return the majority of the same set with probability at least $1 - e^{-\Omega(\delta^3 k)}$.

To show Validity, let \mathcal{C} be the chain of an honest party upon termination of the protocol. Let r denote the greatest round on which a block of $\mathcal{C}^{\lceil k}$ was computed by an honest party. We argue that the rest of the blocks in $\mathcal{C}^{\lceil k}$, that must have been inserted by the adversary, were computed by round $(1 + \frac{\delta}{2})r$. Assume the contrary and let $r' > (1 + \frac{\delta}{2})r$ denote the least round on which an honest player adopted the chain \mathcal{C} (after round $(1 + \frac{\delta}{2})r$). Let X denote the successful rounds from round r to round r' and Z the number of POWs the adversary obtained in these rounds. Lemma 5 implies that the chain of every honest player advanced in length by X blocks at least. By the definition of r' , the adversary inserted all the blocks of \mathcal{C} computed in these $s = r' - r$ rounds. It follows that $Z \geq X$. By Lemma 6 this occurs with probability at most $e^{-\Omega(\delta^2 s)}$. To finish the proof, recall that each block contains the aggregation of all broadcast transactions up to the round it was computed. Thus, $\mathcal{C}^{\lceil k}$ contains POWs computed by honest parties during r rounds and, with high probability, POWs computed by the adversary during at most $(1 + \frac{\delta}{2})r$ rounds. By Lemma 6, the honest parties have computed the majority of the blocks with probability at least $1 - e^{-\Omega(\delta^2 s)}$ and Validity is satisfied. Since $s > \delta r/2$, we need to argue that $r = \Omega(k)$. To see this, note that in $(1 + \frac{\delta}{2})r$ rounds the parties created a chain of length $\text{len}(\mathcal{C}^{\lceil k}) = k$. An application of the Chernoff bound shows that $r = \Omega(k)$ with probability at least $1 - e^{-\Omega(\delta k)}$. Finally note that it is easy to infer from Lemma 5 that the length of chain of all honest parties will reach $2k$ blocks in $O(k)$ rounds with probability $1 - e^{-\Omega(\delta k)}$. \square

Remark 5. Regarding strong validity in the multivalued BA setting, i.e., where the input domain is V and has a constant cardinality strictly larger than 2 we can adapt the above protocol to return the plurality from the values stored in the transactions that are found in the ledger. In order to ensure strong validity by this modification we restrict the hashing power of the adversary to $(1 - \delta)/|V|$ since this will ensure that the adversary's number of transactions cannot overturn the plurality value as defined by the honest parties' inputs (even if those are evenly distributed amongst them).

7 Summary and Directions for Future Work

In this paper we presented a formal treatment of the Bitcoin backbone, the protocol used at the core of Bitcoin’s transaction ledger. We expressed and proved two properties of the backbone protocol — “common prefix” and “chain quality” — and showed how they can be used as foundations for designing Byzantine agreement and robust public transaction ledger protocols. Our results show that an honest majority among the participants suffices, assuming the network synchronizes much faster than the proof of work rate ($f \rightarrow 0$ in our notation) and the proper inputs (e.g., transactions) are available to the honest majority²⁰, while the bound on the adversary for honest parties to reach agreement degenerates as f gets larger.

While these are encouraging results, we have demonstrated deviations that are of concern for the proper operation of Bitcoin. Importantly, we show that as the network ceases to synchronize fast enough compared to the proof-of-work rate (i.e., the worst-case time that takes honest players to “hear” each other becomes substantial compared to the time it takes to solve a proof of work), the honest majority property ceases to hold and the bound offered by our analysis that is required to obtain a robust transaction ledger approaches 0 as f approaches 1. Note that the effects of bad synchronization is in the maintenance of the common prefix property, which is the critical property for showing agreement.

A second important concern is regarding the chain quality property, where our results show that if an adversary controls a hashing power corresponding to β then the ratio of the blocks it can contribute to the blockchain while bounded can be strictly bigger than β . When β gets close to $1/2$, our bounds show that the honest players’ contributions approach 0 in our security model.

The above caveats in the two basic properties of the backbone have repercussions on the Persistence and Liveness properties of the Bitcoin ledger. Firstly, they illustrate that fast information propagation amongst honest players is essential for transaction persistence. Secondly, they show that transaction liveness becomes more fragile as the adversarial power gets close to $1/2$. Note that we achieve Liveness for any adversarial bound less than $1/2$ but we do not assume any upper bound on the number of transactions that may be inserted in a block²¹; it is obvious that the fewer blocks the honest miners get into the blockchain the harder may be for a transaction to get through. Furthermore, the fact that chain quality demonstrably fails to preserve a one-to-one correspondence between a party’s hashing power and the ratio of its contributions to the ledger point to the fact that Bitcoin’s rewarding mechanism is not incentive compatible (cf. [ES14]). Assuming the hashing power of the honest parties γ exceeds the adversary’s hashing power β by a factor λ , we show that the adversary’s contributions to the ledger are bounded by $1/\lambda$ — and this is tight due to the selfish mining strategy we present. In this way our results flesh out the incentive compatibility problems of the Bitcoin backbone, but (on a more positive note) they also point to the fact that honest hashing power majority is sufficient to maintain the public ledger (under favorable network conditions), and hence suggest that Bitcoin can work as long as the majority of the miners *want it to work* (without taking into account the rationality of their decision).

Interesting open questions include the analysis of the Bitcoin backbone protocol in a rational setting as opposed to honest/malicious, its analysis with respect to concurrent/universal composition as opposed to standalone and the development of backbone modifications that improve its characteristics in terms of common prefix and chain quality. In terms of the ledger application, transaction processing times (i.e., reducing the wait time parameter u in the Liveness property) is

²⁰Our formalization is a way to formally express what perhaps was Nakamoto’s intuition when he wrote about Bitcoin that “it takes advantage of the nature of information being easy to spread but hard to stifle” [Nak09].

²¹In the current Bitcoin implementation there is an upper bound of 1MB for blocks, hence the number transactions per block is limited.

also an interesting question with implications to practice (since real world payment systems benefit greatly from fast transaction confirmation and verification). In all these cases, our work offers a formal foundation that allows analyzing the security properties of “tweaks” on the backbone protocol (such as the randomization rule of [ES14] or the “GHOST” rule in [SZ13] used in ethereum²²) towards meeting the above goals.

Another set of interesting directions include the development of other applications that may be built on top of the backbone protocol such as secure multiparty computation with properties such as fairness and guaranteed output delivery (current works in this direction, e.g., [ADMM14, BK14a, BK14b], assume an idealized version of the Bitcoin system).

References

- [ADMM14] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Łukasz Mazurek. Secure multiparty computations on bitcoin. *IEEE Security and Privacy*, 2014.
- [AJK05] James Aspnes, Collin Jackson, and Arvind Krishnamurthy. Exposing computationally-challenged Byzantine impostors. Technical Report YALEU/DCS/TR-1332, Yale University Department of Computer Science, July 2005.
- [Bac97] Adam Back. Hashcash. <http://www.cypherspace.org/hashcash>, 1997.
- [BCL⁺11] Boaz Barak, Ran Canetti, Yehuda Lindell, Rafael Pass, and Tal Rabin. Secure computation without authentication. *J. Cryptology*, 24(4):720–760, 2011.
- [BDOZ12] Moshe Babaioff, Shahar Dobzinski, Sigal Oren, and Aviv Zohar. On bitcoin and red balloons. In Boi Faltings, Kevin Leyton-Brown, and Panos Ipeirotis, editors, *EC*, pages 56–73. ACM, 2012.
- [BG93] Piotr Berman and Juan A. Garay. Randomized distributed agreement revisited. In *Digest of Papers: FTCS-23, The Twenty-Third Annual International Symposium on Fault-Tolerant Computing, Toulouse, France, June 22-24, 1993*, pages 412–419. IEEE Computer Society, 1993.
- [BK14a] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part II*, pages 421–439, 2014.
- [BK14b] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to incentivize correct computations. *ACM CCS 2014*, 2014.
- [BO83] Michael Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols (extended abstract). In Robert L. Probert, Nancy A. Lynch, and Nicola Santoro, editors, *PODC*, pages 27–30. ACM, 1983.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *CCS '93, Proceedings of the 1st ACM Conference on Computer and Communications Security, Fairfax, Virginia, USA, November 3-5, 1993.*, pages 62–73, 1993.
- [BSCG⁺14] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. *IACR Cryptology ePrint Archive*, 2014:349, 2014.
- [Can00] Ran Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
- [Cha82] David Chaum. Blind signatures for untraceable payments. pages 199–203, 1982.
- [Cun13] Cunicula. Why doesn’t bitcoin use a tiebreaking rule when comparing chains of equal length? <https://bitcointalk.org/index.php?topic=355644.0>. 2013.

²²<https://www.ethereum.org/>

- [DN92] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In Ernest F. Brickell, editor, *CRYPTO*, volume 740 of *Lecture Notes in Computer Science*, pages 139–147. Springer, 1992.
- [DW13] Christian Decker and Roger Wattenhofer. Information propagation in the bitcoin network. In *P2P*, pages 1–10. IEEE, 2013.
- [ES14] Ittay Eyal and Emin Gun Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography*, 2014.
- [FG03] Matthias Fitzi and Juan A. Garay. Efficient player-optimal protocols for strong and differential consensus. In Elizabeth Borowsky and Sergio Rajsbaum, editors, *PODC*, pages 211–220. ACM, 2003.
- [FLP85] Michael J. Fischer, Nancy A. Lynch, and Mike Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.
- [FM97] Peasech Feldman and Silvio Micali. An optimal probabilistic protocol for synchronous byzantine agreement. *SIAM J. Comput.*, 26(4):873–933, 1997.
- [GKKZ11] Juan A. Garay, Jonathan Katz, Ranjit Kumaresan, and Hong-Sheng Zhou. Adaptively secure broadcast, revisited. In Cyril Gavoille and Pierre Fraigniaud, editors, *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011*, pages 179–186. ACM, 2011.
- [HZ10] Martin Hirt and Vassilis Zikas. Adaptively secure broadcast. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*, pages 466–485. Springer, 2010.
- [JB99] Ari Juels and John G. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *NDSS*. The Internet Society, 1999.
- [Kin13] Sunny King. Primecoin: Cryptocurrency with prime number proof-of-work. <http://primecoin.io/bin/primecoin-paper.pdf>, July 2013.
- [KK09] Jonathan Katz and Chiu-Yuen Koo. On expected constant-round protocols for byzantine agreement. *Journal of Computer and System Sciences*, 75(2):91 – 112, 2009.
- [LSP82] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- [ML14] Andrew Miller and Joseph J. LaViola. Anonymous byzantine consensus from moderately-hard puzzles: A model for bitcoin. University of Central Florida. Tech Report, CS-TR-14-01, April 2014.
- [Nak08a] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <http://bitcoin.org/bitcoin.pdf>, 2008.
- [Nak08b] Satoshi Nakamoto. “the proof-of-work chain is a solution to the byzantine generals’ problem”. The Cryptography Mailing List, <https://www.mail-archive.com/cryptography@metzdowd.com/msg09997.html>, November 2008.
- [Nak09] Satoshi Nakamoto. Bitcoin open source implementation of p2p currency. <http://p2pfoundation.ning.com/forum/topics/bitcoin-open-source>, February 2009.
- [Nei94] Gil Neiger. Distributed consensus revisited. *Inf. Process. Lett.*, 49(4):195–201, 1994.
- [OB08] Michael Okun and Amnon Barak. Efficient algorithms for anonymous byzantine agreement. *Theor. Comp. Sys.*, 42(2):222–238, January 2008.
- [Oku05a] Michael Okun. Agreement among unacquainted byzantine generals. In Pierre Fraigniaud, editor, *DISC*, volume 3724 of *Lecture Notes in Computer Science*, pages 499–500. Springer, 2005.

- [Oku05b] Michael Okun. Distributed computing among unacquainted processors in the presence of byzantine distributed computing among unacquainted processors in the presence of byzantine failures. Ph.D. Thesis Hebrew University of Jerusalem, 2005.
- [PSL80] Marshall C. Pease, Robert E. Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.
- [Rab83] Michael O. Rabin. Randomized byzantine generals. In *FOCS*, pages 403–409. IEEE Computer Society, 1983.
- [RSW96] R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, Cambridge, MA, USA, 1996.
- [SZ13] Yonatan Sompolinsky and Aviv Zohar. Accelerating bitcoin’s transaction processing. fast money grows on trees, not chains. *IACR Cryptology ePrint Archive*, 2013:881, 2013.

A Model (cont’d)

All entities involved in a protocol execution are described by *interactive Turing machines (ITMs)*; an instance of an execution is described by a collection of instances of these *ITM instances* or *ITIs*. The set of all polynomial-time ITMs is denoted by ITM . We specify our protocols Π in a way similar to Canetti’s synchronous model for “real-world” execution [Can00].

The entities involved in an execution are the set of parties $\mathcal{P} = \{P_1, \dots, P_n\}$ running the protocol Π , the adversary denoted by \mathcal{A} and the environment denoted by \mathcal{Z} that provides inputs to the parties and receives outputs from them. The communication of ITIs is defined in [Can00] as having the sending ITI write directly to a communication input tape of the receiving ITI; the set of potential receivers is not fixed when an ITI is instantiated. This means that we have to include the messages generated by an ITI in its view explicitly. We denote by $\text{RECEIVE}()$ the communication input tape of each party. The communication between parties running the protocol Π will be *non-authenticated* but reliable²³. We capture this by allowing the adversary \mathcal{A} to modify the sender information in any message that is written to the communication input tape of an ITI. The adversary \mathcal{A} is allowed to corrupt parties P_1, \dots, P_n by writing a special message in their input communication tape. When that happens the adversary has complete control of the ITI of the corrupted party. We bound the number of corruptions by $t \leq n$.

Following [Can00] we assume a synchronous mode of interaction between parties. In particular, the protocol execution will be divided in rounds, and in each round all parties will be given the opportunity to act. The order of party activation is controlled by the adversary \mathcal{A} which is adaptive and rushing.

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. The concatenation of the view of all parties $\langle \text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{P_i}(\kappa, z) \rangle_{i=1, \dots, n}$ is denoted by $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}$.

The parties’ inputs are provided by the environment \mathcal{Z} which also receives the parties’ outputs. Parties that receive no input from the environment remain inactive, in the sense that they will not act when their turn comes in each round. The environment may provide input to a party at any

²³This parallels communication over TCP/IP in the Internet where, assuming all intermediate nodes in a communication path are honest and functioning, messages between honest parties are delivered reliably, nevertheless malicious parties may “spoof” the source of a message they transmit and make it appear as originating from an arbitrary party (including an honest party) in the view of the receiver.

round and may also modify that input from round to round. We denote by $\text{INPUT}()$ the input tape of each party.

We will consider ITM's that may have access to an oracle $H(\cdot)$ and perform a number of queries q per round where q is a function of the security parameter κ ; we refer to such parties as q -bounded. The adversary \mathcal{A} is allowed to perform $t \cdot q$ queries per round where $t \leq n$ is the number of corrupted parties. The environment \mathcal{Z} is not permitted any queries to $H(\cdot)$. The rationale for this is that we would like to bound the ‘‘CPU power’’ [Nak08a] of the adversary to be proportionate to the number of parties it controls while making it infeasible for the adversary be aided by external sources in the environment. We express this restriction on the environment, the parties and the adversary as the q -bounded setting. The view of the parties participating in the protocol will be denoted by $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{P, H(\cdot)}(\kappa, q, z)$ and the concatenation of all parties' views by $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{H(\cdot)}(\kappa, q, z)$.

We now define formally a property of an execution which is expressed as a predicate over the view of all parties in \mathcal{P} .

Definition 20. Given a predicate Q and a bound q , we say that *the protocol Π satisfies property Q in the q -bounded setting* provided that for all polynomial-time \mathcal{Z}, \mathcal{A} and strings z the probability that $Q(\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^{H(\cdot)}(\kappa, q, z))$ is false is negligible in κ .

Note that we will only consider properties that are polynomial-time computable predicates.

B Useful Inequalities

We will require the following inequalities.

Fact 1 (Bernoulli's inequality). For $q \geq 1$ and $0 \leq p \leq 1$, $(1 - p)^q \geq 1 - pq$.

Fact 2. For any real $\alpha > 0$, $1 - \alpha < e^{-\alpha} < 1 - \alpha + \frac{\alpha^2}{2}$.

Theorem 21 (Chernoff bounds). Suppose $\{X_i : i \in [n]\}$ are mutually independent Boolean random variables, with $\Pr[X_i = 1] = p$, for all $i \in [n]$. Let $X = \sum_{i=1}^n X_i$ and $\mu = pn$. Then, for any $\delta \in (0, 1]$,

$$\Pr[X \leq (1 - \delta)\mu] \leq e^{-\delta^2 \mu / 2} \quad \text{and} \quad \Pr[X \geq (1 + \delta)\mu] \leq e^{-\delta^2 \mu / 3}.$$