Bitcoin as a Transaction Ledger: A Composable Treatment*

Christian Badertscher[†], Ueli Maurer[†], Daniel Tschudi[†], and Vassilis Zikas[‡**]

[†] {christian.badertscher,maurer,tschudid}@inf.ethz.ch, ETH Zurich
[‡] vzikas@cs.rpi.edu, RPI

Abstract. Bitcoin is one of the most prominent examples of a distributed cryptographic protocol that is extensively used in reality. Nonetheless, existing security proofs are property-based, and as such they do not support composition.
In this work we put forth a universally composable treatment of the Bitcoin protocol. We specify the goal that Bitcoin aims to achieve as a ledger functionality in the (G)UC model of Canetti et al. [TCC'07]. Our ledger functionality is weaker than the one recently proposed by Kiayias, Zhou, and Zikas [EUROCRYPT'16], but unlike the latter suggestion, which is arguably not implementable given the Bitcoin assumptions, we prove that the one proposed here is securely UC realized under standard assumptions by an appropriate abstraction of Bitcoin as a UC protocol. We further show how known property-based approaches can be cast as special instances of our treatment and how their underlying assumptions can be cast in (G)UC without restricting the environment or the adversary.

1 Introduction

Since Nakamoto first proposed Bitcoin as a decentralized cryptocurrency [Nak08], several works have focused on analyzing and/or predicting its behavior under different attack scenarios [BDOZ11, ES14, Eya15, Zoh15, SZ15, KKKT16, PS16]. However, a core question remained:

What security goal does Bitcoin achieve under what assumptions?

An intuitive answer to this question was already given in Nakamoto's original white paper [Nak08]: Bitcoin aims to achieve some form of consensus on a set of valid transactions. The core difference of this consensus mechanism with traditional consensus [LSP82, Lam98, Lam02, Rab83] is that it does not rely on having a known (permissioned) set of participants, but everyone can join and leave at any point in time. This is often referred to as the *permissionless* model. Consensus in this model is achieved by shifting from the traditional assumptions on the fraction of cheating versus honest participants, to assumptions on the collective computing power of the cheating participants compared to the total computing power of the parties that support the consensus mechanism. The core idea is that in order for a party's action to affect the system's behavior, it needs to prove that it is investing sufficient computing resources. In Bitcoin, these resources are measured by means of solutions to a presumably computation-intensive problem.

Although the above idea is implicit in [Nak08], a formal description of Bitcoin's goal had not been proposed or known to be achieved (and under what assumptions) until the recent works of Garay, Kiayias, and Leonardos [GKL15] and Pass, Seeman, and shelat [PSS17]. In a nutshell, these works set forth models of computation and, in these models, an abstraction of Bitcoin as a distributed protocol, and proved that the output of this protocol satisfies certain security properties, for example the *common prefix* [GKL15] or consistency [PSS17] property. This property confirms—under the assumption that not too much of the total computing power of the system is invested in breaking it—a heuristic argument used by the Bitcoin specification: if some block makes it deep enough into the blockchain of an honest party, then it will eventually make it into the blockchain of every honest party and will never be reversed.[1] In addition to the common prefix property, other quality properties of the output of the abstracted blockchain protocol were also defined and proved. A more detailed description of the security properties and a comparison of the assumptions in [GKL15] and [PSS17] is included in Section 4.4.

* This is an extended version (updated February 14, 2018) of a paper that appeared at the 37th International Cryptology Conference (CRYPTO 2017). Proceedings version available at https://doi.org/10.1007/978-3-319-63688-7_11.
** Research supported in part by IOHK.

[1] In the original Bitcoin heuristic "deep enough" is defined as six blocks, whereas in these works it is defined as linear in an appropriate security parameter.

Bitcoin as a service for cryptographic protocols. Evidently, the main use of the Bitcoin protocol is as a decentralized monetary system with a payment mechanism, which is what it was designed for. And although the exact economic forces that guide its sustainability are still being researched, and certain rational models predict it is not a stable solution, it is a fact that Bitcoin has not met any of these pessimistic predictions for several years and it is not clear it ever will do. And even if it does, the research community has produced and is testing several alternative decentralized cryptocurrencies, e.g., [MGGR13, BCG$^+$14, But13], that are more functional and/or resilient to theoretic attacks than Bitcoin. Thus, it is reasonable to assume that decentralized cryptocurrencies are here to stay.

This leads to the natural questions of how one can use this new reality to improve the security and/or efficiency of cryptographic protocols? First answers to this question were given in [ADMM14a, ADMM14b, BK14, KVV16, KB16, KMB15, KB14, AD15] where it was shown how Bitcoin can be used as a punishment mechanism to incentivize honest behavior in higher level cryptographic protocols such as fair lotteries, poker, and general multi-party computation.

But in order to formally define and prove the security of the above constructions in a widely accepted cryptographic framework for multi-party protocols, one needs to define what it means for these protocols to be run in a world that gives them access to the Bitcoin network as a resource to improve their security. In other words, the question now becomes:

What functionality can Bitcoin provide to cryptographic protocols?

To address this question, Bentov and Kumaresan [BK14] introduced a model of computation in which protocols can use a punishment mechanism to incentivize adversaries to adhere to their protocol instructions. As a basis, they use the universal composition framework of Canetti [Can01], but the proposed modifications do not support composition and it is not clear how standard UC cryptographic protocols can be cast as protocols in that model.

In a different direction, Kiayias, Zhou, and Zikas [KZZ16] connected the above question with the original question of Bitcoin's security goal. More concretely, they proposed identifying the resource that Bitcoin (or other decentralized cryptocurrencies) offers to cryptographic protocols as its security goal, and expressing it in a standard language compatible with the existing literature on cryptographic mulit-party protocols. More specifically, they modeled the ideal guarantees as a transaction-ledger functionality in the universal composition framework. To be more precise, the ledger of [KZZ16] is formally a global setup in the (extended) GUC framework of Canetti et al. [CDPW07].

In a nutshell, the ledger proposed by [KZZ16] corresponds to a trusted third party which keeps a state of blocks of transactions and makes it available, upon request, to any party. Furthermore, it accepts messages/transactions from any party and records them as long as they pass an appropriate validation procedure that depends on the above publicly available state as well as other registered messages. Periodically, this ledger puts the transactions that were recently registered into a block and adds them into the state. The state is available to everyone. As proved in [KZZ16], giving multi-party protocols access to such a transaction-ledger functionality allows for formally capturing, within the composable(G)UC framework, the mechanism of leveraging security loss with coins. The proposed ledger functionality guarantees in an ideal manner all properties that one could expect from Bitcoin and encompasses the properties in [GKL15, PSS17]. Therefore, it is natural to postulate that it is a candidate for defining the security goal of Bitcoin (and potentially other decentralized cryptocurrencies). However, the ledger functionality proposed by [KZZ16] was not accompanied by a security proof that any of the known cryptocurrencies implements it.

However, as we show, despite being a step in the right direction, the ledger proposed in [KZZ16] cannot be realized under standard assumptions about the Bitcoin network. On the positive side, we specify a new transaction ledger functionality which still guarantees all properties postulated in [GKL15, PSS17], and prove that a reasonable abstraction of the Bitcoin protocol implements this ledger. In our construction, we describe Bitcoin as a UC protocol which generalizes both the protocols proposed in [GKL15, PSS17]. Along the way we identify the assumptions in each of [GKL15, PSS17] by devising a compound way of capturing such assumptions in UC, which enables us to compare their strengths.

Related Literature. The security of Bitcoin as a cryptographic protocol was previously studied by Garay, Kiayias, and Leonardos [GKL15] and by Pass, Seeman, and shelat [PSS17] who proposed and analyzed an abstraction of the core of the Bitcoin protocol in a *property-based* manner. As such, the treatment of [GKL15, PSS17] does not offer composable security guarantees. More recently, Kiayias, Zhou, and Zikas [KZZ16] proposed capturing the security goal and resource implemented by Bitcoin by means of a shared transaction-ledger functionality in the universal composition with global setup (GUC) framework of Canetti et al. [CDPW07]. However, the proposed ledger-functionality is too strong

to be implementable by Bitcoin. We refer the interested reader to Section A of the appendix for the basic elements of these works, where we also discuss simulation-based security and its advantages. A formal comparison of our treatment with [GKL15, PSS17], which indicates how both these protocols and definitions can be captured as special cases of our security definition, is also given in Section 4.4.

Our Results. We put forth the first universally composable (simulation-based) proof of security of Bitcoin in the (G)UC model of Canetti et at. [CDPW07]. We observe that the ledger functionality proposed by Kiayias et al. [KZZ16] is too strong to be implemented by the Bitcoin protocol—in fact, by any protocol in the permissionless setting, which uses network assumptions similar to Bitcoin. Intuitively, the reason is that the functionality allows too little interference of the simulator with its state, making it impossible to emulate adversarial attacks that result, e.g., in the adversary inserting only transactions coming from parties it wants or that result in parties holding chains of different length.

Therefore, we propose an alternative ledger functionality $\mathcal{G}_{\text{LEDGER}}$ which shares certain design properties with the proposal in [KZZ16] but which can be provably implemented by the Bitcoin protocol. As in [KZZ16], our proposed functionality can be used as a global setup to allow protocols with different sessions to make use of it, thereby enabling the ledger to be cast as shared among any protocol that wants to use it. The ledger is parametrized by a generic transaction validation predicate which enables it to capture decentralized blockchain protocols beyond Bitcoin. Our functionality allows for parties/miners to join and or leave the computation and allows for adaptive corruption.

Having defined our ledger functionality we next prove that for an appropriate validation predicate $\mathcal{G}_{\text{LEDGER}}$ is implemented by Bitcoin assuming that miners which deviate from the Bitcoin protocol do not control a majority of the total hashing power at any point. To this end, we describe an abstraction of the Bitcoin protocol as a synchronous UC protocol. Our protocol construction follows a structure which generalizes both [GKL15, PSS17]—as we argue, the protocols described in these works can be captured as instances of our protocols. The difference between these two instances is the network assumption that is used—more precisely, the assumption about knowledge on the network delay—and the assumption on the number of queries per round. To capture these assumptions in UC, we devise a methodology to formulate functionality wrappers to capture assumptions, and discuss the implications of such a method in preserving universal composability.

We design our protocol to work over a network which basically consists of bounded-delay channels, where similar to the protocol in [PSS17], the miners are not aware of (an upper bound on) the actual delay that the network induces. We argue that such a network is strictly weaker than a network with known bounded delay which, as we argue, is implicit in the synchrony assumptions of [GKL15] (cf. Remark 1.). Notwithstanding, unlike previous works, instead of starting from a complete network that offers multicast, we explain how such a network could be implemented by running the message-diffusion mechanism of the Bitcoin network (which is run over a lower level network of unicast channels). Intuitively, this network is built by every miner, upon joining the system, choosing some existing miners of its choice to use them as relay-nodes.

Our security proof proposes a useful modularization of the Bitcoin protocol. Concretely, we first identify the part of the Bitcoin code which intuitively corresponds to the lottery aspect, provide an ideal UC functionality that reflects this lottery aspect, and prove that this part of the Bitcoin code realizes the proposed functionality. We then analyze the remainder of the protocol in the simpler world where the respective code that implements the lottery aspect is replaced by invocations of the corresponding functionality. Using the UC composition theorem, we can then immediately combine the two parts into a proof of the full protocol.

As is the case with the so-called *backbone* protocol from [GKL15] our above UC protocol description of Bitcoin relies only on proofs of work and not on digital signatures. As a result, it implements a somewhat weaker ledger, which does not guarantee that transactions submitted by honest parties will eventually make it into the blockchain.[2] As a last result, we show that (similarly to [GKL15]) by incorporating public-key cryptography, i.e., taking signatures into account in the validation predicate, we can implement a stronger ledger that ensures that transactions issued by honest users—i.e., users who do not sign contradicting transactions and who keep their signing keys for themselves—are guaranteed to be eventually included into the blockchain. The fact that our protocol is described in UC makes this a straight-forward, modular construction using the proposed transaction ledger as a hybrid. In particular, we do not need to consider the specifics of the Bitcoin protocol in the proof of this step. This also allows us to identify the maximum (worst-case) delay a user needs to wait before being guaranteed to see its transaction on the blockchain and be assured that it will not be inverted.

[2] We formulate a weakened guarantee, which we then amplify using digital signatures.

2 A Composable Model for Blockchain Protocols in the Permissionless Model

In this section we describe our (G)UC-based model of execution for the Bitcoin protocol. We remark that providing such a formal model of execution forces us to make explicit all the implicit assumptions from previous works. As we lay down the theoretical framework, we will also discuss these assumptions along with their strengths and differences.

Bitcoin miners are represented as players—formally Interactive Turing Machine instances (ITIs)—in a multi-party computation. They interact which each other by exchanging messages over an unauthenticated multicast network with eventual delivery (see below) and might make queries to a common random oracle. We will assume a central adversary $\mathcal{A}$ who gets to corrupt miners and might use them to attempt to break the protocol's security. As is common in (G)UC, the resources available to the parties are described as hybrid functionalities. Before we provide the formal specification of such functionalities, we first discuss a delicate issue that relates to the set of parties (ITIs) that might interact with an ideal functionality.

Functionalities with dynamic party sets. In many UC functionalities, the set of parties is defined upon initiation of the functionality and is not subject to change throughout the lifecycle of the execution. Nonetheless, UC does provide support for functionalities in which the set of parties that might interact with the functionality is dynamic. In fact, this dynamic nature is an inherent feature of the Bitcoin protocol—where miners come and go at will. In this work we make this explicit by means of the following mechanism: All the functionalities considered here include the following instructions that allow honest parties to join or leave the set $\mathcal{P}$ of players that the functionality interacts with, and inform the adversary about the current set of registered parties:[3]

- Upon receiving $(\textsc{register}, \mathrm{sid})$ from some party p_i (or from $\mathcal{A}$ on behalf of a corrupted p_i), set $\mathcal{P} = \mathcal{P} \cup \{p_i\}$. Return $(\textsc{register}, \mathrm{sid}, p_i)$ to the caller.
- Upon receiving $(\textsc{de-register}, \mathrm{sid})$ from some party $p_i \in \mathcal{P}$, the functionality sets $\mathcal{P} := \mathcal{P} \setminus \{p_i\}$ and returns $(\textsc{de-register}, \mathrm{sid}, p_i)$ to p_i.
- Upon receiving $(\textsc{is-registered}, \mathrm{sid})$ from some party p_i, return $(\textsc{register}, \mathrm{sid}, b)$ to the caller, where the bit b is 1 if and only if $p_i \in \mathcal{P}$.[4]
- Upon receiving $(\textsc{get-registered}, \mathrm{sid})$ from $\mathcal{A}$, the functionality returns $(\textsc{get-registered}, \mathrm{sid}, \mathcal{P})$ to $\mathcal{A}$.

For simplicity in the description of the functionalities, for a party $p_i \in \mathcal{P}$ we will use p_i to refer to this party's ID.

In addition to the above registration instructions, global setups, i.e., shared functionalities that are available both in the real and in the ideal world and allow parties connected to them to share state [CDPW07], allow also UC functionalities to register with them.[5] Concretely, global setups include, in addition to the above party registration instructions, two registration/de-registration instructions for functionalities:

- Upon receiving $(\textsc{register}, \mathrm{sid}_C)$ from a functionality $\mathcal{F}$, set $F := F \cup \{\mathcal{F}\}$.
- Upon receiving $(\textsc{de-register}, \mathrm{sid}_C)$ from a functionality $\mathcal{F}$, set $F := F \setminus \{\mathcal{F}\}$.
- Upon receiving $(\textsc{get-registered-f}, \mathrm{sid}_C)$ from $\mathcal{A}$, return $(\textsc{get-registered-f}, \mathrm{sid}_C, F)$ to $\mathcal{A}$.

The above four (or seven in case of global setups) instructions will be part of the code of *all* ideal functionalities considered in this work. However, to keep the description simpler we will omit these instructions from the formal descriptions. We are now ready to formally describe each of the available functionalities.

[3] Note that making the set of parties dynamic means that the adversary needs to be informed about which parties are currently in the computation so that he can chose how many (and which) parties to corrupt.

[4] Note that typcially a party knows whether it is registered at a functionality or not. However, it might be useful for another functionality to access this information via the dummy party corresponding to p_i. The exact dynamics of such an information exchange can be found in [CSV16, Section 2].

[5] Although we allow no communication between functionalities, we will allow functionalities to communicate with global setups. (They can use the interface of global setups to additional honest parties, which is anyway open to the environment.)

The Communication Network. In Bitcoin, parties/miners communicate over an incomplete network of asynchronous unauthenticated unidirectional channels. Concretely, every miner chooses a set of other miners as its immediate neighbors—typically by using some public information on IP addresses of existing miners—and uses its neighbors to send messages to all the miners in the Bitcoin network. This corresponds to multicasting the message[6]. This is achieved by a standard diffusion mechanism: The sender sends the message it wishes to multicast to all its neighbors who check that a message with the same content was not received before, and if this is the case forward it to their neighbors, who then do the same check, and so on. We make the following two assumptions about the communication channels in the above diffusion mechanism/protocol:

- They guarantee (reliable) delivery of messages within a delay parameter Δ, but are otherwise specified to be of asynchronous nature (see below) and hence no protocol can rely on timings regarding the delivery of messages. The adversary might delay any message sent through such a channel, but at most by Δ. In particular, the adversary cannot block messages. However, he can induce an arbitrary order on the messages sent to some party.
- The receiver gets no information other than the messages themselves. In particular, a receiver cannot link a message to its sender nor can he observe whether or not two messages were sent from the same sender.
- The channel offers no privacy guarantees. The adversary is given read access to all messages sent on the network.

Our formal description of communication with eventual delivery within the UC framework builds on ideas from [KMTZ13, BHMQU05, CGHZ16]. In particular, we capture such communication by assuming for each miner $p_j \in \mathcal{P}$ a multi-use *unicast* channel $\mathcal{F}_{\text{U-CH}}$ with receiver p_j, to which any miner $p_i \in \mathcal{P}$ can connect and input messages to be delivered to $p_j \in \mathcal{P}$. A miner connecting to the unicast channel with receiver p_j corresponds to the above process of looking up p_j and making him one of its access points. The unicast channel does not provide any information to its receiver about who else is using it. In particular, messages are buffered but the information of who is the sender is deleted; instead, the channel creates unique independent message IDs that are used as handles for the messages. Furthermore, the adversary—who is informed about both the content of the messages and about the handles—is allowed to delay messages by any finite amount, and allowed to deliver them in an arbitrary out-of-order manner.

To ensure that the adversary cannot arbitrarily delay the delivery of messages submitted by honest parties, we use the following idea: We first turn the UC channel-functionality to work in a "fetch message" mode, where the channel delivers the message to its intended recipient p_j if and only if p_j asks to receive it by issuing a special "fetch" command. If the adversary wishes to delay the delivery of some message with message ID mid, he needs to submit to the channel functionality an integer value T_{mid}—the *delay* for message with ID mid. This will have the effect that the channel ignores the next T_{mid} fetch attempts, and only then allows the receipt of the sender's message. Importantly, we require that the channel does not accept more than Δ accumulative delay on any message. To allow the adversary freedom in scheduling the delivery of messages, we allow him to input delays more than once, which are added to the current delay amount. If the adversary wants to deliver the message in the next activation, all he needs to do is submit a negative delay. Furthermore, we allow the adversary to schedule more than one messages to be delivered in the same "fetch" command. Finally, to ensure that the adversary is able to re-order such batches of messages arbitrarily, we allow $\mathcal{A}$ to send special $(\text{SWAP}, \text{mid}, \text{mid}')$ commands that have as an effect to change the order of the corresponding messages. The detailed specification of the described channels, denoted $\mathcal{F}_{\text{U-CH}}$ is given in Appendix B. Note that in the descriptions throughout the paper, for a vector $\vec{M}$ we denote by the symbol $||$ the operation which adds a new element to $\vec{M}$.

From Unicast to Multicast. As already mentioned, the Bitcoin protocol uses the above asynchronous-and-bounded-delay unicast network as a basis to achieve a multicast mechanism. A multicast functionality with bounded delay can be defined similarly to the above unicast channel. The main difference is that once a message is inserted it is recorded $|\mathcal{P}|$ times, once for each possible receiver. The adversary can add delays to any subset of messages, but again for any message the cumulative delay cannot exceed Δ. He is further allowed to do partial and inconsistent multicasts, i.e., where different messages are sent to different parties. This is the main difference of such a multicast network from a broadcast network. The detailed specification of the corresponding functionality $\mathcal{F}_{\text{N-MC}}$ is similar to that of $\mathcal{F}_{\text{U-CH}}$ and can be

[6] In [GKL15] this mechanism is referred to as "broadcast"; here, we use multicast to make explicit the fact that this primitive is different from a standard Byzantine-agreement-type broadcast, in that it does not guarantee any consistency for a malicious sender.

found in Appendix B, where we also show how the simple round-based diffusion mechanism can be used to implement a multicast mechanism from unicast channels as long as the corresponding network among honest parties stays strongly connected. (A network graph is strongly connected if there is a directed path between any two nodes in the network, where the unicast channels are seen as the directed edges from sender to receiver.)

The Random Oracle. As usual in cryptographic proofs, the queries to the hash function are modeled by assuming access to a random oracle (functionality) $\mathcal{F}_{\mathrm{RO}}$. This functionality is specified as follows: upon receiving a query $(\textsc{eval}, \mathrm{sid}, x)$ from a registered party, if x has not been queried before, a value y is chosen uniformly at random from $\{0,1\}^\kappa$ (for security parameter κ) and returned to the party (and the mapping (x, y) is internally stored). If x has been queried before, the corresponding y is returned. For completeness we include the functionality in B.

Synchrony. Katz et al. [KMTZ13], proposed a methodology for casting synchronous protocols in UC by assuming they have access to an ideal functionality $\mathcal{G}_{\textsc{clock}}$, *the clock*, that allows parties to ensure that they proceed in synchronized rounds. Informally, the idea is that the clock keeps track of a round variable whose value the parties can request by sending it $(\textsc{clock-read}, \mathrm{sid}_C)$. This value is updated only once all honest parties sent the clock a $(\textsc{clock-update}, \mathrm{sid}_C)$ command.

Given such a clock, the authors of [KMTZ13] describe how synchronous protocols can maintain their necessary round structure in UC: For every round ρ each party first executes all its round-ρ instructions and then sends the clock a CLOCK-UPDATE command. Subsequently, whenever activated, it sends the clock a CLOCK-READ command and does not advance to round $\rho + 1$ before it sees the clocks variable being updated. This ensures that no honest party will start round $\rho + 1$ before every honest party has completed round ρ. In [KZZ16], this idea was transfered to the (G)UC setting, by assuming that the clock is a global setup. This allows for different protocols to use the same clock and is the model we will also use here. For completeness we include the clock functionality in Appendix B.

As argued in [KMTZ13], in order for an eventual-delivery (aka guaranteed termination) functionality to be UC implementable by a synchronous protocol it needs to keep track of the number of activations that an honest party gets—so that it knows when to generate output for honest parties. This requires that the protocol itself, when described as a UC interactive Turing-machine instance (ITI), has a predictable behavior when it comes to the pattern of activations that it needs before it sends the clock an update command. We capture this property in a generic manner in Definition 1.

In order to make the definition better accessible, we briefly recall the mechanics of activations in UC. In a UC protocol execution, an honest party (ITI) gets activated either by receiving an input from the environment, or by receiving a message from one of its hybrid-functionalities (or from the adversary). Any activation results in the activated ITI performing some computation on its view of the protocol and its local state and ends with either the party sending a message to some of its hybrid functionalities or sending an output to the environment, or not sending any message. In either of this case, the party loses the activation.[7]

For any given protocol execution, we define the *honest-input sequence* $\vec{\mathcal{I}}_H$ to consist of all inputs that the environment gives to honest parties in the given execution (in the order that they were given) along with the identity of the party who received the input. For an execution in which the environment has given m inputs to the honest parties in total, $\vec{\mathcal{I}}_H$ is a vector of the form $((x_1, \mathrm{pid}_1), \ldots, (x_m, \mathrm{pid}_m))$, where x_i is the i-th input that was given in this execution, and pid_i is the corresponding party who received this input. We further define the *timed honest-input sequence,* denoted as $\vec{\mathcal{I}}_H^T$, to be the honest-input sequence augmented with the respective clock time when an input was given. If the timed honest-input sequence of an execution is $\vec{\mathcal{I}}_H^T = ((x_1, \mathrm{pid}_1, \tau_1), \ldots, (x_m, \mathrm{pid}_m, \tau_m))$, this means that $((x_1, \mathrm{pid}_1), \ldots, (x_m, \mathrm{pid}_m))$ is the honest-input sequence corresponding to this execution, and for each $i \in [n]$, τ_i is the time of the global clock when input x_i was handed to pid_i.

Definition 1. *A $\mathcal{G}_{\textsc{clock}}$-hybrid protocol Π has a* predictable synchronization pattern *iff there exist an algorithm* $\mathsf{predict\text{-}time}_\Pi(\cdot)$ *such that for any possible execution of Π (i.e., for any adversary and environment, and any choice of random coins) the following holds: If $\vec{\mathcal{I}}_H^T = ((x_1, \mathrm{pid}_1, \tau_1), \ldots, (x_m, \mathrm{pid}_m, \tau_m))$ is the corresponding timed honest-input sequence for this execution, then for any $i \in [m-1]$:*

$$\mathsf{predict\text{-}time}_\Pi((x_1, \mathrm{pid}_1, \tau_1), \ldots, (x_i, \mathrm{pid}_i, \tau_i)) = \tau_{i+1}.$$

As we argue, all synchronous protocol described in this work are designed to have a predictable synchronization pattern.

[7] In the latter case the activation goes to the environment by default.

Assumptions as UC Functionality Wrappers. In order to prove statements about cryptographic protocols one often makes assumptions about what the environment (or the adversary) can or cannot do. For example, a standard assumption in [GKL15, PSS17] is that in each round the adversary cannot do more calls to the random oracle than what the honest parties (collectively) can do. This can be captured by assuming a restricted environment and adversary which balances the amount of times that the adversary queries the random oracle. In a property-based treatment such as [GKL15, PSS17] this assumptions is typically acceptable.

However, in a simulation-based definition, restricting the class of adversaries and environments in a security statement means that we can no longer generically apply the composition theorem, which dismisses one of the major advantages of using simulation-based security in the first place. Therefore, instead of restricting the class of environments/adversaries, here we take a different approach to capture the fact that the adversary's access to the RO is restricted with respect to that of honest parties. In particular, we capture this assumption by means of a functionality wrapper that wraps the RO functionality and forces the above restrictions on the adversary, for example by assigning to each corrupted party at most q activations per round for a parameter q. To keep track of rounds the functionality registers with the global clock $\mathcal{G}_{\text{CLOCK}}$. For completeness we include the wrapped random oracle functionality $\mathcal{W}^q(\mathcal{F}_{\text{RO}})$ in Appendix B.

Remark 1 (Functionally Black-box Use of the Network (Delay)). A key difference between the models in [GKL15] and [PSS17] is that in the latter the parties do not know any bound on the delay of the network. In particular, although both models are in the synchronous setting, in [PSS17] a party in the protocol does not know when to expect a message which was sent to it in the previous round. Using terminology from [Ros12], the protocol uses the channel in a *functionally black-box* manner. Restricting to such protocols—a restriction which we also adopt in this work—is in fact implying a weaker assumption on the protocol than standard (known) bounded-delay channel. Intuitively the reason is that no such protocol can realize a bounded-delay network with a known upper bound (unless it sacrifices termination) since the protocol cannot decide whether or not the bound has been reached.

3 The Transaction-Ledger Functionality

In this section we describe our ledger functionality, denoted as $\mathcal{G}_{\text{LEDGER}}$, which can, for example, be achieved by (a UC version) of the Bitcoin protocol. As in [KZZ16], our ledger is parametrized by certain algorithms/predicates that allow us to capture a more general version of a ledger which can be instantiated by various cryptocurrencies. Since our abstraction of the Bitcoin protocol is in the synchronous model of computation (this is consistent with known approaches in the cryptographic literature), our ledger is also designed for this synchronous model. Nonetheless, several of our modeling choices are made with the foresight of removing or limiting the use of the clock and leaving room for less synchrony.

At a high level, our ledger $\mathcal{G}_{\text{LEDGER}}$ has a similar structure as the ledger proposed in [KZZ16]. Concretely, anyone (whether an honest miner or the adversary) might submit a transaction which is validated by means of a predicate Validate, and if it is found valid it is added to a buffer `buffer`. The adversary $\mathcal{A}$ is informed that the transaction was received and is given its contents.[8] Informally, this buffer also contains transactions that, although validated, are not yet deep enough in the blockchain to be considered out-of-reach for a adversary.[9] Periodically, $\mathcal{G}_{\text{LEDGER}}$ fetches some of the transactions in the buffer, and using an algorithm Blockify creates a block including these transactions and adds this block to its permanent state `state`, which is a data structure that includes the part of the blockchain the adversary can no longer change. This corresponds to the *common prefix* in [GKL15, PSS17]. Any miner or the adversary is allowed to request a read of the contents of the state.

This sketched specification is simple, but in order to have a ledger that can be implemented by existing blockchain protocols, we need to relax this functionality by giving the adversary more power to interfere with it and influence its behavior. Before sketching the necessary relaxations we discuss the need for a new ledger definition and it potential use as a global setup.

Remark 2 (Impossibility to realize the ledger of [KZZ16]). The main reasons why the ledger functionality in [KZZ16] is not realizable by known protocols under reasonable assumptions are as follows: first, their ledger guarantees that parties always obtain the same common state. Even with strong synchrony

[8] This is inevitable since we assume non-private communication, where the adversary sees any message as soon as it is sent, even if the sender and receiver are honest.

[9] E.g., in [KZZ16] the adversary is allowed to permute the contents of the buffer.

assumptions, this is not realizable since an adversary, who just mined a new block, is not forced to inform each party instantaneously (or at all) and thus could for example make parties observe different lengths of the same prefix. Second, the adversarial influence is restricted to permuting the buffer. This is too optimistic, as in reality the adversary can try to mine a new block and possibly exclude certain transactions. Also, this excludes any possibility to quantify quality. Third, letting the update rate be fixed does not adequately reflect the probabilistic nature of blockchain protocols.

Remark 3 (On the sound usage of a ledger as a global setup). As presented in [KZZ16], a UC ledger functionality $\mathcal{G}_{\text{LEDGER}}$ can be cast as a global setup [CDPW07] which allows different protocols to share state. This fact holds true for any UC functionality as stated in [CDPW07] and [CSV16]. Nonetheless, as pointed out in the recent work of Canetti, Shahaf, and Vald [CSV16], one needs to be extra careful when replacing a global setup by its implementation, e.g., in the case of $\mathcal{G}_{\text{LEDGER}}$ by the UC Bitcoin backbone protocol of Section 4. Indeed, such a replacement does not, in general, preserve a realization proof of some ideal functionality $\mathcal{F}$ that is conducted in a ledger-hybrid world, because the simulator in that proof might rely on specific capabilities that are not available any more after replacement (as the global setup is also replaced in the ideal world). The authors of [CSV16] provide a sufficient condition for such a replacement to be sound. This condition is generally too strong to be satisfied by any natural ledger implementation, which opens the question of devising relaxed sufficient conditions for sound replacements in an MPC context.[10] As this work focuses on the realization of ledger functionalities per se, we can treat $\mathcal{G}_{\text{LEDGER}}$ as a standard UC functionality.

In the following, we first review the necessary relaxations to obtain a realizable ledger. We conclude this section with the specification of our generic ledger functionality.

State-buffer validation. The first relaxation is with respect to the invariant that is enforced by the validation predicate Validate. Concretely, in [KZZ16] it is assumed that the validation predicate enforces that the buffer does not include conflicting transactions, i.e., upon receipt of a transaction, Validate checks that it is not in conflict with the state and the buffer, and if so the transaction is added to the buffer. However, in reality we do not know how to implement such a strong filter, as different miners might be working on different, potentially conflicting sets of transactions.[11] The only time when it becomes clear which of these conflicting transactions will make it into the state is once one of them has been inserted into a block which has made it deep enough into the blockchain (i.e., has become part of `state`). Hence, given that the buffer includes all transactions that might end up in the state, it might at some point include both conflicting transactions.

To enable us for a provably implementable ledger, in this work we take a different approach. The validate predicate will be less restrictive as to which transactions make it into the buffer. Concretely, at the very least, Validate will enforce the invariant that no single transaction in the buffer contradicts the state `state`, while different transactions in `buffer` might contradict each other. Looking ahead, a stronger version that is achievable by employing digital signatures (presented in Section 5), could enforce that no submitted transaction contradicts other submitted transactions. As in [KZZ16], whenever a new transaction x is submitted to $\mathcal{G}_{\text{LEDGER}}$, it is passed to Validate which takes as input a transaction and the current state and decides if x should be added to the buffer. Additionally, as `buffer` might include conflicts, whenever a new block is added to the state, the buffer (i.e., every single transaction in `buffer`) is re-validated using Validate and invalid transactions in `buffer` are removed. To allow for this re-validation to be generic, transactions that are added to the buffer are accompanied by certain metadata, i.e., the identity of the submitter, a unique transaction ID txid[12], or the time τ when x was received.

State update policies and security guarantees. The second relaxation is with respect to the rate and the form and/or origin of transactions that make it into a block. Concretely, instead of assuming that the state is extended in fixed time intervals, we allow the adversary to define when this update occurs. This is done by allowing the adversary, at any point, to propose what we refer to as the next-block candidate

[10] To give an example, a natural condition would be to require that the ideal-world adversary (or simulator) for $\mathcal{F}$ does only use the ledger to submit queries or reading the state, and plays the "dummy adversary" for queries by the environment that request the additional adversarial capabilities. The simulator in [KZZ16] is of this kind.

[11] This will be the case for transactions submitted by the adversary even when signatures are used to authenticate transactions.

[12] In Bitcoin, txidwould be the hash-pointer corresponding to this transaction. Note that the generic ledger can capture explicit guarantees on the ability or disability to link transactions, as this crucially depends on the concrete choice of an ID mechanism.

NxtBC. This is a data structure containing the contents of the next block that $\mathcal{A}$ wants to have inserted into the state. Leaving NxtBC empty can be interpreted as the adversary signaling that it does not want the state to be updated in the current clock tick.

Of course allowing the adversary to always decide what makes it into the state state, or if anything ever does, yields a very weak ledger. Intuitively, this would be a ledger that only guarantees the common prefix property [GKL15] but no liveness or chain quality. Therefore, to enable us to capture also stronger properties of blockchain protocols we parameterize the ledger by an algorithm ExtendPolicy that, informally, enforces a state-update policy restricting the freedom of the adversary to choose the next block and implementing an appropriate compliance-enforcing mechanism in case the adversary does not follow the policy. This enforcing mechanism simply returns a default policy-complying block using the current contents of the buffer. We point out that a good simulator for realizing the ledger will avoid triggering this compliance-enforcing mechanism, as this could result in an uncontrolled update of the state which would yield a potential distinguishing advantage. In other words, a good simulator, i.e., ideal-world adversary, always complies with the policy.

In a nutshell, ExtendPolicy takes the current contents of the buffer buffer, along with the adversary's recommendation NxtBC, and the *block-insertion times vector* $\vec{\tau}_{\texttt{state}}$. The latter is a vector listing the times when each block was inserted into state. The output of ExtendPolicy is a vector including the blocks to be appended to the state during the next state-extend time-slot (where again, ExtendPolicy outputting an empty vector is a signal to not extend). To ensure that ExtendPolicy can also enforce properties that depend on who inserted how many (or which) blocks into the state—e.g. the so-called *chain quality* property from [GKL15]—we also pass to it the timed honest-input sequence $\vec{\mathcal{I}}_H^T$ (cf. Section 2).

Some examples of how ExtendPolicy allows us to define ways that the protocol might restrict the adversary's interference in the state-update include the following properties from [GKL15]:

- *Liveness* corresponds to ExtendPolicy enforcing the following policy: If the state has not been extended for more that a certain number of rounds and the simulator keeps recommending an empty NxtBC, ExtendPolicy can choose some of the transactions in the buffer (e.g., those that have been in the buffer for a long time) and add them to the next block. Note that a good simulator or ideal-world adversary will never allow for this automatic update to happen and will make sure that he keeps the state extend rate within the right amount.
- *Chain quality* corresponds to ExtendPolicy enforcing the following policy: Every block proposal made by the simulator has to be associated with a special flag hFlag, where intuitively hFlag $= 1$ indicates that the proposal is generated using the process that an honest miner would follow. ExtendPolicy enforces two things: first, that block proposal indicating hFlag $= 1$ are frequent enough, and second that such proposals fulfill some specific quality properties (such as including all recent transactions). If these properties are not met, the ledger will define and add a default block to the state. [13] We point out that unlike the original chain-quality property from [GKL15], this policy does not enforce which miner should receive the reward for honest blocks and it is up to the simulator to do so (via the so-called coinbased transaction).[14]

In addition to the above standard properties, ExtendPolicy allows us to also capture additional security properties of various blockchain protocols, e.g., the fact that honest transactions eventually make it into a block and the fact that transactions with higher rewards make it into a block faster than others.

In Section 4 where we prove the security of Bitcoin, we will provide the concrete specification of Validate and ExtendPolicy for which the Bitcoin protocol realizes our ledger.

Output Slackness and Sliding Window of State Blocks. The common prefix property guarantees that blocks which are sufficiently deep in the blockchain of an honest miner will eventually be included in the blockchain of every honest miner. Stated differently, if an honest miner receives as output from the ledger a state state, every honest miner will eventually receive state as its output. However, in reality we cannot guarantee that at any given point in time all honest miners see exactly the same blockchain

[13] More technically, ExtendPolicy looks into the proposed-block sequence and identifies the blocks of state that where proposed by the simulator with hFlag set to 1 to deduce how long ago (in time or block-number) the last proposed block that made it into the chain had hFlag $= 1$.

[14] The actual Bitcoin protocol ensures that at the time when the block was created and circulated in the network the originator of the block was honest. Note that this does not mean that he is still honest when the block makes it into the state *unless* one considers static corruptions only (in which case one can indeed directly argue about the fraction of honest originators in the state). To make this difference is crucial to explicitly see the impact due to adaptive corruptions and was not made explicit in earlier versions of this work.

length; this is especially the case when network delays are incorporated into the model, but it is also true in the zero-delay model of [GKL15]. Thus it is unclear how `state` can be defined so that at any point all parties have the same view on it.

Therefore, to have a ledger implementable by standard assumptions we make the following relaxation: We interpret `state` as the view of the state of the miner with the longest blockchain. And we allow the adversary to define for every honest miner p_i a subchain $\texttt{state}_i$ of `state` of length $|\texttt{state}_i| = \texttt{pt}_i$ that corresponds to what p_i gets as a response when he reads the state of the ledger (formally, the adversary can fix a pointer $\texttt{pt}_i$). For convenience, we denote by $\texttt{state}|_{\texttt{pt}_i}$ the subchain of `state` that finishes in the $\texttt{pt}_i$-th block. Once again, to avoid over-relaxing the functionality to an unuseful setup, our ledger allows the adversary to only move the pointers forward and it forbids the adversary to define pointers for honest miners that are too far apart, i.e., more than `windowSize` state blocks. The parameter $\texttt{windowSize} \in \mathbb{N}$ denotes a core parameter of the ledger. In particular, the parameter `windowSize` reflects the similarity of the blockchain to the dynamics of a so-called *sliding window*, where the window of size `windowSize` contains the possible views of honest miners onto `state` and where the head of the window advances with the head of the `state`. In addition, it is convenient to express security properties of concrete blockchain protocols, including the properties discussed above, as assertions that hold within such a sliding window (for any point in time).

Synchrony. In order to keep the ideal execution indistinguishable from the real execution, the adversary should be unable to use the clock for distinguishing. Since in the ideal world when a dummy party receives a CLOCK-UPDATE-message for $\mathcal{G}_{\text{CLOCK}}$ it will forward it, the ledger needs to be responsible that the clock counter does not advance before all honest parties have received sufficiently many activations. This is achieved by the use of the function $\mathsf{predict\text{-}time}(\vec{\mathcal{I}}_H^T)$ (see Definition 1), which, as we show, is defined for our ledger protocol. This function allows $\mathcal{G}_{\text{LEDGER}}$ to predict when the protocol would update the round and ensure that it only allows the clock to advance if and only if the protocol would. Observe that the ledger can infer all protocol-relevant inputs/activations to honest parties and can therefore easily keep track of the honest inputs sequence $\vec{\mathcal{I}}_H^T$. In particular, in global UC communication between the ledger and the (shared) clock functionality is allowed to access the relevant information (namely via a dummy party as defined in [CSV16]).[15] As the other functions explained above, the function $\mathsf{predict\text{-}time}$ is a parameter of the (general) ledger functionality and hence needs to be instantiated when realizing a specific ledger such as the Bitcoin ledger (which is the topic of the next section).

A final observation is with respect to guarantees that the protocol (and therefore also the ledger) can give to recently registered honest parties, or to registered parties that get de-registered from the clock (temporarily, for instance). We will call miners *de-synchronized* if one of the above properties are fulfilled for this miner. We denote the set of such miners by $\mathcal{P}_{DS}$.

To provide more intuition, consider the following scenario: An honest party registers as miner in round r and waits to receive from honest parties the transactions to mine and the current longest blockchain. In Bitcoin, upon joining, the miner sends out a special request—we denote this here as a special NEW-MINER-message—and as soon as any party receives it, it responds with the set of transactions and longest blockchain it knows. Due to the network delay, the parties might take up to Δ rounds to receive the NEW-MINER notification, and their response might also take up to Δ rounds before it arrives to the new miner. However, because we do not make any assumption on honest parties knowing Δ (see Remark 1) they need to start mining as soon as a message arrives (otherwise they might wait indefinitely). But now the adversary, in the worst case, can make these parties mine on any block he wants and have them accept any valid chain he wants as the current state while they wait for the network's response: simply delay everything sent to these parties by honest miners by the maximum delay Δ, and instead, immediately deliver what you want them to work on. Thus, for the first $\texttt{Delay} := 2\Delta$ rounds[16] (where `Delay` is a parameter of our ledger) these parties are practically in the control of the adversary and their computing power is contributed to his.

The formal specification of our ledger functionality $\mathcal{G}_{\text{LEDGER}}$ is given in the following. Using standard notation, we write $[n]$ to denote the set $\{1, \dots, n\}$.

[15] In order to keep the description below simple, we omit how the ledger exactly infers $\vec{\mathcal{I}}_H^T$, but this is quite straightforward. In particular, the mechanism of [CSV16] allows to assume that the ledger knows whether a party is registered with the clock or not to deduce whether it is synchronized or de-synchronized.

[16] For technical reasons described in Section 4.1, Δ rounds in the protocol correspond to 2Δ clock-ticks.

Functionality $\mathcal{G}_{\text{LEDGER}}$

$\mathcal{G}_{\text{LEDGER}}$ is parametrized by four algorithms, Validate, ExtendPolicy, Blockify, and predict-time, along with two parameters: $\texttt{windowSize}, \texttt{Delay} \in \mathbb{N}$. The functionality manages variables $\texttt{state}$, $\texttt{NxtBC}$, $\texttt{buffer}$, τ_L, and $\vec{\tau}_{\texttt{state}}$, as described above. The variables are initialized as follows: $\texttt{state} := \vec{\tau}_{\texttt{state}} := \texttt{NxtBC} := \varepsilon$, $\texttt{buffer} := \emptyset$, $\tau_L = 0$.

The functionality maintains the set of registered parties $\mathcal{P}$, the (sub-)set of honest parties $\mathcal{H} \subseteq \mathcal{P}$, and the (sub-set) of de-synchronized honest parties $\mathcal{P}_{DS} \subset \mathcal{H}$ (following the definition in the previous paragraph). The sets $\mathcal{P}, \mathcal{H}, \mathcal{P}_{DS}$ are all initially set to $\emptyset$. When a new honest party is registered at the ledger, if it is registered with the clock already then it is added to the party sets $\mathcal{H}$ and $\mathcal{P}$ and the current time of registration is also recorded; if the current time is $\tau_L > 0$, it is also added to $\mathcal{P}_{DS}$. Similarly, when a party is deregistered, it is removed from both $\mathcal{P}$ (and therefore also from $\mathcal{P}_{DS}$ or $\mathcal{H}$). The ledger maintains the invariant that it is registered (as a functionality) to the clock whenever $\mathcal{H} \neq \emptyset$.
For each party $p_i \in \mathcal{P}$ the functionality maintains a pointer $\texttt{pt}_i$ (initially set to 1) and a current-state view $\texttt{state}_i := \varepsilon$ (initially set to empty). The functionality also keeps track of the timed honest-input sequence in a vector $\vec{\mathcal{I}}_H^T$ (initially $\vec{\mathcal{I}}_H^T := \varepsilon$).

Upon receiving any input I from any party or from the adversary, send $(\text{CLOCK-READ}, \text{sid}_C)$ to $\mathcal{G}_{\text{CLOCK}}$ and upon receiving response $(\text{CLOCK-READ}, \text{sid}_C, \tau)$ set $\tau_L := \tau$ and do the following:

1. Let $\widehat{\mathcal{P}} \subseteq \mathcal{P}_{DS}$ denote the set of desynchronized honest parties that have been registered (continuously) since time $\tau' < \tau_L - \texttt{Delay}$ (to both ledger and clock). Set $\mathcal{P}_{DS} := \mathcal{P}_{DS} \setminus \widehat{\mathcal{P}}$.
2. If I was received from an honest party $p_i \in \mathcal{P}$:
   (a) Set $\vec{\mathcal{I}}_H^T := \vec{\mathcal{I}}_H^T || (I, p_i, \tau_L)$;
   (b) Compute $\vec{N} = (\vec{N}_1, \ldots, \vec{N}_\ell) := \mathsf{ExtendPolicy}(\vec{\mathcal{I}}_H^T, \texttt{state}, \texttt{NxtBC}, \texttt{buffer}, \vec{\tau}_{\texttt{state}})$ and if $\vec{N} \neq \varepsilon$ set $\texttt{state} := \texttt{state} || \mathsf{Blockify}(\vec{N}_1) || \ldots || \mathsf{Blockify}(\vec{N}_\ell)$ and $\vec{\tau}_{\texttt{state}} := \vec{\tau}_{\texttt{state}} || \tau_L^\ell$, where $\tau_L^\ell = \tau_L || \ldots, || \tau_L$.
   (c) For each $\texttt{BTX} \in \texttt{buffer}$: if $\mathsf{Validate}(\texttt{BTX}, \texttt{state}, \texttt{buffer}) = 0$ then delete $\texttt{BTX}$ from $\texttt{buffer}$. Also, reset $\texttt{NxtBC} := \varepsilon$.
   (d) If there exists $p_j \in \mathcal{H} \setminus \mathcal{P}_{DS}$ such that $|\texttt{state}| - \texttt{pt}_j > \texttt{windowSize}$ or $\texttt{pt}_j < |\texttt{state}_j|$, then set $\texttt{pt}_k := |\texttt{state}|$ for all $p_k \in \mathcal{H} \setminus \mathcal{P}_{DS}$.
3. Depending on the above input I and its sender's ID, $\mathcal{G}_{\text{LEDGER}}$ executes the corresponding code from the following list:
   - *Submiting a transaction:*
     If $I = (\text{SUBMIT}, \text{sid}, \texttt{tx})$ and is received from a party $p_i \in \mathcal{P}$ or from $\mathcal{A}$ (on behalf of a corrupted party p_i) do the following
     (a) Choose a unique transaction ID txid and set $\texttt{BTX} := (\texttt{tx}, \text{txid}, \tau_L, p_i)$
     (b) If $\mathsf{Validate}(\texttt{BTX}, \texttt{state}, \texttt{buffer}) = 1$, then $\texttt{buffer} := \texttt{buffer} \cup \{\texttt{BTX}\}$.
     (c) Send $(\text{SUBMIT}, \texttt{BTX})$ to $\mathcal{A}$.
   - *Reading the state:*
     If $I = (\text{READ}, \text{sid})$ is received from a party $p_i \in \mathcal{P}$ then set $\texttt{state}_i := \texttt{state}|_{\min\{\texttt{pt}_i, |\texttt{state}|\}}$ and return $(\text{READ}, \text{sid}, \texttt{state}_i)$ to the requestor. If the requestor is $\mathcal{A}$ then send $(\texttt{state}, \texttt{buffer}, \vec{\mathcal{I}}_H^T)$ to $\mathcal{A}$.
   - *Maintaining the ledger state:*
     If $I = (\text{MAINTAIN-LEDGER}, \text{sid}, \text{minerID})$ is received by an honest party $p_i \in \mathcal{P}$ and (after updating $\vec{\mathcal{I}}_H^T$ as above) $\mathsf{predict\text{-}time}(\vec{\mathcal{I}}_H^T) = \widehat{\tau} > \tau_L$ then send $(\text{CLOCK-UPDATE}, \text{sid}_C)$ to $\mathcal{G}_{\text{CLOCK}}$. Else send I to $\mathcal{A}$.
   - *The adversary proposing the next block:*
     If $I = (\text{NEXT-BLOCK}, \text{hFlag}, (\text{txid}_1, \ldots, \text{txid}_\ell))$ is sent from the adversary, update $\texttt{NxtBC}$ as follows:
     (a) Set listOfTxid $\leftarrow \epsilon$
     (b) For $i = 1, \ldots, \ell$ do: if there exists $\texttt{BTX} := (x, \text{txid}, \text{minerID}, \tau_L, p_i) \in \texttt{buffer}$ with ID txid $= \text{txid}_i$ then set listOfTxid := listOfTxid$||\text{txid}_i$.
     (c) Finally, set $\texttt{NxtBC} := \texttt{NxtBC} || (\text{hFlag}, \text{listOfTxid})$ and output $(\text{NEXT-BLOCK}, ok)$ to $\mathcal{A}$.
   - *The adversary setting state-slackness:*
     If $I = (\text{SET-SLACK}, (p_{i_1}, \widehat{\texttt{pt}}_{i_1}), \ldots, (p_{i_\ell}, \widehat{\texttt{pt}}_{i_\ell}))$, with $\{p_{i_1}, \ldots, p_{i_\ell}\} \subseteq \mathcal{H} \setminus \mathcal{P}_{DS}$ is received from the adversary $\mathcal{A}$ do the following:
     (a) If for all $j \in [\ell]$: $|\texttt{state}| - \widehat{\texttt{pt}}_{i_j} \leq \texttt{windowSize}$ and $\widehat{\texttt{pt}}_{i_j} \geq |\texttt{state}_{i_j}|$, set $\texttt{pt}_{i_1} := \widehat{\texttt{pt}}_{i_1}$ for every $j \in [\ell]$ and return $(\text{SET-SLACK}, ok)$ to $\mathcal{A}$.

(b) Otherwise set $\mathtt{pt}_{i_j} := |\mathtt{state}|$ for all $j \in [\ell]$.

- *The adversary setting the state for desychronized parties:*
  If $I = (\text{DESYNC-STATE}, (p_{i_1}, \mathtt{state}'_{i_1}), \ldots, (p_{i_\ell}, \mathtt{state}'_{i_\ell}))$, with $\{p_{i_1}, \ldots, p_{i_\ell}\} \subseteq \mathcal{P}_{DS}$ is received from the adversary $\mathcal{A}$, set $\mathtt{state}_{i_j} := \mathtt{state}'_{i_j}$ for each $j \in [\ell]$ and return $(\text{DESYNC-STATE}, ok)$ to $\mathcal{A}$.

4 Bitcoin as a Transaction Ledger Protocol

In this section we prove our main theorem, namely that, under appropriate assumptions, Bitcoin realizes an instantiation of the ledger functionality from the previous section. More concretely, we cast the Bitcoin protocol as a UC protocol, where consistent with the existing methodology we assume that the protocol is synchronous, i.e., parties can keep track of the current round by using an appropriate global clock functionality. We first describe the UC protocol, denoted Ledger-Protocol, in Section 4.1 which abstracts the components of Bitcoin that are relevant for the construction of such a ledger—similar to how the backbone protocol [GKL15] captures core Bitcoin properties in their respective model of computation.

Later, in Section 4.2, we specify the ledger functionality $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$ that is implemented by the UC ledger protocol as an instance of our general ledger functionality, i.e., by providing appropriate instantiations of algorithms Validate, Blockify, and ExtendPolicy. In fact, for the sake of generality, we specify generic classes of Validate and Blockify and parameterize our Ledger-Protocol with these classes, so that the security statement still stays generic. We then prove our main theorem (Theorem 1) which can be described informally as follows:

Theorem (Informal). *Let* Validate *be the class of predicates that only take into account the current state and a transaction (i.e., no transaction IDs, time, or party IDs), and let* $\mathtt{windowSize} = \omega(\log \kappa)$, κ *being the length of the outputs of the random oracle. Then, for an appropriate* ExtendPolicy *and for any function* Blockify*, the protocol* Ledger-Protocol *instantiated with algorithms* Validate *and* Blockify *securely realizes a ledger functionality* $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$ *(the generic ledger instantiated with the above functions) under the following assumptions on network delays and mining power, where mining power is roughly understood as the ability to find proofs of work via queries to the random oracle (and will be formally defined later):*

- *In any round of the protocol execution, the collective mining power of the adversary, contributed by corrupted and temporarily de-synchronized miners, does not exceed the mining power of honest (and synchronized) parties in that round. The exact relation additionally captures the (negative) impact of network delays on the coordination of mining power of honest parties.*
- *No message can be delayed in the network by more than* $\Delta = O(1)$ *rounds.*

We prove the above theorem via what we believe is a useful modularization of the Bitcoin protocol (cf. Figure 1). Informally, this modularization distills out form the protocol a reactive *state-extend* subprocess which captures the lottery that decides which miner gets to advance the blockchain next and additionally the process of propagating this state to other miners. Lemma 1 shows that the state-extend module/subprocess implements an appropriate reactive UC functionality $\mathcal{F}_{\text{STX}}$. We can then use the UC composition theorem which allows us to argue security of Ledger-Protocol in a simpler hybrid world where, instead of using this subprocess, parties make calls to the functionality $\mathcal{F}_{\text{STX}}$. We conclude this section (Subsection 4.4) by showing how both the GKL and PSs protocols can be cast as special cases of our protocol which provides the basis for comparing the different models and their respective assumptions.

4.1 The Bitcoin ledger as a UC Protocol

In the following we provide the formal description of protocol Ledger-Protocol. The protocol assumes as hybrids the multi-cast network $\mathcal{F}_{\text{N-MC}}$ (recall that we assume that this network does have an upper bound Δ on the delay unknown to the protocol) and a random oracle functionality $\mathcal{F}_{\text{RO}}$. Before providing the detailed specification of our ledger protocol, we establish some useful notation and terminology that we use throughout this section. For compatibility with existing work, wherever it does not overload notation, we use some of the terminology and notation from [GKL15].

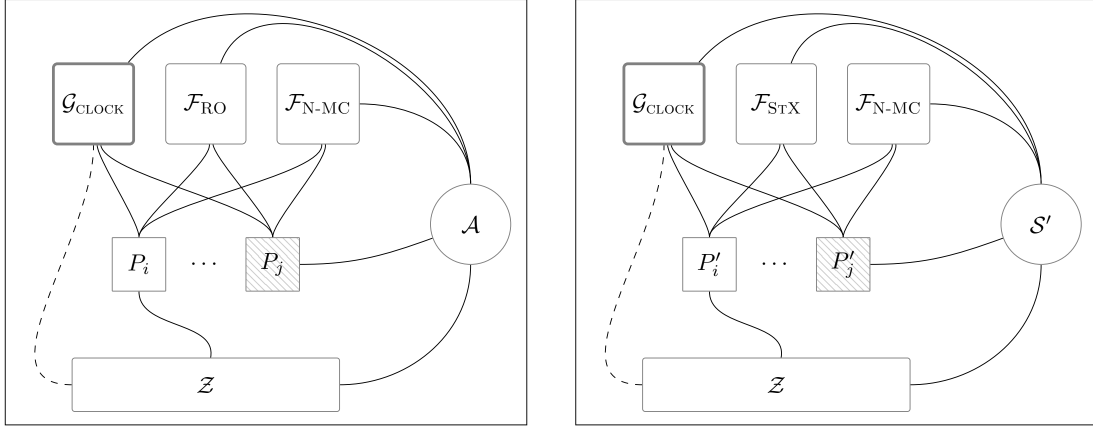

(a) In the real world parties have access to the global clock $\mathcal{G}_{\text{CLOCK}}$, the random oracle $\mathcal{F}_{\text{RO}}$, and network $\mathcal{F}_{\text{N-MC}}$. Here, parties execute the Bitcoin protocol Ledger-Protocol

(b) In the hybrid world parties have access to the state-exchange functionality $\mathcal{F}_{\text{StX}}$ (instead of the random oracle). Here, parties execute the modularized protocol Modular-Ledger-Protocol

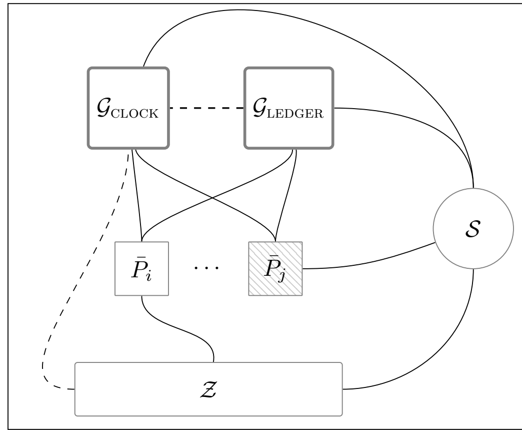

(c) In the ideal world, dummy parties have access to the global clock $\mathcal{G}_{\text{CLOCK}}$ and the ledger $\mathcal{G}_{\text{LEDGER}}$

Fig. 1. Modularization of the Bitcoin protocol.

Blockchain. A *blockchain* $\mathcal{C} = \mathbf{B}_1, \ldots, \mathbf{B}_n$ is a (finite) sequence of blocks where each *block* $\mathbf{B}_i = \langle \mathtt{s}_i, \mathtt{st}_i, \mathtt{n}_i \rangle$ is a triple consisting of the *pointer* $\mathtt{s}_i$, the *state block* $\mathtt{st}_i$, and the *nonce* $\mathtt{n}_i$. A special block is the *genesis block* $\mathbf{G} = \langle \bot, \mathtt{gen}, \bot \rangle$ which contains the genesis state $\mathtt{gen} := \varepsilon$. The *head* of chain $\mathcal{C}$ is the block $\mathtt{head}(\mathcal{C}) := \mathbf{B}_n$ and the *length* $\mathtt{length}(\mathcal{C})$ of the chain is the number of blocks, i.e., $\mathtt{length}(\mathcal{C}) = n$. The chain $\mathcal{C}^{\lceil k}$ is the (potentially empty) sequence of the first $\mathtt{length}(\mathcal{C}) - k$ blocks of $\mathcal{C}$. The *state* $\vec{\mathtt{st}}$ *encoded* in $\mathcal{C}$ is defined as a sequence of the corresponding state blocks, i.e., $\vec{\mathtt{st}} := \mathtt{st}_1 || \ldots || \mathtt{st}_n$. In other words, one should think of the blockchain $\mathcal{C}$ as an encoding of its underlying state $\vec{\mathtt{st}}$; such an encoding might, e.g., organize $\mathcal{C}$ is an efficient searchable data structure as is the case in the Bitcoin protocol where a blockchain is a linked list implemented with hash-pointers.

In the protocol, the blockchain is the data structure storing a sequence of entries, often referred to as transactions. Furthermore, as in [KZZ16], in order to capture blockchains with syntactically different state encoding, we use an algorithm $\mathsf{blockify}_{\mathcal{B}}$ to map a vector of transactions into a state block. Thus, each block $\mathtt{st} \in \vec{\mathtt{st}}$ (except the genesis state) of the state encoded in the blockchain has the form $\mathtt{st} = \mathsf{Blockify}(\vec{N})$ where $\vec{N}$ is a vector of transactions.

For a blockchain $\mathcal{C}$ to be considered a valid blockchain, it needs to satisfy certain conditions. Concretely, the validity of a blockchain $\mathcal{C} = \mathbf{B}_1, \ldots, \mathbf{B}_n$ where $\mathbf{B}_i = \langle \mathtt{s}_i, \mathtt{st}_i, \mathtt{n}_i \rangle$ depends on two aspects: *chain-level* validity, also referred to as syntactic validity, and a *state-level* validity also referred to as semantic validity. Syntactic validity is defined with respect to a difficulty parameter $\mathtt{D} \in [2^\kappa]$, where κ is the security parameter, and a given hash function $H(\cdot) : \{0,1\}^* \to \{0,1\}^\kappa$; it requires that, for each

$i > 1$, the value $\mathtt{s}_i$ contained in $\mathbf{B}_i$ satisfies $\mathtt{s}_i = H[\mathbf{B}_{i-1}]$ and that additionally $H[\mathbf{B}_i] < \mathtt{D}$ holds, where we interpret the output of the hash-function as an integer in this comparison.

The semantic validity on the other hand is defined on the state $\vec{\mathtt{st}}$ encoded in the blockchain $\mathcal{C}$ and specifies whether this content is valid (which might depend on a particular application). The validation predicate Validate defined in the ledger functionality (cf. Section 3) plays a similar role. In fact, the semantic validity of the blockchain can be defined using an algorithm that we denote isvalidstate which is builds upon the Validate predicate. The idea is that for any choice of Validate, the blockchain protocol using isvalidstate for semantic validation of the chain implements the ledger parametrized with Validate. More specifically, algorithm isvalidstate checks that a given blockchain state can be built in an iterative manner, such that each contained transaction is considered valid according to Validate upon insertion. It further ensures that the state starts with the genesis state and that state blocks contain a special *coin-base* transaction $\mathtt{tx}^{\text{coin-base}}_{\text{minerID}}$ which assigns them to a miner. We remark that this only works for predicates Validate which ignore all information other than the state and transaction that is being validated.[17] To avoid confusion, throughout this section we use $\mathsf{ValidTx}_{\text{B}}$ to refer to the validate predicate with the above restriction. The pseudo-code of the algorithm isvalidstate which builds upon $\mathsf{ValidTx}_{\text{B}}$ is provided in Appendix C.

We succinctly denote by $\mathsf{isvalidchain}_D(\mathcal{C})$ the predicate that returns true iff chain $\mathcal{C}$ satisfies syntactic and semantic validity as defined above.

The Ledger Protocol. We can now formally define our blockchain protocol $\mathsf{Ledger\text{-}Protocol}_{q,\mathtt{D},T}$ (we usually omit the parameters when clear from the context). The protocol allows an arbitrary number of parties/miners to communicate by means of a multicast network $\mathcal{F}_{\text{N-MC}}$. Note that this means that the adversary can send different messages to different parties. New miners might dynamically joint or leave the protocol by means of the registration/de-registration commands: when they join they register with all associated functionalities and when they leave they deregister.[18]

Each party maintains a local blockchain which initially consists of the genesis block. The chains of honest parties might differ (but as we will prove, it will have a common prefix which will define the ledger state). New transactions are added in a 'mining process'. First, a party collects valid transactions (according to $\mathsf{ValidTx}_{\text{B}}$) and creates a new state block $\mathtt{st}$ using $\mathsf{blockify}_{\text{B}}$. Next, the party attempts to mine a new block which can be validly added to their local blockchain. The mining is done using the $\mathsf{extendchain}_D$ algorithm which takes as inputs a chain $\mathcal{C}$, a state block $\mathtt{st}$, and the number q of attempts. The core idea of the algorithm is to find a proof-of-work which allows to extend $\mathcal{C}$ by a block which encodes $\mathtt{st}$. The pseudo-code of this algorithm is provided in Appendix C.2. After each mining attempt parties will multicast their current chain. A party will replace its local chain if it receives a longer chain. When queried to output the state of the ledger, Ledger-Protocol outputs the state of its longest chain, where it first chops-off the most recent T blocks (or ε if the state has less than T blocks). This behavior will ensure that all honest parties output a consistent ledger state.

As already mentioned, our Bitcoin-Ledger protocol proceeds in rounds which are implemented by using a global synchronization clock $\mathcal{G}_{\text{CLOCK}}$. For formal reasons that have to do with how activations are handled in UC, we have each round correspond to two sub-rounds (also known as mini-rounds). To avoid confusion we refer to clock rounds as *clock-ticks*. We say that a protocol is in round r if the current time of the clock is $\tau \in \{2r-1, 2r\}$. In fact, having two clock-ticks per round is the way to ensure in synchronous UC that messages (e.g., a block) sent within a round are delivered at the beginning of the next round. The idea is that each round is divided into two mini-rounds, where each mini-round corresponds to a clock tick, and treat the first mini-round as a *working mini-round* where parties might mine new blocks and submit them to the multicast network for delivery, and in the second *reading mini-round* they simply fetch messages from the network to obtain messages sent in the previous round. The pseudo-code of this UC blockchain protocol, denoted as Ledger-Protocol, is provided in Appendix C.3 where we also argue that the protocol satisfies Definition 1, i.e., there is a concrete function $\mathsf{predict\text{-}time}_{BC}$ that predicts the synchronization pattern of our synchronous UC Bitcoin backbone protocol

[17] Recall that in the general ledger description, Validate might depend on some associated metadata; although this might be useful to capture alternative blockchains, it is not the case for Bitcoin.

[18] Note that when a party registers to a local functionality such as the network or the random oracle it does not lose its activation token. This is a subtle point to ensure that the real and ideal worlds are in-sync regarding activations.

4.2 The Bitcoin Ledger

We next show how to instantiate the ledger functionality from Section 3 with appropriate parameters so that it is implemented by protocol Ledger-Protocol. To define this Bitcoin ledger $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$, we give the specific instantiations of the relevant functions Validate, Blockify, ExtendPolicy, and predict-time. First, predict-time is defined to be predict-time$_{BC}$ to reflect the synchronization pattern of the Bitcoin backbone protocol.

Similarly, in case of Validate we use the same predicate as the protocol uses to validate the states: For a given transaction tx and a given state state, the predicate decides whether this transaction is valid with respect to state. Given such a validation predicate, the ledger validation predicate takes a specific simple form which, excludes dependency on anything other than the transaction tx and the state state, i.e., for any values of txid, τ_L, p_i, and buffer:

$$\mathsf{Validate}((\mathtt{tx}, \text{txid}, \tau_L, p_i), \mathtt{state}, \mathtt{buffer}) := \mathsf{ValidTx}_{\text{B}}(\mathtt{tx}, \mathtt{state}).$$

Blockify can be an arbitrary algorithm, and if the same algorithm is used in Ledger-Protocol the security proof will go through. However, as discussed below (in Definition 2), a meaningful Blockify should be in certain relation with the ledger's Validate predicate. (This relation is satisfied by the Bitcoin protocol.)

Finally, we define ExtendPolicy. At a high level, upon receiving a list of possible candidate blocks which should go into the state of the ledger, ExtendPolicy does the following: for each block it first verifies that the blocks are valid with respect to the state they extend. (Only valid blocks might be added to the state.) Moreover, ExtendPolicy ensures the following property:

1. The speed of the ledger is not too slow. This is implemented by defining an upper bound $\mathtt{maxTime}_{\mathtt{window}}$ on the time interval (number of clock-ticks) within which at least windowSize state blocks have to be added. This is known as minimal chain-growth.
2. The speed of the ledger is not too fast. This is implemented by defining a lower bound $\mathtt{minTime}_{\mathtt{window}}$ on the time interval (number of clock-ticks), such that the adversary is not allowed to propose new blocks if windowSize or more blocks have already been added during that time interval.
3. The adversary cannot create too many blocks with arbitrary (but valid) contents. This is formally enforced by defining an upper bound η on the number of these so-called adversarial blocks within a sequence of state blocks. This is known as chain quality. Formally, this is enforced by requiring that a certain fraction of blocks need to satisfy higher quality standards (to model blocks that are honestly generated).
4. Last but not least, ExtendPolicy guarantees that if a transaction is "old enough", and still valid with respect to the actual state, then it is included into the state. This is a weak form of guaranteeing that a transaction will make it into the state unless it is in conflict. As we show in Section 5, this guarantee can be amplified by using digital signatures.

In order to enforce these policies, ExtendPolicy first defines alternative blocks which satisfy all of the above criteria in an ideal way, and whenever it catches the adversary in trying to propose blocks that do not obey the policies, it punishes the adversary by proposing its own generated blocks. In particular, if the adversary violates the policy regarding minimal chain-growth, the ExtendPolicy will directly propose a sequence of complying blocks. The formal description of the extend policy (as pseudo-code) for $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$ is given in Appendix C.4.

On the relation between Blockify and Validate. As already discussed above, ExtendPolicy guarantees that the adversary cannot block the extension of the state indefinitely, and that occasionally an honest miner will create a block. These are implications of the chain-growth and chain-quality properties from [GKL15]. However, our generic ExtendPolicy makes explicit that a priori, we cannot exclude that the chain always extends with blocks that include, for example, only a coin-base transaction, i.e., any submitted transaction is ignored and never inserted into a new block. This issue is an orthogonal one to ensuring that honest transactions are not invalidated by adversarial interaction—which, as argued in [GKL15], is achieved by adding digital signatures.

To see where this could be problematic in general, consider a blockify that, at a certain point, creates a block that renders all possible future transactions invalid. Observe that this does not mean that our protocol is insecure and that this is as well possible for the protocols of [GKL15, PSS17]; indeed our proof shows that the protocol will give exactly the same guarantees as an $\mathcal{G}_{\text{LEDGER}}$ parametrized with such an algorithm Blockify.

Nonetheless, a look in reality indicates that this situation never occurs with Bitcoin. To capture that this is the case, Validate and Blockify need to be in a certain relation with each other. Informally,

this relation should ensure that the above sketched situation never occurs. A way to ensure this, which is already implemented by the Bitcoin protocol, is by restricting Blockify to only make an invertible manipulation of the blocks when they are inserted into the state—e.g., be an encoding function of a code—and define Validate to depend on the inverse of Blockify. This is captured in the following definition.

Definition 2. *A co-design of Blockify and Validate is* non-self-disqualifying *if there exists an efficiently computable function Dec mapping outputs of Blockify to vectors $\vec{N}$ such that there exists a validate predicate Validate′ for which the following properties hold for any possible state* $\texttt{state} = \texttt{st}_1||\ldots||\texttt{st}_\ell$*, buffer* $\texttt{buffer}$ *vectors* $\vec{N} := (\texttt{tx}_1, \ldots, \texttt{tx}_m)$*, and transaction* $\texttt{tx}$*:*

1. $\mathsf{Validate}(\texttt{tx}, \texttt{state}, \texttt{buffer}) = \mathsf{Validate}'(\texttt{tx}, \mathsf{Dec}(\texttt{st}_1)||\ldots||\mathsf{Dec}(\texttt{st}_\ell), \texttt{buffer})$
2. $\mathsf{Validate}(\texttt{tx}, \texttt{state}||\mathsf{Blockify}(\vec{N}), \texttt{buffer}) = \mathsf{Validate}'(\texttt{tx}, \mathsf{Dec}(\texttt{st}_1)||\ldots||\mathsf{Dec}(\texttt{st}_\ell)||\vec{N}, \texttt{buffer})$

We remark that the actual validation of Bitcoin does satisfy the above definition, since a transaction is only rendered invalid with respect to the state if the coins it is trying to spend have already been spent, and this only depends on the transactions in the state and not the metadata added by Blockify. Hence, in the following, we assume that $\mathsf{ValidTx}_{B}$ and $\mathsf{blockify}_{B}$ satisfy the relation in Definition 2.

4.3 Security Analysis

We next turn to the security analysis of our protocol. As already mentioned, we argue security in two step. In a first step, we distill out from the protocol Ledger-Protocol a state-extend module/subprocess, denoted as StateExchange-Protocol, and devise an alternative, modular description of the Ledger-Protocol protocol in which every party makes invocations of this subprocess. We denote this modularized protocol by Modular-Ledger-Protocol. By a game-hopping argument, we prove that the original protocol Ledger-Protocol and the modularized protocol Modular-Ledger-Protocol are in fact functionally equivalent. The advantage of having such a modular description is that we are now able to define an appropriate ideal functionality $\mathcal{F}_{\text{STX}}$ that is realized by StateExchange-Protocol. Using the universal composition theorem we can deduce that Ledger-Protocol UC emulates Modular-Ledger-Protocol where invocations of StateExchange-Protocol are replaced by invocations of $\mathcal{F}_{\text{STX}}$. The second step of the proof consists of proving that, under appropriate assumptions, Modular-Ledger-Protocol, where invocations of StateExchange-Protocol are replaced by invocations of $\mathcal{F}_{\text{STX}}$, implements the Bitcoin ledger described in Section 4.2.

Step 1. The state-exchange functionality $\mathcal{F}_{\text{STX}}$ allows parties to submit ledger states which are accepted with a certain probability. Accepted states are then multicast to all parties. Informally, it can be seen as as lottery on which (valid) states are exchanged among the participants. Parties can use $\mathcal{F}_{\text{STX}}$ to multicast a valid state, but instead of accepting any submitted state and sending it to all (registered) parties, $\mathcal{F}_{\text{STX}}$ keeps track of all states that it ever saw, and implements the following mechanism upon submission of a new ledger state $\vec{\texttt{st}}$ and a state block $\texttt{st}$ from any party: If $\vec{\texttt{st}}$ was previously submitted to $\mathcal{F}_{\text{STX}}$ and $\vec{\texttt{st}}||\texttt{st}$ is a valid state, then $\mathcal{F}_{\text{STX}}$ accepts $\vec{\texttt{st}}||\texttt{st}$ with probability p_H (resp. p_A for dishonest parties); accepted states are then sent to all registered parties. The formal specification follows:

Functionality $\mathcal{F}_{\text{STX}}^{\Delta, p_H, p_A}$

The functionality is parametrized with a set of parties $\mathcal{P}$. Any newly registered (resp. deregistered) party is added to (resp. deleted from) $\mathcal{P}$. For each party $p \in \mathcal{P}$ the functionality manages a tree $\mathcal{T}_p$ where each rooted path corresponds to a valid state the party has received. Initially each tree contains the genesis state. Finally, it manages a buffer $\vec{M}$ which contains successfully submitted states which have not yet been delivered to (some) parties in $\mathcal{P}$.

Submit/receive new states:

- Upon receiving $(\textsc{submit-new}, \text{sid}, \vec{\texttt{st}}, \texttt{st})$ from some participant $p_s \in \mathcal{P}$, if $\mathsf{isvalidstate}(\vec{\texttt{st}}||\texttt{st}) = 1$ and $\vec{\texttt{st}} \in \mathcal{T}_p$ do the following:
  1. Sample B according to a Bernoulli-Distribution with parameter p_H (or p_A if p_s is dishonest).
  2. If $B = 1$, set $\vec{\texttt{st}}_{new} \leftarrow \vec{\texttt{st}}||\texttt{st}$ and add $\vec{\texttt{st}}_{new}$ to $\mathcal{T}_{p_s}$. Else set $\vec{\texttt{st}}_{new} \leftarrow \vec{\texttt{st}}$.
  3. Output $(\textsc{success}, \text{sid}, B)$ to p_s.

4. On response (CONTINUE, sid) where $\mathcal{P} = \{p_1, \dots, p_n\}$ choose n new unique message-IDs $\text{mid}_1, \dots, \text{mid}_n$, initialize n new variables $D_{\text{mid}_1} := D^{MAX}_{\text{mid}_1} := \dots := D_{\text{mid}_n} := D^{MAX}_{\text{mid}_n} := 1$ set $\vec{M} := \vec{M}||(\vec{\mathtt{st}}_{new}, \text{mid}_1, D_{\text{mid}_1}, p_1)||\dots||(\vec{\mathtt{st}}_{new}, \text{mid}_n, D_{\text{mid}_n}, p_n)$, and send (SUBMIT-NEW, sid, $\vec{\mathtt{st}}_{new}, p_s, (p_1, \text{mid}_1), \dots, (p_n, \text{mid}_n)$) to the adversary.

- Upon receiving (FETCH-NEW, sid) from a party $p \in \mathcal{P}$ or $\mathcal{A}$ (on behalf of p), do the following:
  1. For all tuples $(\vec{\mathtt{st}}, \text{mid}, D_{\text{mid}}, p) \in \vec{M}$ set $D_{\text{mid}} := D_{\text{mid}} - 1$.
  2. Let $\vec{M}^p_0$ denote the subvector of $\vec{M}$ including all tuples of the form $(\vec{\mathtt{st}}, \text{mid}, D_{\text{mid}}, p)$ where $D_{\text{mid}} = 0$ (in the same order as they appear in $\vec{M}$). For each tuple $(\vec{\mathtt{st}}, \text{mid}, D_{\text{mid}}, p) \in \vec{M}^p_0$ add $\vec{\mathtt{st}}$ to $\mathcal{T}_p$. Delete all entries in $\vec{M}^p_0$ from $\vec{M}$ and send $\vec{M}^p_0$ to p.
- Upon receiving (SEND, sid, $\vec{\mathtt{st}}, p'$) from $\mathcal{A}$ on behalf some *corrupted* $p \in \mathcal{P}$, if $p' \in \mathcal{P}$ and $\vec{\mathtt{st}} \in \mathcal{T}_p$, choose a new unique message-ID mid, initialize $D := 1$, add $(\vec{\mathtt{st}}, \text{mid}, D_{\text{mid}}, p')$ to $\vec{M}$, and return (SEND, sid, $\vec{\mathtt{st}}, p'$, mid) to $\mathcal{A}$.

Further adversarial influence on the network:

- Upon receiving (SWAP, sid, mid, mid′) from $\mathcal{A}$, if mid and mid′ are message-IDs registered in the current $\vec{M}$, swap the corresponding tuples in $\vec{M}$. Return (SWAP, sid) to $\mathcal{A}$.
- Upon receiving (DELAY, sid, T, mid) from $\mathcal{A}$, if T is a valid delay, mid is a message-ID for a tuple $(\vec{\mathtt{st}}, \text{mid}, D_{\text{mid}}, p)$ in the current $\vec{M}$ and $D^{MAX}_{\text{mid}} + T \leq \Delta$, set $D_{\text{mid}} := D_{\text{mid}} + T$ and set $D^{MAX}_{\text{mid}} := D^{MAX}_{\text{mid}} + T$.

The Modular-Ledger-Protocol uses the same hybrids as Ledger-Protocol but abstracts the lottery implemented by the mining process by making calls to the above state exchange functionality $\mathcal{F}^{\Delta,p_H,p_A}_{\text{STX}}$. The detailed specification of the Modular-Ledger-Protocol protocol can be found in Appendix D.1. Note that the only remaining parameter of Modular-Ledger-Protocol is the chop-off parameter T, the rest is part of $\mathcal{F}^{\Delta,p_H,p_A}_{\text{STX}}$. The following Lemma states that our Bitcoin protocol implements the above modular ledger protocol. The proof appears in Appendix D.2.

Lemma 1. *The blockchain protocol* Ledger-Protocol$_{q,D,T}$ *UC emulates protocol* Modular-Ledger-Protocol$_T$ *that runs in a hybrid world with access to the functionality* $\mathcal{F}^{\Delta,p_H,p_A}_{\text{STX}}$ *with* $p_A := \frac{D}{2^\kappa}$ *and* $p_H = 1-(1-p_A)^q$, *and where* Δ *denotes the upper bound on the network delay.*

Step 2. We are now ready to complete the proof of our main theorem. Before providing the formal statement it is useful to discuss some of the key properties used in both, the statement and the proof. The security of the Bitcoin protocol depends on various key properties of an execution. This means that its security depends on the number of random oracle queries (or, in the $\mathcal{F}_{\text{STX}}$ hybrid world, the number of submit-queries) by the pool of corrupted miners. Therefore it is important to capture the relevant properties of such a UC execution. In the following we denote by upper-case R the number of rounds of a given protocol execution.

Capturing query power in an execution. In an execution, we measure the query power per logical round r, which can be conveniently captured as a function $\mathtt{T}_{qp}(r)$. We observe that in an interval of, say, t_{rc} rounds, the total number of queries is

$$Q^{r'}_{t_{rc}} = \sum_{r=r'}^{r'+t_{rc}-1} \mathtt{T}_{qp}(r).$$

In each round $r \in [R]$, each honest miner gets a certain number $q^{(r)}_i$ of activations from the environment to maintain the ledger (i.e., to try to extend the state). Let

$$q^{(r)}_H := \sum_{p_i \text{ honest in round r}} q^{(r)}_i.$$

Also, the adversary makes a certain number $q^{(r)}_A$ of queries to $\mathcal{F}_{\text{STX}}$. We get

$$\mathtt{T}_{qp}(r) = q^{(r)}_A + q^{(r)}_H.$$

Quantifying total mining power in an execution. Mining power is the expected number of successful state extensions, i.e., the number of times a new state block is successfully mined. The mining power of round r is therefore

$$\mathtt{T}_{mp}(r) := q_A^{(r)} \cdot p_A + q_H^{(r)} \cdot p_H,$$

Recall that p_H is the success probability per query of an honest miner and p_A is the success probability per query of a corrupted miner. If $p_A = p$ and $p_H = 1-(1-p)^q$, it is convenient to consider $(q_A^{(r)} + q \cdot q_H^{(r)}) \cdot p$ as the total mining power (by applying Bernoulli's inequality). Within an interval of t_{rc} rounds, we can for example quantify the overall expectation by $\mathtt{T}_{mp}^{\text{total}}(t_{rc}) := \sum_{r=1}^{t_{rc}} \mathtt{T}_{mp}(r)$. This allows to formulate the goal of a re-calibration of the difficulty parameter as requiring that this value should be 2016 blocks for t_{rc} corresponding a desired time bound (such as roughly two weeks), which is part of future work.

Quantifying adversarial mining power in an execution. The *adversarial mining power* $\mathtt{mp}_A(r)$ per round is made up of two parts: first, queries by corrupted parties, and second, queries by honest, but de-synchronized miners (recall that the latter are either active honest parties that are de-registered from the clock or not yet registered with the clock for long enough and thus still out of sync).

$$\mathtt{mp}_A(r) := p_A \cdot q_A^{(r)} \ + \ p_H \cdot \sum_{p_i \text{ is de-sync}} q_i^{(r)}.$$

Recall that a party is considered desynchronized for 2Δ rounds after its registration.

It is convenient to measure the adversary's contribution to the mining power as the fraction of the overall mining power. In particular, we assume there is a parameter $\rho \in (0,1)$ such that in any round r, the relation $\mathtt{mp}_A(r) \leq \rho \cdot \mathtt{T}_{mp}(r)$ holds. We then define $\beta_r := \rho \cdot \mathtt{T}_{mp}(r)$. Looking ahead, if a model is flat, then the fraction $(1-\rho)$ corresponds to the fraction of users that are honest and synchronized.

Quantifying honest and synchronized mining power in an execution. In each round $r \in [R]$, each honest miner gets a certain number $q_{i,r}$ of activations from the environment, where it can submit one new state to $\mathcal{F}_{\text{STX}}$. This state is accepted with probability p_H. We define the vector $\vec{q}_r$ such that for any honest miner p_i in round r, $\vec{q}_r[i] = q_{i,r}$. The probability that a miner p_i is successful to extend the state by at least one block is $\alpha_{i,r} := 1-(1-p_H)^{q_{i,r}}$ and the probability that at least one *registered and synchronized*, uncorrupted miner successfully queries $\mathcal{F}_{\text{STX}}$ to extend its local longest state is

$$\alpha_r := 1 - \prod_{\text{honest sync } p_i} (1-\alpha_{i,r}) = 1 - \prod_{\text{honest sync } p_i} (1-p_H)^{q_{i,r}}.$$

Looking ahead, in existing flat models of Bitcoin, parties are expected to be synchronized and are otherwise counted as dishonest and the quantity $(1-\rho)$ is the fraction of honest and synchronized miners.

Worst-Case Analysis. We analyze Bitcoin in a worst-case fashion. Let us assume that the protocol runs for $[R]$ rounds (e.g., $R = t_{rc}$ if we do not take re-calibration into account), then

$$\alpha := \min \{\alpha_r\}_{r \in [R]}, \text{ and } \beta := \max \{\beta_r\}_{r \in [R]}.$$

Remark 4. This view on Bitcoin gives already a glimpse for the relevance of the re-calibration sub-protocol which is considered as part of future work. Ideally, we would like the variation among the values α_r and among the values β_r to be small, which needs an additional assumption on the increase of computing power per round. Thanks to the re-calibration phase, such a bound can exist at all. If no re-calibration phase would happen, any strictly positive gradient of the computing power development would eventually provoke Bitcoin failing, as the value β (as a fraction of the total mining power) could not be reasonably bounded.

We are now ready to state the main theorem. The proof of the theorem can be found in Appendix D.3.

Theorem 1. *Let the functions* $\mathsf{ValidTx}_{\mathcal{B}}$, $\mathsf{blockify}_{\mathcal{B}}$, *and* $\mathsf{ExtendPolicy}$ *be as defined above. Let* $p \in (0,1)$, *integer* $q \geq 1$, $p_H = 1-(1-p)^q$, *and* $p_A = p$. *Let* $\Delta \geq 1$ *be the upper bound on the network delay. Consider* $\mathsf{Modular\text{-}Ledger\text{-}Protocol}_T$ *in the* $(\mathcal{G}_{\text{CLOCK}}, \mathcal{F}_{\text{STX}}^{\Delta,p_H,p_A}, \mathcal{F}_{\text{N-MC}}^{\Delta})$*-hybrid world. If, for some* $\lambda > 1$, *the relation*

$$\alpha \cdot (1 - 2 \cdot (\Delta+1) \cdot \alpha) \geq \lambda \cdot \beta \tag{1}$$

is satisfied in any real-world execution, where α and β are defined as above, then the protocol $\mathsf{Modular\text{-}Ledger\text{-}Protocol}_T$ *UC-realizes* $\mathcal{G}_{\text{LEDGER}}^{\not B}$ *for any ledger parameters (which are positive and integer-valued) in the range*

$$\mathtt{windowSize} = T \quad \text{and} \quad \mathtt{Delay} = 4\Delta,$$

$$\mathtt{maxTime}_{\mathtt{window}} \geq \frac{2 \cdot \mathtt{windowSize}}{(1-\delta)\cdot\gamma} \quad \text{and} \quad \mathtt{minTime}_{\mathtt{window}} \leq \frac{2 \cdot \mathtt{windowSize}}{(1+\delta)\cdot \max_r \mathtt{T}_{mp}(r)},$$

$$\eta > (1+\delta)\cdot \mathtt{windowSize}\cdot \frac{\beta}{\gamma},$$

where $\gamma := \frac{\alpha}{1+\Delta\alpha}$ and $\delta > 0$ is an arbitrary constant. In particular, the realization is perfect except with probability $R \cdot negl(T)$, where R denotes the upper bound on the number of rounds (consisting of two clock-ticks), and $negl(T)$ denotes a negligible function in T.

Remark 5. It is worth noting the implications of Equation 1. In practice, typically p is small such that α (and thus γ) can be approximated using Bernoulli's inequality to be $(1-\rho)mp$, where m is the estimated number of hash queries in the Bitcoin network per round. Hence, by canceling out the term mp and letting p be sufficiently small (compared to $\frac{1}{\Delta m}$), Equation 1 collapses roughly to the condition that $(1-\rho)(1-\epsilon) \geq (1+\delta)\rho$, which basically relates the fractions of adversarial vs. honest mining power. Also, as pointed out by [PSS17], for too large values of p in the order of $p > \frac{1}{mp}$, Equation 1 is violated for any constant fraction ρ of corrupted miners and they present an attack in this case.

Proof (Overview). To show the theorem we specify a simulator for the ideal world that internally runs the round-based mining procedure of every honest party. Whenever the real world parties complete a working round, then the simulator has to assemble the views of all honest (and synchronized) miners that it simulates and determine their common prefix of states, i.e., the longest state stored or received by each simulated party when chopping off T blocks. The adversary will then propose a new block candidate, i.e., a list of transactions, to the ledger to announce that the common prefix has increased. To reflect that not all parties have the same view on this common prefix, the simulator can adjust the state pointers accordingly. This simulation is perfect and corresponds to an emulation of real-world processes. What possibly prevents a perfect simulation is the requirement of a consistent prefix and the restrictions imposed by $\mathsf{ExtendPolicy}$. In order to show that these restrictions do not forbid a proper simulation, we have to justify why the choice of the parameters in the theorem statement is acceptable. To this end, we analyze the real-world execution to bound the corresponding bad events that prevent a perfect simulation. This can be done following the detailed analysis provided by Pass, Seeman, and shelat [PSS17] which includes the necessary claims for lower and upper on chain growth, chain quality, and prefix consistency. From these claims, it follows that our simulator can simulate the real-world, since the restrictions imposed by the ledger prohibit a prefect simulation only with probability $R \cdot \text{negl}(T)$. This is an upper bound on the distinguishing advantage of the real and ideal world. The detailed proof is found in Appendix D.3 ☐

Note that the theorem statement a-priori holds for any environment (but simply yields a void statement if the conditions are not met). In order to turn this into a composable statement, we follow the approach proposed in Section 2 and model restrictions as wrapper functionalities to ensure the condition of the theorem. We review two particular choices in 4.4. The general conceptual principle behind this is the following: For the hybrid world, that consists of a network $\mathcal{F}_{\text{N-MC}}$, a clock $\mathcal{G}_{\text{CLOCK}}$ and a random oracle $\mathcal{F}_{\text{RO}}$ with output length κ (or alternatively the state-exchange functionality $\mathcal{F}_{\text{StX}}$ instead of the random oracle), define a wrapper functionality $\mathcal{W}$ which ensures the condition in Equation 1 and (possibly) additional conditions on minimal (honest) and maximal (dishonest) mining power. This can be done by enforcing appropriate restrictions along the lines of the basic example in Section 2 (e.g., imposing an upper bound on parties, or RO queries per round etc.). We provide the details and the specification of such a general random-oracle wrapper $\mathcal{W}_{\alpha,\beta,\mathtt{D}}^{\Delta,\lambda,\mathtt{T}_{mp}}(\mathcal{F}_{\text{RO}})$ with its parameters[19] in Appendix B.

For this wrapper we have the following desired corollary to Theorem 1 and Lemma 1. This statement is guaranteed to compose according to the UC composition theorem.

Corollary 1. *The protocol $\mathsf{Ledger\text{-}Protocol}_{q,D,T}$ that is executed in the $(\mathcal{G}_{\text{CLOCK}}, \mathcal{F}_{\text{N-MC}}^{\Delta}, \mathcal{W}_{\alpha,\beta,D}^{\Delta,\lambda,T_{mp}}(\mathcal{F}_{\text{RO}}))$-hybrid world, UC-realizes functionality $\mathcal{G}_{\text{LEDGER}}^{\not B}$ (with the respective parameters assured by Theorem 1).*

[19] The parameters are the ones introduced in this section: a lower bound on honest mining power (per round) α, an upper bound on adversarial mining power (per round) β, the total mining power (per round) $\mathtt{T}_{mp}$, the network delay Δ, the difficulty parameter $\mathtt{D}$ (that influences the probability of a successful PoW), and finally a value $\lambda > 1$ describing the required gap between honest and dishonest mining power.

4.4 Comparison with Existing Work

We demonstrate how the protocols, assumptions, and results from the two existing works analyzing security of Bitcoin (in a property based manner) can be cast as special cases of our construction.

We start with the result in [GKL15], which is the so-called flat and synchronous model[20] with instant delivery and a constant number of parties n (i.e., Bitcoin is seen as an n-party MPC protocol). [21] Consider the concrete values for α and β as follows:

- Let n denote the number of parties. Each corrupted party gets at most q activations to query the $\mathcal{F}_{\text{StX}}$ per round. Each honest party is activated exactly once per round.
- In the model of GKL, we have $q \geq 1$. Thus, we get $p_H = 1-(1-p)^q$ and $p_A = p$. We can further conclude that $\mathtt{T}_{mp}^{\text{GKL}}(r) \leq p \cdot q \cdot n$.
- The adversary gets (at most) q queries per corrupted party with probability $p_A = p$ and one query per honest but desynchronized party with success probability $p_H = 1-(1-p)^q$. If t_r denotes the number of corrupted or desynchronized parties in round r, we get $\mathtt{mp}_A^{\text{GKL}}(r) \leq t_r \cdot q \cdot p$ and thus $\beta_r^{\text{GKL}} = p \cdot q \cdot (\rho \cdot n)$, where ρn is the (assumed) upper bound on the number of miners contributing to the adversarial mining power (independent of r).
- Each honest and synchronized miner gets exactly one activation per round, i.e., $q_{i,r} := 1$, with $p_H = 1-(1-p)^q \in (0,1)$, for some integer $q > 0$. Inserting it into the general equation yields $\alpha_r^{\text{GKL}} = 1-(1-p)^{q(1-\rho)\cdot n}$ (independent of r). Note that since n is assumed to be fixed in their model, $q(1-\rho)\cdot n$ is in fact a lower bound on the honest and synchronized hashing power.

We can now easily specify a wrapper $\mathcal{W}_{\text{GKL}}$ as special case of the above general wrapper. In the hybrid world $(\mathcal{G}_{\text{CLOCK}}, \mathcal{W}_{\text{GKL}}(\mathcal{F}_{\text{StX}}^{\Delta,p_H,p_A}), \mathcal{F}_{\text{N-MC}}^{\Delta})$ this ensures the condition of Theorem 1and we arrive at the following composable statement:

Corollary 2. *The ledger protocol Modular-Ledger-Protocol$_T$ UC-realizes the functionality $\mathcal{G}_{\text{LEDGER}}^{B}$ in the $(\mathcal{G}_{\text{CLOCK}}, \mathcal{W}_{\text{GKL}}(\mathcal{F}_{\text{StX}}^{1,p_H,p_A}), \mathcal{F}_{\text{N-MC}}^{1})$-hybrid model (setting delay $\Delta = 1$) for the parameters assured by Theorem 1 for the above choice:*

$$\alpha^{\text{GKL}} = 1-(1-p)^{(1-\rho)\cdot q\cdot n} \quad \text{and} \quad \beta^{\text{GKL}} = p \cdot q \cdot (\rho \cdot n).$$

Similarly, we can instantiate the above values with the assumptions of [PSS17]:

- For each corrupted (and desynchronized) party, the adversary gets at most one query per round. Each honest miner makes exactly one query per round. This means that $q_A^{(r)} + q_H^{(r)} = n_r$.
- In the PSs model, $p_H = p_A = p$ and hence $\mathtt{T}_{mp}^{\text{PSs}}(r) \leq p \cdot n_r = p \cdot n$, where n is as above. With these values we get $\mathtt{mp}_A^{\text{PSs}}(r) = p \cdot n_r^{\text{corr}}$ and consequently $\beta_r^{\text{PSs}} = p \cdot (\rho \cdot n)$, where ρn denotes the upper bound on corrupted parties in any round. Putting things together, we also have $\alpha_r^{\text{PSs}} = 1-(1-p)^{(1-\rho)\cdot n}$. Note that since n is assumed to be fixed in their model, $(1-\rho)\cdot n$ is in fact a lower bound on the honest and synchronized hashing power.

We can again specify a wrapper functionality $\mathcal{W}_{\text{PSs}}$ as above (where the restriction is 1 query per corrupted instead of q). We again have that the hybrid world $(\mathcal{G}_{\text{CLOCK}}, \mathcal{W}_{\text{PSs}}(\mathcal{F}_{\text{StX}}^{\Delta,p,p}), \mathcal{F}_{\text{N-MC}}^{\Delta})$ will ensure the condition of the theorem and directly yields the following composable statement.

Corollary 3. *The protocol Modular-Ledger-Protocol$_T$ UC-realizes the ledger functionality $\mathcal{G}_{\text{LEDGER}}^{B}$ in the $(\mathcal{G}_{\text{CLOCK}}, \mathcal{W}(\mathcal{F}_{\text{StX}}^{\Delta,p,p}), \mathcal{F}_{\text{N-MC}}^{\Delta})$-hybrid model (with network delay $\Delta \geq 1$) for the parameters assured by Theorem 1 for the above choice:*

$$\alpha^{\text{PSs}} = 1-(1-p)^{(1-\rho)\cdot n} \quad \text{and} \quad \beta^{\text{PSs}} = p \cdot (\rho \cdot n).$$

[20] The flat model means that every party gets the same number of hash queries in every round.

[21] In a recent paper, the authors of [GKL15] propose an analysis of Bitcoin for a variable number of parties. Capturing the appropriate assumptions for this case, as a wrapper in our composable setting, is part of future work.

5 Implementing a Stronger Ledger

As already observed in [GKL15], the Bitcoin protocol makes use of digital signatures to protect transactions which allows it to achieve stronger guarantees. Informally, the stronger guarantee ensures that every transaction submitted by an honest miner will eventually make it into the state. Using our terminology, this means that by employing digital signatures, Bitcoin implements a stronger ledger. In this section we present this stronger ledger and show how such an implementation can be captured as a UC protocol which makes black-box use of the Ledger-Protocol to implement this ledger. The UC composition theorem makes such a proof immediate, as we do not need to think about the specifics of the invoked ledger protocol, and we can instead argue security in a world where this protocol is replaced by $\mathcal{G}_{\text{LEDGER}}^{\text{B}}$.

Protection of transactions using accounts. In Bitcoin, a miner creates an account ID AccountID by generating a signature key pair and hashing the public key. Any transaction of this party includes this account ID, i.e., $\texttt{tx} = (\mathsf{AccountID}, \texttt{tx}')$. An important property is that a transaction of a certain account cannot be invalidated by a transaction with a different account ID. Hence, to protect the validity of a transaction, upon submitting tx, party p_i has to sign it, append the signature and verification key to get a transaction $((\mathsf{AccountID}, \texttt{tx}'), vk, \sigma)$. The validation predicate now additionally has to check that the account ID is the hash of the public key and that the signature σ is valid with respect to the verification key vk. Roughly, an adversary can invalidate tx, only by either forging a signature relative to vk, or by possessing key pair whose hash of the public key collides with the account ID of the honest party. The details of the protocol and the validate predicate as pseudo-code are provided in Appendix E.

Realized ledger. The realized ledger abstraction, denoted by $\mathcal{G}_{\text{LEDGER}}^{\text{B}+}$, is formally specified in Appendix E. Roughly, it is a ledger functionality as the one from the previous section, but which additionally allows parties to create unique accounts. Upon receiving a transaction from party p_i, $\mathcal{G}_{\text{LEDGER}}^{\text{B}+}$ only accepts a transaction containing the AccountID that was previously associated to p_i and ensures that parties are restricted to issue transactions using their own accounts.

Amplification of transaction liveness. In Bitcoin a given transaction can only be invalidated due to another one with the same account. By definition of the enhanced ledger, this means that no other party can make a transaction of p_i not enter the state. The liveness guarantee for transactions specified by ExtendPolicy in the previous chapter captures that if a valid transaction is in the buffer for long enough then it eventually enters the state. For $\mathcal{G}_{\text{LEDGER}}^{\text{B}+}$, this implies that if p_i submits a single transaction which is valid according to the current state, then this transaction will eventually be contained in the state. More precisely, we can conclude that this essentially happens within the next $3 \cdot \texttt{windowSize}$ new state blocks in the worst case (neglecting the offset in the beginning). Relative to the current view of p_i this is no more than within the next $4 \cdot \texttt{windowSize}$ blocks as argued in Appendix E.

References

[AD15] Marcin Andrychowicz and Stefan Dziembowski. PoW-based distributed cryptography with no trusted setup. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 379–399. Springer, Heidelberg, August 2015.

[ADMM14a] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Fair two-party computations via bitcoin deposits. In Rainer Böhme, Michael Brenner, Tyler Moore, and Matthew Smith, editors, *FC 2014 Workshops*, volume 8438 of *LNCS*, pages 105–121. Springer, Heidelberg, March 2014.

[ADMM14b] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 443–458. IEEE Computer Society Press, May 2014.

[BCG+14] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014.

[BDOZ11] Moshe Babaioff, Shahar Dobzinski, Sigal Oren, and Aviv Zohar. On bitcoin and red balloons. *SIGecom Exchanges*, 10(3):5–9, 2011.

[BHMQU05] Michael Backes, Dennis Hofheinz, Jörn Müller-Quade, and Dominique Unruh. On fairness in simulatability-based cryptographic systems. FMSE '05, 2005.

[BK14] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 421–439. Springer, Heidelberg, August 2014.

[But13] Vitalik Buterin. A next-generation smart contract and decentralized application platform, 2013. `https://github.com/ethereum/wiki/wiki/White-Paper`.

[Can00] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13(1):143–202, 2000.

[Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.

[CDPW07] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 61–85. Springer, Heidelberg, February 2007.

[CGHZ16] Sandro Coretti, Juan A. Garay, Martin Hirt, and Vassilis Zikas. Constant-round asynchronous multi-party computation based on one-way functions. In *ASIACRYPT*, 2016.

[CSV16] Ran Canetti, Daniel Shahaf, and Margarita Vald. Universally composable authentication and key-exchange with global PKI. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part II*, volume 9615 of *LNCS*, pages 265–296. Springer, Heidelberg, March 2016.

[ES14] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In Nicolas Christin and Reihaneh Safavi-Naini, editors, *FC 2014*, volume 8437 of *LNCS*, pages 436–454. Springer, Heidelberg, March 2014.

[Eya15] Ittay Eyal. The miner's dilemma. In *2015 IEEE Symposium on Security and Privacy*, pages 89–103. IEEE Computer Society Press, May 2015.

[GKL15] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 281–310. Springer, Heidelberg, April 2015.

[GKL16] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol with chains of variable difficulty. Cryptology ePrint Archive, Report 2016/1048, 2016.

[KB14] Ranjit Kumaresan and Iddo Bentov. How to use bitcoin to incentivize correct computations. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 14*, pages 30–41. ACM Press, November 2014.

[KB16] Ranjit Kumaresan and Iddo Bentov. Amortizing secure computation with penalties. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 16*, pages 418–429. ACM Press, October 2016.

[KKKT16] Aggelos Kiayias, Elias Koutsoupias, Maria Kyropoulou, and Yiannis Tselekounis. Blockchain mining games. In *EC*, 2016.

[KMB15] Ranjit Kumaresan, Tal Moran, and Iddo Bentov. How to use bitcoin to play decentralized poker. In Indrajit Ray, Ninghui Li, and Christopher Kruegel:, editors, *ACM CCS 15*, pages 195–206. ACM Press, October 2015.

[KMTZ13] Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Universally composable synchronous computation. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 477–498. Springer, Heidelberg, March 2013.

[KVV16] Ranjit Kumaresan, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. Improvements to secure computation with penalties. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 16*, pages 406–417. ACM Press, October 2016.

[KZZ16] Aggelos Kiayias, Hong-Sheng Zhou, and Vassilis Zikas. Fair and robust multi-party computation using a global transaction ledger. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 705–734. Springer, Heidelberg, May 2016.

[Lam98] Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, 1998.

[Lam02] Leslie Lamport. Paxos made simple, fast, and byzantine. In *OPODIS*, 2002.

[LSP82] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.

[MGGR13] Ian Miers, Christina Garman, Matthew Green, and Aviel D. Rubin. Zerocoin: Anonymous distributed E-cash from Bitcoin. In *2013 IEEE Symposium on Security and Privacy*, pages 397–411. IEEE Computer Society Press, May 2013.

[Nak08] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. `http://bitcoin.org/bitcoin.pdf`.

[PS16] Rafael Pass and Elaine Shi. FruitChains: A fair blockchain. Cryptology ePrint Archive, Report 2016/916, 2016. `http://eprint.iacr.org/2016/916`.

[PSS17] Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part II*, pages 643–673. Springer, 2017.

[Rab83] Michael O. Rabin. Randomized byzantine generals. In *FOCS*, 1983.

[Ros12] Mike Rosulek. Must you know the code of f to securely compute f? In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 87–104. Springer, Heidelberg, August 2012.

[SZ15] Yonatan Sompolinsky and Aviv Zohar. Secure high-rate transaction processing in bitcoin. In Rainer Böhme and Tatsuaki Okamoto, editors, *FC 2015*, volume 8975 of *LNCS*, pages 507–527. Springer, Heidelberg, January 2015.

[Zoh15] Aviv Zohar. Bitcoin: under the hood. *Commun. ACM*, 58(9):104–113, 2015.

A Related Literature

At a high level, the Bitcoin works as follows: The parties (also referred to as *miners*) collect and circulate messages (transactions) from users of the network, check that they satisfy some commonly agreed validity property, put the valid transactions into a block, and them try to find appropriate metadata such that the hash of the block-contents and this metadata is of a specific form—concretely that, parsed as a binary string, it has a sufficient number of leading zeros. This is often referred to as a solving a mining puzzle. If we assume that the hash function in this experiment is fully unpredictable (i.e., before computing it one has probability $2^{-\kappa}$ of predicting its output, where κ is the length of the hash), then the best strategy for finding such metadata is by trial-and-error. Thus, informally, the probability that some party finds appropriate metadata increases proportional to the number of times some party attempts a hash computation. And the more leading zeros we require from a correct puzzle solution the harder it is to find one.

Intuitively, a successful solution can be seen as a *proof of work* that testifies to the fact that the miner presenting is tried a large number of hash queries. Once a miner finds such a solution, he puts it into a block and sends it to the other miners. The miners who receive it check that it satisfies some validity property (see below) and if so create new metadata using the hash of this (newly minted) block and put this metadata together with transactions that are still valid into a new block and start working on solving the puzzle induced by this block. This creates for each miner a sequence of seen miner blocks. Moreover, a block is rendered valid by any miner only if it includes a hash-pointer the last valid block that thus miner has seen. Thus, the sequence of valid blocks forms a linked list which is often referred to as the *blockchain*.

The works of Garay, Kiayias, and Leonardos [GKL15] and that of Pass, Seeman, and shelat [PSS17] include a formal specification and security proof of the Bitcoin protocol.[22] However, the proved security in these works is property-based, i.e., informally, they prove that conditioned on the largest part of the network following the Bitcoin protocol (in fact and abstraction and generalization thereof), the output of bitcoin satisfies certain properties.

These properties are as follows:

- The *common prefix property* from [GKL15] is a property of the linked list of blocks commonly referred to as the *blockchain* that is created by executing Bitcoin. It requires that if a block $\mathbf{B}$ is "deep enough" in a blockchain $\mathcal{C}$ that is considered valid by some honest party, then the prefix-subchain $\mathcal{C}|_{\mathbf{B}}$ of $\mathcal{C}$ that results by ignoring in $\mathcal{C}$ all blocks that are after $\mathbf{B}$ will eventually become a prefix-subchain of all honest parties. In [PSS17] this property was refined (and augmented) by requiring a similar *consistency* property which in addition to the above mandates that for every honest party, cutting his chain at a deep enough block, yields a prefix-chain that will be prefix for ever. These properties are satisfied with overwhelming probability (in a security parameter κ), where deep enough is defined as $O(\kappa)$.
- The *chain growth*, which was implicitly defined in [GKL15] and posted as a required property in [PSS17], mandates, informally, that the blockchain of honest parties will increase with time. More concretely, it postulates a lower bound on the speed in which blocks are added in the blockchain. In fact, [PSS17] introduced also the notion of a corresponding upper bound which, as the argue, might be useful for synchronization purposes.We observer that although [PSS17] adopts a round-based model of execution (see below) which makes it unclear where or how such extra synchronization would help, the proposed property would indeed be useful to have if one was to remove that assumed synchronous structure from the model.[23]
- The *chain quality* [GKL15, PSS17] property postulates that the honest miners get to insert a fraction of the blocks that eventually make it into the common prefix of honest parties.[24] This, means that the adversary is not able to monopolize the blocks that are inserted in the blockchain, and every so often an honest miner is allowed to do so. This abstracts the property which was already postulated in the original blockchain whitepaper that the Bitcoin implements a lottery which decides (according to the computing power that each party contributes—alas not necessarily in a fair allocation manner [PS16]) who will insert that next block to the chain.

[22] An extension of [GKL15] to allow capturing dynamically evolving miner-set was also recently posted on the IACR eprint [GKL16].

[23] Such an adaptation would however harm several of the arguments and would require a new analysis and protocol.

[24] More precisely, a state block is considered honest if the respective miner that proposed the block was honest at the time of the proposal.

Both [GKL15] and [PSS17] assume a multicast network—i.e., a network where a party sends messages to arbitrary other parties[25]—and abstract the hash as a random oracle. Furthermore, they both have an explicit round-based model of execution where parties proceed in rounds. The main differences between the two are: (1) that in [GKL15] every party make q of hash-queries (i.e., q hash calls) in each round as opposed to [PSS17] where every party exactly one hash-query per round; and (2) in [GKL15] message sent in some round r are guaranteed to be delivered at the beginning of round $r+1$, whereas in [PSS17] the adversary might choose to delay message delivery but the statements are proved assuming no message is delayed by more than Δ rounds. We note that Property 1 implies that both models assume that parties are fully synchronized, since in both models every party can simply count the number of queries it has made to the random oracle and decide in what round it is (and by assumption this decision will ensure that no party goes to round $r+1$ before all parties have finished round r). Notwithstanding, as we argue in Section 4 the network assumption in [PSS17] in combination with the design of the protocol yields *strictly* weaker setting than the one in [GKL15], i.e., one can achieve strictly more with the assumptions in [GKL15] even assuming an external synchronization mechanism. We will discuss these difference further when we describe these network assumptions and in particular in Section 4 where we argue how these two models are captured in UC and how the corresponding protocols can be cast as special case of our UC protocol.

Property-based vs simulation-based security. Proving that Bitcoin satisfies the above properties has been an essential step into the direction of understanding the security goals of Bitcoin. But as argued above, this does not offer the tool to be able to argue security of cryptographic protocols that use Bitcoin—e.g., to achieve an improved fairness notion [ADMM14a, ADMM14b, BK14, KVV16, KB16, KMB15, KB14, AD15]—without the need to always look at the Bitcoin specifics. In other words, such property based security definitions do not support composition. The standard way to allow for such a generic use of blockchain protocols as a cryptographic resource, is to prove that it implements an ideal functionality in a one of the composable simulation-based frameworks, e.g., [Can01, Can00, CDPW07]. In such a framework, security of a protocol is defined as follows: First we specify the goal the protocol is supposed to achieve by means of a trusted third party, usually referred to as the *ideal functionality* $\mathcal{F}$. Then we prove that the protocol implements this functionality $\mathcal{F}$ which means that for any adversary $\mathcal{A}$ attacking the protocol, there exists an ideal adversary (aka a *simulator*) $\mathcal{S}$ that attacks an ideal invocation of $\mathcal{F}$ and emulates the attach that $\mathcal{A}$ launches to the protocol. We assume that the reader has some familiarity with simulation-based security, and in particular with the UC and GUC frameworks of Canetti et al. [Can01, CDPW07].

The advantage of simulation-based security is that it often comes with a composition theorem which, intuitively, states that we can replace calls to a functionality with invocation of a protocol implementing it without worrying about the protocol's internals. Thus, in our case, a ledger functionality that is implemented by Bitcoin allows us to use this functionality in any protocol that wants to use Bitcoin as a resource, and the composition theorem will then imply that replacing the functionality with Bitcoin does not compromise security. Note that [KZZ16] already includes an attempt to define such a functionality, but as we argue here, the proposed functionality is too strong to be implemented from Bitcoin under standard assumptions.

B Model (Cont'd)

This section includes complementary material for Section 2.

B.1 The Unicast Channel

Functionality $\mathcal{F}_{\text{U-CH}}^{\Delta}$

The functionality is parametrized with a receiver p_R, and and upper bound Δ on the delay of any channel. I keeps track of the set of possible senders $\mathcal{P}$. Any newly registered (resp. deregistered) party is added to (resp. deleted from) $\mathcal{P}$.

[25] Unlike [GKL15] where this operation is referred to as broadcast, we choose to call it multicast here to avoid confusion with the standard broadcast primitive in the Byzantine agreement literature that offers stronger consistency guarantees.

- Upon receiving a message (SEND, m) from some $p_s \in \mathcal{P}$ or from the adversary $\mathcal{A}$, choose a new unique message-ID midfor m, initialize variables $D_{\text{mid}} := 1$ and $D_{\text{mid}}^{MAX} = 1$, set $\vec{M} := \vec{M}||(m, \text{mid}, D_{\text{mid}})$, and send $(m, \text{mid}, D_{\text{mid}})$ to the adversary.
- Upon receiving a message (FETCH) from p_R:
  1. For all regired mids, set $D_{\text{mid}} := D_{\text{mid}} - 1$.
  2. Let $\vec{M}_0$ denote the subvector $\vec{M}$ including all triples $(m, \text{mid}, D_{\text{mid}})$ with $D_{\text{mid}} = 0$ (in the same order as they appear in $\vec{M}$). Delete all entries in $\vec{M}_0$ from $\vec{M}$ and send $\vec{M}_0$ to p_R.
- Upon receiving a message $(\text{DELAY}, T_{\text{mid}}, \text{mid})$ from the adversary, if $D_{\text{mid}}^{MAX} + T_{\text{mid}} \leq \Delta$ and mid is a message-ID registered in the current $\vec{M}$, set $D_{\text{mid}} := D_{\text{mid}} + T_{\text{mid}}$ and $D_{\text{mid}}^{MAX} := D_{\text{mid}}^{MAX} + T_{\text{mid}}$; otherwise, ignore the message.
- Upon receiving a message $(\text{SWAP}, \text{mid}, \text{mid}')$ from the adversary, if mid and mid′ are message-IDs registered in the current $\vec{M}$, then swap the triples $(m, \text{mid}, D_{\text{mid}})$ and $(m, \text{mid}', D_{\text{mid}'})$ in $\vec{M}$. Return (SWAP-OK) to the adversary.

B.2 The Multicast Network

The multicast network with bounded delay is described in the following.

Functionality $\mathcal{F}_{\text{N-MC}}^{\Delta}$

The functionality is parametrized with a set possible senders and receivers $\mathcal{P}$. Any newly registered (resp. deregistered) party is added to (resp. deleted from) $\mathcal{P}$.

- *Honest sender multicast:*
  Upon receiving a message $(\text{MULTICAST}, \text{sid}, m)$ from some $p_s \in \mathcal{P}$, where $\mathcal{P} = \{p_1, \ldots, p_n\}$ denotes the current party set, choose n new unique message-IDs $\text{mid}_1, \ldots, \text{mid}_n$, initialize $2n$ new variables $D_{\text{mid}_1} := D_{\text{mid}_1}^{MAX} \ldots := D_{\text{mid}_n} := D_{\text{mid}_n}^{MAX} := 1$, set $\vec{M} := \vec{M}||(m, \text{mid}_1, D_{\text{mid}_1}, p_1)|| \ldots ||(m, \text{mid}_n, D_{\text{mid}_n}, p_n)$, and send $(\text{MULTICAST}, \text{sid}, m, p_s, (p_1, \text{mid}_1), \ldots, (p_n, \text{mid}_n))$ to the adversary.
- *Adversarial sender (partial) multicast:*
  Upon receiving a message $(\text{MULTICAST}, \text{sid}, (m_{i_1}, p_{i_1}), \ldots, (m_{i_\ell}, p_{i_\ell})$ from the adversary with $\{p_{i_1}, \ldots, p_{i_\ell}\} \subseteq \mathcal{P}$, choose ℓ new unique message-IDs $\text{mid}_{i_1}, \ldots, \text{mid}_{i_\ell}$, initialize ℓ new variables $D_{\text{mid}_{i_1}} := D_{\text{mid}_{i_1}}^{MAX} := \ldots := D_{\text{mid}_{i_\ell}} := D_{\text{mid}_{i_\ell}}^{MAX} := 1$, set $\vec{M} := \vec{M}||(m_{i_1}, \text{mid}_{i_1}, D_{\text{mid}_{i_1}}, p_{i_1})|| \ldots ||(m_{i_\ell}, \text{mid}_{i_\ell}, D_{\text{mid}_{i_\ell}}, p_{i_\ell})$, and send $((\text{MULTICAST}, \text{sid}, (m_{i_1}, p_{i_1}, \text{mid}_{i_1}), \ldots, (m_{i_\ell}, p_{i_\ell}, \text{mid}_{i_\ell})$ to the adversary.
- *Honest party fetching:*
  Upon receiving a message $(\text{FETCH}, \text{sid})$ from $p_i \in \mathcal{P}$ (or from $\mathcal{A}$ on behalf of p_i if p_i is corrupted):
  1. For all tuples $(m, \text{mid}, D_{\text{mid}}, p_i) \in \vec{M}$, set $D_{\text{mid}} := D_{\text{mid}} - 1$.
  2. Let $\vec{M}_0^{p_i}$ denote the subvector $\vec{M}$ including all tuples of the form $(m, \text{mid}, D_{\text{mid}}, p_i)$ with $D_{\text{mid}} = 0$ (in the same order as they appear in $\vec{M}$). Delete all entries in $\vec{M}_0^{p_i}$ from $\vec{M}$, and send $\vec{M}_0^{p_i}$ to p_i.
- *Adding adversarial delays:*
  Upon receiving a message $(\text{DELAYS}, \text{sid}, (T_{\text{mid}_{i_1}}, \text{mid}_{i_1}), \ldots, (T_{\text{mid}_{i_\ell}}, \text{mid}_{i_\ell}))$ do the following for each pair $(T_{\text{mid}_{i_j}}, \text{mid}_{i_j})$ in this message:
  If $D_{\text{mid}_{i_j}}^{MAX} + T_{\text{mid}_{i_j}} \leq \Delta$ and mid is a message-ID registered in the current $\vec{M}$, set $D_{\text{mid}_{i_j}} := D_{\text{mid}_{i_j}} + T_{\text{mid}_{i_j}}$ and set $D_{\text{mid}_{i_j}}^{MAX} := D_{\text{mid}_{i_j}}^{MAX} + T_{\text{mid}_{i_j}}$; otherwise, ignore this pair.
- *Adversarially reordering messages:*
  Upon receiving a message $(\text{SWAP}, \text{sid}, \text{mid}, \text{mid}')$ from the adversary, if mid and mid′ are message-IDs registered in the current $\vec{M}$, then swap the triples $(m, \text{mid}, D_{\text{mid}}, \cdot)$ and $(m, \text{mid}', D_{\text{mid}'}, \cdot)$ in $\vec{M}$. Return $(\text{SWAP}, \text{sid})$ to the adversary.

On implementing a multicast network. We briefly sketch how to realize such a multicast network, in particular its synchronized version along the lines of [KMTZ13], by means of a synchronized message-diffusion protocol over a network of unicast channels (and implicitly assuming a local clock to obtain the round structure). The core of this diffusion protocol are the assumed and known (e.g., by a common list of IP addresses) relay-nodes to which parties thus can connect and which forward in each round all new messages they received (either from registered parties or other relay nodes) in the previous round

to all the unicast channels they are connected to as senders.[26] Let $G = (V, E)$ denote the (dynamically updatable) directed graph whose vertices V are the parties and the relay-nodes which are currently participating in the execution and an edge (p_i, p_j) is in E iff p_i is one of the senders of the multicast channel with receiver p_j. It is straightforward to verify that provided that G restricted to the honest parties (i.e., when corrupted parties and the edges that use them are deleted from G) remains *strongly connected* (i.e., there is a directed path between any two honest parties, in either direction), then the diffusion mechanism executed over unicast channels with delay at most Δ security realizes a multicast network with delay Δd where d is an upper bound of the diameter of G. Indeed, the simulator, which is given any message submitted to any unicast channel and enough activations when the dummy parties themselves get activated (note that it is essentially a synchronous computation among the relay-nodes), needs to simply simulate when the respective parties would see a message and schedule the corresponding deliveries by using the delays submitted by the adversary. The fact that each channel has at most Δ delay means that it will take delay at most ΔL rounds for it to travel through an honest path of length L. Last but not least, in order to receive messages from the network established this way, when a party joins the network, it has to multicast a special message to the relay-nodes that has to contain its identifier such that the relay-nodes can start sending messages to that party. This induces at most a delay of Δ rounds until the party is guaranteed to receive the messages sent over the network. For simplicity and for the sake of presentation, we ignore this additional delay incurred by the registration to the network, and we therefore also omit it in our specification of the multicast functionality.[27]

B.3 The Clock

The global clock functionality, i.e., a shared clock that may interact with more than one protocol session, is denoted by $\bar{\mathcal{G}}_{\text{CLOCK}}$. The standard UC functionality is denoted by $\mathcal{G}_{\text{CLOCK}}$ (without the bar).

Functionality $\mathcal{G}_{\text{CLOCK}}$

The functionality is available to all participants. The functionality is parametrized with variable τ, a set of parties $\mathcal{P}'$, and a set F of functionalities. For each party $p \in \mathcal{P}'$ it manages variable d_p. For each $\mathcal{F} \in F$ it manages variable $d_{\mathcal{F}}$

Initially, $\tau := 0$, $\mathcal{P}' := \emptyset$ and $F := \emptyset$.

Synchronization:

- Upon receiving $(\text{CLOCK-UPDATE}, \text{sid}_C)$ from some party $p \in \mathcal{P}'$ set $d_p := 1$; execute *Round-Update* and forward $(\text{CLOCK-UPDATE}, \text{sid}_C, p)$ to $\mathcal{A}$.
- Upon receiving $(\text{CLOCK-UPDATE}, \text{sid}_C)$ from some functionality $\mathcal{F} \in F$ set $d_{\mathcal{F}} := 1$, execute *Round-Update* and return $(\text{CLOCK-UPDATE}, \text{sid}_C, \mathcal{F})$ to $\mathcal{F}$.
- Upon receiving $(\text{CLOCK-READ}, \text{sid}_C)$ from any participant (including the environment, the adversary, or any ideal—shared or local—functionality) return $(\text{CLOCK-READ}, \text{sid}_C, \tau)$ to the requestor.

Procedure Round-Update:
If $d_{\mathcal{F}} := 1$ for all $\mathcal{F} \in F$ and $d_p = 1$ for all honest p in $\mathcal{P}'$, then set $\tau := \tau + 1$ and reset $d_{\mathcal{F}} := 0$ and $d_p := 0$ for all parties in $\mathcal{P}'$.

B.4 The Random Oracle Functionality

Functionality $\mathcal{F}_{\text{RO}}$

The functionality is parametrized by a security parameter κ. It maintains a set of registered parties/miners $\mathcal{P}$ (initially set to $\emptyset$) and a (dynamically updatable) function table $\mathcal{T}$ (initially $\mathcal{T} = \emptyset$). For simplicity we write $T[x] = \perp$ to denote the fact that no pair of the form $(x, \cdot)$ is in $\mathcal{T}$.

[26] In order to ensure that parties can send some messages twice, a nonce is attached to each input message that is to be multicasted. The relayers do not add another nonce to the message they relay.

[27] Formally, one would have to define an additional party set $\mathcal{P}'$ which contains all parties that have joined (and not yet left) the network at least Δ rounds ago. All the guarantees would then hold the same, but with respect to the set $\mathcal{P}'$ instead of the party set $\mathcal{P}$ in the above multicast specification.

- Upon receiving $(\text{EVAL}, \text{sid}, x)$ from some party $p \in \mathcal{P}$ (or from $\mathcal{A}$ on behalf of a corrupted p), do the following:
  1. If $H[x] = \bot$ sample a value y uniformly at random from $\{0,1\}^\kappa$, set $H[x] \leftarrow y$ and add $(x, T[x])$ to $\mathcal{T}$.
  2. Return $(\text{EVAL}, \text{sid}, x, H[x])$ to the requestor.

B.5 Wrapping Functionalities

Basic example. We first provide a basic introductory example that illustrates how one can employ a wrapper functionality to enforce an upper bound on the number of adversarial queries to the random oracle (per round).

Functionality $\mathcal{W}^q(\mathcal{F}_{\text{RO}})$

The wrapper functionality is parametrized by an upper bound q which restricts the $\mathcal{F}$-evaluations of each corrupted party per round. (To keep track of rounds the functionality registers with the global clock $\mathcal{G}_{\text{CLOCK}}$.) The functionality manages the variable `counter` and the current set of corrupted miners $\mathcal{P}'$. For each party $p \in \mathcal{P}'$ it manages variables $\texttt{count}_p$.

Initially, $\mathcal{P}' = \emptyset$ and $\texttt{counter} = 0$.

General:

- The wrapper stops the interaction with the adversary as soon as the adversary tries to exceed its budget of q queries per corrupted party.

Relaying inputs to the random oracle:

- Upon receiving $(\text{EVAL}, \text{sid}, x)$ from $\mathcal{A}$ on behalf of a corrupted party $p \in \mathcal{P}'$, then first execute *Round Reset*. Then, set $\texttt{count}_p \leftarrow \texttt{count}_p + 1$ and only if $\texttt{count}_p \leq q$ forward the request to $\mathcal{F}_{\text{RO}}$ and return to $\mathcal{A}$ whatever $\mathcal{F}_{\text{RO}}$ returns.
- Any other request from any participant or the adversary is simply relayed to the underlying functionality without any further action and the output is given to the destination specified by the hybrid functionality.

Standard UC Corruption Handling:

- Upon receiving $(\text{CORRUPT}, \text{sid}, p)$ from the adversary, set $\mathcal{P}' \leftarrow \mathcal{P}' \cup \{p\}$. If p has already issued $t > 0$ random oracle queries in this round, set $\texttt{count}_p \leftarrow t$. Otherwise set $\texttt{count}_p \leftarrow 0$.

Procedure Round-Reset:
Send $(\text{CLOCK-READ}, \text{sid}_C)$ to $\bar{\mathcal{G}}_{\text{CLOCK}}$ and receive $(\text{CLOCK-READ}, \text{sid}_C, \tau)$ from $\bar{\mathcal{G}}_{\text{CLOCK}}$. If $\left|\tau - \texttt{counter}\right| > 0$ and the new time τ is even (i.e., a new round started), then set $\texttt{count}_p := 0$ for each participant $p \in \mathcal{P}'$ and set $\texttt{counter} \leftarrow \tau$.

Bitcoin Assumptions as a wrapper. We provide here the more elaborate version of a wrapped random-oracle functionality which controls the amount of hashing power of honest parties relative to the adversary. To be more flexible, we also allow to specify an upper and lower bound on the computing power that the environment has to provide. Otherwise, the execution is not allowed to advance.

Functionality $\mathcal{W}_{\alpha,\beta,\texttt{D}}^{\Delta,\lambda,\texttt{T}_{mp}}(\mathcal{F}_{\text{RO}})$

The wrapper functionality is parametrized by the parameters for network delay, lower bound on honest mining power, upper bound on adversarial mining power, the difficulty, the upper bound $\texttt{T}_{mp}$ on the total mining power per round (which thereby also implies an upper bound on the total number of RO-queries per round), and a value $\lambda > 1$ (the parameter that describes the gap between the honest and adversarial mining power). The wrapper is assumed to be registered with the global clock $\mathcal{G}_{\text{CLOCK}}$. The functionality manages the variable `counter` and is aware of set of registered parties, and the set of corrupted parties.

Initially, $\mathcal{P}' = \emptyset$ and $\texttt{counter} = 0$, $q_A = 0$ and $q_H = 0$. Define $p := \frac{\texttt{D}}{2^\kappa}$ (where κ is the output length of the underlying random oracle).

General:

– The wrapper stops the interaction with the adversary as soon as the adversary tries to exceed its allowed budget of hashing power.

Relaying inputs to the random oracle:

– Upon receiving ($\textsc{eval}, \mathrm{sid}, x$) from $\mathcal{A}$ on behalf of a party P which is corrupted or registered but de-synchronized, then first execute *Round Reset*. Then do the following:

  $q_A \leftarrow q_A + 1$; $\beta_{\mathsf{counter}} \leftarrow q_A \cdot p$
  if $(q_A + q_H) \cdot p \leq \mathtt{T}_{mp}$ **then**
    if $\beta_{\mathsf{counter}} \leq \beta \ \wedge\ \alpha_{\mathsf{counter}} \cdot (1 - 2 \cdot (\Delta + 1) \cdot \alpha_{\mathsf{counter}}) \geq \lambda \cdot \beta_{\mathsf{counter}}$ **then**
      Forward the request to $\mathcal{F}_{\mathrm{RO}}$ and return to $\mathcal{A}$ whatever $\mathcal{F}_{\mathrm{RO}}$ returns.
    end if
  end if

– Upon receiving ($\textsc{eval}, \mathrm{sid}, x$) from an uncorrupted, registered and synchronized party P, then first execute *Round Reset*. Then do the following:

  $q_H \leftarrow q_H + 1$; $\alpha_{\mathsf{counter}} \leftarrow 1 - (1 - p)^{q_H}$
  if $(q_A + q_H) \cdot p \leq \mathtt{T}_{mp}$ **then**
    if $\alpha_{\mathsf{counter}} \cdot (1 - 2 \cdot (\Delta + 1) \cdot \alpha_{\mathsf{counter}}) \geq \lambda \cdot \beta_{\mathsf{counter}}$ **then**
      Forward the request to $\mathcal{F}_{\mathrm{RO}}$ and return to P whatever $\mathcal{F}_{\mathrm{RO}}$ returns.
    end if
  end if
  if $\alpha_{\mathsf{counter}} \geq \alpha$ **then**
    Send ($\textsc{clock-update}, \mathrm{sid}_C$) to $\mathcal{G}_{\textsc{clock}}$ // Release the clock if lower bound is reached.
  end if

– Any other request is relayed to the underlying functionality (and recorded by the wrapper) and the corresponding output is given to the destination specified by the underlying functionality.

Standard UC Corruption Handling:

– Upon receiving ($\textsc{corrupt}, \mathrm{sid}, P$) from the adversary, set $\mathcal{P}' \leftarrow \mathcal{P}' \cup \{P\}$.

Procedure Round-Reset:
Send ($\textsc{clock-read}, \mathrm{sid}_C$) to $\bar{\mathcal{G}}_{\textsc{clock}}$ and receive ($\textsc{clock-read}, \mathrm{sid}_C, \tau$) from $\bar{\mathcal{G}}_{\textsc{clock}}$. If $\left|\tau - \mathtt{counter}\right| > 0$ and the new time τ is even (i.e., a new round started), then set $\mathtt{counter} \leftarrow \tau$ and set $q_A \leftarrow 0$ and $q_H \leftarrow 0$.

C Bitcoin Protocol (Cont'd)

This section includes complementary material for Section 4.1.

C.1 Algorithm isvalidstate

The algorithm isvalidstate takes a state $\vec{\mathtt{st}}$ as input and checks whether $\vec{\mathtt{st}}$ is valid with respect to $\mathsf{ValidTx}_{\text{B}}$, i.e., whether it was built by adding only valid transaction (note that this validity will be ensured by the ledger protocol below).

Algorithm isvalidstate($\vec{\mathtt{st}}$)

Let $\vec{\mathtt{st}} := \mathtt{st}_1 || \dots || \mathtt{st}_n$
for each $\mathtt{st}_i$ **do**
  Extract the transaction sequence $\vec{\mathtt{tx}}_i \leftarrow \mathtt{tx}_{i,1}, \dots, \mathtt{tx}_{i,n_i}$ contained in $\mathtt{st}_i$
end for
$\vec{\mathtt{st}}' \leftarrow \mathtt{gen}$ // Initialize the genesis state
for $i = 1$ to n **do**
  if the first transaction in $\vec{\mathtt{tx}}_i$ is not a coin-base transaction **return** false
  $\vec{N}_i \leftarrow \mathtt{tx}_{i,1}$
  for $j = 2$ to $|\vec{\mathtt{tx}}_i|$ **do**
    $\mathtt{st} \leftarrow \mathsf{blockify}_{\text{B}}(\vec{N}_i)$
    if $\mathsf{ValidTx}_{\text{B}}(\mathtt{tx}_{i,j}, \vec{\mathtt{st}}' || \mathtt{st}) = 0$ **return** false
    $\vec{N}_i \leftarrow \vec{N}_i || \mathtt{tx}_{i,j}$

end for
$\vec{\mathtt{st}}' \leftarrow \vec{\mathtt{st}}'||\mathtt{st}_i$
end for
return true

C.2 Algorithm extendchain

The algorithm takes a chain $\mathcal{C}$, a state block $\mathtt{st}$ and the number of attempts q as inputs. It tries to find a proof-of-work which allows to extend the $\mathcal{C}$ by a block which encodes $\mathtt{st}$.

Algorithm $\mathsf{extendchain}_D(\mathcal{C}, \mathtt{st}, q)$

Require: Chain $\mathcal{C}$ is valid with state $\vec{\mathtt{st}}$. The state $\vec{\mathtt{st}}||\mathtt{st}$ is valid.
Set $\mathbf{B} \leftarrow \bot$
$\mathtt{s} \leftarrow H[\mathtt{head}(\mathcal{C})]$ // Compute the pointer $\mathtt{s}$ of the new block
for $i \in \{1, \ldots, q\}$ **do**
Choose nonce $\mathtt{n}$ uniformly at random from $\{0,1\}^\kappa$ and set $\mathbf{B} \leftarrow \langle \mathtt{s}, \mathtt{st}, \mathtt{n} \rangle$.
if $H[\mathbf{B}] < \mathtt{D}$ **then**
break
end if
end for
if $\mathbf{B} \neq \bot$ **then**
$\mathcal{C} \leftarrow \mathcal{C}||\mathbf{B}$
end if
return $\mathcal{C}$

C.3 Ledger-Protocol

The Bitcoin ledger protocol is described in the following. It assumes as hybrids a random oracle $\mathcal{F}_{\text{RO}}$, a network $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$ for blockchains, a network $\mathcal{F}^{\mathsf{tx}}_{\text{N-MC}}$ for transcations, and clock $\mathcal{G}_{\text{CLOCK}}$. Note that the two networks are (named) instances of $\mathcal{F}^{\Delta}_{\text{N-MC}}$ and can be realized from a single network $\mathcal{F}^{\Delta}_{\text{N-MC}}$ using different message-IDs. The protocol is parametrized by q, $\mathtt{D}$, T where q is the number of mining attempts per round, $\mathtt{D}$ is the difficulty of the proof-of-work, and T is the number of blocks chopped off to get the ledger state.

For the sake of simplicity, we omit the registration part from the explicit protocol description and describe the general structure here: The registration process in the protocol works as follows. If a party receives (REGISTER, sid) from the environment it registers at the random oracle and the network. Since the clock is a shared functionality, the registrations are fully controlled by the environment and thus the protocol relays such registration queries to the clock. Only if a party is registered to the clock already, it reacts to such REGISTER queries and otherwise stays idle. Once registration has succeeded the party returns activation to the environment. Upon the next activation to maintain the ledger (MAINTAIN-LEDGER), the party initializes its local variables and executes the main mining procedure. In this first execution of the mining procedure, the party multicasts a special NEW-PARTY message over the network (in addition to the round messages). Consequently, in the mining procedure parties would additionally check if they have received such a NEW-PARTY message in the current round and if so, they multicast their transaction buffer in addition to their other round messages (such as longest chains) to share also these transactions.

De-registering from the ledger (via a query (DE-REGISTER, sid)) from the environment) works analogously, upon which the party erases all its state and becomes idle until its is freshly invoked with a REGISTER-query.

Protocol $\mathsf{Ledger\text{-}Protocol}_{q,\mathtt{D},T}(p)$

Initialization:

We assume that the party p is registered to $\mathcal{F}_{\text{RO}}$, $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$, $\mathcal{F}^{\mathsf{tx}}_{\text{N-MC}}$, and $\mathcal{G}_{\text{CLOCK}}$ (note that a non-idle party is otherwise considered de-synchronized). A party is considered not initialized after registration (define $\mathtt{isInit} \leftarrow$ false).
The protocol stores a local (working) chain $\mathcal{C}_{loc}$ which initially contains the genesis block, i.e., $\mathcal{C}_{loc} \leftarrow (\mathbf{G})$.
It additionally manages a separate chain $\mathcal{C}_{\exp}$ to store the current chain whose encoded state $\vec{\mathtt{st}}$ is exported as the ledger state (initially this chain contains the genesis block).

The protocol also manages an initially empty buffer `buffer` which contains the received transactions. The buffer forms the information base to build new chain blocks. Finally, set $t \leftarrow 0$.
Once the initialization is complete after the next MAINTAIN-LEDGER-command as described above (including requesting previous transactions), set `isInit` ← `true`

Clock:

- Upon receiving $(\textsc{clock-read}, \mathrm{sid}_C)$ forward the query to $\mathcal{G}_{\textsc{clock}}$ and return whatever is received as answer from $\mathcal{G}_{\textsc{clock}}$
- Upon receiving $(\textsc{clock-update}, \mathrm{sid}_C)$, do the following: remember that a clock-update was received in the current mini-round.

Ledger:

- Upon receiving $(\textsc{submit}, \mathrm{sid}, \mathtt{tx})$, set $\mathtt{buffer} \leftarrow \mathtt{buffer}||\mathtt{tx}$, and send $(\textsc{multicast}, \mathrm{sid}, \mathtt{tx})$ to $\mathcal{F}^{\mathsf{tx}}_{\text{N-MC}}$.
- Upon receiving $(\textsc{read}, \mathrm{sid})$ send $(\textsc{clock-read}, \mathrm{sid}_C)$ to $\mathcal{G}_{\textsc{clock}}$, receive as answer $(\textsc{clock-read}, \mathrm{sid}_C, \tau)$ and proceed as follows.

    // Fetch new states in update mini-rounds
    if τ corresponds to an update mini-round and $t < \tau$ and `isInit` **then**
        Execute **FetchInformation** and set $t \leftarrow \tau$.
    end if
    // Return current ledger state
    Let $\vec{\mathtt{st}}$ be the encoded state in $\mathcal{C}_{\mathrm{exp}}$
    Return $(\textsc{read}, \mathrm{sid}, \vec{\mathtt{st}}^{\lceil T})$.

- Upon receiving $(\textsc{maintain-ledger}, \mathrm{sid}, \mathrm{minerID})$ execute **MiningProcedure**

MiningProcedure:

Step 1: If a $(\textsc{clock-update}, \mathrm{sid}_C)$ has been received during this update mini-round then send $(\textsc{clock-update}, \mathrm{sid}_C)$ to $\mathcal{G}_{\textsc{clock}}$ (if it hasn't been sent already in the current mini-round), and in the next activation go to the next step. Else in the next activation repeat this step.

Step 2: Send $(\textsc{clock-read}, \mathrm{sid}_C)$ to $\mathcal{G}_{\textsc{clock}}$, receive as answer $(\textsc{clock-read}, \mathrm{sid}_C, \tau)$, and proceed as follows.

    if τ corresponds to a working mini-round **then**
        // Generate a new block: extract transactions and form a state-block and append
        Let $\vec{\mathtt{st}}$ be the encoded state in $\mathcal{C}_{loc}$
        Set $\mathtt{buffer}' \leftarrow \mathtt{buffer}$
        Parse $\mathtt{buffer}'$ as sequence $(\mathtt{tx}_1, \ldots, \mathtt{tx}_n)$
        Set $\vec{N} \leftarrow \mathtt{tx}^{\text{coin-base}}_{\text{minerID}}$
        Set $\mathtt{st} \leftarrow \mathsf{blockify}_{\text{B}}(\vec{N})$
        repeat
            Let $(\mathtt{tx}_1, \ldots, \mathtt{tx}_n)$ be the current list of (remaining) transactions in $\mathtt{buffer}'$
            for $i = 1$ to n **do**
                if $\mathsf{ValidTx}_{\text{B}}(\mathtt{tx}_i, \vec{\mathtt{st}}||\mathtt{st}) = 1$ **then**
                    $\vec{N} \leftarrow \vec{N}||\mathtt{tx}_i$
                    Remove `tx` from $\mathtt{buffer}'$
                    Set $\mathtt{st} \leftarrow \mathsf{blockify}_{\text{B}}(\vec{N})$
                end if
            end for
        until $\vec{N}$ does not increase anymore
        Execute **ExtendState(st)** and go to step 3 in the next activation.
    else
        Go to the beginning of step 2 in the next activation.
    end if

Step 3:
If a $(\textsc{clock-update}, \mathrm{sid}_C)$ has been received during this working round then send $(\textsc{clock-update}, \mathrm{sid}_C)$ to $\mathcal{G}_{\textsc{clock}}$, and in the next activation go to the next step. Else in the next activation repeat this step.

Step 4:
Send $(\textsc{clock-read}, \mathrm{sid}_C)$ to $\mathcal{G}_{\textsc{clock}}$, receive as answer $(\textsc{clock-read}, \mathrm{sid}_C, \tau)$, and proceed as follows.

    if τ corresponds to an update mini-round **then**
        If $t < \tau$ execute **FetchInformation** and set $t \leftarrow \tau$.
        Go to step 1 in the next activation.

else
Go to the beginning of step 4 in the next activation.
end if

ExtendState(st):
$\mathcal{C}_{\text{new}} \leftarrow \mathsf{extendchain}_{\mathsf{D}}(\mathcal{C}_{loc}, \mathtt{st}, q)$
if $\mathcal{C}_{\text{new}} \neq \mathcal{C}_{loc}$ **then**
Update the local chain, i.e., $\mathcal{C}_{loc} \leftarrow \mathcal{C}_{\text{new}}$.
end if
// Broadcast current chain
Send $(\text{MULTICAST}, \text{sid}, \mathcal{C}_{loc})$ to $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$

FetchInformation:
// Fetch new chains and update the local state
Send $(\text{FETCH}, \text{sid})$ to $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$; denote the response from $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$ by $(\text{FETCH}, \text{sid}, b)$.
Extract valid chains $\mathcal{C}_1, \ldots, \mathcal{C}_k$ from b.
Set both $\mathcal{C}_{loc}$, $\mathcal{C}_{\text{exp}}$ to the longest valid chain in $\mathcal{C}_{loc}, \mathcal{C}_{\text{exp}}, \mathcal{C}_1, \ldots, \mathcal{C}_k$ (to resolve ties the ordering decides).
// Fetch new transactions and add them to the buffer
Send $(\text{FETCH}, \text{sid})$ to $\mathcal{F}^{\mathsf{tx}}_{\text{N-MC}}$; denote the response from $\mathcal{F}^{\mathsf{tx}}_{\text{N-MC}}$ by $(\text{FETCH}, \text{sid}, b)$.
Extract received transactions $(\mathtt{tx}_1, \ldots, \mathtt{tx}_k)$ from b.
Set $\mathtt{buffer} \leftarrow \mathtt{buffer}||(\mathtt{tx}_1, \ldots, \mathtt{tx}_k)$.
Remove all transactions from `buffer` which are invalid with respect to $\vec{\mathtt{st}}^{\lceil T}$

We now show that the ledger protocol has a predictable synchronization pattern according to Definition 1.

Lemma 2. *The protocol* $\mathsf{Ledger\text{-}Protocol}_{q,\mathcal{D},T}$ *satisfies Definition 1.*

Proof (Sketch). Recall that we have to argue that there exists an algorithm $\mathsf{predict\text{-}time}_{\Pi}(\cdot)$ such that for any possible execution of Π (i.e., for any adversary and environment, and any choice of random coins) it holds the following: If $\vec{\mathcal{I}}_H^T = ((x_1, \text{pid}_1, \tau_1), \ldots, (x_m, \text{pid}_m, \tau_m))$ is the corresponding timed honest-inputs sequence for this execution, then for any $i \in [m-1]$:

$$\mathsf{predict\text{-}time}_{\Pi}((x_1, \text{pid}_1, \tau_1), \ldots, (x_i, \text{pid}_i, \tau_i)) = \tau_{i+1}.$$

This is straightforward to see for our ledger protocol (and all protocols that share the same structure) in all the respective hybrid worlds they are executed. Roughly, the predicate predict-time can be implemented as follows: browse through the entire sequence $\vec{\mathcal{I}}_H^T$ and determine how many times the clock advances. The clock advances for the first time, when all miners got a maintain command, followed by a clock-update command. By definition of Ledger-Protocol, in this each party will then send a clock-update to the clock. If every party has done that, the clock advances. By an inductive argument, whenever the clock has ticked, the check when the clock advances the next time is checked exactly the same way. Overall, this allows to check whether the next activation of an honest party, given the history of activations will provoke a clock update. Note that only an activation of an honest parties can make the clock advance. □

C.4 The Bitcoin Ledger (Cont'd)

This section includes complementary material for Section 4.2. We here give the formal description of the Extend Policy for $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$. It is easy to observe that the computation performed by this algorithm is well-defined for any definition of Validate and Blockify.

Compared to previous versions of this work, this policy implements the updated chain quality definition. In addition, the policy makes the initial bootstrapping time of the chain now explicit, where by bootstrapping time we mean the time it takes for the first state block to be inserted into the state.

Algorithm ExtendPolicy for $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$

function $\text{EXTENDPOLICY}(\vec{\mathcal{I}}_H^T, \mathtt{state}, \mathtt{NxtBC}, \mathtt{buffer}, \vec{\tau}_{\mathtt{state}})$
Let τ_L be current ledger time (computed from $\vec{\mathcal{I}}_H^T$)
// The function must not have side-effects: Only modify copies of relevant values.
Create local copies of the values `buffer`, `state`, and $\vec{\tau}_{\mathtt{state}}$.
// First, create a default honest client block as alternative:
Set $\vec{N}_{\mathtt{df}} \leftarrow \mathtt{tx}^{\text{coin-base}}_{\text{minerID}}$ of an honest miner

Sort buffer according to time stamps and let $\vec{\texttt{tx}} = (\texttt{tx}_1, \ldots, \texttt{tx}_n)$ be the transactions in buffer
Set $\texttt{st} \leftarrow \mathsf{blockify}_{\text{B}}(\vec{N}_{\texttt{df}})$
repeat
  Let $\vec{\texttt{tx}} = (\texttt{tx}_1, \ldots, \texttt{tx}_n)$ be the current list of (remaining) transactions
  for $i = 1$ to n **do**
    if $\mathsf{ValidTx}_{\text{B}}(\texttt{tx}_i, \texttt{state}||\texttt{st}) = 1$ **then**
      $\vec{N}_{\texttt{df}} \leftarrow \vec{N}_{\texttt{df}}||\texttt{tx}_i$
      Remove $\texttt{tx}_i$ from $\vec{\texttt{tx}}$
      Set $\texttt{st} \leftarrow \mathsf{blockify}_{\text{B}}(\vec{N}_{\texttt{df}})$
    end if
  end for
until $\vec{N}_{\texttt{df}}$ does not increase anymore
// Possibly more than one block have to be added
// Let $\tau_{\texttt{low}}$ be the time of the block which is windowSize $-$ 1 blocks behind the head of the state.
if $|\texttt{state}| + 1 \geq \texttt{windowSize}$ **then**
  Set $\tau_{\texttt{low}} \leftarrow \vec{\tau}_{\texttt{state}}[|\texttt{state}| - \texttt{windowSize} + 2]$
else
  Set $\tau_{\texttt{low}} \leftarrow 0$
end if
$c \leftarrow 1$
while $\tau_L - \tau_{\texttt{low}} > \texttt{maxTime}_{\texttt{window}}$ **do**
  Set $\vec{N}_c \leftarrow \texttt{tx}^{\text{coin-base}}_{\text{minerID}}$ of an honest miner
  $\vec{N}_{\texttt{df}} \leftarrow \vec{N}_{\texttt{df}}||\vec{N}_c$
  $c \leftarrow c + 1$
  // Update $\tau_{\texttt{low}}$ to the time of the state block which is windowSize $- c$ blocks behind the head.
  if $|\texttt{state}| + c \geq \texttt{windowSize}$ **then**
    Set $\tau_{\texttt{low}} \leftarrow \vec{\tau}_{\texttt{state}}[|\texttt{state}| - \texttt{windowSize} + c + 1]$
  else
    Set $\tau_{\texttt{low}} \leftarrow 0$
  end if
end while
// Now, parse the proposed block by the adversary
// Possibly more than one block should be added
Parse NxtBC as a vector $((\text{hFlag}_1, \texttt{NxtBC}_1), \cdots, (\text{hFlag}_n, \texttt{NxtBC}_n))$
$\vec{N} \leftarrow \varepsilon$ // Initialize Result
// Determine the time of the state block which is windowSize blocks behind the head of the state
if $|\texttt{state}| \geq \texttt{windowSize}$ **then**
  Set $\tau_{\texttt{low}} \leftarrow \vec{\tau}_{\texttt{state}}[|\texttt{state}| - \texttt{windowSize} + 1]$
else
  Set $\tau_{\texttt{low}} \leftarrow 0$
end if
oldValidTxMissing $\leftarrow$ false // Flag to keep track whether old enough, valid transactions are inserted.
for each list $\texttt{NxtBC}_i$ of transaction IDs **do**
  // Compute the next state block
  $\vec{N}_i \leftarrow \varepsilon$
  // Verify validity of $\texttt{NxtBC}_i$ and compute content
  Use the txid contained in $\texttt{NxtBC}_i$ to determine the list of transactions
  Let $\vec{\texttt{tx}} = (\texttt{tx}_1, \ldots, \texttt{tx}_{|\texttt{NxtBC}_i|})$ denote the transactions of $\texttt{NxtBC}_i$
  if $\texttt{tx}_1$ is not a coin-base transaction **then**
    return $\vec{N}_{\texttt{df}}$
  else
    $\vec{N}_i \leftarrow \texttt{tx}_1$
    for $j = 2$ to $|\texttt{NxtBC}_i|$ **do**
      Set $\texttt{st}_i \leftarrow \mathsf{blockify}_{\text{B}}(\vec{N}_i)$
      if $\mathsf{ValidTx}_{\text{B}}(\texttt{tx}_j, \texttt{state}||\texttt{st}_i) = 0$ **then**
        return $\vec{N}_{\texttt{df}}$
      end if
      $\vec{N}_i \leftarrow \vec{N}_i||\texttt{tx}_j$
    end for
    Set $\texttt{st}_i \leftarrow \mathsf{blockify}_{\text{B}}(\vec{N}_i)$
  end if

```
            // Test that all old valid transaction are included
            if the proposal is declared to be an honest block, i.e., hFlag_i = 1 then
                for each BTX = (tx, txid, τ', p_i) ∈ buffer of an honest party p_i with time τ' < τ_low − Delay/2 do
                    if ValidTx_B(tx, state||st_i) = 1 but tx ∉ N⃗_i then
                        oldValidTxMissing ← true
                    end if
                end for
            end if
            N⃗ ← N⃗||N⃗_i
            state ← state||st_i
            τ⃗_state ← τ⃗_state||τ_L
            // Must not proceed with too many adversarial blocks
            Determine the most recent honest block st_j in state (last proposal added with hFlag = 1)
            if |state| − j ≥ η then
                return N⃗_df
            end if
            // Update τ_low: the time of the state block which is windowSize blocks behind the head of the
            // current, potentially already extended state
            if |state| ≥ windowSize then
                Set τ_low ← τ⃗_state[|state| − windowSize + 1]
            else
                Set τ_low ← 0
            end if
        end for
        // Final checks (if policy is violated, it is enforced by the ledger):
        // Must not proceed too fast, too slow, or with missing transactions.
        if τ_L − τ_low < minTime_window then
            return ε
        else if τ_low > 0 and τ_L − τ_low > maxTime_window then // A sequence of blocks cannot take too much time.
            return N⃗_df
        else if τ_low = 0 and τ_L − τ_low > 2 · maxTime_window then // Bootstrapping cannot take too much time.
            return N⃗_df
        else if oldValidTxMissing then // If not all old enough, valid transactions have been included.
            return N⃗_df
        end if
        return N⃗
    end function
```

D Security Analysis (Cont'd)

This section includes complementary material for Section 4.3.

D.1 Modular-Ledger-Protocol

The Modular-Ledger-Protocol uses the same hybrids as our original protocol Ledger-Protocol but abstracts the lottery implemented by the mining process by making calls to the above state exchange functionality $\mathcal{F}_{\text{STX}}^{p_H,p_A}$. The protocol is parameterized by T which is the number of blocks chopped off to get the ledger state. The registration process work as in the Ledger-Protocol.

Protocol Modular-Ledger-Protocol$_T(p)$

Initialization:

We assume that the party p is registered to $\mathcal{F}_{\text{STX}}^{p_H,p_A}$, $\mathcal{F}_{\text{N-MC}}^{\text{tx}}$, and $\mathcal{G}_{\text{CLOCK}}$. A party is considered not initialized after registration (define `isInit` ← false).
The protocol manages the exported ledger state $\vec{\texttt{st}}_{\text{exp}}$ which initially is the genesis state, i.e. , $\vec{\texttt{st}} \leftarrow (\texttt{gen})$. It also manages a local (working) state $\vec{\texttt{st}}$ (initially also the genesis state).
The protocol also manages an initially empty buffer `buffer` which contains the received transactions. The buffer forms the information base to build new state blocks. Finally, set $t \leftarrow 0$.
Once the initialization is complete after the next MAINTAIN-LEDGER-command as described above (including requesting previous transactions), set `isInit` ← true

Clock:

- Upon receiving ($\textsc{clock-read}, \mathrm{sid}_C$) forward the query to $\mathcal{G}_{\textsc{clock}}$ and return whatever is received as answer from $\mathcal{G}_{\textsc{clock}}$
- Upon receiving ($\textsc{clock-update}, \mathrm{sid}_C$), do the following: remember that a clock-update was received in the current mini-round.

Ledger:

- Upon receiving ($\textsc{submit}, \mathrm{sid}, \mathtt{tx}$), set $\mathtt{buffer} \leftarrow \mathtt{buffer}||\mathtt{tx}$, and send ($\textsc{multicast}, \mathrm{sid}, \mathtt{tx}$) to $\mathcal{F}^{\mathsf{tx}}_{\text{N-MC}}$.
- Upon receiving ($\textsc{read}, \mathrm{sid}$) send ($\textsc{clock-read}, \mathrm{sid}_C$) to $\mathcal{G}_{\textsc{clock}}$, receive as answer ($\textsc{clock-read}, \mathrm{sid}_C, \tau$) and proceed as follows.

    // Fetch new states in update mini-rounds
    if τ corresponds to an update mini-round and $t < \tau$ and $\mathtt{isInit}$ **then**
        Execute **FetchInformation** and set $t \leftarrow \tau$.
    end if
    // Return current ledger state
    Return ($\textsc{read}, \mathrm{sid}, \vec{\mathtt{st}}^{\lceil T}_{\mathrm{exp}}$)

- Upon receiving ($\textsc{maintain-ledger}, \mathrm{sid}, \mathrm{minerID}$) execute **MiningProcedure**

MiningProcedure:

Step 1:
    If a ($\textsc{clock-update}, \mathrm{sid}_C$) has been received during this update mini-round then send ($\textsc{clock-update}, \mathrm{sid}_C$) to $\mathcal{G}_{\textsc{clock}}$ (if it hasn't been sent already in the current mini-round), and in the next activation go to the next step. Else in the next activation repeat this step.

Step 2:
    Send ($\textsc{clock-read}, \mathrm{sid}_C$) to $\mathcal{G}_{\textsc{clock}}$, receive as answer ($\textsc{clock-read}, \mathrm{sid}_C, \tau$), and proceed as follows.
    if τ corresponds to a working mini-round **then**
        // Generate a new block: extract transactions and form a state-block
        Set $\mathtt{buffer}' \leftarrow \mathtt{buffer}$
        Parse $\mathtt{buffer}'$ as sequence ($\mathtt{tx}_1, \ldots, \mathtt{tx}_n$)
        Set $\vec{N} \leftarrow \mathtt{tx}^{\text{coin-base}}_{\text{minerID}}$
        Set $\mathtt{st} \leftarrow \mathsf{blockify}_{\text{₿}}(\vec{N})$
        repeat
            Let ($\mathtt{tx}_1, \ldots, \mathtt{tx}_n$) be the current list of (remaining) transactions in $\mathtt{buffer}'$
            for $i = 1$ to n **do**
                if $\mathsf{ValidTx}_{\text{₿}}(\mathtt{tx}, \vec{\mathtt{st}}||\mathtt{st}) = 1$ **then**
                    $\vec{N} \leftarrow \vec{N}||\mathtt{tx}$
                    Remove $\mathtt{tx}$ from $\mathtt{buffer}'$
                    Set $\mathtt{st} \leftarrow \mathsf{blockify}_{\text{₿}}(\vec{N})$
                end if
            end for
        until $\vec{N}$ does not increase anymore
        Execute **ExtendState(st)** and go to step 3 in the next activation.
    else
        Go to the beginning of step 2 in the next activation.
    end if

Step 3:
    If a ($\textsc{clock-update}, \mathrm{sid}_C$) has been received during this working round then send ($\textsc{clock-update}, \mathrm{sid}_C$) to $\mathcal{G}_{\textsc{clock}}$, and in the next activation go to the next step. Else in the next activation repeat this step.

Step 4:
    Send ($\textsc{clock-read}, \mathrm{sid}_C$) to $\mathcal{G}_{\textsc{clock}}$, receive as answer ($\textsc{clock-read}, \mathrm{sid}_C, \tau$), and proceed as follows.
    if τ corresponds to an update mini-round **then**
        if $t < \tau$ **then**
            Execute **FetchInformation** and set $t \leftarrow \tau$.
        end if
        Go to step 1 in the next activation.
    else
        Go to the beginning of step 4 in the next activation.
    end if

ExtendState(st):

Send (SUBMIT-NEW, sid, $\vec{\texttt{st}}, \texttt{st}$) to $\mathcal{F}_{\text{StX}}$.
Denote the response by (SUCCESS, sid, B) of $\mathcal{F}_{\text{StX}}$.
if $B = 1$ **then**
Update the local state, i.e., $\vec{\texttt{st}} \leftarrow \vec{\texttt{st}}||\texttt{st}$.
end if
// Broadcast current state using $\mathcal{F}_{\text{StX}}$.
Send (CONTINUE, sid) to $\mathcal{F}_{\text{StX}}$

FetchInformation:
// Fetch new states and update the local state
Send (FETCH-NEW, sid) to $\mathcal{F}_{\text{StX}}$.
Denote the response from $\mathcal{F}_{\text{StX}}$ by (FETCH-NEW, sid, $(\vec{\texttt{st}}_1, \dots, \vec{\texttt{st}}_k)$).
Set both $\vec{\texttt{st}}$,$\vec{\texttt{st}}_{\text{exp}}$ to the longest state in $\vec{\texttt{st}}, \vec{\texttt{st}}_{\text{exp}}, \vec{\texttt{st}}_1, \dots, \vec{\texttt{st}}_k$ (to resolve ties the ordering decides).
// Fetch new transactions and add them to the buffer
Send (FETCH, sid) to $\mathcal{F}^{\text{tx}}_{\text{N-MC}}$; denote the response from $\mathcal{F}^{\text{tx}}_{\text{N-MC}}$ by (FETCH, sid, b).
Extract received transactions $(\texttt{tx}_1, \dots, \texttt{tx}_k)$ from b.
Set $\texttt{buffer} \leftarrow \texttt{buffer}||(\texttt{tx}_1, \dots, \texttt{tx}_k)$.
Remove all transactions from `buffer` which are invalid with respect to $\vec{\texttt{st}}^{\lceil T}$

D.2 Proof of Lemma 1

In the following we provide a proof for Lemma 1.

Lemma (1). *The protocol* $\mathsf{Ledger\text{-}Protocol}_{q,D,T}$ *UC emulates the protocol* $\mathsf{Modular\text{-}Ledger\text{-}Protocol}_T$ *that runs in a hybrid world with access to the functionality* $\mathcal{F}^{\Delta,p_H,p_A}_{\text{StX}}$ *with* $p_A := \frac{D}{2^\kappa}$ *and* $p_H = 1 - (1 - p_A)^q$*, and where* Δ *denotes upper bound on the network delay.*

In a first step, we describe protocol StateExchange-Protocol and show that it UC-realizes the $\mathcal{F}_{\text{StX}}$ functionality in the $\mathcal{F}_{\text{RO}}$, $\mathcal{F}^{\text{bc}}_{\text{N-MC}}$ hybrid world. Note that $\mathcal{F}^{\text{bc}}_{\text{N-MC}}$ is a (named) instance of the $\mathcal{F}^{\Delta}_{\text{N-MC}}$ functionality. The protocol is parametrized by q and D where q is the number of mining attempts per submission attempt and D is the difficulty of the proof-of-work.

Protocol $\mathsf{StateExchange\text{-}Protocol}_{q,\mathtt{D}}(p)$

Initialization:
We assume that the party p is registered to $\mathcal{F}_{\text{RO}}$ and $\mathcal{F}^{\text{bc}}_{\text{N-MC}}$.
The protocol maintains a tree $\mathcal{T}$ of all valid chains. Initially it contains the genesis chain (**G**).

Message Exchange:

– Upon receiving (SUBMIT-NEW, sid, $\vec{\texttt{st}}, \texttt{st}$) do
  // Check if there exists a chain in $\mathcal{T}$ which contains the state $\vec{\texttt{st}}$
  if $\mathsf{isvalidstate}(\vec{\texttt{st}}||\texttt{st}) = 1$ **then**
    if there exists $\mathcal{C} \in \mathcal{T}$ with $\vec{\texttt{st}}$ **then**
      // Try to extend the chain
      $\mathcal{C}_{\text{new}} \leftarrow \mathsf{extendchain}_{\mathtt{D}}(\mathcal{C}, \texttt{st}, q)$
      if $\mathcal{C}_{\text{new}} \neq \mathcal{C}$ **then**
        Update the local tree, i.e., add $\mathcal{C}_{\text{new}}$ to $\mathcal{T}$
        Output (SUCCESS, sid, 1) to p.
      else
        Output (SUCCESS, sid, 0) to p.
      end if
      // Broadcast current chain
      On response (CONTINUE, sid) send (MULTICAST, sid, $\mathcal{C}_{\text{new}}$) to $\mathcal{F}^{\text{bc}}_{\text{N-MC}}$.
    end if
  end if
– Upon receiving (FETCH-NEW, sid) if do the following:
  Send (FETCH, sid) to $\mathcal{F}^{\text{bc}}_{\text{N-MC}}$ and denote the response by (FETCH, sid, b).
  Extract all valid chains $\mathcal{C}_1, \dots, \mathcal{C}_k$ from b and add them to $\mathcal{T}$.
  Extract states $\vec{\texttt{st}}_1, \dots, \vec{\texttt{st}}_k$ from $\mathcal{C}_1, \dots, \mathcal{C}_k$ and output them.

Lemma 3. *Let* $p := \frac{D}{2^\kappa}$*. The protocol* $\mathsf{StateExchange\text{-}Protocol}_{q,D}$ *UC-realizes functionality* $\mathcal{F}^{\Delta,p_H,p_A}_{\text{StX}}$ *in the* $(\mathcal{F}_{\text{RO}}, \mathcal{F}^{\Delta}_{\text{N-MC}})$*-hybrid model where* $p_A := p$ *and* $p_H := 1 - (1 - p)^q$*.*

Proof. We consider the following simulator.

Simulator $\mathcal{S}_{\mathrm{stx}}$

Initialization:

Set up a tree of valid chains $\mathcal{T} \leftarrow \{(\mathbf{G})\}$ and an empty network buffer $\vec{M}$.
Set up an empty random oracle table H and set $H[\mathbf{G}]$ to a uniform random value in $\{0,1\}^\kappa$. If the simulator ever tries to add a colliding entry to H, abort with COLLISION-ERROR.
The simulator manages a set $\mathcal{P}_{RO}$ of parties registered to the random oracle and a set of parties $\mathcal{P}_{net}$ registered to the network.

Random Oracle:

- Upon receiving $(\textsc{eval}, \mathrm{sid}, v)$ for $\mathcal{F}_{\mathrm{RO}}$ from $\mathcal{A}$ on behalf of corrupted $p \in \mathcal{P}_{RO}$ do the following.
  1. If $H[v]$ is already defined, output $(\textsc{eval}, \mathrm{sid}, v, H[v])$.
  2. If v is of the form $(\mathtt{s}, \mathtt{st}, \mathtt{n})$ and there exists[a] a chain $\mathcal{C} = \mathbf{B}_1, \ldots, \mathbf{B}_n$ such that $H[\mathbf{B}_n] = \mathtt{s}$ proceed as follows. If $\mathcal{C} \notin \mathcal{T}$ abort with TREE-ERROR. Otherwise continue. Extract the state $\vec{\mathtt{st}}$ from $\mathcal{C}$ and extract the state block $\mathtt{st}$ from v. Send $(\textsc{submit-new}, \mathrm{sid}, \vec{\mathtt{st}}, \mathtt{st})$ to $\mathcal{F}_{\mathrm{STX}}$ and denote by $(\textsc{success}, B)$ the output of $\mathcal{F}_{\mathrm{STX}}$. If $B = 1$ set $H[v]$ to a uniform random value in $\{0,1\}^\kappa$ strictly smaller[b] than $\mathtt{D}$. Add $\mathcal{C}||v$ to $\mathcal{T}$. Otherwise set $H[v]$ to a uniform random value in $\{0,1\}^\kappa$ larger than $\mathtt{D}$. Output $(\textsc{eval}, \mathrm{sid}, v, H[v])$.
  3. Otherwise set v to a uniform random value in $\{0,1\}^\kappa$ and output $(\textsc{eval}, \mathrm{sid}, v, H[v])$.

Network:

- Upon receiving $(\textsc{multicast}, \mathrm{sid}, (m_{i_1}, p_{i_1}), \ldots, (m_{i_\ell}, p_{i_\ell}))$ for $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$ from $\mathcal{A}$ on behalf of corrupted $p \in \mathcal{P}_{net}$ with $\{p_{i_1}, \ldots, p_{i_\ell}\} \subseteq \mathcal{P}_{net}$ proceed as follows.
  1. Choose ℓ new unique message-IDs $\mathrm{mid}_{i_1}, \ldots, \mathrm{mid}_{i_\ell}$, initialize ℓ new variables $D_{\mathrm{mid}_{i_1}} := \ldots := D_{\mathrm{mid}_{i_\ell}} := 1$, set $\vec{M} := \vec{M}||(m_{i_1}, \mathrm{mid}_{i_1}, D_{\mathrm{mid}_{i_1}}, p_{i_1})|| \ldots ||(m_{i_\ell}, \mathrm{mid}_{i_\ell}, D_{\mathrm{mid}_{i_\ell}}, p_\ell)$.
  2. For each (m_{i_j}, p_{i_j}) where m_{i_j} is a chain $\mathcal{T}$ extract the state $\vec{\mathtt{st}}_{i_j}$ from m_{i_j}, and send $(\textsc{send}, \mathrm{sid}, \vec{\mathtt{st}}, p_{i_j})$ to $\mathcal{F}_{\mathrm{STX}}$. Store the message-ID $\widehat{\mathrm{mid}_{i_j}}$ returned by $\mathcal{F}_{\mathrm{STX}}$ with mid_{i_j}.
  3. Output $(\textsc{multicast}, \mathrm{sid}, (m_{i_1}, p_{i_1}, \mathrm{mid}_{i_1}), \ldots, (m_{i_\ell}, p_{i_\ell}, \mathrm{mid}_{i_\ell})$ to $\mathcal{A}$.
- Upon receiving $(\textsc{fetch}, \mathrm{sid})$ for $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$ from $\mathcal{A}$ on behalf of corrupted $p \in \mathcal{P}_{net}$ proceed as follows.
  1. For all tuples $(m, \mathrm{mid}, D_{\mathrm{mid}}, p) \in \vec{M}$, set $D_{\mathrm{mid}} := D_{\mathrm{mid}} - 1$.
  2. Let $\vec{M}_0^p$ denote the subvector $\vec{M}$ including all tuples of the form $(m, \mathrm{mid}, D_{\mathrm{mid}}, p)$ with $D_{\mathrm{mid}} = 0$ (in the same order as they appear in $\vec{M}$). Delete all entries in $\vec{M}_0^p$ from $\vec{M}$, and send $\vec{M}_0^p$ to $\mathcal{A}$.
- Upon receiving a message $(\textsc{delays}, \mathrm{sid}, (T_{\mathrm{mid}_{i_1}}, \mathrm{mid}_{i_1}), \ldots, (T_{\mathrm{mid}_{i_\ell}}, \mathrm{mid}_{i_\ell}))$ do the following for each pair $(T_{\mathrm{mid}}, \mathrm{mid})$ in this message:
  1. If T_{mid} is a valid delay (i.e., it encodes an integer in unary notation) and mid is a message-ID registered in the current $\vec{M}$, set $D_{\mathrm{mid}} := \max\{1, D_{\mathrm{mid}} + T_{\mathrm{mid}}\}$; otherwise, ignore this tuple.
  2. If the simulator knows a corresponding $\mathcal{F}_{\mathrm{STX}}$-message-ID $\widehat{\mathrm{mid}}$ for mid send $(\textsc{delay}, \mathrm{sid}, T_{\mathrm{mid}}, \widehat{\mathrm{mid}})$ to $\mathcal{F}_{\mathrm{STX}}$.
- Upon receiving a message $(\textsc{swap}, \mathrm{sid}, \mathrm{mid}_1, \mathrm{mid}_2)$ from the adversary do the following:
  1. If mid_1 and mid_2 are message-IDs registered in the current $\vec{M}$, then swap the corresponding tuples in $\vec{M}$.
  2. If the simulator knows for both mid_1 and mid_2 $\mathcal{F}_{\mathrm{STX}}$-message-IDs $\widehat{\mathrm{mid}_1}$ and $\widehat{\mathrm{mid}_2}$ send $(\textsc{swap}, \mathrm{sid}, \widehat{\mathrm{mid}_1}, \widehat{\mathrm{mid}_2})$ to $\mathcal{F}_{\mathrm{STX}}$.
  3. Output $(\textsc{swap}, \mathrm{sid})$ to $\mathcal{A}$.

State Exchange Functionality :

- Upon receiving $(\textsc{submit-new}, \mathrm{sid}, \vec{\mathtt{st}}, p_s, (p_1, \widehat{\mathrm{mid}}_1), \ldots, (p_n, \widehat{\mathrm{mid}}_n))$ from $\mathcal{F}_{\mathrm{STX}}$ where $\vec{\mathtt{st}} = \mathtt{st}_1, \ldots, \mathtt{st}_k$ and $\{p_1, \ldots, p_n\} := \mathcal{P}_{net}$ proceed as follows
  1. If there exist a chain $\mathcal{C} \in \mathcal{T}$ with state $\vec{\mathtt{st}}$ generate new unique message-IDs $\mathrm{mid}_1, \ldots, \mathrm{mid}_n$, initialize $D_1 := \cdots := D_n = 1$, set $\vec{M}||(\mathcal{C}, \mathrm{mid}_{i1}, D_{\mathrm{mid}_1}, p_1)|| \ldots ||(\mathcal{C}, \mathrm{mid}_n, D_{\mathrm{mid}_n}, p_n)$, and store the message-IDs $\widehat{\mathrm{mid}}_i$ along the message-IDs mid_i.
  Output $(\textsc{multicast}, \mathrm{sid}, \mathcal{C}, p_s, (p_1, \mathrm{mid}_1), \ldots, (p_n, \mathrm{mid}_n))$ to the adversary.
  2. Otherwise find a chain $\mathcal{C}'$ in $\mathcal{T}$ with state $\mathtt{st}_1, \ldots, \mathtt{st}_{k-1}$[c]. Choose a random nonce $\mathtt{n}$ and set $\mathbf{B}_k = (H[\mathbf{B}_{k-1}], \mathtt{st}_k, \mathtt{n})$ and set $H[\mathbf{B}_k]$ to a uniform random value in $\{0,1\}^\kappa$ strictly smaller than $\mathtt{D}$. Add the chain $\mathcal{C} = \mathcal{C}'||\mathbf{B}_k$ to $\mathcal{T}$.
  Generate new unique message-IDs $\mathrm{mid}_1, \ldots, \mathrm{mid}_n$, initialize $D_1 := \cdots := D_n = 1$, set $\vec{M}||(\mathcal{C}, \mathrm{mid}_{i1}, D_{\mathrm{mid}_1}, p_1)|| \ldots ||(\mathcal{C}, \mathrm{mid}_n, D_{\mathrm{mid}_n}, p_n)$, and store the message-IDs $\widehat{\mathrm{mid}}_i$ along the message-IDs mid_i. Output $(\textsc{multicast}, \mathrm{sid}, \mathcal{C}, p_s, (p_1, \mathrm{mid}_1), \ldots, (p_n, \mathrm{mid}_n))$ to the adversary.

[a] This can be checked efficiently using H under the assumption that there are no collisions.
[b] Can be done efficiently using rejection sampling.
[c] Such a chain must exist as $\texttt{st}_1, \ldots, \texttt{st}_{k-1}$ is a successfully submitted state in $\mathcal{F}_{\mathrm{STX}}$ in which case the simulator knows a corresponding chain.

The proof works similar as the one for Lemma 5.1 in [PSS17].Denote by $\mathrm{EXEC}_{Real,\mathcal{A},\mathcal{Z}}$ the joint view of all parties in the execution of StateExchange-Protocol for adversary $\mathcal{A}$ and environment $\mathcal{Z}$. Denote by $\mathrm{EXEC}_{Ideal,\mathcal{A},\mathcal{Z}}$ the joint view of all parties for $\mathcal{F}_{\mathrm{STX}}$ with simulator $\mathcal{S}_{\mathrm{stx}}$.

Define $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ the same as $\mathrm{EXEC}_{Real,\mathcal{A},\mathcal{Z}}$ except that the random oracle aborts on collisions with COLLISION-ERROR and that adversarial oracle queries are emulated as in $\mathcal{S}_{\mathrm{stx}}$. The only difference is thus that in $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ we may abort with COLLISION-ERROR or TREE-ERROR.

Let event1 be the event that some parties query two different values v, v' such that $H[v] = H[v']$, i.e. the event that a hash-collision occurs. For any two queries the probability that they return the same hash value is $2^{-\kappa}$. By a union bound over all queries we have that event1 happens with probability at most $\mathsf{poly}(\kappa) \cdot 2^{-\kappa}$ in both worlds. Note that if event1 does not happen $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ will not abort with COLLISION-ERROR.

Let event2 be the event that some party makes a query $H[(\texttt{s}, \cdot, \cdot)]$ where no v exists such that $H[v] = \texttt{s}$, but later some party makes a query v' such that $H[v'] = \texttt{s}$. The probability that any query $H[(\texttt{s}, \cdot, \cdot)]$ a later query returns $\texttt{s}$ is $2^{-\kappa}$ in both worlds By a union bound over all queries we have that event2 happens with probability at most $\mathsf{poly}(\kappa) \cdot 2^{-\kappa}$ in both worlds.

Next, we show that the TREE-ERROR abort does not occur in $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ conditioned under event1 and event2 not happening. Assume for contradiction that $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ aborts with TREE-ERROR with event1 and event2 not happening. Let $\mathcal{C} = \mathbf{B}_1, \ldots, \mathbf{B}_n$ be the shortest valid chain created in the experiment $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ such that $\mathbf{B}_1, \ldots, \mathbf{B}_{n-1} \in \mathcal{T}$ but $\mathbf{B}_1, \ldots, \mathbf{B}_n \notin \mathcal{T}$. Let $\mathbf{B}_i = (\texttt{s}_i, \texttt{st}_i, \texttt{n}_i)$. Since $\mathcal{C}$ is a valid chain we have $H[(\texttt{s}_n, \texttt{st}_n \texttt{n}_n)] < D$. But at the time $\mathbf{B}_n$ was added to H no valid chain existed where the last block has hash value $\texttt{s}_n$ (otherwise $\mathcal{C}$ would be in $\mathcal{T}$). This implies that no earlier query to H could have returned $\texttt{s}_n$, since if the query was $\mathbf{B}_{n-1}$ $\mathcal{C}$ would not be the shortest chain with the above property and if the query was not $\mathbf{B}_{n-1}$ the event event1 must have happened. This implies that event2 must have happened, which is a contradiction.

This implies that conditioned under event1 and event2 not happening $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ proceeds the same as $\mathrm{EXEC}_{Real,\mathcal{A},\mathcal{Z}}$. It follows that $\mathrm{EXEC}_{Real,\mathcal{A},\mathcal{Z}}$ and $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ are statistically close.

Now we compare $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ and $\mathrm{EXEC}_{Ideal,\mathcal{A},\mathcal{Z}}$. Consider the event where a honest miner queries a block $(\texttt{s}, \texttt{st}, \texttt{n})$ and fails, i.e. where $H[(\texttt{s}, \texttt{st}, \texttt{n})] > \texttt{D}$. In $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ this query is stored in the random oracle table while the simulator in $\mathrm{EXEC}_{Ideal,\mathcal{A},\mathcal{Z}}$ does not store the query in the random oracle table. Under the condition that such failed queries are never queried again $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ and $\mathrm{EXEC}_{Ideal,\mathcal{A},\mathcal{Z}}$ are identically distributed as the network simulation in $\mathcal{S}_{\mathrm{stx}}$ is perfect. Note that the nonce $\texttt{n}$ in a 'failed' query $(\texttt{s}, \texttt{st}, \texttt{n})$ is choosen uniform at random from $\{0,1\}^{\kappa}$. This implies that with probability $\mathsf{poly}(\kappa) \cdot 2^{-\kappa}$ it was never queried before. As honest miner discard 'failed' queries it also follows that except with probability $\mathsf{poly}(\kappa) \cdot 2^{-\kappa}$ the query will not be queried again (by honest parties or $\mathcal{A}$). By a union bound over all failed queries we have that failed queries are never queried twice except with probability $\mathsf{poly}(\kappa) \cdot 2^{-\kappa}$. Thus $\mathrm{HYB}_{\mathcal{A},\mathcal{Z}}$ and $\mathrm{EXEC}_{Ideal,\mathcal{A},\mathcal{Z}}$ are statistically close.

We conclude the proof for Lemma 1, by giving a game-hopping argument to show that Ledger-Protocol UC emulates the protocol Modular-Ledger-Protocol. We start with the original Ledger-Protocol and consider the protocol part below where will alter the protocol.

Protocol Original Protocol Part

Initialization:

The protocol stores a local (working) chain $\mathcal{C}_{loc}$ which initially contains the genesis block, i.e., $\mathcal{C}_{loc} \leftarrow (\mathbf{G})$.
[...]

ExtendState(st):

$\mathcal{C}_{\mathrm{new}} \leftarrow \mathsf{extendchain}_{\texttt{D}}(\mathcal{C}_{loc}, \texttt{st}, q)$
if $\mathcal{C}_{\mathrm{new}} \neq \mathcal{C}_{loc}$ **then**
Update the local chain, i.e., $\mathcal{C}_{loc} \leftarrow \mathcal{C}_{\mathrm{new}}$.

end if
// Broadcast current chain
Send $(\text{MULTICAST}, \text{sid}, \mathcal{C}_{loc})$ to $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$

FetchInformation:
// Update the local state
Send $(\text{FETCH}, \text{sid})$ to $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$; denote the response from $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$ by $(\text{FETCH}, \text{sid}, b)$.
Extract valid chains $\mathcal{C}_1, \ldots, \mathcal{C}_k$ from b.
Set both $\mathcal{C}_{loc}, \mathcal{C}_{\text{exp}}$ to the longest valid chain in $\mathcal{C}_{loc}, \mathcal{C}_{\text{exp}}, \mathcal{C}_1, \ldots, \mathcal{C}_k$ (to resolve ties the ordering decides).
// Fetch new transactions and add them to the buffer
...

The first modification of the protocol (see below) proceeds as Ledger-Protocol except (a) it stores a history of all valid chains in a tree $\mathcal{T}$ and (b) in the **ExtendState(st)** procedure it checks that $\vec{\mathtt{st}}||\mathtt{st}$ is a valid state and that there exists a chain in $\mathcal{T}$ which encodes the state $\vec{\mathtt{st}}$. We observe that the protocol calls **ExtendState(st)** only with $\mathtt{st}$ where $\vec{\mathtt{st}}||\mathtt{st}$ is a valid state. This implies that the first check is always satisfied. Moreover, not that the current local chain $\mathcal{C}_{loc}$ which encodes state $\vec{\mathtt{st}}$ is at any time stored in the tree $\mathcal{T}$. The second check is therefore also always satisfied. Hence, the modified protocol has the same input/output behavior has Ledger-Protocol.

Protocol Modification 1

Initialization:
The protocol stores a local (working) chain $\mathcal{C}_{loc}$ which initially contains the genesis block, i.e., $\mathcal{C}_{loc} \leftarrow (\mathbf{G})$.
[...]
The protocol additionally maintains a tree $\mathcal{T}$ of valid chains which initially contains the (genesis) chain $(\mathbf{G})$.

ExtendState(st):
if $\mathsf{isvalidstate}(\vec{\mathtt{st}}||\mathtt{st}) = 1$ **then**
  if there exists $\mathcal{C} \in \mathcal{T}$ which encodes $\vec{\mathtt{st}}$ **then**
    $\mathcal{C}_{\text{new}} \leftarrow \mathsf{extendchain}_{\mathsf{D}}(\mathcal{C}_{loc}, \mathtt{st}, q)$
    if $\mathcal{C}_{\text{new}} \neq \mathcal{C}_{loc}$ **then**
      Update the local chain, i.e., $\mathcal{C}_{loc} \leftarrow \mathcal{C}_{\text{new}}$.
      Add $\mathcal{C}_{loc}$ to $\mathcal{T}$
    end if
    // Broadcast current chain
    Send $(\text{MULTICAST}, \text{sid}, \mathcal{C})$ to $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$
  end if
end if

FetchInformation:
// Update the local state
Send $(\text{FETCH}, \text{sid})$ to $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$; denote the response from $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$ by $(\text{FETCH}, \text{sid}, b)$.
Extract all valid chains $\mathcal{C}_1, \ldots, \mathcal{C}_k$ from b and add them to $\mathcal{T}$.
Set both $\mathcal{C}_{loc}, \mathcal{C}_{\text{exp}}$ to the longest valid chain in $\mathcal{C}_{loc}, \mathcal{C}_{\text{exp}}, \mathcal{C}_1, \ldots, \mathcal{C}_k$ (to resolve ties the ordering decides).
// Fetch new transactions and add them to the buffer
...

In Modification 2 (see below) the local state $\vec{\mathtt{st}}$ is stored directly instead of being encoded in chain $\mathcal{C}_{loc}$. The procedures **ExtendState(st)** and **FetchInformation** are modified to accommodate this change. Note that the $\mathcal{C}_{loc}$ is stored in $\mathcal{T}$ as we have seen in the first modification. This implies that the behavior of **ExtendState(st)** remains the same as in the first modification.

Protocol Modification 2

Initialization:
The protocol manages [...] a local (working) state $\vec{\mathtt{st}}$ (initially also the genesis state).[...]
The protocol additionally maintains a tree $\mathcal{T}$ of valid chains which initially contains the genesis chain $(\mathbf{G})$.

ExtendState(st):

if isvalidstate($\vec{\mathtt{st}}||\mathtt{st}$) $= 1$ **then**
  if there exists $\mathcal{C} \in \mathcal{T}$ which encodes $\vec{\mathtt{st}}$ **then**
    $\mathcal{C}_{\text{new}} \leftarrow \mathsf{extendchain}_{\mathsf{D}}(\mathcal{C}, \mathtt{st}, q)$
    if $\mathcal{C}_{\text{new}} \neq \mathcal{C}$ **then**
      Add $\mathcal{C}$ to $\mathcal{T}$
      Update the local state, i.e., $\vec{\mathtt{st}} \leftarrow \vec{\mathtt{st}}||\mathtt{st}$.
    end if
    // Broadcast current chain
    Send (MULTICAST, sid, $\mathcal{C}$) to $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$
  end if
end if

FetchInformation:

// Update the local state
Send (FETCH, sid) to $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$; denote the response from $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$ by (FETCH, sid, b).
Extract all valid chains $\mathcal{C}_1, \ldots, \mathcal{C}_k$ from b and add them to $\mathcal{T}$.
Extract all state $\vec{\mathtt{st}}_1, \ldots, \vec{\mathtt{st}}_k$ from chains $\mathcal{C}_1, \ldots, \mathcal{C}_k$.
Set both $\vec{\mathtt{st}}, \vec{\mathtt{st}}_{\text{exp}}$ to the longest state in $\vec{\mathtt{st}}, \vec{\mathtt{st}}_{\text{exp}}, \vec{\mathtt{st}}_1, \ldots, \vec{\mathtt{st}}_k$ (to resolve ties the ordering decides).
// Fetch new transactions and add them to the buffer
...

In Modification 3 (see below) parts of the procedures **ExtendState(st)** and **FetchInformation** are split off into separate sub-procedures. Otherwise the protocol remains the same. As there are no changes to the program logic the protocol still has the same behavior as the original protocol.

Protocol Modification 3

Initialization:

The protocol manages [...] a local (working) state $\vec{\mathtt{st}}$ (initially also the genesis state).[...]
The protocol additionally maintains a tree $\mathcal{T}$ of valid chains which initially contains the (genesis) chain ($\mathbf{G}$).

ExtendState(st):

$B \leftarrow$ SUBMIT-NEW($\vec{\mathtt{st}}, \mathtt{st}$)
if $B = 1$ **then**
  Update the local state, i.e., $\vec{\mathtt{st}} \leftarrow \vec{\mathtt{st}}||\mathtt{st}$.
end if
// Broadcast current chain
Execute CONTINUE.

Procedure SUBMIT-NEW($\vec{\mathtt{st}}, \mathtt{st}$):

if isvalidstate($\vec{\mathtt{st}}||\mathtt{st}$) $= 1$ **then**
  if there exists $\mathcal{C}' \in \mathcal{T}$ which encodes $\vec{\mathtt{st}}$ **then**
    Set $\mathcal{C} \leftarrow \mathcal{C}'$. // $\mathcal{C}$ is assumed to be a global variable
    $\mathcal{C}_{\text{new}} \leftarrow \mathsf{extendchain}_{\mathsf{D}}(\mathcal{C}, \mathtt{st}, q)$
    if $\mathcal{C}_{\text{new}} \neq \mathcal{C}$ **then**
      Add $\mathcal{C}$ to $\mathcal{T}$
      return 1
    end if
    return 0
  end if
end if

Procedure CONTINUE:

Send (MULTICAST, sid, $\mathcal{C}$) to $\mathcal{F}^{\mathsf{bc}}_{\text{N-MC}}$

FetchInformation:

// Update the local state
$(\vec{\mathtt{st}}_1, \ldots, \vec{\mathtt{st}}_k) \leftarrow$ FETCH-NEW
Set both $\vec{\mathtt{st}}, \vec{\mathtt{st}}_{\text{exp}}$ to the longest state in $\vec{\mathtt{st}}, \vec{\mathtt{st}}_{\text{exp}}, \vec{\mathtt{st}}_1, \ldots, \vec{\mathtt{st}}_k$ (to resolve ties the ordering decides).

// Fetch new transactions and add them to the buffer
...

Procedure FETCH-NEW:
Send (FETCH, sid) to $\mathcal{F}_{\text{N-MC}}^{\text{bc}}$; denote the response from $\mathcal{F}_{\text{N-MC}}^{\text{bc}}$ by (FETCH, sid, b).
Extract all valid chains $\mathcal{C}_1, \ldots, \mathcal{C}_k$ from b and add them to $\mathcal{T}$.
Extract states $\vec{\mathtt{st}}_1, \ldots, \vec{\mathtt{st}}_s$ from $\mathcal{C}_1, \ldots, \mathcal{C}_k$ and output them.

Finally consider the part of Modular-Ledger-Protocol below which is the same as Modification 3 except that the chain storage $\mathcal{T}$ and the calls to sub-procedures SUBMIT-NEW, CONTINUE, and FETCH-NEW are replaced by the calls to $\mathcal{F}_{\text{StX}}$. Lemma 3 implies that the behavior of **Modification** 3 and the protocol Modular-Ledger-Protocol is the same. This concludes the game-hopping argument and the proof. □

Protocol Modular-Ledger-Protocol Part

Initialization:
The protocol manages [...] a local (working) state $\vec{\mathtt{st}}$ (initially also the genesis state). [...]

ExtendState(st):
Send (SUBMIT-NEW, sid, $\vec{\mathtt{st}}$, $\mathtt{st}$) to $\mathcal{F}_{\text{StX}}$.
Denote the response by (SUCCESS, sid, B) of $\mathcal{F}_{\text{StX}}$.
if $B = 1$ **then**
  Update the local state, i.e., $\vec{\mathtt{st}} \leftarrow \vec{\mathtt{st}}||\mathtt{st}$.
end if
// Broadcast current state using $\mathcal{F}_{\text{StX}}$.
Send (CONTINUE, sid) to $\mathcal{F}_{\text{StX}}$

FetchInformation:
// Fetch new states and update the local state
Send (FETCH-NEW, sid) to $\mathcal{F}_{\text{StX}}$.
Denote the response from $\mathcal{F}_{\text{StX}}$ by (FETCH-NEW, sid, $(\vec{\mathtt{st}}_1, \ldots, \vec{\mathtt{st}}_k)$).
Set $\vec{\mathtt{st}}, \vec{\mathtt{st}}_{\text{exp}}$ to the longest state in $\vec{\mathtt{st}}, \vec{\mathtt{st}}_{\text{exp}}, \vec{\mathtt{st}}_1, \ldots, \vec{\mathtt{st}}_k$ (to resolve ties the ordering decides).
// Fetch new transactions and add them to the buffer
...

D.3 Proof of Theorem 1

Theorem (1). *Let the functions* ValidTx$_{\mathcal{B}}$, blockify$_{\mathcal{B}}$, *and* ExtendPolicy *be as defined above. Let $p \in (0,1)$, integer $q \geq 1$, $p_H = 1 - (1-p)^q$, and $p_A = p$. Let $\Delta \geq 1$ be the upper bound on the network delay. Consider* Modular-Ledger-Protocol$_T$ *in the* $(\mathcal{G}_{\text{CLOCK}}, \mathcal{F}_{\text{StX}}^{\Delta, p_H, p_A}, \mathcal{F}_{\text{N-MC}}^{\Delta})$*-hybrid world. If, for some $\lambda > 1$, the relation*

$$\alpha \cdot (1 - 2 \cdot (\Delta + 1) \cdot \alpha) \geq \lambda \cdot \beta \tag{1}$$

is satisfied in any real-world execution, where α and β are defined as above, then the protocol Modular-Ledger-Protocol$_T$ *UC-realizes* $\mathcal{G}_{\text{LEDGER}}^{\mathcal{B}}$ *for any ledger parameters (which are positive and integer-valued) in the range*

$$\mathtt{windowSize} = T \quad \text{and} \quad \mathtt{Delay} = 4\Delta,$$

$$\mathtt{maxTime}_{\mathtt{window}} \geq \frac{2 \cdot \mathtt{windowSize}}{(1-\delta) \cdot \gamma} \quad \text{and} \quad \mathtt{minTime}_{\mathtt{window}} \leq \frac{2 \cdot \mathtt{windowSize}}{(1+\delta) \cdot \max_r \mathtt{T}_{mp}(r)},$$

$$\eta > (1+\delta) \cdot \mathtt{windowSize} \cdot \frac{\beta}{\gamma},$$

where $\gamma := \frac{\alpha}{1+\Delta\alpha}$ and $\delta > 0$ is an arbitrary constant. In particular, the realization is perfect except with probability $R \cdot negl(T)$, where R denotes the upper bound on the number of rounds, and $negl(T)$ denotes a negligible function in T.

Proof. In order to show the theorem we specify the simulator for the ideal world $\mathcal{S}_{\text{ledg}}$. $\mathcal{S}_{\text{ledg}}$ is specified below as pseudo-code. Let us explain the general structure: the simulator internally runs the round-based

mining procedure of every honest party. Whenever a working mini-round is over, i.e., whenever the real world parties have issued their queries to $\mathcal{F}_{\text{STX}}$, then the simulator will assemble the views of its simulated honest miners and determine their common prefix of states, which is the longest state stored or received by each simulated party when chopping off T blocks. The adversary will then propose a new block candidate, i.e., a list of transactions, to the ledger to announce that the common prefix has increased (procedure ExtendLedgerState). The ledger will apply the Blockify on this list of transactions and add it to the state. Note that since Blockify does not depend on time, the current time of the ledger has no influence on this output. To reflect that not all parties have the same view on this common prefix, the simulator can adjust the state pointers accordingly (procedure AdjustView). The simulation inside the simulator is perfect and is simply the emulation of real-world processes. What restricts a perfect simulation is the requirement of a consistent prefix and the restrictions imposed by ExtendPolicy. In order to show that these restrictions are not forbidding a proper simulation, we have to justify, why the choice of the parameters in the theorem are sufficient to guarantee that (except with negligible probability). To this end, we analyze the real-world execution to bound the corresponding bad events that prevent a perfect simulation.

We start by analyzing the real-world execution $\text{EXEC}_{\mathsf{Modular\text{-}Ledger\text{-}Protocol},\mathcal{A},\mathcal{Z}}$. We follow the detailed analysis provided by Pass, Seeman, and shelat to analyze the real-world execution. The analysis is divided into six different claims about the real-world execution. They include properties such as a lower-bound on the chain growth (Claim 2.), the chain quality (Claim 3.), or an upper-bound on the chain growth (Claim 6.). For completeness, we prove the Claims 1.-6. below, which represent the core arguments (with some slight adaptations) also in our setting. These Claims prove that our simulator can simulate the real-world view perfectly, since the restrictions imposed by the ledger prohibit that only with negligible probability, where the distinguishing advantage is upper bounded by $R \cdot \text{negl}(T)$, where R denotes the number of rounds the protocol is running and $\text{negl}(\cdot)$ denotes a negligible function in the parameter T.

By abusing a bit of notation, use the term $\text{EXEC}_{Real,\mathcal{A},\mathcal{Z}}$ to refer to the entire random experiment defined by the UC execution of a protocol with adversary $\mathcal{A}$ and environment $\mathcal{Z}$ (instead of only to the output of $\mathcal{Z}$). We then denote by $f(\text{EXEC}_{Real,\mathcal{A},\mathcal{Z}})$ the induced random variable where f is a function on the entire view of an execution. For notational simplicity, for Claims 1 to 6, if we speak of *honest* miners, we always mean *honest and synchronized* miners. In addition, recall that each round consists of two time-ticks. Hence, if a statement is expressed with respect to a certain number r of rounds, it can equivalently be expressed with respect to $2r$ clock-ticks. [28]

Claim. 1. State dissemination: *Let p_i and p_j be miners, and let $r \geq 0$. Assume p_i is honest in round r, and the longest state received or stored by p_j has length ℓ. For any honest miner p_j in round $r + \Delta$, it holds that the longest state received or stored by p_j has length at least ℓ.*

Proof: By assumption, all messages, and in particular transmitted states of honest miners, are delayed maximally by Δ rounds. Thus, if an honest miner receives a state of length ℓ, then any other honest miner will receive this state within the next Δ rounds. Additionally, if a honest miner successfully mines a block, this new state will arrive at any honest miner at latest after Δ rounds. By then, any honest miner will have adopted a chain of length at least ℓ. ■

Claim. 2. Minimal number of mined blocks: *Let p_i be a miner, and let $r \geq 0$. Assume p_i is honest in round r, and the longest state received or stored by p_i in round r has length ℓ. Then, in round $r + t$, it holds that for any $\delta > 0$, except with probability $R \cdot negl(T)$, the length of the longest state (received or stored) of any honest miner p_j in that round has length at least $\ell + T$ if $t \geq \frac{T}{(1-\delta)\cdot\gamma}$.*

Proof: We first prove that for any real-world adversary $\mathcal{A}$, there is an adversary $\mathcal{A}'$ that, starting at the given round r, maximally delays messages starting and prove that in a real-world execution with $\mathcal{A}'$ the expected state length of an honest miner in round $r + t$, where the expectation is taken over the randomness of the adversarial strategy, is no larger than with adversary $\mathcal{A}$ in round $r + t$.

[28] The theorem statement is expressed with respect to clock ticks since the clock is the main reference point of time. However, for the proof, it is easier to think in rounds.

Construction of $\mathcal{A}'$. Given adversary $\mathcal{A}$, the adversary $\mathcal{A}'$ works as follows. It internally runs $\mathcal{A}$ until round r without any modifications. At round r, $\mathcal{A}'$ first delays all current messages in the network to the maximally possible delay. Also, from round r onward, whenever an honest party sends a message containing a state, $\mathcal{A}'$ sets the maximal delay Δ for this message. Message delays defined by $\mathcal{A}$ for messages that contain valid states of honest parties are ignored. The adversary further ignores any message sent by $\mathcal{A}$ on behalf of corrupted parties starting from round r.

We define the function $\text{LEN}_i^r(\text{EXEC}_{Real,\mathcal{A}(\sigma),\mathcal{Z}})$ to be the length of the longest chain (honest) miner i in round r in the real world experiment. Further, let $\mathcal{A}(\sigma)$ denote the behavior of $\mathcal{A}$ on (internal) randomness σ, further, denote by $Real(\sigma')$ the real-system, where the internal randomness of $\mathcal{F}_{\text{STX}}$ is fixed to σ'. We give an inductive proof to show that for any $r > 0$, $\text{LEN}_i^{r+t}(\text{EXEC}_{Real(\sigma'),\mathcal{A}(\sigma),\mathcal{Z}}) \geq \text{LEN}_i^{r+t}(\text{EXEC}_{Real(\sigma'),\mathcal{A}'(\mathcal{A}(\sigma)),\mathcal{Z}})$.

Base Case, $t = 0$: Since adversary $\mathcal{A}'$ and $\mathcal{A}'$ behave identical up to and including round $r-1$, the length of the longest state known or received by any party is the same. The reason is that $\mathcal{A}'$ and $\mathcal{A}$ play exactly the same strategy when σ is fixed. Furthermore, when the randomness σ' of $\mathcal{F}_{\text{STX}}$ is fixed, a miner i in any round r' is successful, if and only if it is successful in round r' with adversary $\mathcal{A}'$. Thus, $\text{LEN}_i^r(\text{EXEC}_{Real(\sigma'),\mathcal{A}(\sigma),\mathcal{Z}}) < \text{LEN}_i^r(\text{EXEC}_{Real(\sigma'),\mathcal{A}'(\mathcal{A}(\sigma)),\mathcal{Z}})$ only if player i receives a longer state in round r. Since $\mathcal{A}'$ additionally maximally delays messages, if any state arrives in round r in the real execution with $\mathcal{A}'$, then it arrives no later than r in the real execution with $\mathcal{A}$. This concludes the base case.

Induction Step: $t \to t+1$: By the induction hypothesis, we have

$$\text{LEN}_i^{r+t}(\text{EXEC}_{Real(\sigma'),\mathcal{A}(\sigma),\mathcal{Z}}) \geq \text{LEN}_i^{r+t}(\text{EXEC}_{Real(\sigma'),\mathcal{A}'(\mathcal{A}(\sigma)),\mathcal{Z}}).$$

We argue that $\text{LEN}_i^{r+t+1}(\text{EXEC}_{Real(\sigma'),\mathcal{A}(\sigma),\mathcal{Z}}) \geq \text{LEN}_i^{r+t+1}(\text{EXEC}_{Real(\sigma'),\mathcal{A}'(\mathcal{A}(\sigma)),\mathcal{Z}})$ holds as well. Assume not, then, by the above reasoning, it can only be due to miner i receiving a state in round $r+t+1$ that increases $\text{LEN}_i^{r+t+1}(\text{EXEC}_{Real(\sigma'),\mathcal{A}'(\mathcal{A}(\sigma)),\mathcal{Z}})$ but not $\text{LEN}_i^{r+t+1}(\text{EXEC}_{Real(\sigma'),\mathcal{A}(\sigma),\mathcal{Z}})$ (since the success of miner i in round $r+t+1$ is fixed given σ'. By the same reasoning as above, since $\mathcal{A}'$ maximally delays delivery of new states, if any state arrives in round r in the real execution with $\mathcal{A}'$, then it arrives no later than r in the real execution with $\mathcal{A}$. This concludes the induction proof. Taking the expectation over the randomness σ and σ', we conclude that for any round r, any $c \geq 0$, for the real UC-execution, it holds that

$$\begin{aligned}&\Pr[\text{LEN}_i^{r+t}(\text{EXEC}_{Real,\mathcal{A},\mathcal{Z}}) \leq \text{LEN}_i^{r}(\text{EXEC}_{Real,\mathcal{A},\mathcal{Z}}) + c]\\ &\leq \Pr[\text{LEN}_i^{r+t}(\text{EXEC}_{Real,\mathcal{A}',\mathcal{Z}}) \leq \text{LEN}_i^{r}(\text{EXEC}_{Real,\mathcal{A}',\mathcal{Z}}) + c].\end{aligned}$$

We say a round r' is *uniform* if $\text{LEN}_i^{r'}(\text{EXEC}_{Real,\mathcal{A}',\mathcal{Z}}) = \text{LEN}_j^{r'}(\text{EXEC}_{Real,\mathcal{A}',\mathcal{Z}})$ for all honest miners i and j. Recall that adversary $\mathcal{A}'$ does not broadcast adversarially generated states and any new state is delayed by exactly Δ rounds. The slowest progress of the overall maximal state length known to an honest party occurs in case uniform rounds are the only successful rounds (if at all). Otherwise, the honest miner with the longest state could be successful and broadcast a longer state at round r', which would be guaranteed to arrive to any other honest miner in $r + \Delta$.

Fix some round r. If in round $s = r + t$, the length increase of the overall longest state of an honest miner is less than c blocks, then at most $c \cdot \Delta$ non-uniform rounds occurred. Hence, there were at least $t - c \cdot \Delta$ uniform rounds. The probability that less than c new blocks are mined by honest miners (i.e., that less than c successful queries by honest miners to $\mathcal{F}_{\text{STX}}$ extended the state by one block) is thus $\Pr[\sum_{i=1}^{t-c\Delta} X_i < c]$, where X_i is a boolean random variable with mean α. Let $X := \sum_{i=1}^{t-c\Delta} X_i$ with mean $E[X] = \alpha \cdot t - a \cdot c \cdot \Delta$. To get an appropriate tail-estimate, we set $c = \frac{\alpha t}{1+\alpha\Delta}$ to obtain $E[X] = c$ and can apply the Chernoff bound to get

$$\begin{aligned}&\Pr[\text{LEN}_i^{s}(\text{EXEC}_{Real,\mathcal{A}',\mathcal{Z}}) \leq \text{LEN}_i^{r}(\text{EXEC}_{Real,\mathcal{A}',\mathcal{Z}}) + (1-\delta)c]\\ &\leq \Pr[\text{LEN}^{s}(\text{EXEC}_{Real,\mathcal{A}',\mathcal{Z}}) \leq \text{LEN}^{r}(\text{EXEC}_{Real,\mathcal{A}',\mathcal{Z}}) + (1-\delta)c]\\ &\leq \Pr[X < (1-\delta)c] = \Pr[X < (1-\delta)\gamma t] \leq \exp\Big(\frac{-\delta^2}{2} \cdot \gamma t\Big).\end{aligned}$$

where we plugged-in the definition of γ.

By Claim 1 (chain dissemination), we know that if some honest party has a state in round r, then any honest miner will have a state of at least this length in round $r + \Delta$. Hence, for realizing the ledger,

we see that state extend happens if at least one honest miner has a new state which happens at a rate at least γt (except with probability $\mathrm{negl}(\gamma t) = \mathrm{negl}(T)$ by the lower bound on t).

By Claim 1, we see that if an honest miner knows some state, then within Δ rounds, every other honest miner will be aware of that. A similar calculation shows that the lower bound on the time to have a state increase by T blocks by all honest parties follows the same law (and hence the perceived ledger speed is the same). By requiring $s = r + t - \Delta$ above (such that in round $r + t$ all honest miners have a state that increased by at least $(1 - \delta)\gamma t$ since round r), and letting $Y := \sum_{i=1}^{t-\Delta-c\Delta} X_i$ (for X_i as above), we get

$$\Pr[\textsc{Len}^s(\mathrm{EXEC}_{Real,\mathcal{A}',\mathcal{Z}}) \leq \textsc{Len}^r(\mathrm{EXEC}_{Real,\mathcal{A}',\mathcal{Z}}) + (1-\delta)c]$$
$$\leq \Pr[Y < (1-\delta)\gamma(t-\Delta)] \leq \exp\Big(\frac{-\delta^2}{2} \cdot \gamma(t-\Delta)\Big). \quad (2)$$

Since $\gamma t - 1 < \gamma t - \gamma\Delta < \gamma t$, this implies that $\Pr[Y < (1-\delta')\gamma t] \leq \exp\Big(\frac{-\delta^2}{2} \cdot \gamma t\Big)$ for any δ' by choosing a sufficiently small constant δ in Equation 2, yielding a negligible function in $\gamma t \geq T$. Finally, since α (and thus γ) is the lower bound for any round r, taking the union bound over the polynomial number of rounds yields Claim 2. ■

The following claims use the fact that Equation 1 implies that there exists $0 < \widehat{\delta} < 1$ such that 1.) $\alpha > \gamma > (1 + \widehat{\delta})\beta$, and 2.) $(1 - \widehat{\delta}) > (\Delta + 1)\alpha$. This is proven in [PSS17] (Claim 6.12).

Claim. 3. Fraction of honest blocks: *Let p_i be a miner, and let $r \geq 0$. Assume p_i is honest in round r, and the length of the longest state received or stored is $\ell \geq T$. The fraction of adversarially mined blocks within a sequence of T blocks in the state is at most* $\min\{1, (1+\delta)\cdot\frac{\beta}{\gamma}\}$ *except with probability $R \cdot negl(T)$ for any $\delta > 0$.*

Proof: Let us assume that at round r', the state of miner p_i is $\vec{\mathtt{st}}_{r'} = \mathtt{st}_0||\ldots||\mathtt{st}_k$. We show that in any sub-sequence of T state blocks $\mathtt{st}_{j+1},\ldots,\mathtt{st}_{j+T}$ in $\vec{\mathtt{st}}_r$, the fraction of adversarially mined blocks is bounded. Without loss of generality, one can assume that the state $\vec{\mathtt{st}}^{<j} := \mathtt{st}_0||\ldots||\mathtt{st}_j$ as well as the state $\vec{\mathtt{st}}^{>j+T} := \mathtt{st}_0||\ldots||\mathtt{st}_{j+T+1}$ are mined by honest miners (unless $j+T$ is the maximum length of any state known to $\mathcal{F}^{\Delta}_{\text{STX-BD}}$). Otherwise, one can enlarge T to meet this condition, as any state is finite and starts with the genesis block. We further assume that $\vec{\mathtt{st}}^{<j}$ is mined at round r, and that in round $r+t$, the state $\vec{\mathtt{st}}^{>j+T}$ appears for the first time as the state, or the prefix of a state, of at least one honest miner. We conclude that if an adversary successfully extended the state during some round by a new state block $\mathtt{st}_{j+s}$ of the above sequence $\mathtt{st}_{j+1},\ldots,\mathtt{st}_{j+T}$, then this happens in a round between r and $r+t$.

We now relate the number t of rounds to the number T of blocks. Since t is assumed to be the minimal number of rounds until the first honest miner adopted a state containing $\mathtt{st}_{j+1}$, we can invoke Claim 2, to conclude that the probability that the condition $t > \frac{T}{(1-\delta')\gamma}$ occurs in such an execution is at most $\mathrm{negl}(T)$. We hence have $t \leq \frac{T}{(1-\delta')\gamma}$ with overwhelming probability in T.

On the other hand, we can lower bound the number of rounds needed to generate a state increase by T blocks by standard Chernoff bound: using the relation of $p_H = 1-(1-p)^q$ our assumed $\mathcal{F}_{\text{STX}}$, we have $\mathtt{T}_{mp}(r) \leq (q_A^{(r)} + q \cdot q_H^{(r)}) \cdot p$. Let $\mathtt{T}_{mp} := \max_{r\in[R]} \mathtt{T}_{mp}(r)$ and denote by $q_{\max}$ be the corresponding upper bound on the query power. The event that during t rounds, more than $T = (1+\delta)\cdot q_{\max}\cdot p\cdot t$ times a successful state extension happens (via a query to $\mathcal{F}_{\text{STX}}$), occurs with probability at most $\exp\Big(-\frac{\delta^2}{3}q_{\max}pt\Big) \in \mathrm{negl}(T)$ only (for any $0 < \delta < 1$). Stated differently, with overwhelming probability, for the assumed sequence of T state blocks, the number of rounds t needed to mine the new state is at least $\frac{T}{(1+\delta)q_{\max}p}$.

Additionally, we know that the upper bound on the expected number of adversarial successes to extend a state in one round is β, and the upper bound on the expected number of successes (i.e., newly minded state blocks) within t rounds by the adversary, denoted as the random variable N_A^t, is thus $t\cdot\beta$ by linearity of expectation. The random variable N_A^t is hence the sum of $q'_{\max}$ binary random variables (where $q'_{\max}$ is the maximum number of queries contributed by corrupted as well as by honest but de-synchronized miners) each being successful with probability p. By the Chernoff bound, we get a tail-estimate of

$$\Pr[N_A^t > (1+\delta)t\cdot\beta] \leq \exp\Big(\frac{-\delta^2}{3}t\beta\Big),$$

again for any $0 < \delta < 1$. By the above lower bound on t and since β is the maximum expected value over all rounds, and recall that β is a ρ fraction of the maximum mining power $\mathtt{T}_{mp}$ of a round within this t-round interval, we conclude that $\frac{\rho T}{1+\delta} \leq \beta t$ holds with overwhelming probability. Hence, the above function is indeed a negligible function in T. Therefore, except with negligible probability in T, the number of times the adversary was successful in extending the state by one block is upper bounded by

$$N_A^{\frac{T}{(1-\delta')\gamma}} \leq \frac{1+\delta}{1-\delta'} \cdot T \cdot \frac{\beta}{\gamma}.$$

Hence, the fraction of adversarial blocks within T consecutive blocks cannot be more than $f = \min\{1, (1+\delta'')\frac{\beta}{\gamma}\}$ for any δ'' and sufficiently small constants $\delta, \delta' > 0$, except with negligible probability in the length T of the sequence. By Equation 1, we even have $f < 1$ (for an appropriate choice of δ''). Since β is the maximum expected value in any round, the proof is concluded by taking the union bound over the number of such sequences (which is in the order of number of rounds). ■

Claim. 4. Withholding of adversarial block: *For any round r, the event that an honest miner accepts a new (i.e., longer) state at round r and at least one block in the extension was mined before round $r - \omega t$, happens with probability negl(βt), for any $0 < \omega < 1$.*

Proof: Let us define $\vec{\mathtt{st}}_r = \mathtt{st}_0 || \dots || \mathtt{st}_k$ to be the longest state known to $\mathcal{F}_{\text{STX}}$ at round r. Let $\vec{\mathtt{st}}_{r'}$ be the longest prefix of $\vec{\mathtt{st}}_r$ such that $\vec{\mathtt{st}}_{r'}$ was mined by an honest miner or it is the genesis block. This state was known to at least one honest party by round $r' \leq r$. Now, assume that $r - r' \geq \omega t$ (as otherwise, by Claim 1, the extension trivially contains a more recent block from an honest party). By the chain growth lowerbound, we know that except with probability negl(γt), $|\vec{\mathtt{st}}_r| - |\vec{\mathtt{st}}_{r'}| \geq (1-\delta) \cdot \gamma \omega t$. Since $\gamma t > \beta t$, we have that this holds except with probability negl(βt).

Analogously to the previous claim, the number of new states mined by the adversary is upper bounded by $(1+\delta') \cdot \beta \omega t$ (except with probability negl(βt)). Since all $|\vec{\mathtt{st}}_r| - |\vec{\mathtt{st}}_{r'}|$ blocks are mined by the adversary we have $|\vec{\mathtt{st}}_r| - |\vec{\mathtt{st}}_{r'}| \leq (1+\delta') \cdot \beta \omega t$. We get

$$(1-\delta) \cdot \gamma \omega t \leq (1+\delta') \cdot \beta \omega t,$$

which, for sufficiently small δ, δ' implies that $\gamma < (1+\delta'')\beta$ for any δ''. This contradicts Equation 1 and the claim follows. ■

Claim. 5. Consistent states: *Let p_i and p_j be miners, and let $r' \geq r \geq 0$. Assume p_i is honest in round r, and p_j is honest in round r'. Assume that the length of the longest state received or stored by p_i in round r is $\ell \geq T$. Then, the $\ell - T$-prefix of that longest state of p_i in round r is identical to the $\ell - T$-prefix of the state of p_j stored or received in round r' except with probability $R \cdot$ negl(T).*

Proof: We again follow the exposition in [PSS17]. Since an inconsistency at round r implies an inconsistency at round $r' > r$, if the claim is proven for the case that $r \leq r' \leq r+1$, then by an inductive argument, the claim holds for any $r' \geq r$.

The protocol for the honest miners truncates the T newest blocks from the current respective state of each miner. Thus, we need to argue that the block which is $T+1$ far away from the head will be part of any state output by any honest miner. Suppose we are at round r' in the protocol, then the time it takes to generate the last T blocks is at least $t \geq \frac{T}{(1+\delta) q_{\max} p}$ except with negligible probability as argued in the previous claim, where $q_{\max}$ is the maximum number of queries per round of all registered miners (maximum over all rounds) and any $0 < \delta < 1$. We can thus follow the argument by [PSS17] to conclude that the probability that the states of two honest miners diverge at round $s := r - t$ is a negligible function negl(βt) and thus also a negligible function in T. The last step follows as in the previous claim by observing that $\beta t \geq \frac{\rho}{(1+\delta)} T$, where β is the maximum expected value of adversarially mined blocks (over all rounds).

In the interval between s and r', the expected number of rounds, where at lest one honest miner is successful, is at least αt. Thus, by a standard Chernoff bound, the probability that the number of these successful rounds is smaller than $\bar{q}_{min} := (1-\delta') \cdot \alpha t$ is no more than $\exp\Big(-\frac{\delta'^2}{2}\alpha t\Big)$ in the real-world UC random experiment. Since Equation 1 implies that there exists a $\widehat{\delta} \in (0,1)$ such that $\alpha > \beta \cdot (1+\widehat{\delta})$, this probability is upper bounded by negl(βt).

Unfortunately, the "race" between the good guys and the bad guys is not yet conclusively analyzed, since the mere superiority of honestly mined blocks does not imply that the honest parties will reach agreement. In particular, not all of the, in expectation, αt, blocks are equally useful to obtain a so-called convergence opportunity. In particular, we need to know how many of the honestly mined blocks happen in isolated, sufficiently silent intervals.

Formally, let us introduce the random variable R_i that measures the number of elapsed round between successful round $i-1$ and successful round i in the real-world UC execution, where R_1 measures the number of elapsed rounds to the first successful round. Based on R_i, the random variable X_i is defined as follows: $X_i = 1$ if and only if $R_i > \Delta$ and exactly one honest miner mines a new state (i.e., successfully appends a new block to the state) in the ith successful round. Let E_1^i be the event that there is at least one successful round in the interval of Δ rounds starting after successful round $i-1$ (or at the onset of the experiment). Let E_2^i be the event that strictly more than one miner is successful in the following successful round i.

Overall, our goal is to suitably bound the number of blocks that prevent those events of "success & silence" (i.e., bound the probability of the event $X_i = 0$) in an interval of t rounds. We call these the undesirable blocks. Clearly, since we do a worst-case analysis, all (roughly βt) adversarial blocks will contribute to the number of undesirable blocks.

We now conduct the following thought-experiment: given the minimal set of synchronized parties and their associated mining power α, we study their production of undesirable blocks if only they ran the protocol alone (as in Claim 2). This thought-experiment is sound: all hash queries $q^{(r)}$ in a round r either contribute to the power of the synchronized parties or the adversary, i.e., $q^{(r)} = q_H^{(r)} + q_A^{(r)}$, and each query to $\mathcal{F}_{\mathrm{STX}}$ yields an i.i.d. trial of extending a state. Now, let x be the number of queries such that $\alpha = 1-(1-p)^x$ holds. Since by (worst-case) definition $\beta = \max_{r\in[R]} \rho \cdot \mathtt{T}_{mp}(r) = \rho \cdot p \cdot \max_{r\in[R]} q^{(r)}$, we can always assign the difference $(q_H^{(r)} - x)$ to the adversary's budget (and the condition $\alpha > \beta$ still holds). In particular, we have for integers $x, y > 1$,

$$\begin{aligned}\alpha_r - \alpha &= (1-(1-p)^{x+y}) - (1-(1-p)^x) = (1-p)^x - (1-p)^{x+y} \\ &= (1-p)^x \cdot (1-(1-p)^y) \leq (1-(1-p)^y) \leq (1-(1-y\cdot p)) \\ &= y \cdot p,\end{aligned}$$

where the last inequality is a consequence of Bernoulli's inequality. Hence, considering the worst-case adversary with mining power $\beta = p \cdot \max_{r\in[R]} q^{(r)}$ over t rounds allows us to study the number of undesirable blocks that the minimally honest parties produce alone according to our thought-experiment (and think of the additive term $(q_H^{(r)} - x)\cdot p$ as being part of the adversarial mining power β that anyway contributes to the production of undesirable blocks). In this thought experiment, we need to suitably bound the occurrence of the above two bad events E_j^i in our thought experiment. By a union bound we directly have that $\Pr[X_i = 0] = \Pr[E_1^i \cup E_2^i] \leq \Delta\alpha + \alpha$, hence, on the positive side, $\Pr[X_i = 1] \geq 1 - \alpha(\Delta+1)$.

Let $X := \sum_{i=1}^{\bar{q}_{min}} X_i$, where $E[X] = (1-\alpha(\Delta+1)) \cdot \bar{q}_{min}$. Let us define $\bar{q}'_{min} := (1-\delta'') \cdot (1-\alpha(\Delta+1)) \cdot \bar{q}_{min}$. Since the random variables X_i defined in our thought-experiment are i.i.d., it follows that $\Pr[X \leq \bar{q}'_{min}] \leq \exp\left(-\frac{\delta''^2}{2}(1-\alpha(\Delta+1))\cdot\bar{q}_{min}\right)$. Aside of $\alpha > \beta$, Equation 1 implies that there exists some $0 < \widehat{\delta} < 1$ such that $\alpha(\Delta+1) < 1-\widehat{\delta}$. We conclude that $(1-\alpha(\Delta+1))\cdot\bar{q}_{min} > \widehat{\delta}(1-\delta')\cdot\alpha t > \widehat{\delta}(1-\delta')\cdot\beta t$. Thus, $\Pr[X \leq \bar{q}'_{min}] \leq \exp\left(-\frac{\delta''^2}{2}\widehat{\delta}(1-\delta')\cdot\beta t\right) \in \mathrm{negl}(\beta t)$. And hence, we have a (high-probability) lower bound on the number of silent patterns.

We are actually interested in the number of times that $X_i = X_{i+1} = 1$. This situation, as already introduced above, means that we have a situation, in which for Δ rounds, no honest miner is successful, then exactly one honest miner is successful, and afterwards, we again have Δ rounds where no honest miner is successful. This is denoted in [PSS17] as a *convergence opportunity*. For example, a convergence opportunity has the desirable property, that *at the end* of such an opportunity, if the adversary is unable to provide a longer state to the honest miners during this period, all honest miners will reach an agreement on the current longest state. Thus, in order to prevent this, an adversary needs to be successful in mining roughly at the rate of the number of convergence opportunities within t rounds.

We have already seen that with overwhelming probability, there are at least $\bar{q}_{min}$ successful rounds, and among which $(\bar{q}_{min}-\bar{q}'_{min})$ could prevent convergence opportunities. Thus, the number of convergence opportunities C can, except with probability $\mathrm{negl}(\beta t)$, be (generously) lower bounded by $C \geq \bar{q}_{min} - 2(\bar{q}_{min} - \bar{q}'_{min}) = 2\bar{q}'_{min} - \bar{q}_{min} > (1-\epsilon)(1-2\alpha(\Delta+1))\alpha t$, for any constant ϵ (by picking δ' and δ'' sufficiently small).

The final argument is a counting argument. Let us denote by $\mathcal{S}_r$ the set of maximal states known to $\mathcal{F}^{\Delta}_{\text{STX-BD}}$ at round r' (i.e., any path from the root to some a leaf) of length at least $\ell + C$, where ℓ is the length of the longest state known to at least one honest miner at round s. Note that $\mathcal{S}_r^{\ell+C}$ is non-empty: since each convergence opportunity increases the length by at least one, and before each successful round, there is a period of Δ rounds where no honest miner mines a new state, there has to exist at least one state with length at least $\ell + C$ at round r'.

Assume that the number of successful state extensions made by the adversary (to $\mathcal{F}^{\Delta}_{\text{STX-BD}}$) between round s and r' is $T_{\mathcal{A}} < C$. Then, by the pigeonhole principle, for all $\vec{\mathtt{st}} \in \mathcal{S}_r$, it holds that there is at least one block $\mathtt{st}_k$, such that functionality $\mathcal{F}_{\text{STX}}$ is successfully queried by exactly one miner P in round i to extend the state to length $k+1$, but no query by the adversary to extend a state of length k to a state of length $k+1$ has been successful up to and including round r'. Even more, $T_{\mathcal{A}} < C$ implies that such an i has to exist that also constitutes a convergence opportunity.

After this convergence opportunity at round i, all honest miners have a state whose first $k+1$ blocks are $\vec{\mathtt{st}}_i = \mathtt{st}_0 \ldots, \mathtt{st}_k$. Unless the adversary provides an alternative state with a prefix $\vec{\mathtt{st}}_i'$ of length $k+1$, such that $\mathtt{st}_l' \neq \mathtt{st}_l$ for at least one index $0 < l \leq k$, no honest miner will ever mine on a state whose first $k+1$ blocks do not agree with $\vec{\mathtt{st}}_i$.

The existence of an alternative prefix $\vec{\mathtt{st}}_i'$ of length $k+1$ for any such i and for all states $\vec{\mathtt{st}} \in \mathcal{S}_r^{\ell+C}$ implies $T_{\mathcal{A}} \geq C$ and therefore contradicts the assumption that $T_{\mathcal{A}} < C$.

What is left to prove is that for any such interval of size t (from round s to round r'), the probability that $T_{\mathcal{A}} < C$ holds in any real-world execution except with negligible probability in βt.

First, for any $\omega > 0$, the probability that any new state accepted by an honest miner during the period of $(t+1)$ rounds is mined before round $s - \omega t$ is at most $\text{negl}(\beta t)$. Analogously to the previous claim, the number of adversarial blocks (i.e., successful state extensions by $\mathcal{A}$) generated within $(1+\omega)(t+1)$ rounds is (except with probability $\text{negl}(\beta t)$) upper bounded by $T_{\mathcal{A}} \leq (1+\delta)(1+\omega)(t+1)\beta \leq \frac{(1+\delta)(1+\omega)}{\lambda}(t+1)\alpha \cdot (1 - 2\alpha \cdot (\Delta + 1))$, where the last inequality follows from Equation 1. By picking the constants δ and ω, and ϵ sufficiently small relative to λ, we hence get $T_{\mathcal{A}} < C$ except with probability $\text{negl}(\beta t)$. Since β is the maximal value over all β_r for any round, we again apply the union bound over the number of rounds and get the desired claim. ■

Claim. 6. Maximal number of mined blocks *Let p_i be a miner, and let $r \geq 0$. Assume p_i is honest in round r, and the longest state received or stored by p_i in round r has length ℓ. Then, in round $r+t$, it holds, except with probability $R \cdot \text{negl}(T)$, that the length of the longest state (received or stored) of at least one honest miner p_j in that round has length at most $\ell + T$ if $t \leq \frac{T}{(1+\delta) \cdot p \cdot q_{\max}}$, where $q_{\max}$ is the maximum query power in any round of this interval.*

Proof: To upper bound the number of accepted blocks, we have to combine two observations made above: we have seen that the time it takes to generate T blocks is at least $t \geq \frac{T}{(1+\delta) q_{\max} p}$ except with probability $\text{negl}(\beta t)$. Hence, with overwhelming probability, in $t \leq \frac{T}{(1+\delta) \cdot p \cdot q_{\max}}$, no more than T blocks are mined. Furthermore, we have seen in Claim 4 that for any round r, the probability that a new state is accepted by an honest miner but this state was mined before round $r - \omega t$ happens with probability $\text{negl}(\beta t)$, for any $0 < \omega < 1$. Thus, in the interval between r and t, for t as bound in the Claim statement, the state can increase by at most T state blocks except with probability $\text{negl}(T)$, since we have again, $\beta t \geq \frac{\rho}{(1+\delta)} T$. Since β is the maximum expected value of adversarially mined blocks (over all rounds), the claim follows by taking the union bound over all rounds. ■

Finally, we conclude the proof by noting that after a delay of Δ rounds, all honest transactions are known to all honest miners, so, as soon as an honest miner mines the next state block, he for sure puts all these transactions into his next blocks if they are valid. In the simulation, the simulator also does it in the ideal world and hence will never propose blocks of honest parties that do not comply with the conditions of the defined ExtendPolicy of $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$. Further, the synchronization of a party takes at most $\mathtt{Delay} = 4\Delta$ clock ticks: if p_j joins the network, his knowledge of the longest chain and the set of valid transactions relative to that state, which is known to at least one honest and synchronized miner is only reliable after 2Δ rounds (4Δ clock ticks) since it takes at most Δ rounds to multicast the initial message that the miner has joined the network, and additional Δ rounds until the replies are received. During this 2Δ round the new miner will also have received all messages sent at or after he joined the network, and in particular all transactions that are more than Δ rounds ($2\Delta = \frac{\mathtt{Delay}}{2}$) old and potentially valid. Finally, setting $\mathtt{windowSize} \geq T$, follows from the arguments of Claim 4 and 5, since the adversary cannot have mined a state that is T blocks larger than some chain from a honest and synchronized party. But this would be a necessary condition to provoke a slackness that exceeds T. This concludes the proof. □

It follows the formal specification of the simulator.

Simulator $\mathcal{S}_{\text{ledg}}$

Initialization:

The simulator manages internally a simulated state-exchange functionality $\mathcal{F}_{\text{STX}}$, a simulated network $\mathcal{F}_{\text{N-MC}}$. An honest miner p registered to $\mathcal{G}^{B}_{\text{LEDGER}}$ is assumed to be registered in all simulated functionalities. Moreover, the simulator maintains the local state $\vec{\mathtt{st}}_p$ and the buffer of transactions $\mathtt{buffer}_p$ of such a party. Upon any activation, the simulator will query the current party set from the ledger (and simulate the corresponding message they send out to the network in the first maintain-ledger activation after registration), query all activations from honest parties $\vec{\mathcal{I}}_H^T$, and read the current clock value to learn the time. In particular, the simulator knows which parties are honest and synchronized and which parties are de-synchronized.

General Structure:

The simulator internally runs adversary $\mathcal{A}$ in a black-box way and simulates the interaction between $\mathcal{A}$ and the (emulated) real-world hybrid functionalities. The inputs from $\mathcal{A}$ to the clock are simply relayed (and the replies given back to $\mathcal{A}$). The ideal world consists of the ledger functionality and the clock.

Messages from the Clock:

– Upon receiving ($\text{CLOCK-UPDATE}, \text{sid}_C, p$) from $\mathcal{G}_{\text{CLOCK}}$, if p is an honest registered party, then remember that this party has received such a clock update (and the environment gets an activation). Otherwise, send ($\text{CLOCK-UPDATE}, \text{sid}_C, p$) to $\mathcal{A}$. In addition (before releasing the activation token), the simulator checks whether the clock advances. If so, and if this was a working mini-round (and hence all maintain commands have already been submitted by honest and synchronized parties), then execute ExtendLedgerState before giving the activation to $\mathcal{A}$.

Messages from the Ledger:

– Upon receiving ($\text{SUBMIT}, \mathtt{BTX}$) from $\mathcal{G}^{B}_{\text{LEDGER}}$ where $\mathtt{BTX} := (\mathtt{tx}, \text{txid}, \tau, p)$ forward ($\text{MULTICAST}, \text{sid}, \mathtt{tx}$) to the simulated network $\mathcal{F}_{\text{N-MC}}$ in the name of p. Output the answer of $\mathcal{F}_{\text{N-MC}}$ to the adversary.
– Upon receiving ($\text{MAINTAIN-LEDGER}, \text{sid}, \text{minerID}$) from $\mathcal{G}^{B}_{\text{LEDGER}}$, extract from $\vec{\mathcal{I}}_H^T$ the party p_i that issued this query. If p_i has already done its instructions for the current mini-round, then ignore the request. Otherwise, do:
  1. Execute SimulateMining(p_{minerID}, τ) and if this was the last maintain command in a working mini-round and the round will advance, then execute ExtendLedgerState before giving the activation to $\mathcal{A}$.
  2. In addition, remember that party p_i is done with mining in the current mini-round.
– Upon any further activation of the simulator, the simulator inspects the entire sequence of inputs by honest parties to the ledger $\vec{\mathcal{I}}_H^T$ and does the following:
  1. For any input, $I = (\text{READ}, \text{sid})$ of party P, if the current round is an update mini-round, then execute Step 4 of the mining procedure as below in SimulateMining
  2. Remember that the update for party P is done for this round.

Simulation of the State Exchange Functionality:

– Upon receiving ($\text{SUBMIT-NEW}, \text{sid}, \vec{\mathtt{st}}, \mathtt{st}$) from $\mathcal{A}$ on behalf of a corrupted $p \in \mathcal{P}_{stx}$, then relay it to the simulated $\mathcal{F}_{\text{STX}}$ and do the following:
  1. If $\mathcal{F}_{\text{STX}}$ returns ($\text{SUCCESS}, B$) give this reply to $\mathcal{A}$
  2. If $\mathcal{A}$ replies with ($\text{CONTINUE}, \text{sid}$), input ($\text{CONTINUE}, \text{sid}$) to the simulated $\mathcal{F}_{\text{STX}}$
  3. If the current mini-round is an update mini-round, then execute ExtendLedgerState
– Upon receiving ($\text{FETCH-NEW}, \text{sid}$) from $\mathcal{A}$ (on behalf of a corrupted p) forward the request to the simulated $\mathcal{F}_{\text{STX}}$ and return whatever is returned to $\mathcal{A}$.
– Upon receiving ($\text{SEND}, \text{sid}, s, p'$) from $\mathcal{A}$ on behalf some *corrupted* party P, do the following:
  1. Forward the request to the simulated $\mathcal{F}_{\text{STX}}$.
  2. If the current mini-round is an update mini-round, then execute ExtendLedgerState
  3. Return to $\mathcal{A}$ the return value from $\mathcal{F}_{\text{STX}}$.
– Upon receiving ($\text{SWAP}, \text{sid}, \text{mid}, \text{mid}'$) from $\mathcal{A}$, forward the request to the simulated $\mathcal{F}_{\text{STX}}$ and return whatever is returned to $\mathcal{A}$.
– Upon receiving ($\text{DELAY}, \text{sid}, T, \text{mid}$) from $\mathcal{A}$ forward the request to the simulated $\mathcal{F}_{\text{STX}}$ and do the following:
  1. Query the ledger state $\mathtt{state}$

2. Execute $\textsc{AdjustView}(\texttt{state})$
3. Return to $\mathcal{A}$ the output of $\mathcal{F}_{\text{StX}}$

Simulation of the Network (over which transactionss are sent) :

– Upon receiving ($\textsc{multicast}, \text{sid}, (m_{i_1}, p_{i_1}), \ldots, (m_{i_\ell}, p_{i_\ell})$ with list of transactions from $\mathcal{A}$ on behalf some *corrupted* $P \in \mathcal{P}_{net}$, then do the following:
  1. Submit the transactions to the ledger on behalf of this corrupted party, and receive for each transaction the transaction id txid
  2. Forward the request to the internally simulated $\mathcal{F}_{\text{N-MC}}$, which replies for each message with a message-ID mid
  3. Remember the association between each mid and the corresponding txid
  4. Provide $\mathcal{A}$ with whatever the network outputs.
– Upon receiving (an ordinary input) ($\textsc{multicast}, \text{sid}, m$) from $\mathcal{A}$ on behalf of some *corrupted* $P \in \mathcal{P}_{net}$, then execute the corresponding steps 1. to 4. as above.
– Upon receiving ($\textsc{fetch}, \text{sid}$) from $\mathcal{A}$ on behalf some *corrupted* $P \in \mathcal{P}_{net}$ forward the request to the simulated $\mathcal{F}_{\text{N-MC}}$ and return whatever is returned to $\mathcal{A}$.
– Upon receiving ($\textsc{delays}, \text{sid}, (T_{\text{mid}_{i_1}}, \text{mid}_{i_1}), \ldots, (T_{\text{mid}_{i_\ell}}, \text{mid}_{i_\ell})$) from $\mathcal{A}$ forward the request to the simulated $\mathcal{F}_{\text{N-MC}}$ and return whatever is returned to $\mathcal{A}$.
– Upon receiving ($\textsc{swap}, \text{sid}, \text{mid}, \text{mid}'$) from $\mathcal{A}$ forward the request to the simulated $\mathcal{F}_{\text{N-MC}}$ and return whatever is returned to $\mathcal{A}$.

procedure $\textsc{SimulateMining}(P, \tau)$
  Simulate the mining procedure of P of the protocol:
  if time-tick τ corresponds to a working sub-round **then**
    Execute Step 2 of the mining protocol. This includes:
      -Define the next state block $\texttt{st}$ using the transaction set Txs_P
      -Send ($\textsc{submit-new}, \text{sid}, \vec{\texttt{st}}_P, \texttt{st}$) to simulated functionality $\mathcal{F}_{\text{StX}}$.
      -If successful, store $\vec{\texttt{st}}_P || \texttt{st}$ as the new $\vec{\texttt{st}}_P$
      -If successful, distribute the new state via $\mathcal{F}_{\text{StX}}$.
  else if time-tick τ corresponds to an update sub-round **then**
    Execute Step 4 of the mining protocol. This means that if the new information has not been fetched in this round already, then the following is executed:
      -Fetch transactions $(\texttt{tx}_1, \ldots, \texttt{tx}_u)$ (on behalf of P) from simulated $\mathcal{F}_{\text{N-MC}}$ and add them to Txs_P.
      -Fetch states $\vec{\texttt{st}}_1, \ldots, \vec{\texttt{st}}_s$ (on behalf of P) from the simulated $\mathcal{F}_{\text{StX}}$ and update $\vec{\texttt{st}}_P$ to the largest state among $\vec{\texttt{st}}_P$ and $\vec{\texttt{st}}_i$.
  end if
end procedure

procedure $\textsc{ExtendLedgerState}$
  Consider all honest and synchronized players P:
    - Let $\vec{\texttt{st}}$ be the longest state among all states $\vec{\texttt{st}}_p$ or states contained in a receiver buffer $\vec{M}_P$ with delay 1 (and hence is a potential output in the next round)
  Compare $\vec{\texttt{st}}^{\lceil T}$ with the current state $\texttt{state}$ of the ledger
  if $|\texttt{state}| > |\vec{\texttt{st}}^{\lceil T}|$ **then**
    Execute $\textsc{AdjustViiew}(\texttt{state})$
  end if
  if $\texttt{state}$ is not a prefix of $\vec{\texttt{st}}^{\lceil T}$ **then**
    Abort the simulation (due to inconsistency)
  end if
  Define the difference diff to be the block sequence s.t. $\texttt{state}||\mathsf{diff} = \vec{\texttt{st}}^{\lceil T}$.
  Let $n \leftarrow |\mathsf{diff}|$
  for each block diff_j, $j = 1$ to n **do**
    Map each transaction $\texttt{tx}$ in this block to its unique transaction ID txid
    If a transaction does not yet have an txid, then submit it to the ledger and receive the corresponding txid from $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$
    Let $\mathsf{list}_j = (\text{txid}_{j,1}, \ldots, \text{txid}_{j,\ell_j})$ be the corresponding list for this block.
    if coinbase $\text{txid}_{j,1}$ specifies a party that was honest at block creation time **then**

hFlag$_j \leftarrow 1$
else
hFlag$_j \leftarrow 0$
end if
Output (NEXT-BLOCK, hFlag$_j$, list_j) to $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$ (receiving (NEXT-BLOCK, ok) as an immediate answer)
end for
Execute ADJUSTVIEW($\mathtt{state}||\mathsf{diff}$)
end procedure

procedure ADJUSTVIEW($\mathtt{state}$)
$\mathsf{pointers} \leftarrow \varepsilon$
for each honest and synchronized party p_i **do**
Using the simulated functionality $\mathcal{F}_{\text{STX}}$ do the following:
- Let $\vec{\mathtt{st}}$ be the longest state among $\vec{\mathtt{st}}_{p_i}$ and those contained in the receiver buffer $\vec{M}_{p_i}$ with delay 1
Determine the pointer $\mathtt{pt}_i$ s.t. $\vec{\mathtt{st}}^{\lceil T} = \mathsf{state}|_{\mathtt{pt}_i}$
if such a pointer value does not exist **then**
Abort simulation (due to inconsistency)
end if
if Party p_i has not executed step 4 of the mining protocol in this current mini-round **then**
$\mathsf{pointers} \leftarrow \mathsf{pointers}||(p_i, \mathtt{pt}_i)$
end if // As otherwise, the new state is only fetched in the next round
end for
Output (SET-SLACK, $\mathsf{pointers}$) to $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$
$\mathsf{pointers} \leftarrow \varepsilon$
$\mathsf{desyncStates} \leftarrow \varepsilon$
for each honest but de-synchronized party p_i **do**
Using the simulated functionality $\mathcal{F}_{\text{STX}}$ do the following:
- Let $\vec{\mathtt{st}}$ be the longest state among $\vec{\mathtt{st}}_{p_i}$ and those contained in the receiver buffer $\vec{M}_{p_i}$ with delay 1
if Party p_i has not executed step 4 of the mining protocol in this current mini-round **then**
Set the pointer $\mathtt{pt}_i$ to be $|\vec{\mathtt{st}}^{\lceil T}|$
$\mathsf{pointers} \leftarrow \mathsf{pointers}||(p_i, \mathtt{pt}_i)$
$\mathsf{desyncStates} \leftarrow \mathsf{desyncState}||(p_i, \vec{\mathtt{st}}^{\lceil T})$
end if // As otherwise, the new state is only fetched in the next round
Output (SET-SLACK, $\mathsf{pointers}$) to $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$
Output (DESYNC-STATE, $\mathsf{desyncStates}$) to $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$
end for
end procedure

E Implementing a Stronger Ledger (Cont'd)

To achieve stronger guarantees than our original Bitcoin ledger, a party issues transactions relative to an account. More abstractly speaking, a transaction contains an identifier, AccountID, which can be seen as the abstract identity that claims ownership of the transaction. More specifically, we can represent this situation by having transactions $\mathtt{tx}$ be pairs $(\mathsf{AccountID}, \mathtt{tx}')$ with the above meaning. Signatures enter the picture at this level: an honest participant of the Bitcoin network will issue only signed transactions on the network. In order to link verification key to the account, AccountID is the hash of the verification keys, where we require collision resistance. More concretely, whenever a miner is supposed to submit a transaction $\mathtt{tx}$, it signs it and appends the signature and its verification key. This bundle is distributed into the Bitcoin network. The validation consists now of three parts. First, it is verified that the public key matches the account, second, the signature is verified, and third, its validated whether the actual transaction $(\mathsf{AccountID}, \mathtt{tx}')$ is valid, with respect to a separate validation predicate $\mathsf{ValidTx}_{\text{B}}$ on states and transactions $\mathtt{tx}$ of the above format. Only if all three tests succeed, the transactions is valid.

Looking ahead, the goal of this is the following: Assume that for the validation predicate $\mathsf{ValidTx}_{\text{B}}$ it holds that if a transaction $(\mathsf{AccountID}, \mathtt{tx})$ is valid relative to a state, then the only reason why it can get invalid is due to the presence of another transaction with the same account. If we think of wallets, if a miner can spend his coins at current time, then only another transaction by himself can

invalidate that (by spending the same coins, which the Bitcoin network will refuse). In combination with the unforgeability of signatures, no adversary can ever render a valid transaction invalid. Due to the guarantee of the liveness guarantee in Bitcoin, that guarantees that if a transaction too old, but valid relative to the state, then it will enter the state.

We now show how to implement this account management in the $\mathcal{G}^{\text{B}}_{\text{LEDGER}}$ hybrid world to achieve a stronger ledger that formalizes account management in an ideal manner. Our protocol makes use of an existentially unforgeable digital signatures scheme DSS $:= (K, S, V)$.

Definition 3. *A* digital signature scheme DSS $:= (K, S, V)$ *for a message space* $\mathcal{M}$ *and signature space* Ω *consists of a (probabilistic) key generation algorithm* K *that returns a key pair* (sk, vk)*, a (possibly probabilistic) signing algorithm* S*, that given a message* $m \in \mathcal{M}$ *and the signing key sk returns a signature* $s \leftarrow S_{sk}(m)$*, and a (possibly probabilistic, but usually deterministic) verification algorithm* V*, that given a message* $m \in \mathcal{M}$*, a candidate signature* $s' \in \Omega$*, and the verification key vk returns a bit* $V_{vk}(m, s')$*. The bit* 1 *is interpreted as a successful verification and* 0 *as a failed verification. It is required that* $V_{vk}(m, S_{sk}(m)) = 1$ *for all* m *and all* (vk, sk) *in the support of* K*. We call a* DSS *secure if it existentially unforgeable under chosen message attacks.*

Definition 4. *A digital signatures scheme is existentially unforgeable under chosen message attacks if no efficient adversary* A *can win the following game* $\mathbf{G}^{\mathsf{EU-CMA}}$ *better than with negligible probability.* $\mathbf{G}^{\mathsf{EU-CMA}}$ *first chooses a key pair* $(sk, vk) \leftarrow K$*. Then it acts as a signing oracle, receiving messages* $m \in \mathcal{M}$ *at its interface and responding with* $S_{sk}(m)$*. At any point,* A *can undertake a forging attempt by providing a message* m' *and a candidate signature* s' *to* $\mathbf{G}^{\mathsf{EU-CMA}}$*. The game is won if and only if* $V_{vk}(m', s') = 1$ *and* m' *was never queried before by* A*.*

E.1 The protocol for Account Management

Hybrid ledger functionality Let $\mathsf{ValidTx}_{\text{B}}$ and $\mathsf{blockify}_{\text{B}}$ be as in the previous section but with the following additional property: each transaction is a pair $\mathtt{tx} = (\mathsf{AccountID}, \mathtt{tx}')$ where the first part is bitstring of fixed length and the second part is an arbitrary transaction. In addition we require the following property: for any state $\mathtt{state}$ and any transaction $\mathtt{tx}$ it holds that $\mathsf{ValidTx}_{\text{B}}(\mathtt{tx}, \mathtt{state}) = 1$ implies, for any state extension $\mathtt{state}||\mathtt{st}'$, that $\mathsf{ValidTx}_{\text{B}}(\mathtt{tx}, \mathtt{state}||\mathtt{st}') = 1$, if $\mathtt{st}'$ does not contain a transaction with the same identifier $\mathsf{AccountID}$. Recall that we assume that Definition 2 is satisfied.

We assume the Bitcoin ledger functionality with the following validation predicate, which is defined relative to a collision-resistant hash function H, and a signature scheme DSS.

Algorithm to describe the assumed validation predicate

function $\mathsf{Val}_{\text{DSS}}(\mathtt{BTX}, \mathtt{state}, \mathtt{buffer})$
    Let $\mathtt{BTX} = (\mathtt{tx}, \text{txid}, \tau_L, p_i)$
    Parse $\mathtt{tx}$ as $((\mathsf{AccountID}, \mathtt{tx}'), vk, \sigma)$ (Return 0 in case of a wrong format)
    if $\mathsf{AccountID} = H(vk)$ **and** $V_{vk}(\mathtt{tx}, \sigma) = 1$ **then**
        return $\mathsf{ValidTx}_{\text{B}}(\mathtt{tx}, \mathtt{state})$
    else
        return 0
    end if
end function

The protocol. The protocol is straightforward: whenever the protocol is given an input of the form $(\mathsf{AccountID}, \mathtt{tx})$ it first checks that it is the party associated with this account ID. Then, it receives the newest state from the ledger and checks, whether this input is valid with respect to the current state. If this is the case, the party signs the input and submits it to the ledger.

Protocol $\mathsf{accountMgmt}(p)$

Initialization:

This protocol talks to the $\mathcal{G}_{\text{LEDGER}}$, but only changes the behavior of on read or submit-queries to the ledger. Any other command is simply relayed to $\mathcal{G}_{\text{LEDGER}}$ and the corresponding output is given to the environment $\mathcal{Z}$.
The protocol keeps a counter i and a vector $\mathsf{submitted}$ of inputs submitted to the ledger which are not yet contained in the state of the ledger.

Account Management:

- Upon receiving ($\textsc{CreateAccount}$, sid), execute $(sk, vk) \leftarrow K$, update $i \leftarrow i + 1$ and set $\mathsf{AccountID}_i \leftarrow H(vk)$. Return ($\textsc{CreateAccount}$, sid, $\mathsf{AccountID}_i$)

Ledger Read and Write:

- Upon receiving ($\textsc{read}$, sid) send ($\textsc{read}$, sid) to $\mathcal{G}_{\textsc{ledger}}$ and receive as answer the current $\mathtt{state} = \mathtt{st}_1 || \dots || \mathtt{st}_n$. Then do the following:

  $\mathtt{state}' \leftarrow \mathtt{st}_1$ // Genesis state
  for $i = 2$ to n **do**
  From state block $\mathtt{st}_i$, extract the contents $(\mathtt{tx}_1, vk_1, \sigma_1) || \dots || (\mathtt{tx}_n, vk_n, \sigma_n)$
  Define new block-content $\vec{x}' \leftarrow \mathtt{tx}_1 || \dots || \mathtt{tx}_n$
  $\mathtt{state}' \leftarrow \mathtt{state} || \mathsf{Blockify}(\vec{x})$
  end for

  Return ($\textsc{read}$, sid, $\mathtt{state}'$)

- Upon receiving ($\textsc{submit}$, sid, $\mathtt{tx}$), check that $\mathtt{tx} = (\mathsf{AccountID}, \mathtt{tx}')$ for $\mathsf{AccountID} \in \{\mathsf{AccountID}_1, \dots \mathsf{AccountID}_i\}$. If the check fails, ignore the input. Otherwise, do the following:
  1. Read the state $\mathtt{state}$ from $\mathcal{G}_{\textsc{ledger}}$ as above.
  2. If $\mathsf{ValidTx}_{\text{₿}}(\mathtt{tx}, \mathtt{state}) = 1$, then sign the input by $\sigma \leftarrow S_{sk}(\mathtt{tx})$ and send ($\textsc{submit}$, sid, $(\mathtt{tx}, vk, \sigma)$)

The enhanced ledger functionality. We present an enhanced ledger functionality with a validation predicate that enforces that an adversarial transaction cannot prevent a transaction by an honest party to eventually make it into the stable state of the ledger. In particular, we get the following enhanced functionality:

Functionality $\mathcal{G}_{\textsc{ledger}}^{\text{₿}+}$

$\mathcal{G}_{\textsc{ledger}}^{\text{₿}+}$ is identical to $\mathcal{G}_{\textsc{ledger}}^{\text{₿}}$ except with the following additional capabilities:

Difference to standard Ledger:

- Upon receiving ($\textsc{CreateAccount}$, sid) from party p_i (or the adversary on behalf of a party p_i), send ($\textsc{AccountReq}$, sid, p_i) to $\mathcal{A}$ and upon receiving a reply ($\textsc{AccountReq}$, sid, p_i, $\mathsf{AccountID}$) do the following:
  1. If $\mathsf{AccountID}$ is not yet associated to any party, store the pair ($\mathsf{AccountID}$, p_i) internally and return ($\textsc{CreateAccount}$, sid, $\mathsf{AccountID}$) to p_i.
  2. If $\mathsf{AccountID}$ is already associated to a party, then output ($\textsc{CreateAccount}$, sid, Fail) to p_i.

Standard Bitcoin Ledger:

- Identical to $\mathcal{G}_{\textsc{ledger}}$ with validation predicate $\mathsf{Val}_{\text{strong}}$ and thus omitted from this description.

The following validation predicate is used within $\mathcal{G}_{\textsc{ledger}}^{\text{₿}+}$.

Algorithm to define the strong validation

```
function Val_strong(BTX, state, buffer)
    Let BTX = (tx, txid, τ_L, p_i)
    if tx = (AccountID, tx') and AccountID is associated with p_i then
        return Validate(tx, state)
    else
        return 0
    end if
end function
```

The stronger guarantee for honestly submitted transactions stems from two facts. First, by Definition 2, the state blocks contain transactions beyond coin-base transactions. Second, since a transaction of a party is associated with its account, and cannot be invalidated by another transaction with a different account, this implies that the transaction remains valid relative to $\mathtt{state}$ (unless the honest party itself

issues a transaction that contradicts a previous transaction for one of its accounts, but we neglect this here). As an example, assume an honest party submits a single transaction for one of its accounts, and assume this transaction is valid relative to the state $\texttt{state}$. Then, by the defined enforcing mechanism of $\mathsf{ExtendPolicy}$, this transaction is guaranteed to enter the state after staying in the the buffer for long enough, and when an honest party mines a subsequent block after this delay. This means that after that delay has passed, the transaction has to appear within the subsequent window of $\texttt{windowSize}$ blocks.

A brief worst-case calculation. Looking at the ledger abstraction, we can directly compute the following worst-case upper bound for any miner (we neglect here the offset at the beginning of the execution for simplicity): after submitting the transaction, the transaction will appear (relative to the view of the submitting party) within the next $4 \cdot \texttt{windowSize}$ blocks after submitting the transaction (except with negligible probability). The reason is that upon submitting, (1) the view of the miner submitting the transaction could be $\texttt{windowSize}$ blocks behind the head of the state of the ledger and (2) by the definition of $\mathsf{ExtendPolicy}$ (and observing that $\frac{\texttt{Delay}}{2} < \texttt{minTime}_{\texttt{window}}$), at most $2 \cdot \texttt{windowSize}$ blocks can be added to the state while the transaction is staying in the buffer before the ledger starts enforcing that the transaction be part of the subsequent next honest state block. Recall that the ledger guarantees that such an honest block has to appear within a window of $\texttt{windowSize}$ state blocks.

Recall that in the real Bitcoin, the expected time until a transaction appears in the confirmed part of the state, i.e., in any block which is $T = 6$ blocks behind the head of the state, is approximately one hour. By composition, the correspondence of the ledger parameter $\texttt{windowSize}$ with the protocol parameter T (as proven in the previous section), our analysis suggests that four hours is a worst-case bound for Bitcoin given that transactions are correctly signed and are not invalidated due to other transactions with the same account.

We conclude this section by stating the following lemma:

Lemma 4. *Let* DSS *be a secure digital signature scheme and let H be a collision resistant hash function. Then the protocol* $\mathsf{accountMgmt}$ *in the $\mathcal{G}^{B}_{\text{LEDGER}}$-hybrid world UC-realizes ledger $\mathcal{G}^{B+}_{\text{LEDGER}}$, where the functionalities are instantiated as described above.*