

# Etherless Ethereum Tokens: Simulating Native Tokens in Ethereum

John Andrews<sup>1</sup>, Michele Ciampi<sup>2</sup>, and Vassilis Zikas<sup>\*3</sup>

<sup>1</sup>Sunday Group, [jandrews@sundaygroupinc.com](mailto:jandrews@sundaygroupinc.com)

<sup>2</sup>The University of Edinburgh, [michele.ciampi@ed.ac.uk](mailto:michele.ciampi@ed.ac.uk)

<sup>3</sup>Purdue University, [vzikas@cs.purdue.edu](mailto:vzikas@cs.purdue.edu)

## Abstract

Standardized Ethereum tokens, e.g., ERC-20 tokens, have become the norm in fundraising (through ICOs) and kicking off blockchain-based DeFi applications. However, they require the user’s wallet to hold both tokens and ether to pay the gas fee for making a transaction. This makes for a cumbersome and counterintuitive—at least for less tech-savvy users—user experience, especially when the token creator intends to switch to their own blockchain down the line, or wishes the flexibility of transferring the token to a different smart-contract enabled blockchain. We formalize, instantiate, and analyze in a composable manner a system that we call *Etherless Ethereum Tokens* (in short, EETs), which allows the token creator to allow its users to transact in a closed-economy manner, i.e., having only tokens on their wallet and paying any transaction fees in token units rather than gas. In the process, we devise a methodology for capturing Ethereum token-contracts in the Universal Composability (UC) framework, which can be of independent interest. We have implemented and benchmarked our system and compared it to another solution for obtaining similar functionality in Ethereum, i.e., the Gas Station Networks (GSN); in addition to being the first system with a rigorous security analysis, we demonstrate that EETs are not only far easier to deploy, but are also far less gas intensive than the GSN.

## 1 Introduction

As applications of smart contracts, e.g., Decentralized Finance (DeFi) and Non-Fungible Tokens (NFTs), become mainstream, there is a need to make them as independent from the Ethereum chain as possible. The paradigm in which such a need is most prominent is in the creation of Ethereum tokens (e.g., ERC-20 tokens [VB15]). The usual implementation of such tokens requires that, for a token-holder to use them (exchange them with other tokens, or transact with other holders of the same token,) they need to also hold Ether, to be used to fuel the Ethereum transaction. This poses a challenge both to the token creator and the users: on the one hand, token creators need to provide a wallet which supports both their token and Ethereum, making it more challenging to transition to their own blockchain or switch token platforms while offering a smooth user experience. On the other hand, users need to make sure that they hold not only the token but also Ether, which

---

\*Work supported in part by NSF grant no. 2055599 and by Sunday Group.

makes it more challenging to expand this technology to less tech-savvy audiences, thereby hindering wider societal adoption.

The easiest way to conceptualize the relevant bottleneck is through considering the life cycle of an ETH-based initial coin offering (ICO): in a first stage, the token creator solicits investment (typically in different cryptocurrencies), under the promise of a certain (prearranged) amount of tokens once the token launches.<sup>1</sup> In a second phase, the token creator initializes the promised new token by launching a token smart contract (e.g. an ERC 20 token) on the Ethereum chain. In order to hold their promise to the investors, the token creator would then have the investors create a new token-specific Ethereum address where the promised tokens can be transferred. This can be done by means of a wallet that offers generic support for Ethereum tokens.

Often, however, ICO-funded applications launch tokens which have the ultimate goal of eventually being disconnected from the main Ethereum blockchain, and/or which aim to create an ecosystem independent of Ethereum. In such cases, the token creator would typically also offer its users a token-specific wallet application. However, in order for anyone to use this application to transfer his tokens, the token-specific wallet needs to also support Ether as a currency. This leads to confusion for less tech-savvy investors, and makes the user experience of migrating the token to a different smart contract platform — e.g. a different smart-contract-enabled blockchain or a blockchain developed by the token creator — less intuitive. We note that such migration is becoming more relevant as more smart-contract-enabled blockchains are released, and as the gas price for Ethereum smart contracts rises to a point where its use makes the corresponding tokens less attractive.

In this work, we propose a design methodology and formal treatment of Ethereum tokens which allow their creator to provide the option to its users of making transfers without the need to hold Ether in their wallet, a mechanism which we term *Etherless Ethereum Tokens (in short, EETs)*. The high-level idea is simple: allow the token creator to take on the cost (i.e., gas) for the token transaction, and have the token contract perform an on-the-fly exchange of token-to-ether at a pre-agreed rate, giving the user the experience of a native token. As one might expect, properly specifying, implementing, and proving such a protocol secure is a challenging task; in particular, it requires a model for token-enabled ledgers, which we provide and believe can be a result of independent interest. We remark that, as a concept, etherless transactions have been frequently discussed within the Ethereum community for several years, often under the term *meta transactions* [Gri18, AB20, gsn]. However, to our knowledge, our work is the first formal treatment and security analysis of the concept.

At a less technical level, we believe that in addition to offering a more intuitive, closed-economy user experience, such a mechanism also provides assurance to the original ICO investors that the token creator indeed expects value on the token, as he is willing to make marginal exchanges. Nonetheless, the study of this market effect is outside the scope of our current work. We note in passing that despite being explicitly implemented on the Ethereum blockchain, our design is generic and can be ported to any smart-contract-enabled blockchain platform, and thus can enable transferring the tokens from one blockchain to another.

We have implemented our EET design, and we demonstrate how it outperforms existing generic systems that enable etherless transactions, such as the Gas Station Network (GSN) [gsn], both in terms of simplicity of deployment and in terms of gas usage. We also compare such a deployment

---

<sup>1</sup>There are a number of legal issues regarding ICO's — in particular, how to hold the token creator to his promise and how to avoid scamming attacks — and there are technological advances that allow us to circumvent them; these topics are outside the scope of this paper.

with how a native token could perform on Ethereum and demonstrate that the overhead makes the flexibility which is offered by smart-contract-based tokens a reasonable compromise for the moderate increase in the required gas it incurs over what a native token would require.

## 2 Our Contributions and Related Work

Our contribution is threefold:

- A. A universally composable (UC) [Can01] treatment of ledgers supporting a broad class of smart contracts, which includes token contracts (e.g. ERC 20).
- B. A design and UC security analysis of EETs.
- C. An implementation of our EET, benchmarks, and comparison with alternative approaches.

In the following, we expand on the key components of the above contributions, and put our results in perspective with existing literature and systems.

### 2.1 Smart-Contract-Enabled Transaction Ledgers

The first analyses of blockchain protocols showed that they satisfy certain desirable properties, such as common-prefix (also referred to as safety or consistency), chain-growth (also referred to as liveness), chain quality, etc [BMTZ17, BGK<sup>+</sup>18, GKL15, PSs17, GKL17, DPS19, PS17]. Badertscher *et al.* [BMTZ17] put forth the first universally composable treatment of the bitcoin backbone (i.e. consensus layer) by introducing a UC functionality, called  $\mathcal{F}_{\text{LEDGER}}$ , which captures the interface that Bitcoin offers to external applications, rather than the way in which this interface is implemented. At a very high level,  $\mathcal{F}_{\text{LEDGER}}$  takes as input transactions which are validated by means of a validation predicate `Validate`. All valid transactions are then stored into a data structure denoted as `state`. The adversary has full control over the order in which transactions appear in `state`, and can define (in a limited way) the portion of the state that each party can access. However, once something is added to the state, it cannot be removed (not even by the adversary).

We note that the advantage of proving security in UC is that it enables use of the ledger as an ideal primitive, and ensures that replacing this ideal ledger primitive by its implementation—the corresponding blockchain—does not compromise the security of primitives that make ideal calls to the ledger; nor does it affect the security of systems and protocols that run alongside the ledger. This property is often referred to as universal composability, and it allows for a constructive approach to cryptographic/security protocols, analogous to how programming uses libraries with fixed APIs without worrying about their implementation. Following that work, a number of papers on the design and analysis of blockchains have adopted UC as the model to prove their security and have devised systems implementing variants of the above ledger [BGK<sup>+</sup>18, KKKZ19].

UC [BMTZ17] has also been leveraged to describe how  $\mathcal{F}_{\text{LEDGER}}$  may be used together with a digital signature scheme to derive a *transaction ledger*, abstracting the cryptocurrency aspects of Bitcoin in addition to its backbone guarantrees.<sup>2</sup> This was done by relying on digital signatures where, to ensure composability, the ideal adversary is allowed to choose the signing and verification keys (cf. [KZZ16]).

---

<sup>2</sup>Unlike transaction ledgers, the bare  $\mathcal{F}_{\text{LEDGER}}$  captures the consensus layer, and does not interpret its contents as transactions which need to be verified with respect to whether or not they are spending some already spent coin.

### 2.1.1 The Transaction Ledger

In this paper we consider a simpler, more UC-friendly approach that abstracts away the public-key infrastructure (PKI), analogous to how the UC signatures functionality [Can03] would. In a nutshell, instead of having `Validate` rely on a specific signature scheme, we define a new transaction ledger  $\mathcal{F}_{\text{T-LEDGER}}$  that internally runs  $\mathcal{F}_{\text{LEDGER}}$  and also emulates existentially unforgeable signatures, similar to [Can03].  $\mathcal{F}_{\text{T-LEDGER}}$  accepts transactions with the format  $\text{tx} := (v, \text{addr}_i, \text{addr}_j, \text{fee})$  where  $v$  represents the number of coins involved in the transaction,  $\text{fee}$  is the fee that the issuer of the transaction is willing to pay, and  $\text{addr}_i$  and  $\text{addr}_j$  represent the wallet addresses of the sender the receiver respectively. Upon receiving a transaction,  $\mathcal{F}_{\text{T-LEDGER}}$  checks the state of  $\mathcal{F}_{\text{LEDGER}}$  to ensure that the wallet address  $\text{addr}_i$  has at least  $v + \text{fee}$  coins and that the  $\text{fee}$  is sufficient, i.e. that  $\text{fee} \geq f(\text{tx})$ , where  $f$  is function specified in the description of  $\mathcal{F}_{\text{T-LEDGER}}$  that determines the fee that needs to be payed for the input transaction.

We note that it is straightforward to adapt the analysis of the transaction ledger [BMTZ17]—using a specific existentially-unforgeable signatures scheme—to prove security of our ledger for a standard Bitcoin-style blockchain protocol, such as bitcoin or the proof-of-work-based version of Ethereum. Nonetheless, as we shall see, this makes it more intuitive to add cryptocurrency-relevant features to the ledger—such as etherless tokens.

### 2.1.2 Adding Smart Contracts

The functionality  $\mathcal{F}_{\text{T-LEDGER}}$  is sufficient to capture the base functionality of cryptocurrencies, but it does not support smart contracts. To achieve that, we define an augmented functionality, which we denote  $\mathcal{F}_{\text{TSC-LEDGER}}$ .  $\mathcal{F}_{\text{TSC-LEDGER}}$  internally manages  $\mathcal{F}_{\text{T-LEDGER}}$  and a functionality  $\mathcal{F}_{\text{SC}}$  that abstracts a smart contract:  $\mathcal{F}_{\text{SC}}$  maintains its own state  $\text{cstate}$ —corresponding to the state of a (virtual) machine VM<sup>3</sup>—and is parametrized by a function  $f_{\text{CFee}}$ , that takes as input the query to the contract (which contains also the fee that the caller is willing to pay to run the contract), and checks whether or not the fee is enough for the VM to process the input and update its state.

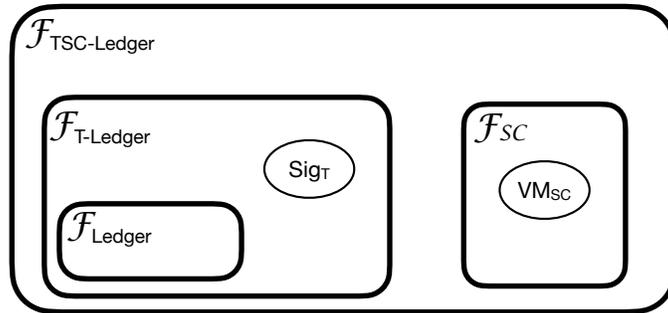


Figure 1: The Smart-Contract-Enabled Transaction Ledger Functionality  $\mathcal{F}_{\text{TSC-LEDGER}}$

The construction of  $\mathcal{F}_{\text{TSC-LEDGER}}$  from its components is illustrated in Figure 1.  $\mathcal{F}_{\text{TSC-LEDGER}}$  accepts either standard transactions in the native currency  $\mathbb{E}$  (that are forwarded to  $\mathcal{F}_{\text{T-LEDGER}}$ ) or

<sup>3</sup>We do not specify a model of computation for describing the VM; one can use any such model, e.g. Turing machines, RAMs, etc.

inputs/transactions that are intended as queries to the contract  $\mathcal{F}_{\text{SC}}$ . Upon receiving such a query for the smart contract,  $\mathcal{F}_{\text{TSC-LEDGER}}$  forwards the query to  $\mathcal{F}_{\text{SC}}$ , which checks if the fee specified in the query is sufficient to update its state, and if so it updates `cstate` by running the VM on input the given transaction and the state of  $\mathcal{F}_{\text{T-LEDGER}}$  (which is handed to  $\mathcal{F}_{\text{SC}}$  by  $\mathcal{F}_{\text{TSC-LEDGER}}$ )<sup>4</sup>, and returns the updated state (including the received input) to  $\mathcal{F}_{\text{TSC-LEDGER}}$ .  $\mathcal{F}_{\text{TSC-LEDGER}}$  then pushes the query and the updated state `cstate` to the state of  $\mathcal{F}_{\text{T-LEDGER}}$  (by submitting it as a transaction).

Consistently with the Ethereum smart contract mechanism,  $\mathcal{F}_{\text{SC}}$  charges the contract caller only for the fee that is required to update its state, even if the contract’s caller specified a higher fee. Moreover, if a contract caller did not specify a fee high enough to conclude an update on the contract’s state, the fee will be deducted from the caller account, and the input used to query the contract will appear in the state of  $\mathcal{F}_{\text{T-LEDGER}}$ , though no change to the contract’s state will be committed.

### 2.1.3 Tokens as Smart Contracts

Given the above smart-contract-enabled ledger, it is straightforward to capture a smart contract for creating a standard (e.g. ERC 20 [VB15]) Ethereum token by instantiating  $\mathcal{F}_{\text{TSC-LEDGER}}$  with contract functionality that stores and updates the state (balances for different addresses) of such a token. Note that this results in a token-enabled transaction ledger  $\mathcal{F}_{\text{LEDGER}}^{\text{Token}}$  which allows parties both to issue transactions in the native coin **E**, and to exchange tokens **T**.

In more detail,  $\mathcal{F}_{\text{LEDGER}}^{\text{Token}}$  instantiates  $\mathcal{F}_{\text{TSC-LEDGER}}$  with a token-contract  $\mathcal{F}_{\text{SC}}^{\text{T}}$  which works as follows:  $\mathcal{F}_{\text{SC}}^{\text{T}}$  collects all token transactions, and upon receiving a read-request returns only the *valid* token transactions. Similarly to the way the ledger  $\mathcal{F}_{\text{T-LEDGER}}$  deals with native transactions, a token transaction consists of the components  $(v, \text{addr}_i^{\text{T}}, \text{addr}_j^{\text{T}})$ , where  $v$  is the number of tokens involved in the transaction, and  $\text{addr}_i$  and  $\text{addr}_j$  represent the token wallet addresses of the sender and the receiver respectively. Furthermore,  $\mathcal{F}_{\text{SC}}^{\text{T}}$  internally emulates an existentially-unforgeable signature scheme related to the token which is independent of the one that is used in  $\mathcal{F}_{\text{T-LEDGER}}$ .<sup>5</sup>

We observe that there is no fee appearing in the description of the token transaction. The reason is that the fee will be part of the query to the contract, and it is expressed in the native currency **E**. Indeed, the issuer of the token transaction, in order to query the contract  $\mathcal{F}_{\text{SC}}^{\text{T}}$ , needs to possess coins of type **E**.

### 2.1.4 The EET Functionality

As discussed in the introduction, the above contract implementation of tokens—which has become a standard for Ethereum—has the undesirable property that a party who wants to send tokens requires coins of type **E** to do so, coins which they might not have. In this work, we introduce EETs to allow the token creator to offer, as a service, to take on the cost of the token transaction, in exchange for tokens at a pre-agreed **E**-to-**T** rate. This is captured by tweaking the token-enabled ledger  $\mathcal{F}_{\text{LEDGER}}^{\text{Token}}$  toward an EET-enabled ledger, denoted as  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , which supports an additional input called `SUBMIT-DELEGATION`. Upon receiving `SUBMIT-DELEGATION`,  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  allows the user to issue a token transaction which pays a fee, in **T**, to a special party, called *intermediary* (that we

<sup>4</sup>Note that  $\mathcal{F}_{\text{TSC-LEDGER}}$  also keeps track of the history of the state of  $\mathcal{F}_{\text{T-LEDGER}}$ .

<sup>5</sup>Note that we cannot generically use the same signature emulator procedure of  $\mathcal{F}_{\text{T-LEDGER}}$ , as a token address is typically overloaded to also be an Ethereum address.

denote with  $M$ ), in exchange for the intermediary submitting the token transaction to  $\mathcal{F}_{SC}^T$  and paying the  $E$  needed for the token contract to process the transaction.

## 2.2 EET Construction and Analysis

To realize  $\mathcal{F}_{LEDGER}^{EET}$  we rely only on  $\mathcal{F}_{T-LEDGER}$  and signatures. In particular, any party that wants to issue a token transaction and has enough coins of type  $E$  to cover for the fee can issue a transaction  $\mathbf{tx} = (0, \mathbf{addr}_i, 0^\lambda, (\mathbf{aux}, \sigma), \mathbf{fee})$ , where  $\mathbf{aux} = (v, \mathbf{addr}_i^T, \mathbf{addr}_j^T)$  and  $\sigma$  is a signature of  $\mathbf{aux}$  that verifies under  $\mathbf{addr}_i^T$ .<sup>6</sup>

In a nutshell,  $\mathbf{tx}$  is a standard transaction for  $\mathcal{F}_{T-LEDGER}$  that contains in its payload the information related to the token transaction properly signed by the sender. By definition, if the fee  $\mathbf{fee}$  is high enough, then  $\mathbf{tx}$  will become part of  $\mathcal{F}_{T-LEDGER}$ 's state. Let  $\mathbf{addr}_M^T$  be the token wallet address of  $M$ . To delegate a transaction, the sender  $P_i$  creates a special token transaction  $\mathbf{aux} = ([v, \mathbf{del-fee}], \mathbf{addr}_i^T, [\mathbf{addr}_j^T, \mathbf{addr}_M^T])$  (where  $\mathbf{del-fee}$  is a fee expressed in  $T$  that parametrizes  $\mathcal{F}_{LEDGER}^{EET}$ ) and signs it to obtain  $\sigma$ .  $\mathbf{aux}$  is the atomic representation of two token transactions: the first moves  $v$  tokens from  $\mathbf{addr}_i^T$  to  $\mathbf{addr}_j^T$ , and the second moves  $\mathbf{del-fee}$  from  $\mathbf{addr}_i^T$  to  $\mathbf{addr}_M^T$ .  $M$ , upon receiving  $(\mathbf{aux}, \sigma)$  submits a transaction to  $\mathcal{F}_{T-LEDGER}$  that contains  $(\mathbf{aux}, \sigma)$  in its payload. If a party wants to obtain only the valid token transaction, they need to filter out the payload of the transactions stored in  $\mathcal{F}_{T-LEDGER}$ 's state, and output only the valid transactions. Similarly to what we have described above, a token transaction  $(v, \mathbf{addr}_i^T, \mathbf{addr}_j^T)$  is valid if the sum of tokens with receiver address  $\mathbf{addr}_j^T$  minus the sum of tokens in the state with sender address  $\mathbf{addr}_i$  (including the fees) is greater than or equal to  $v$ .

## 2.3 Implementation, Benchmarks, and Comparisons

The Gas Station Network (GSN) is a relatively recent development in the Ethereum community that shares some of our goals, but a broader scope. In particular, the GSN aims to create a decentralized, trustless network of *relay servers* which can pick up the transaction fees for any GSN-enabled contract.

The GSN is built around a RelayHub smart contract that:

1. Records available relay servers and their service fees,
2. Keeps ether deposits from GSN-enabled contracts for repayment of relay servers,
3. Facilitates the interaction between relays and GSN-enabled contracts, and punishes any detected bad actors.

This is in contrast to our mechanism, in which there is no separate smart contract to manage the delegation of transactions. Additionally, each GSN-enabled contract must interact with a separate *paymaster* contract, which is responsible for performing any action needed to extract or verify payment from users. Paymaster contracts may be written generically and shared between multiple contracts, or purpose-written for particular contracts.

The outward functionality of the GSN is similar to our mechanism: a gasless user submits a transaction to an intermediary relay server instead of directly to the blockchain, and the relay submits the transaction on the user's behalf, receiving an ether repayment from the target contract.

---

<sup>6</sup>In the protocol, the addresses become verification keys for a signature scheme.

The target contract, in turn, is allowed to extract any payment it wishes from the user, e.g. tokens. The primary difference is in the complexity of implementation and development; where the GSN aims to be fully generic and decentralized, and admits a great deal of complexity in service of that aim, we have endeavored to keep our efforts very self-contained in order to ease implementation, simplify formal analysis, and keep operational costs manageable.

As is common in designs that aim for maximally generic functionality, the GSN pays for its genericity with increased complexity. This complexity manifests both in development effort — anecdotally, we found setting up a testing environment for a GSN-enabled contract to be significantly more cumbersome than for other contracts — and in gas consumption. Our experiments indicate a 4-5x overhead in gas consumption when using the GSN as opposed to using our EET contract. (Note that gas is pretty much the only relevant measurable unit of comparison. Other metrics — e.g. running time, settlement time, etc. — are either very difficult to test in a controlled way, are irrelevant for a contract which aims only to facilitate token exchange, or are negligible compared to other confounding factors.) We note in passing that, to our knowledge, there is no formal security analysis of the GSN, making our work the first rigorous treatment of the etherless token paradigm.

*Remark (Contract-based vs Native Tokens).* Recently, the blockchain/cryptocurrency community has been entertaining the idea of making tokens native to the cryptocurrency chain. For example, [Kia] outlines a plan to introduce such a mechanism, which would allow a user to post a token transaction along with a token-to-native exchange rate he is willing to pay; any miner/minter who agrees with the rate can pick this up and create a transaction in which they “foot the bill” in terms of native-currency fees, in exchange for tokens at the proposed exchange rate. This clearly yields an advantage in terms of fees needed for the transaction, but it does come at a cost: (1) The token functionality is limited to what is hardwired on the token chain, and is therefore far less flexible than a smart-contract-based solution. For example, it is unclear if or how such a solution would allow the use of amortization/batching to save on bulk transactions. (2) If one adopts the natural “pay-per-use” principle for fees — i.e. you pay more for a more complex transaction — as Ethereum does, then adding this functionality would increase the cost of all transactions, including those that only involve the native cryptocurrency. Although this increase is expected to be minimal, it is unclear how the implicit auction for the submitted token transaction created by such a mechanism would affect fees. In Appendix A, we have included an attempt to estimate the overhead this might incur in a hypothetical implementation on Ethereum, and compare it with using a smart contract. We note that in the absence of a (platform or blockchain supporting an) actual implementation of native tokens, the relevant experiments are somewhat artificial and speculative. Thus, we do not consider these experiments an important part of our contributions (and we defer them to the appendix). Nonetheless, we do believe they give an interesting perspective to the discussion on native tokens, and a pointer for experiments once such a functionality is implemented on a mainstream blockchain.

### 3 Preliminaries

We use “=” to denote equality of two different elements (i.e.  $a = b$  then...) and “ $\leftarrow$ ” as the assignment operator (e.g. to assign to  $a$  the value of  $b$  we write  $a \leftarrow b$ ). A randomized assignment is denoted with  $a \stackrel{\$}{\leftarrow} A$ , where  $A$  is a randomized algorithm and the randomness used by  $A$  is not explicit. We call a function  $\nu : \mathbb{N} \rightarrow \mathbb{R}^+$  *negligible* if for every positive polynomial  $p(\kappa)$ , there exists a  $\kappa_0 \in \mathbb{N}$  such that for all  $\kappa > \kappa_0 : \nu(\kappa) < 1/p(\kappa)$ .

### 3.1 Signatures

**Definition 3.1** (Signature scheme [Can03]). A triple of PPT algorithms  $(\text{Kgen}, \text{Sign}, \text{Ver})$  is called a *signature scheme* if it satisfies the following properties.

**Completeness:** For every pair  $(s, v) \xleftarrow{\$} \text{Kgen}(1^\lambda)$ , and every  $m \in \{0, 1\}^\lambda$ , we have that  $\Pr[\text{Ver}(v, m, \text{Sign}(s, m)) = 0] < \nu(\lambda)$ .

**Consistency (non-repudiation):** For any  $m$ , the probability that  $\text{Kgen}(1^\lambda)$  generates  $(s, v)$  and  $\text{Ver}(v, m, \sigma)$  generates two different outputs in two independent invocations is smaller than  $\nu(\lambda)$ .

**Unforgeability:** For every PPT  $\mathcal{A}$ , there exists a negligible function  $\nu$ , such that for all auxiliary input  $z \in \{0, 1\}^*$  it holds that:

$$\Pr[(s, v) \xleftarrow{\$} \text{Kgen}(1^\lambda); (m, \sigma) \xleftarrow{\$} \mathcal{A}^{\text{Sign}(s, \cdot)}(z, v) \wedge \text{Ver}(v, m, \sigma) = 1 \wedge m \notin Q] < \nu(\lambda)$$

where  $Q$  denotes the set of messages whose signatures were requested by  $\mathcal{A}$  from the oracle  $\text{Sign}(s, \cdot)$ .

## 4 The Model

Following the recent line of works proving composable security of blockchain ledgers [BMTZ17, BGK<sup>+</sup>18], we provide our protocols and security proofs in Canetti’s universal composition (UC) framework [Can01]. In this section we discuss the main components of our real-world model (including the associated hybrids). We assume that the reader is familiar with simulation-based security and has basic knowledge of the UC framework. We review all the aspects of the execution model that are needed for our protocols and proof, but omit some of the low-level details and refer the interested reader to relevant works wherever appropriate.

We now 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, sending an output to the environment, or not sending any message at all. In any of these cases, the party loses the activation.<sup>7</sup> We denote the identities of parties by  $P_i$ , i.e.  $P_i = (\text{pid}_i, \text{sid}_i)$ , and call  $P_i$  a party for short. The index  $i$  is used to distinguish two identifiers, i.e.,  $P_i \neq P_j$ , and otherwise carries no meaning. 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.

Our protocols are synchronous (G)UC protocols [BMTZ17, KMTZ13]: parties have access to a (global) clock setup, denoted by  $\mathcal{F}_{\text{clock}}$ , and can communicate over a network of authenticated multicast channels.

<sup>7</sup>In the latter case the activation goes to the environment by default.

We adopt the *dynamic availability* model implicit in [BMTZ17] which was fleshed out in [BGK<sup>+</sup>18]. We next sketch its main components: All functionalities, protocols, and setups have a dynamic party set. I.e., they all include special instructions allowing parties to register and deregister, and allow the adversary to learn the current set of registered parties. Additionally, global setups allow any other setup (or functionality) to register and deregister with them, and also allow other setups to learn their set of registered parties (we refer to [Appendix B](#) for the formal treatment).

We conclude this section by elaborating on the main hybrid functionality used in our paper. For self containment we have included formal descriptions of the ideal functionalities we consider in [Appendices C and D](#).

#### 4.1 The functionality $\mathcal{F}_{\text{ledger}}$ .

The main functionality (in fact, a global setup) we rely on is a cryptographic distributed transaction ledger. We use the (backbone) ledgers proposed in the recent literature [BMTZ17, BGK<sup>+</sup>18] in order to describe a transaction ledger and its properties. As proved in [BMTZ17, BGK<sup>+</sup>18], such a ledger is implemented by known permissionless blockchains based on either proof-of-work (PoW), e.g. Bitcoin, or poof-of-stake (PoS), e.g. Ouroboros Genesis. The ledger stores an immutable sequence of blocks called *state*—each block containing several messages typically referred to as *transactions* and denoted by  $\text{tx}$ —which is accessible from the parties under some restrictions discussed below. It enforces the following basic properties that are inspired by [GKL15, PSs17]:

- *Ledger growth.* The size of the ledger’s state should grow—new blocks should be added—as the rounds advance.
- *Chain quality.* It is guaranteed that a percentage of honest blocks are created in a sufficiently long sequence of blocks.
- *Transaction liveness.* Old enough (valid) transactions are included in the next block added to the ledger state.

We next give a brief overview of the ledger functionality  $\mathcal{F}_{\text{LEDGER}}$  proposed in [BMTZ17, BGK<sup>+</sup>18], focusing on the properties of  $\mathcal{F}_{\text{LEDGER}}$  that are relevant for the understanding our results. Along the way we also introduce some useful notation and terminology. We refer the reader interested in the low-level details of the ledger functionality and its UC implementation to [Figure 11](#) in [Appendix D](#) and [BMTZ17, BGK<sup>+</sup>18]. We note that with minor differences related to the nature of the resource used to implement the ledger, PoW vs PoS, the ledgers proposed in these works are identical.

The functionality  $\mathcal{F}_{\text{LEDGER}}$  is parametrized by three main functions `Validate`, `ExtendPolicy` and `Blockify`. At a high level, anyone (honest miner or the adversary) may submit a transaction to  $\mathcal{F}_{\text{LEDGER}}$ . The transaction is validated by means of a filtering predicate `Validate`, and if it is found to be valid it is added to a *buffer* that we denote `buffer`. Taking a peak at the actual implementation of the ledger, this buffer contains transactions that, although validated, are either not yet inserted into a valid block, or are in a block which is not yet deep enough in the blockchain to be considered immutable for an adversary. The adversary  $\mathcal{A}$  is informed that the transaction was received and is given its contents. Periodically,  $\mathcal{F}_{\text{LEDGER}}$  does the following: 1) fetches some of the transactions in the buffer under the influence of the adversary (more on this will follow), 2) modifies them by means of a procedure `Blockify`, 3) creates a block including the output of `Blockify`, and 4) adds this block to its permanent state, denoted as `state`. `state` is a data structure that includes the sequences

of blocks that the adversary can no longer change. (In [GKL15, PSs17] this corresponds to the *common prefix*.) Any miner or the adversary is allowed to request a read of the contents of the state and, every honest miner will eventually receive `state` as its output.<sup>8</sup>

To enforce transaction liveness and chain-quality,  $\mathcal{F}_{\text{LEDGER}}$  relies on the function `ExtendPolicy`. At a high level, `ExtendPolicy` makes sure that the adversary cannot create too many blocks with arbitrary (but valid) contents (chain quality) and that if a transaction is old enough, and still valid with respect to the actual state, then it is included into the state. In more detail, `ExtendPolicy` takes the current contents of the buffer, along with the adversary’s recommendation `NxtBC`, and the block-insertion times vector  $\tau_{\text{state}}$ . The latter is a vector listing the times when each block was inserted into the state. The output of `ExtendPolicy` is a vector including the blocks to be appended to the state during the next state-extend time-slot. Each of these blocks is then given as input to `Blockify`.

We conclude the discussion by providing a high-level description of the main input command of  $\mathcal{F}_{\text{LEDGER}}$  used in our protocols/definitions, and refer to [Appendix D](#) for a detailed description of the functionality.

- The input `(READ, sid)` is used to request the content of the ledger’s `state`. Concretely, upon receiving `(READ, sid)` from some party (or the adversary on behalf of a corrupted party), the ledger returns (a prefix of) `state` to the caller.
- The input `(SUBMIT, sid, tx)` is used to request that a transaction `tx` be added to the buffer. That is, upon receiving a `(SUBMIT, sid, tx)` message from any party (or the adversary), the ledger adds the transaction `tx` to the buffer `buffer`. If the validation predicate `Validate`, on input `state, buffer, tx` outputs 1, then `tx` will be included in `state`.<sup>9</sup> The time required for the transaction to be part of `state` and visible to all honest parties who query  $\mathcal{F}_{\text{LEDGER}}$  depends on the transaction liveness parameter defined in `ExtendPolicy`.

## 5 Define and Instantiate a New Cryptocurrency from $\mathcal{F}_{\text{ledger}}$

The ledger  $\mathcal{F}_{\text{LEDGER}}$  does not itself realize a cryptocurrency. We use `E` to denote the symbol of the coins that will be maintained by our ledger  $\mathcal{F}_{\text{T-LEDGER}}$ , and show how to realize  $\mathcal{F}_{\text{T-LEDGER}}$  from  $\mathcal{F}_{\text{LEDGER}}$

The validation predicate of  $\mathcal{F}_{\text{LEDGER}}$ , in this case, is defined to always output 1, and it is  $\mathcal{F}_{\text{T-LEDGER}}$ ’s responsibility to make sure that only valid transactions are submitted to  $\mathcal{F}_{\text{LEDGER}}$ .  $\mathcal{F}_{\text{T-LEDGER}}$  also generates and manages the *wallets* of the parties. A transaction supported by  $\mathcal{F}_{\text{T-LEDGER}}$  consists of five main components  $(v, \text{addr}_i, \text{addr}_j, \text{aux}, \text{fee})$ , where  $v$  represents the amount of coins of type `E`,  $\text{addr}_i$  is the sender’s wallet address,  $\text{addr}_j$  is the receiver’s wallet address,  $\text{aux}$  is a payload, and  $\text{fee}$  represents the fee. At a high level, a transaction is valid if the fee  $\text{fee}$  is high enough and if the amount of coins stored in the wallet with address  $\text{addr}_i$  is at least  $v + \text{fee}$ . How high the fee should be in order for the transaction to be considered is specified by a function  $f$

<sup>8</sup>As observed in [BMTZ17], it is not possible to guarantee with existing constructions that at any given point in time all honest parties see exactly the same state (blockchain) length, so each party may have a different view of the state which is defined by the adversary. However, the adversary can restrict the view of the honest parties only by a bounded number of blocks. The parameter that defines such a bound is called `windowSize`.

<sup>9</sup>We have the guarantee that any transaction (either generated by a malicious or honest party) that manages to go in `buffer` will eventually be included in `state`.

that is part of the description of  $\mathcal{F}_{\text{T-LEDGER}}$ .  $f$  takes as input the transaction  $\mathbf{tx}$  and computes the required fee. In the case where the output of  $f$  is greater than fee, the transaction is immediately discarded. Otherwise,  $\mathcal{F}_{\text{T-LEDGER}}$  replaces fee with the output of the function and submits it. This captures the fact that  $\mathcal{F}_{\text{T-LEDGER}}$  charges the issuer of the transaction only for the cost of processing the transaction, even if the transaction specifies a higher fee.

In more detail, each party has an associated wallet address, and different parties have different wallet addresses.  $\mathcal{F}_{\text{T-LEDGER}}$  manages a table  $\mathcal{T}$  that, for each party  $P_i$ , stores  $P_i$ 's wallet address  $\text{addr}_i$ . We initialize  $\mathcal{F}_{\text{T-LEDGER}}$  with a party  $P_0$  which initially holds all the coins (e.g.,  $V$  coins) of type  $\mathbf{E}^{10}$ . To do so,  $\mathcal{F}_{\text{T-LEDGER}}$  generates an address  $\text{addr}_0$  and sends  $(\text{SUBMIT}, \text{sid}, \mathbf{tx})$  to the wrapped  $\mathcal{F}_{\text{LEDGER}}$  with  $\mathbf{tx} := (V, 0^\lambda, \text{addr}_0, \perp, 0)$ , where  $V$  is the initial amount of coins held by  $P_i$  and  $0^\lambda$  is a special address used only for the initialization.

Upon receiving a registration request from a party  $P_i$ ,  $\mathcal{F}_{\text{T-LEDGER}}$  creates a new wallet address  $\text{addr}_i$  and adds  $(\text{addr}_i, P_i)$  to the table  $\mathcal{T}$ .

$\mathcal{F}_{\text{T-LEDGER}}$ , upon receiving  $(\text{SUBMIT}, \text{sid}, \mathbf{tx})$  from a party  $P_i$ , performs the following steps.

- Parse  $\mathbf{tx}$  as  $(v, \text{addr}_i, \text{addr}_j, \text{aux}, \text{fee})$  and continue if and only if  $(P_i, \text{addr}_i) \in \mathcal{T}$  and  $\text{fee} \geq f(\mathbf{tx})$ .
- Get **state** and **buffer** of  $\mathcal{F}_{\text{LEDGER}}$  and check that the balance of transactions to/from the wallet address  $\text{addr}_i$  is at least  $v' \geq v + f(\mathbf{tx})$  coins. That is, the sum of coins with receiver address  $\text{addr}_i$  minus the sum of coins in the state with sender address  $\text{addr}_i$  (including the fees) is greater than or equal to  $v + f(\mathbf{tx})$ . If this is not the case, deem the transaction invalid; otherwise, submit  $\mathbf{tx}$  to  $\mathcal{F}_{\text{LEDGER}}$  with the fee  $f(\mathbf{tx})$ .

$\mathcal{F}_{\text{T-LEDGER}}$  is also parametrized with the identifier of an ideal functionality  $\mathcal{F}_{\text{trap}}$ . Whenever  $\mathcal{F}_{\text{T-LEDGER}}$  receives the command  $(\text{SUBMIT-TRAPDOOR}, \text{sid}, \mathbf{tx}, P_i)$  from  $\mathcal{F}_{\text{trap}}$ , it forwards the transaction  $\mathbf{tx}$  on behalf of  $P_i$  to  $\mathcal{F}_{\text{LEDGER}}$  without checking anything about  $\mathbf{tx}$  in terms of balances and fees. This simple mechanism allows  $\mathcal{F}_{\text{T-LEDGER}}$  to interact with other ideal functionalities when required. This becomes particularly helpful when we want to enhance the behavior of  $\mathcal{F}_{\text{T-LEDGER}}$  with smart contracts, and in the next section we show how to do that.

For all the other input commands,  $\mathcal{F}_{\text{T-LEDGER}}$  just acts as a proxy between  $\mathcal{F}_{\text{LEDGER}}$  and its external interface.

To conclude the description of  $\mathcal{F}_{\text{T-LEDGER}}$ , we need to specify how **Blockify** works. **Blockify** is a simple procedure that takes as input the next block to be added to the state, and outputs a concatenation of the transactions contained in the block. This means that the state of  $\mathcal{F}_{\text{LEDGER}}$  (which will correspond also to the state of  $\mathcal{F}_{\text{T-LEDGER}}$ ) is represented by just list of transactions. We do not specify how **ExtendPolicy** works, as any realization of **ExtendPolicy** can be used in our formalization. We provide a formal description of  $\mathcal{F}_{\text{T-LEDGER}}$  in [Figure 2](#).

We note that  $\mathcal{F}_{\text{T-LEDGER}}$  does not specify who gets the fee, but this would not be difficult to do since  $\mathcal{F}_{\text{LEDGER}}$  keeps track of the party that generated each block. Hence, it would be easy to modify  $\mathcal{F}_{\text{T-LEDGER}}$  to keep track of which party gets the fees of the transactions that constitute a block. Another simplification we make is to consider fixed relation between the cost required to execute a transaction (or call a contract as we will see) and the complexity of the transaction (or the contract call). In system like Ethereum this is not the case, as the fee that a party pays depends on the complexity of the transaction (which determines the amount of gas) and on the gas price.

---

<sup>10</sup>It is easy to initialize the functionality with an arbitrary number of parties that hold an initial amount of coin. To simplify the description on the functionality, we decided to use only one party in this phase.

## $\mathcal{F}_{\text{T-Ledger}}$

### Initialization

1. Parameters: the trapdoor functionality  $\mathcal{F}_{\text{trap}}$  and the fee function  $f$ .
2. Send (REGISTER,  $P_0$ ) to  $\mathcal{A}$ .
3. Upon receiving  $\text{addr}_0$  from  $\mathcal{A}$ , if  $\text{addr}_0 = 0^\lambda$ , then ignore the command and stop, else add  $(P_0, \text{addr}_0)$  to  $\mathcal{T}$ .
4. Initialize the functionality  $\mathcal{F}_{\text{LEDGER}}$  with a registered party  $P_0$ .

### Registration

- Upon receiving (REGISTER) from a party  $P_i$ , send (REGISTER,  $P_i$ ) to  $\mathcal{A}$ . Upon receiving  $\text{addr}_i$  from  $\mathcal{A}$ , if there is already an entry  $(P_j, \text{addr}_i) \in \mathcal{T}$  for some  $P_j \in \mathcal{P}$ , then ignore the command, else add  $(P_i, \text{addr}_i)$  to  $\mathcal{T}$ , register  $P_i$  to  $\mathcal{F}_{\text{LEDGER}}$ , and send ( $\text{addr}_i$ ) to  $P_i$ .

### Transactions

- Upon receiving (SUBMIT,  $\text{sid}, \text{tx}$ ) from a party  $P_i$ , parse  $\text{tx}$  as  $(v, \text{addr}_i, \text{addr}_j, \text{aux}, \text{fee})$ . If there exists an entry  $(P_i, \text{addr}_i)$  in  $\mathcal{T}$  and  $\text{fee} \geq f(\text{tx})$ , then continue with the following steps, else ignore the command.
  - Get **state** and **buffer** from  $\mathcal{F}_{\text{LEDGER}}$ , initialize  $\text{balance} \leftarrow 0$  and for each  $\text{tx}^*$  in **buffer** and in **state**.
    - If  $\text{tx}^* = (v^*, \text{addr}_i, \text{addr}, \text{aux}^*, \text{fee}^*)$ , then compute  $\text{balance} \leftarrow \text{balance} - v^* - \text{fee}^*$ .
    - If  $\text{tx}^* = (v^*, \text{addr}^*, \text{addr}_i, \text{aux}^*, \text{fee}^*)$ , then compute  $\text{balance} \leftarrow \text{balance} + v^*$ .
  - If  $\text{balance} \geq v + f(\text{fee})$ , then send (SUBMIT,  $\text{sid}, (v, \text{addr}_i, \text{addr}_j, \text{aux}, f(\text{fee}))$ ) to  $\mathcal{F}_{\text{LEDGER}}$  on behalf of  $P_i$ .

**Trapdoor input** Upon receiving (SUBMIT-TRAPDOOR,  $\text{sid}, \text{tx}, P_i$ ) from  $\mathcal{F}_{\text{trap}}$ , send (SUBMIT,  $\text{sid}, \text{tx}$ ) to  $\mathcal{F}_{\text{LEDGER}}$  on behalf of  $P_i$ .

### Getting state and other commands

- Upon receiving (READ,  $\text{sid}$ ) from  $P_i$ , send (READ,  $\text{sid}$ ) to  $\mathcal{F}_{\text{LEDGER}}$ . Upon receiving (READ,  $\text{sid}, \text{state}$ ), forward (READ,  $\text{sid}, \text{state}$ ) to  $P_i$
- Upon receiving any other input from an honest party  $P_i \in \mathcal{P}$  (resp. from  $\mathcal{A}$ ), forward it to  $\mathcal{F}_{\text{LEDGER}}$  on behalf of  $P_i$  (resp.  $\mathcal{A}$ ). Upon receiving a reply to a command sent on behalf of a party  $P_i$  (resp. from  $\mathcal{A}$ ), forward it to  $P_i$  (resp.  $\mathcal{A}$ ).

Figure 2: This ledger allows exchanging coins of type **E** between parties.

This means that how fast and if a transaction will be executed depends on the product of gas price and amount of required gas. We could modify  $\mathcal{F}_{\text{T-LEDGER}}$  (and the other functionalities we will consider) to accommodate for an additional mechanism that allows the adversary communicating to the functionality the average gas price, in such a way that we can use this gas cost to decide whether to accept or reject a transaction.

However, since these aspects are not relevant for our results, to simplify the description of our already involved ideal functionalities, we have decided to not include such mechanisms in our model.

## 6 How to handle smart-contracts

In this section we define the functionality  $\mathcal{F}_{\text{TSC-LEDGER}}$  that, in addition to  $\mathcal{F}_{\text{T-LEDGER}}$ , captures a ledger that enables a large class of smart contracts.  $\mathcal{F}_{\text{TSC-LEDGER}}$  internally runs  $\mathcal{F}_{\text{T-LEDGER}}$  and a smart contract (formally defined by means of an additional ideal functionality). The contract has a state that can be updated by any party that can afford to pay a fee (that depends on the contract and on the input). After any valid update, the new contract state is pushed onto the  $\mathcal{F}_{\text{T-LEDGER}}$ 's state. As we have alluded, in order for the contract to freely interact with  $\mathcal{F}_{\text{T-LEDGER}}$ , the parameter  $\mathcal{F}_{\text{trap}}$  of  $\mathcal{F}_{\text{T-LEDGER}}$  is set to be equal to the identity of  $\mathcal{F}_{\text{TSC-LEDGER}}$ , which will act as a bridge between the contract functionality and  $\mathcal{F}_{\text{T-LEDGER}}$ . To simplify the description of the

functionality, we describe the case where only one smart contract is running; however, it is easy to extend the functionality to the case where multiple smart contracts are running at the same time.

A smart contract  $\mathcal{F}_{\text{SC}}$  is a small functionality managed by  $\mathcal{F}_{\text{TSC-LEDGER}}$  that maintains its own state  $\text{cstate}$ . The behavior of  $\mathcal{F}_{\text{SC}}$  is fully determined by three procedures:  $f_{\text{CFee}}$ ,  $f_{\text{filter}}$  and  $f_{\text{trans}}$ .

- $f_{\text{CFee}}$  (the contract fee function) takes as input the contract state  $\text{cstate}$ , the ledger state of  $\mathcal{F}_{\text{T-LEDGER}}$ , a transaction, (which represents the input received by the contract's caller) and the fee specified in the input transaction. If the fee indicated is sufficient to update the contract state, then  $f_{\text{CFee}}$  returns the actual fee required to run the contract (which could be less than the fee indicated by the contract's caller). If the submitted fee is not sufficient, then the function returns  $\perp$ .
- $f_{\text{trans}}$  (the state transition function) takes as input the payload of the input transaction,  $\mathcal{F}_{\text{T-LEDGER}}$ 's state, and the contract state  $\text{cstate}$ , and returns a new contract state updated according to its inputs.
- $f_{\text{filter}}$  (the filtering function) takes as input 1) the view that the contract's caller has of  $\mathcal{F}_{\text{T-LEDGER}}$ 's state  $\text{state}_i$  and 2) the contract state, and returns an arbitrary sub-set of the information contained in  $\text{state}_i$ .

The functionality  $\mathcal{F}_{\text{TSC-LEDGER}}$  is also parametrized by  $\text{Fee}$ , which represents the minimum fee that a party should pay in order to query a contract (to update the contract the fee might be higher). In more detail,  $\mathcal{F}_{\text{TSC-LEDGER}}$  accepts transactions with the following format:  $\text{tx} := (v, \text{addr}_i^{\text{E}}, \text{addr}_j^{\text{E}}, \text{aux}, \text{fee}, \text{type})$ , where  $\text{type} \in \{\text{E}, \text{SC}\}$  denotes whether the transaction should be treated as a normal transaction or as a call to the contract. In particular,  $\mathcal{F}_{\text{TSC-LEDGER}}$  checks whether  $\text{type} = \text{E}$  or  $\text{type} = \text{SC}$ . In the former case,  $\mathcal{F}_{\text{TSC-LEDGER}}$  removes the field  $\text{type}$  from the transaction and forwards it to  $\mathcal{F}_{\text{T-LEDGER}}$ . In the latter,  $\mathcal{F}_{\text{TSC-LEDGER}}$  checks that  $\text{fee} \geq \text{Fee}$  and that the issuer of the transaction has at least  $\text{fee}$  coins of type  $\text{E}$  in its wallet<sup>11</sup>. If this check is successful, then  $\mathcal{F}_{\text{TSC-LEDGER}}$  forwards the transaction and the current ledger state to  $\mathcal{F}_{\text{SC}}$ , which does the following: It uses  $f_{\text{CFee}}$  to check whether the fee specified in  $\text{tx}$  minus the fee required to query the contract (denoted with  $\text{Fee}$ ) would be sufficient to update the contract state using the input  $\text{aux}$ . If  $f_{\text{CFee}}$  returns  $\perp$ , then the contract returns  $(\text{ko}, \text{cstate}, \text{fee})$ .

Else, if  $f_{\text{CFee}}$  returns  $\text{fee}^{\text{SC}}$ ,  $\mathcal{F}_{\text{SC}}$  computes the updated contract state  $\text{cstate}$  by running  $f_{\text{trans}}$  on input the payload of  $\text{tx}$  (denoted with  $\text{aux}$ ), the ledger state, and the contract state, and returns  $(\text{ok}, \text{cstate}, \text{fee}^{\text{SC}} + \text{Fee})$ .

$\mathcal{F}_{\text{TSC-LEDGER}}$  upon receiving  $(\text{Flag}_{\text{C}}, \text{cstate}, \text{actualfee})$  from  $\mathcal{F}_{\text{SC}}$ , constructs and sends to  $\mathcal{F}_{\text{T-LEDGER}}$  the transaction  $\text{tx}^{\text{E}} := (0, \text{addr}_i^{\text{E}}, 0^\lambda, (\text{Flag}_{\text{C}}, \text{aux}, \text{cstate}, \mathcal{F}_{\text{SC}}.\text{id}), \text{actualfee})$  using the command `SUBMIT-TRAPDOOR`, where we recall that  $\text{aux}$  is the payload of  $\text{tx}$ ,  $\text{Flag}_{\text{C}} \in \{\text{ok}, \text{ko}\}$ , and  $\mathcal{F}_{\text{SC}}.\text{id}$  is the identifier of  $\text{SC}$ .

We note that the transaction  $\text{tx}^{\text{E}}$  is a standard  $\mathcal{F}_{\text{T-LEDGER}}$  transaction that contains in its payload the updated state of the contract (or the old state if the fee was not sufficient), the input used to eventually update the contract's state, and the fee  $\text{actualfee}$  such that:

- if  $\text{Flag}_{\text{C}} = \text{ko}$  (i.e. the fee specified by the contract's caller was not sufficient to update the contract state) then  $\text{actualfee} = \text{fee}$

<sup>11</sup> $\mathcal{F}_{\text{TSC-LEDGER}}$  can do this check since it has full access to  $\mathcal{F}_{\text{T-LEDGER}}$ 's state and buffer.

- if  $\text{Flag}_C = \text{ok}$  (i.e. fee was sufficient to update the contract's state) then  $\text{actualfee} \leq \text{fee}$ .

Note that it might be that  $\text{actualfee} < \text{fee}$  in the case where the fee required to update the contract state is less than  $\text{fee}$ . That is,  $\mathcal{F}_{\text{TSC-LEDGER}}$  only charges the contract caller exactly for the fee required to run the contract. When  $\text{fee}$  is insufficient to complete execution of the contract, the issuer of the transaction pays the full amount of  $\text{fee}$  even though no change to the contract state is committed. (This is consistent with Ethereum and other blockchains that support Turing-complete smart contracts.) We refer to Figure 3 for the formal description of  $\mathcal{F}_{\text{TSC-LEDGER}}$  and for the abstraction of  $\mathcal{F}_{\text{SC}}$ .

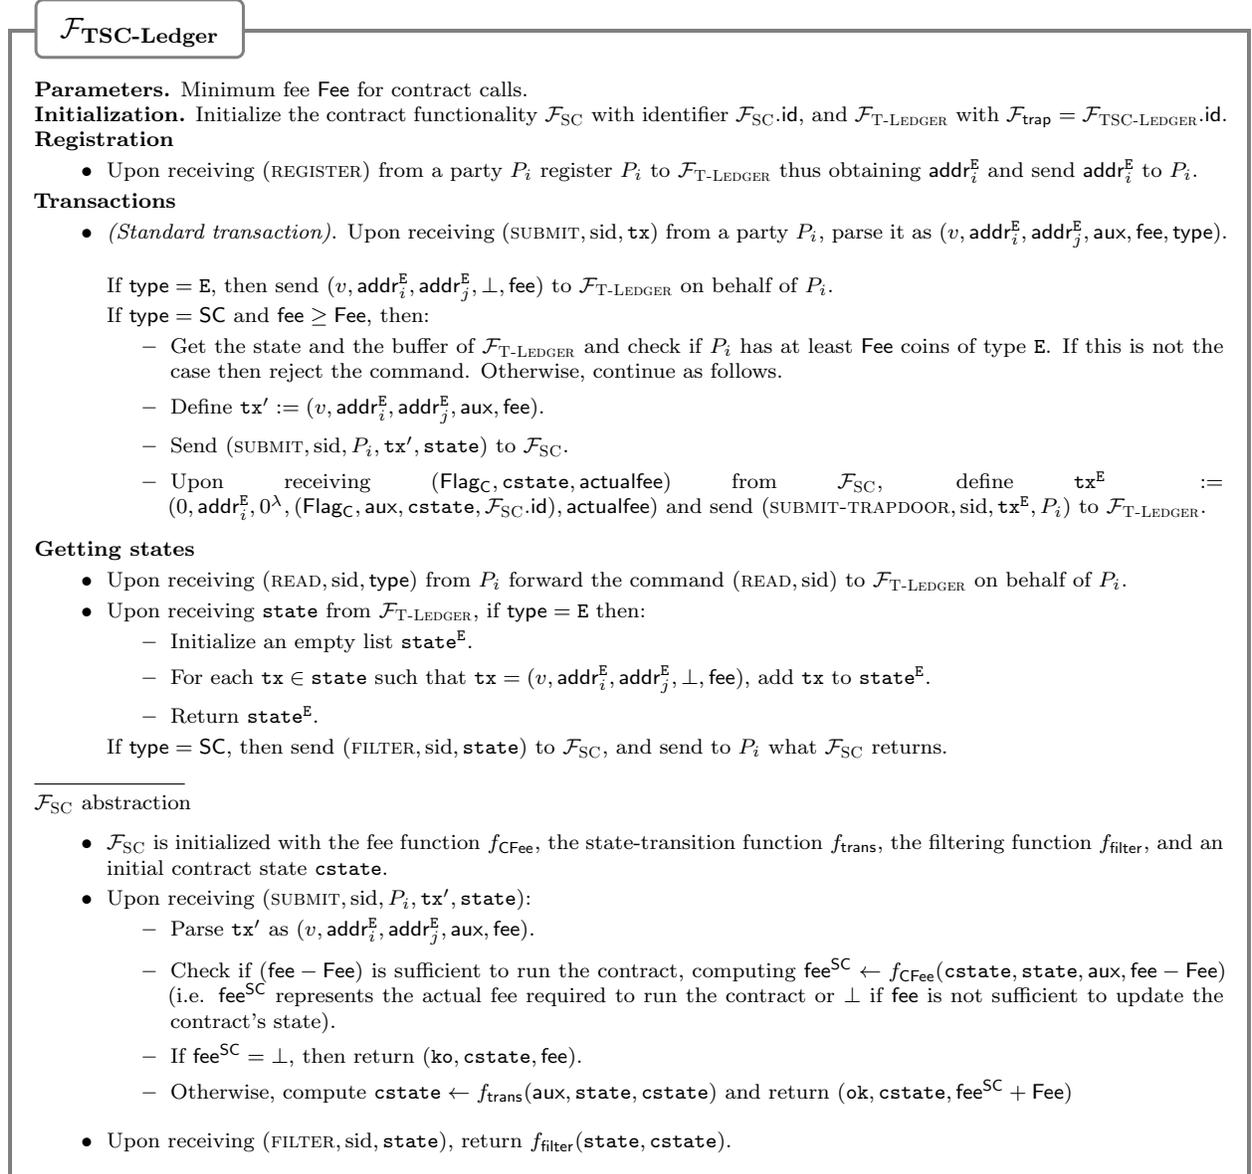


Figure 3: This ledger tolerates any type of contract abstracted by  $\mathcal{F}_{\text{SC}}$ .

## $\mathcal{F}_{\text{SC}}^{\text{T}}$ functions

**Initialization.** The contract is parametrized by the functions  $f_{\text{CFee}}^{\text{T}}$ ,  $f_{\text{trans}}$ , and  $f_{\text{filter}}^{\text{T}}$  described below.  $\text{cstate}$  consists of the following components:

- Constants:  $y$ ,  $\text{addr}_0^{\text{T}}$ ,  $\text{Fee}$ , identifier  $\mathcal{F}_{\text{SC}}^{\text{T}}.\text{id}$ .
- Empty set  $\text{token-set}$ .

### Functions and helper procedures

$f_{\text{trans}}(\text{aux}^{\text{E}}, \text{state}, \text{cstate})$

- Add  $\text{aux}^{\text{E}}$  to  $\text{token-set}$  of  $\text{cstate}$  and return the updated  $\text{cstate}$ .

$f_{\text{CFee}}(\text{cstate}, \text{state}, \text{tx}, \text{fee})$

- Define and initialize  $\text{tfee} \leftarrow 0$ .
- Parse  $\text{tx} := (0, \text{addr}_i^{\text{E}}, 0^\lambda, \text{aux}^{\text{E}}, \text{fee}^{\text{E}})$ .
- Parse  $\text{tx}^{\text{T}}$  as  $(v, \text{addr}_i^{\text{T}}, \text{addrs}, \text{id})$  and compute  $\text{tfee} \leftarrow \text{tfee} + \text{Fee}|v|$ .
- If  $\text{tfee} > \text{fee}$ , then return  $\perp$ ; otherwise, return  $\text{tfee}$ .

$f_{\text{filter}}(\text{cstate}, \text{state})$

- Initialize the empty list  $\text{state}^{\text{T}}$  and set  $\text{temp-buffer} \leftarrow \text{token-set}$ .
- For each  $\text{tx}^{\text{E}} = (0^\lambda, \text{addr}_i^{\text{E}}, 0^\lambda, \text{aux}^{\text{E}}, \text{fee}^{\text{E}})$  in  $\text{state}$  where  $\text{aux}^{\text{E}} = (\text{ok}, (v^*, \text{addr}_i^*, \text{addrs}^*, \text{id}_i^*), \text{cstate}^*, \mathcal{F}_{\text{SC}}^{\text{T}}.\text{id})$ :
  - Define  $\text{aux} := (v^*, \text{addr}_i^*, \text{addrs}^*, \text{id}_i^*)$ .
  - If  $\text{verify}(\text{aux}, \text{state}^{\text{T}}) = 1$  and  $\text{aux} \in \text{temp-buffer}$ , then add  $\text{aux}$  to  $\text{state}^{\text{T}}$  and remove  $\text{aux}$  from  $\text{temp-buffer}$ .
- Return  $(\text{READ}, \text{sid}, \text{state})$ .

$\text{verify}(\text{aux}^{\text{E}}, \text{state}^{\text{T}})$

- Parse  $\text{aux}^{\text{E}}$  as  $(v^{\text{T}}, \text{addr}_i^{\text{T}}, \text{addrs}, \text{id}_i^{\text{T}})$
- If  $\text{addr}_0^{\text{T}} = \text{addr}_i^{\text{T}}$ , then initialize  $\text{balance} \leftarrow y$ ; otherwise,  $\text{balance} \leftarrow 0$ .
- For each  $\text{tx}^* = (v^*, \text{addr}_i^*, \text{addrs}^*, \text{id}_i^*)$  in  $\text{state}^{\text{T}}$ :
  - If  $\text{addr}^* = \text{addr}_i^{\text{T}}$ , then compute  $\text{balance} \leftarrow -\sum_k v^*[k]$
  - For each  $k$  such that  $\text{addrs}^*[k] = \text{addr}_i^{\text{T}}$ , compute  $\text{balance} \leftarrow +v^*[k]$ .
- If  $\sum_k v^{\text{T}}[k] \geq \text{balance}$  then return 1 else return 0.

Figure 4: Smart contract for the creation of a new token. The fee required to run the contract is computed by multiplying the number of token transactions encoded in the payload of the input  $\text{tx}$  and  $\text{Fee}$ .

## 7 Our ideal functionality $\mathcal{F}_{\text{Ledger}}^{\text{EET}}$

In this section we can finally define the functionality  $\mathcal{F}_{\text{Ledger}}^{\text{EET}}$ .  $\mathcal{F}_{\text{Ledger}}^{\text{EET}}$  internally runs  $\mathcal{F}_{\text{TSC-Ledger}}$ , parametrized by a contract  $\mathcal{F}_{\text{SC}}^{\text{T}}$ .  $\mathcal{F}_{\text{SC}}^{\text{T}}$  maintains a token  $\text{T}$ , and allows parties to issue transactions with respect to such a token. At a high level, any party that has some tokens can send it to another party by querying the contract  $\mathcal{F}_{\text{SC}}^{\text{T}}$ . However, invoking the contract requires payment of a fee in the native currency  $\text{E}$ , even if the transaction involves only tokens. To mitigate this problem, our functionality allows a sender  $P_i$  to send tokens to another party  $P_j$ , even if  $P_i$  does not have native coins. In particular, the sender will pay a fee of at least  $\text{del-fee}$  tokens  $\text{T}$  to a special party  $\text{M}$ , called the intermediary, and  $\text{M}$  will pay the fee in  $\text{E}$  on the behalf of the sender ( $\text{del-fee}$  is a fixed amount of tokens that parametrizes our functionality). The functionality guarantees that either the transaction by  $P_i$  becomes part of the ledger state *and*  $\text{M}$  gets a fixed amount of tokens  $\text{del-fee}$ , or nothing happens. We propose the formal description of  $\mathcal{F}_{\text{Ledger}}^{\text{EET}}$  in Figure 5 and the description of

## $\mathcal{F}_{\text{Ledger}}^{\text{EET}}$

### Initialization

- Initialize an empty set  $\mathcal{T}^T$ .
- Send (REGISTER,  $P_0$ ) to  $\mathcal{A}$ .
- Upon receiving  $\text{addr}_0^T$  from  $\mathcal{A}$ , add  $(P_0, \text{addr}_0^T)$  to  $\mathcal{T}^T$ , run the initialization procedure of  $\mathcal{F}_{\text{SC}}^T$ , and initialize the wrapped functionality  $\mathcal{F}_{\text{TSC-LEDGER}}$  with the contract  $\mathcal{F}_{\text{SC}}^T$ , using identifier  $\mathcal{F}_{\text{SC}}^T.\text{id}$
- Send (REGISTER,  $M$ ) to  $\mathcal{A}$ .
- Upon receiving  $\text{addr}_M^T$  from  $\mathcal{A}$ , add  $(M, \text{addr}_M^T)$  to  $\mathcal{T}^T$ .

### Registration

- Upon receiving (REGISTER) from a party  $P_i$ , send (REGISTER,  $P_i$ ) to  $\mathcal{A}$ .
- Upon receiving  $\text{addr}_i^T$  from  $\mathcal{A}$ , if there is already an entry  $(P_j, \text{addr}_j^T) \in \mathcal{T}^T$  for some  $P_j \in \mathcal{P}$ , then ignore the command; otherwise, add  $(P_i, \text{addr}_i^T)$  to  $\mathcal{T}^T$  and register  $P_i$  to  $\mathcal{F}_{\text{TSC-LEDGER}}$ , thus obtaining  $\text{addr}_i^E$ , and send  $(\text{addr}_i^E, \text{addr}_i^T)$  to  $P_i$ .

### Transactions

- (*Standard transaction*). Upon receiving (SUBMIT,  $\text{sid}, \text{tx}^T$ ) from a party  $P_i$ , parse  $\text{tx}^T$  as  $(v, \text{addr}_i^E, \text{addr}_i^T, \text{addr}_j^E, \text{addr}_j^T, \text{fee}, \text{Coin})$ .

If  $\text{Coin} = T$ , and there exists an entry  $(P_i, \text{addr}_i^T)$  in  $\mathcal{T}^T$ , and  $\text{fee} \geq 2\text{Fee}$ , then do the following :

- Send (REQ-TRX,  $P_i, \text{tx}^T$ ) to  $\mathcal{A}$  and, upon receiving  $\text{id}_i$ , define  $\text{aux} := (v, \text{addr}_i^T, \text{addr}_j^T, \text{id}_i)$  and  $\text{tx} = (0, \text{addr}_i^E, 0^\lambda, \text{aux}, \text{fee}, \text{SC})$ .

If  $\text{Coin} = E$  and  $\text{fee} \geq \text{Fee}$ , then:

Define  $\text{tx} := (v, \text{addr}_i^E, \text{addr}_j^E, \perp, \text{fee}, E)$

- Send (SUBMIT,  $\text{sid}, \text{tx}$ ) to  $\mathcal{F}_{\text{TSC-LEDGER}}$  on the behalf of  $P_i$ .
- (*Delegatable transaction*). Upon receiving (SUBMIT-DELEGATION,  $\text{sid}, \text{tx}^T$ ) from a party  $P_i$ , parse  $\text{tx}^T$  as  $(v, \text{addr}_i^T, \text{addr}_j^T, \text{fee}^T)$ . If there exists an entry  $(P_i, \text{addr}_i^T)$  in  $\mathcal{T}^T$  and  $\text{fee}^T \geq \text{del-fee}$ , then do the following, ignoring the command otherwise:
  - Send (REQ-TRX-DEL,  $P_i, \text{tx}^T$ ) to  $\mathcal{A}$  and, upon receiving  $\text{id}_i$ , define  $\text{aux} := ([v, \text{fee}^T], \text{addr}_i^T, [\text{addr}_j^T, \text{addr}_M^T], \text{id}_i)$ .
  - If  $M$  is honest, then define  $\text{tx} := (0, \text{addr}_M, 0^\lambda, \text{aux}, 3\text{Fee}, \text{SC})$  and send  $\text{tx}$  to  $\mathcal{F}_{\text{TSC-LEDGER}}$  on behalf of  $M$ . If  $M$  is corrupted, do the following:
    - Send (DELEGATE,  $\text{aux}, P_i$ ) to  $M$ .
    - If  $M$  replies with (REJECT,  $P_i$ ), then send REJECT to  $P_i$ . If  $M$  replies with (ACCEPT,  $P_i, \text{fee}$ ), then define  $\text{tx} := (0, \text{addr}_M, 0^\lambda, \text{aux}, \text{fee})$  and send  $\text{tx}$  to  $\mathcal{F}_{\text{TSC-LEDGER}}$  on behalf of  $M$ .

### Getting states

- Upon receiving (READ,  $\text{sid}, \text{Coin}$ ) from  $P_i$ , forward the command to  $\mathcal{F}_{\text{TSC-LEDGER}}$  on behalf of  $P_i$ .
- Upon receiving (READ,  $\text{sid}, \text{state}$ ), forward it to  $P_i$ .

### Forwarding queries to $\mathcal{F}_{\text{TSC-Ledger}}$

- Upon receiving (INNER-INPUT,  $\text{sid}, m, P_i$ ) from  $\mathcal{A}$ , if  $P_i$  is an honest party, then ignore the command. Otherwise, if  $P_i$  is corrupted, send  $m$  to  $\mathcal{F}_{\text{TSC-LEDGER}}$  on behalf of  $P_i$ .
- Upon receiving any other input from an honest party  $P_i \in \mathcal{P}$  (resp. from  $\mathcal{A}$ ), forward it to  $\mathcal{F}_{\text{TSC-LEDGER}}$  on behalf of  $P_i$ .
- Upon receiving a reply to a command sent on behalf of a party  $P_i \in \mathcal{P}$  (resp. from  $\mathcal{A}$ ), forward it to  $P_i$  (resp.  $\mathcal{A}$ ).

Figure 5: This ledger allows parties with no coins of type E to post transactions using tokens of type T (we call this transaction a delegated transaction). In the case where  $M$  is honest and has enough coins of type E to pay the fee, the delegated transactions are always included in the ledger state.

$\mathcal{F}_{\text{SC}}^T$  in Figure 4, and provide a high level description of those below.

The functionality  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , interacts with a set of parties, with the adversary, and with a special party that we denote with  $M$  (the intermediary), and manages the *token wallet* addresses of the

registered parties. We assume that a party  $P_0$  initially holds all of the available tokens<sup>12</sup>. We denote the token wallet addresses of  $P_0$  and  $M$  with  $\text{addr}_0^T$  and  $\text{addr}_M^T$  respectively. Any time  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  receives a registration command from a party  $P_i$ , it registers  $P_i$  to the ledger  $\mathcal{F}_{\text{TSC-LEDGER}}$ , thus obtaining  $\text{addr}_i^E$ . It then generates a token wallet address  $\text{addr}_i^T$  and returns  $(\text{addr}_i^E, \text{addr}_i^T)$  to  $P_i$ .  $(\text{addr}_i^E, \text{addr}_i^T)$  represents respectively the wallet addresses for the native currency  $E$  and for the token  $T$ .  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  tolerates two types of transactions: *standard* and *delegated* transactions. Any registered party  $P_i$  can issue a standard transaction  $\text{tx}^T := (v, \text{addr}_i^E, \text{addr}_i^T, \text{addr}_j^T, \text{fee})$ , where  $v$  denotes the amount of tokens,  $(\text{addr}_i^E, \text{addr}_i^T)$  are the addresses of the sender,  $\text{addr}_j^T$  is the token wallet address of the receiver, and  $\text{fee}$  is the fee expressed in coins of type  $E$ .

$\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  takes  $\text{tx}^T$  and creates a transaction  $\text{tx}^E$  for the ledger  $\mathcal{F}_{\text{TSC-LEDGER}}$  that 1) has as a sender address  $\text{addr}_i^E$ , 2) has a fee  $\text{fee}$ , and 3) calls the contract  $\mathcal{F}_{\text{SC}}^T$  and includes in its payload what we call a *token transaction*  $\text{tx}' := (v, \text{addr}_i^T, \text{addr}_j^T)$ .<sup>13</sup>  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  then forwards  $\text{tx}^E$  to the ledger  $\mathcal{F}_{\text{TSC-LEDGER}}$  on behalf of  $P_i$ .

The contract  $\mathcal{F}_{\text{SC}}^T$  maintains a set **token-set** as part of its state, and if the fee specified in  $\text{tx}^E$  is sufficient, it updates its state by adding  $\text{tx}'$  to **token-set** and returns  $(\text{ok}, \text{cstate}, \text{actualfee})$ . Note that this means that the  $\text{tx}'$  is part of the contract state and appears in the  $\mathcal{F}_{\text{TSC-LEDGER}}$ 's state by definition.

To complete this first part of the description of  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , it remains to specify the function  $f_{\text{filter}}$  (and  $f_{\text{CFee}}$ , which we describe later in this section) of  $\mathcal{F}_{\text{SC}}^T$ .  $f_{\text{filter}}$  receives as input the contract state and the state of  $\mathcal{F}_{\text{TSC-LEDGER}}$  (which we denote **state**) and, for each transaction  $\text{tx}$  in **state** such that  $\text{tx}^E := (0^\lambda, \text{addr}_i^E, 0^\lambda, \text{aux}^E, \text{fee}^E)$  (where  $\text{aux}^E = (\text{ok}, \text{tx}', \text{cstate}^*, \mathcal{F}_{\text{SC}}^T.\text{id})$ ), adds  $\text{tx}'$  to **state**<sup>T</sup> if and only if:

1.  $\text{tx}'$  appears in **token-set** (which is part of the token state).
2.  $\text{tx}' := (v, \text{addr}_i^T, \text{addr}_j^T)$  and the sum of tokens in the token transactions stored so far in **state**<sup>T</sup> with receiver address  $\text{addr}_j^T$ , minus the sum of coins in the state with sender address  $\text{addr}_i^T$ , is greater than or equal to  $v$ .

$\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  captures the main characteristics of a token, relying on the smart contract to filter out invalid transactions. Unfortunately, the mechanism that we have discussed so far has a major drawback: if a party wants to issue a token transaction, they must have the required amount of coins of type  $E$  to query the contract.

To get rid of this requirement,  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  admits what we call *delegated transactions*. A party that wants to issue a delegated transaction submits  $\text{tx}^T := (v, \text{addr}_i^T, \text{addr}_j^T, \text{fee}^T)$  to  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , which in turns asks the special party denoted  $M$  to pay the fee in  $E$  in exchange of (at least) **del-fee** tokens  $T$ , which will be taken from  $P_i$ 's account. If  $M$  is honest and  $\text{fee}^T \geq \text{del-fee}$ , (where we recall that **del-fee** is the minimum fee required for the delegation to be considered,) then  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  submits a call to the contract  $\mathcal{F}_{\text{SC}}^T$  on behalf of  $M$  with the input (the payload of the transaction)  $\text{aux} := (([v, \text{fee}^T], \text{addr}_i^T, [\text{addr}_j^T, \text{addr}_M^T]))$ . If  $M$  has enough coins of type  $E$  to afford the call to  $\mathcal{F}_{\text{SC}}^T$ , then  $\text{aux}$  will become part of the contract state. To accommodate for this special input, we modify the filtering function  $f_{\text{filter}}$  of  $\mathcal{F}_{\text{SC}}^T$  in such a way that the value  $\text{aux}$  can also be understood as two atomic token transactions: the first moves  $v$  tokens from the wallet address  $\text{addr}_i^T$  to the wallet

<sup>12</sup>As before, we could have multiple addresses having different amounts of tokens, but for simplicity, we assume that only one party initially holds tokens.

<sup>13</sup>The payload also includes an identifier chosen by the adversary, which we omit in this informal description.

address  $\text{addr}_j^T$ , and the second moves  $\text{fee}^T$  from the wallet address  $\text{addr}_i^T$  to the wallet address  $\text{addr}_M^T$ . It remains to specify how the contract computes the fee. The function  $f_{\text{CFee}}$  charges  $\text{Fee}$  coins of type  $\mathbf{E}$  for each token transaction encoded in  $\text{aux}$  (the input that is used to update the contract state). Hence, for a non-delegated token transaction,  $f_{\text{CFee}}$  would return  $\text{Fee}$ , and for a delegated token transaction, it would return  $2\text{Fee}$ . In addition to this fee, we need to consider the fee required simply to query the contract. Hence, the total cost of a non-delegated transaction would be of  $2\text{FeeE}$ , and the total cost of a delegated transaction would be  $3\text{FeeE}$ . We stress that this is a simplified method of computing the fee, and that a more fine-grained calculation could be used to capture what actually happens in the real world. We use this mechanism only to simplify the description of the functionality, and later the security proof. In the [Section 9](#), we provide experimental results to estimate the cost of executing these transactions in a real world realization of  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ .

## 8 Our Protocol: how to realize $\mathcal{F}_{\text{Ledger}}^{\text{EET}}$

Our protocol is described in the  $\mathcal{F}_{\text{T-LEDGER}}$ -hybrid world, where  $\mathcal{F}_{\text{T-LEDGER}}$  is parametrized by  $\mathcal{F}_{\text{trap}} = \perp$ , and the fee function  $f$  which, upon receiving an input transaction  $\text{tx}^E$ , does the following:

- Parse  $\text{tx}$  as  $(v, \text{addr}_i, \text{addr}_j, \text{aux}, \text{fee})$ .
- If  $\text{aux} = \perp$ , then return  $\text{Fee}$ .
- Otherwise, return  $\text{Fee} + |\text{aux}|/\kappa\text{Fee}$ .

In a nutshell, the fee required for a transaction to settle in the  $\mathcal{F}_{\text{T-LEDGER}}$ 's state is  $\text{Fee}$ , plus and additional  $\text{Fee}$  for each  $\kappa$  bits contained in the payload, where  $\text{Fee}$  and  $\kappa$  are part of the description of  $f$ .

We provide the formal description of our protocol in [Figure 6](#). At a very high level, the protocol works as follows: Each party registers with  $\mathcal{F}_{\text{T-LEDGER}}$  and runs  $\text{Kgen}(1^\lambda)$  to obtain  $(\text{sk}_i^T, \text{addr}_i^T)$ , where  $\text{addr}_i^T$  represents the token wallet address. A party  $P_i$  that wants to send  $vT$  to  $P_j$  and has at least  $2\text{Fee}$  coins of type  $\mathbf{E}$  can do so by issuing a transaction for  $\mathcal{F}_{\text{T-LEDGER}}$  that contains in its payload  $\text{aux} := (v, \text{addr}_i^T, \text{addr}_j^T, \text{id}, \sigma_i^T)$ , where  $\text{id}$  is a random value, and  $\sigma_i^T$  is a signature of  $(v, \text{addr}_i^T, \text{addr}_j^T, \text{id})$  that verifies under the verification key  $\text{addr}_i^T$ . We require  $P_i$  to pay a fee of at least  $2\text{Fee}$  because we assume that, in this case,  $|\text{aux}| = \kappa$ .

When an honest party  $P_i$  receives the command  $(\text{READ}, \text{sid}, T)$ , they shall retrieve  $\mathcal{F}_{\text{T-LEDGER}}$ 's state, filter out the payload of each transaction (thus obtaining only the information related to token transactions), and output only the *valid* token transactions. A token transaction  $(v, \text{addr}_i^T, \text{addr}_j^T, \text{id}, \sigma_i^T)$  is valid if  $\text{addr}_i^T$  has received at least  $v$  tokens,  $\sigma_i^T$  is a signature of  $(v, \text{addr}_i^T, \text{addr}_j^T, \text{id})$  that verifies under the verification key  $\text{addr}_i^T$ , and there does not exist any other token transaction with the same sender address and identifier  $\text{id}$ .

Our protocol allows any party  $P_i$  that does not have coins of type  $\mathbf{E}$  to delegate the payment of the fee to  $M$ , paying  $M$  with at least  $\text{del-fee}$  tokens  $T$ . To do so,  $P_i$  creates  $m := ([v, \text{del-fee}], \text{addr}_i^T, [\text{addr}_j^T, \text{addr}_M^T], \text{id})$  and signs it, thus obtaining  $\sigma_i^T$ .  $P_i$  then sends  $(m, \sigma_i^T)$  to  $M$ . The honest  $M$  then creates a transaction for  $\mathcal{F}_{\text{T-LEDGER}}$  that includes  $(m, \sigma_i^T)$  in its payload and has a fee of at least  $3\text{Fee}$ , and submits it. We require  $M$  to pay a fee of at least  $3\text{FeeE}$  because we assume that, in this case, the payload of the transaction is  $2\kappa$  bits (as, indeed, the payload of this type of transaction contains more information). The honest  $M$  would immediately create

## Protocol $\Pi^{\text{Token}}$

### Initialization

- The issuer  $P_0$  does the following:
  1. Register to  $\mathcal{F}_{\text{T-LEDGER}}$ , thus obtaining  $\text{addr}_0$ .
  2. Compute  $(\text{sk}_0^{\text{T}}, \text{addr}_0^{\text{T}}) \xleftarrow{\$} \text{Kgen}(1^\lambda)$ .
- The intermediary  $M$  registers to  $\mathcal{F}_{\text{T-LEDGER}}$ , thus obtaining  $\text{addr}_M$ , and computes  $(\text{sk}_M^{\text{T}}, \text{addr}_M^{\text{T}}) \xleftarrow{\$} \text{Kgen}(1^\lambda)$ .

### Registration

- Upon receiving  $(\text{REGISTER}, \text{sid})$ , the party  $P_i$  sends  $(\text{REGISTER}, \text{sid})$  to  $\mathcal{F}_{\text{T-LEDGER}}$ , thus obtaining  $\text{addr}_i$ , and computes  $(\text{sk}_i^{\text{T}}, \text{addr}_i^{\text{T}}) \xleftarrow{\$} \text{Kgen}(1^\lambda)$ .

### Transactions

- $P_i$ , upon receiving  $(\text{SUBMIT-DELEGATION}, \text{sid}, \text{tx}^{\text{T}})$ , parses  $\text{tx}^{\text{T}}$  as  $(v, \text{addr}_i^{\text{T}}, \text{addr}_j^{\text{T}}, \text{fee}^{\text{T}})$  and does the following:
  1. If  $\text{fee}^{\text{T}} < \text{del-fee}$ , then ignore the command. Otherwise, continue.
  2. Sample  $\text{id} \xleftarrow{\$} \{0, 1\}^\lambda$  and define  $m := (([v, \text{fee}], \text{addr}_i^{\text{T}}, [\text{addr}_j^{\text{T}}, \text{addr}_M^{\text{T}}]), \text{id})$ .
  3. Compute  $\sigma_i^{\text{T}} \xleftarrow{\$} \text{Sign}(\text{sk}_i^{\text{T}}, m)$ .
  4. Send  $(\text{delegate}, m, \sigma_i^{\text{T}})$  to  $M$ .
- $M$ , upon receiving  $(\text{delegate}, m, \sigma_i^{\text{T}})$  from  $P_i$ , does the following:
  1. Parse  $m$  as  $(([v, \text{fee}], \text{addr}_i^{\text{T}}, [\text{addr}_j^{\text{T}}, \text{addr}_M^{\text{T}}]), \text{id})$ .
  2. If  $\text{Ver}(\text{addr}_i^{\text{T}}, m, \sigma_i^{\text{T}}) = 0$  or  $\text{fee}^{\text{T}} < \text{del-fee}$ , then ignore the message. Otherwise, continue.
  3. Define  $\text{tx}_M = (0, \text{addr}_M, 0^\lambda, (m, \sigma_i^{\text{T}}), 3\text{Fee})$ .
  4. Send  $(\text{ACCEPT}, P_i)$  to  $P_i$  and  $(\text{SUBMIT}, \text{sid}, \text{tx}_M)$  to  $\mathcal{F}_{\text{T-LEDGER}}$ .
- $P_i$ , upon receiving  $(\text{SUBMIT}, \text{sid}, \text{tx}^{\text{T}})$ , parses  $\text{tx}^{\text{T}}$  as  $(v, \text{addr}_i, \text{addr}_i^{\text{T}}, \text{addr}_j, \text{addr}_j^{\text{T}}, \text{fee}, \text{Coin})$  and does the following:
  - If  $\text{Coin} = \text{T}$  and  $\text{fee} \geq 2\text{Fee}$  then:
    - \* Sample  $\text{id} \xleftarrow{\$} \{0, 1\}^\lambda$ , define  $m := ((v, \text{addr}_i^{\text{T}}, \text{addr}_j^{\text{T}}), \text{id})$  and compute  $\sigma_i^{\text{T}} \xleftarrow{\$} \text{Sign}(\text{sk}, m)$ .
    - \* Define  $\text{tx} = (0, \text{addr}_i, 0^\lambda, (m, \sigma_i^{\text{T}}), \text{fee})$ .
  - If  $\text{Coin} = \text{E}$  and  $\text{fee} \geq \text{Fee}$  then
    - Define  $\text{tx} := (v, \text{addr}_i, \text{addr}_j, \perp, \text{fee})$ .
  - Send  $I = (\text{SUBMIT}, \text{sid}, \text{tx})$  to  $\mathcal{F}_{\text{T-LEDGER}}$ .

### Getting states

- Upon receiving  $(\text{READ}, \text{sid}, \text{type})$ ,  $P$  forwards the command  $(\text{READ}, \text{sid})$  to  $\mathcal{F}_{\text{T-LEDGER}}$ .
- Upon receiving **state** from  $\mathcal{F}_{\text{T-LEDGER}}$ , if  $\text{type} = \text{E}$ , then  $P$  does the following:
  - Initialize an empty list  $\text{state}^{\text{E}}$ .
  - For each  $\text{tx} \in \text{state}$  such that  $\text{tx} = (v, \text{addr}_i^{\text{E}}, \text{addr}_j^{\text{E}}, \perp, \text{fee})$ , add  $\text{tx}$  to  $\text{state}^{\text{E}}$ .
  - Return  $(\text{READ}, \text{sid}, \text{state}^{\text{E}})$ .
- Otherwise,  $P$  does the following:
  - Initialize the the list  $\text{state}^{\text{T}}$  with  $(y, 0^\lambda, \text{addr}_0, 0)$  and, for each  $\text{tx}^{\text{E}} = (0, \text{addr}_i^{\text{E}}, 0^\lambda, \text{aux}^{\text{E}}, \text{fee})$  in  $\text{state}$  where  $\text{aux}^{\text{E}} = (v^{\text{T}}, \text{addr}_i^{\text{T}}, \text{addrs}, \text{id}^{\text{T}}, \sigma_i^{\text{T}})$ , do the following:
    - If  $\text{checkvalidity}(\text{aux}^{\text{E}}, \text{state}^{\text{T}}) = 1$ , then add  $(v^{\text{T}}, \text{addr}_i^{\text{T}}, \text{addrs}, \text{id}^{\text{T}})$  to  $\text{state}^{\text{T}}$ .
  - Return  $(\text{READ}, \text{sid}, \text{state}^{\text{T}})$ .

### $\text{checkvalidity}(\text{aux}^{\text{E}}, \text{state}^{\text{T}})$

- Parse  $\text{aux}^{\text{E}}$  as  $(v^{\text{T}}, \text{addr}_i^{\text{T}}, \text{addrs}, \text{id}_i^{\text{T}}, \sigma_i^{\text{T}})$  and define  $m := (v^{\text{T}}, \text{addr}_i^{\text{T}}, \text{addrs}, \text{id}_i^{\text{T}})$ .
- If  $\text{Ver}(\text{addr}_i^{\text{T}}, m, \sigma_i^{\text{T}}) = 0$ , then return 0. Otherwise, continue.
- Initialize  $\text{balance} \leftarrow 0$ .
- For each  $\text{tx}^* = (v^*, \text{addr}_i^*, \text{addrs}^*, \text{id}_i^*)$  in  $\text{state}^{\text{T}}$ :
  - If  $\text{addr}^* = \text{addr}_i^{\text{T}}$  and  $\text{id}_i^* = \text{id}_i^{\text{T}}$ , then return 0.
  - If  $\text{addr}^* = \text{addr}_i^{\text{T}}$ , then compute  $\text{balance} \leftarrow -\sum_k v^*[k]$ .
  - For each  $k$  such that  $\text{addrs}^*[k] = \text{addr}_i^{\text{T}}$ , compute  $\text{balance} \leftarrow +v^*[k]$ .
- If  $\sum_k v^{\text{T}}[k] \geq \text{balance}$ , then return 1. Otherwise, return 0.

Figure 6: Our protocol.

and submit such a transaction, whereas the corrupted  $M$  might decide when (and if) to create the transaction. We require each token transaction to contain a random identifier in order to avoid replay attacks; without such an identifier, the adversary could take the payload of any transaction from  $\mathcal{F}_{\text{T-LEDGER}}$ 's state, (for instance, the payload of a transaction that moves  $v$  tokens from the address  $\text{addr}_i^{\text{T}}$  of an honest party to some potentially adversarial address,) copy this payload, and use it to generate a new transaction for  $\mathcal{F}_{\text{T-LEDGER}}$ . In this way, the adversary could empty the token wallet of the honest party without their knowledge. The other advantage of using identifiers is that an honest party that has delegated a transaction to a malicious intermediary can at any point decide to withdraw the delegation. Indeed, if  $M$  is not responding to a party that has delegated the transaction  $m := ([v, \text{del-fee}], \text{addr}_i^{\text{T}}, [\text{addr}_j^{\text{T}}, \text{addr}_M^{\text{T}}], \text{id})$  for a long time, and  $m$  does not appear in the payload of any transaction that appears in the ledger's state, then  $P_i$  can withdraw the delegation by submitting (or delegating) a token transaction with the same identifier; then, at most one of these transactions will be valid and accepted by the functionality.

**Theorem 8.1.** *The protocol  $\Pi^{\text{Token}}$  realizes  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  in the  $\mathcal{F}_{\text{T-LEDGER}}$ -hybrid model.*

*Proof.* In our proof we consider the more involved case where  $M$  is colluding with an arbitrary set of parties  $S \subset \mathcal{P}$  and  $P_0$  is honest; in any other case the proof follows from similar arguments. The simulator  $\text{Sim}_M$  internally runs the corrupted parties, emulating for them the functionality  $\mathcal{F}_{\text{T-LEDGER}}$ .

Moreover,  $\text{Sim}_M$  reacts as follows on the inputs it receives:

1. Upon receiving  $(\text{REGISTER}, P_i)$  from  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , compute  $(\text{sk}_i^{\text{T}}, \text{addr}_i^{\text{T}}) \xleftarrow{\$} \text{Kgen}(1^\lambda)$  and send  $(P_i, \text{addr}_i^{\text{T}})$  to  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ .
2. Upon receiving  $(\text{DELEGATE}, \text{aux}, P_i)$  from  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , compute  $\sigma_i^{\text{T}} \xleftarrow{\$} \text{Sign}(\text{sk}_i^{\text{T}}, \text{aux})$  and send  $(\text{delegate}, \text{aux}, \sigma_i^{\text{T}})$  to  $M$ .
3. Upon receiving  $(\text{REQ-TRX-DEL}, P_i, \text{tx}^{\text{T}})$  from  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , sample  $\text{id}_i \xleftarrow{\$} \{0, 1\}^\lambda$  and send  $\text{id}_i$  to  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ .
4. Upon receiving  $(\text{REQ-TRX}, P_i, \text{tx}^{\text{T}})$  from  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , parse  $\text{tx}^{\text{T}}$  as  $(v, \text{addr}_i, \text{addr}_i^{\text{T}}, \text{addr}_j, \text{addr}_j^{\text{T}}, \text{fee}, \text{Coin})$  and do the following:
  - If  $\text{Coin} = \text{T}$  and  $\text{fee} \geq 2\text{Fee}$ :
    - Sample  $\text{id} \xleftarrow{\$} \{0, 1\}^\lambda$ , define  $m := ((v, \text{addr}_i^{\text{T}}, \text{addr}_j^{\text{T}}), \text{id})$ , and compute  $\sigma_i^{\text{T}} \xleftarrow{\$} \text{Sign}(\text{sk}, m)$ .
    - Define  $\text{tx} = (0, \text{addr}_i, 0^\lambda, (m, \sigma_i^{\text{T}}), \text{fee})$ .
  - If  $\text{Coin} = \text{E}$  and  $\text{fee} \geq \text{Fee}$ , define  $\text{tx} := (v, \text{addr}_i, \text{addr}_j, \perp, \text{fee})$ .
  - Send  $I = (\text{SUBMIT}, \text{id}, \text{tx})$  to  $\mathcal{F}_{\text{T-LEDGER}}$  and return  $\text{id}$  to  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ .
5. If  $\mathcal{F}_{\text{T-LEDGER}}$  (emulated in the real-world) receives the input  $(\text{SUBMIT}, \text{id}, \text{tx}_M)$  from  $M$ , and all of the following are true:
  - $\text{tx}_M = (0, \text{addr}_M, 0^\lambda, (m, \sigma_i^{\text{T}}), \text{fee})$ , where  $m := (([v, \text{fee}^{\text{T}}], \text{addr}_i^{\text{T}}, [\text{addr}_j^{\text{T}}, \text{addr}_M^{\text{T}}], \text{id})$ ,
  - An honest party  $P_i$  has queried  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  with  $(\text{SUBMIT-DELEGATION}, \text{id}, \text{tx}^{\text{T}})$ , where  $\text{tx}^{\text{T}} = (v, \text{addr}_i^{\text{T}}, \text{addr}_j^{\text{T}}, \text{del-fee})$  and  $\text{Ver}(\text{addr}_i^{\text{T}}, ([v, \text{fee}^{\text{T}}], \text{addr}_i^{\text{T}}, [\text{addr}_j^{\text{T}}, \text{addr}_M^{\text{T}}], \text{id}), \sigma_i^{\text{T}})$ ,

- (ACCEPT,  $P_i$ , fee) has not been sent yet to  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ ,

then send (ACCEPT,  $P_i$ , fee) to  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ .

6. Upon receiving any other command  $m$  on the adversarial interface of  $\mathcal{F}_{\text{T-LEDGER}}$ ,  $\text{Sim}_M$  forwards it to  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ .

We now argue that if the real and the ideal world are distinguishable, then there exists an adversary that breaks either the security of the signature scheme or the security of  $\mathcal{F}_{\text{T-LEDGER}}$ . There are few cases that could make the simulator to fail, thus allowing the environment to distinguish the two worlds. We now summarize these cases and argue that none of them occur with more than negligible probability.

- In step 4, the simulator does not have the secret key  $\text{sk}_i^T$  to generate the signature  $\sigma_i^T$ . We recall that  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  forwards the command REQ-TRX to  $\mathcal{A}$  only if a party  $P_i$  is querying  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  with either (SUBMIT-DELEGATION, sid,  $\text{tx}^T$ ) or (SUBMIT, sid,  $\text{tx}_M$ ), with  $(P_i, \text{addr}_i) \in \mathcal{T}^T$ .

If that is not the case, (implying that the calling party is malicious and has generated his own address,) then we do not have the secret key. However, we do not need it, as  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  would ignore the message for the reason above.

- The simulator computes an address  $\text{addr}$  (upon receiving the command (REGISTER,  $P_i$ )), and there already exists an entry  $(P_j, \text{addr})$  in  $\mathcal{T}^T$ . If this happens with non-negligible probability, then we can construct a reduction to the signature scheme. The reduction takes as input a signature verification key  $\text{addr}$  for  $\Sigma$ , and computes  $(s, v) \stackrel{\$}{\leftarrow} \text{Kgen}(1^\lambda)$ . If  $\text{addr} \neq v$ , the reduction aborts, and otherwise computes and outputs  $\text{Sign}(s, 0^\lambda)$ .

Since, by contradiction,  $\text{addr} = v$  with non-negligible probability, and the reduction has never queried the signing oracle, then we can claim that we have broken the security of  $\Sigma$ .

- Let  $\text{state}^{\text{ideal}}$  be the output obtained when an honest party  $P_i$  queries  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$  with the input (READ, sid, SC). Let  $\text{state}^{\text{real}}$  be the output that an honest  $P_i$  would obtain when running  $\Pi^{\text{Token}}$  with the same input. The simulation fails if  $\text{state}^{\text{real}}$  contains  $\text{tx} = (v, \text{addr}_i^T, \text{addr}_j^T, \text{id}^T)$  and  $\text{state}^{\text{ideal}}$  does not. Note that if  $\text{tx} = (v, \text{addr}_i^T, \text{addr}_j^T, \text{id}^T)$  appears in  $\text{state}^{\text{real}}$ , it means that the state of  $\mathcal{F}_{\text{T-LEDGER}}$  contains a transaction whose payload contains  $(\text{tx}, \sigma_i^T)$  such that  $\text{Ver}(\text{addr}_i^T, \text{tx}, \sigma_i^T) = 1$ .

This can happen in the following two scenarios:

- The simulator has never signed  $\text{tx}$ . In this case, we make a reduction to the security of  $\Sigma$ , since the reduction would not need to query the signing oracle with the input  $\text{tx}$ .
- The simulator has signed  $\text{tx}$ . This means that  $\text{tx}$  was submitted by  $P_i$  through  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ 's interface. By construction of  $\mathcal{F}_{\text{T-LEDGER}}$  and  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , if a transaction is accepted by  $\mathcal{F}_{\text{LEDGER}}^{\text{EET}}$ , it must be accepted by  $\mathcal{F}_{\text{T-LEDGER}}$  as well. If this is not the case, then we can make a reduction to the security of  $\mathcal{F}_{\text{T-LEDGER}}$ .

The opposite scenario could also occur; that is,  $\text{state}^{\text{ideal}}$  contains  $\text{tx} = (v, \text{addr}_i^T, \text{addr}_j^T, \text{id}^T)$  and  $\text{state}^{\text{real}}$  does not. This can happen in the case where the simulator has already sampled the identifier  $\text{id}^T$  for a previous transaction. However, the probability that this event occurs is negligible. As before, the only other reason we might have such a situation is if the security

of  $\mathcal{F}_{\text{T-LEDGER}}$  has been compromised. In this case, we could again make a reduction to the security of  $\mathcal{F}_{\text{T-LEDGER}}$ .

□

## 9 Implementation, Benchmarks, and Comparisons

We implemented our EET via an Ethereum smart contract, measured its gas consumption, and compared it with other approaches. Our EET conforms to the ERC-20 standard. In its testing mode, our contract has the added functionality of allowing unlimited minting of new tokens by any account. This feature is included for ease of testing rather than for actual use, and should be disabled when the EET is deployed in production. We wrote our contracts in Solidity 0.6.10, and tested them using Hardhat 2.0.8, a Javascript and TypeScript framework for Ethereum smart contract development and testing. We tested only on a locally-running test network, not on any live public network; however, since the Hardhat test node is a faithful implementation of the Ethereum protocol and virtual machine, this should not affect the amount of gas used on any given contract invocation.

To compare the gas usage of our EET with that of the Gas Station Network (GSN), we deployed the GSN infrastructure contracts (most notably the `RelayHub` contract) to our local test network and ran a local relay server. The OpenGSN project provides a testing infrastructure that automatically deploys the required contracts and runs a local relay server; our tests used version 2.1.0 of the OpenGSN repository. The experiment code itself is written in TypeScript, using the Ethers 5.0.26 library for blockchain and contract interaction.

For each of our evaluation and comparison experiments below, we first select sender, receiver, and other relevant addresses randomly (without replacement) from a pool of 20 addresses. After executing the relevant transaction, we record the amount of gas consumed by the transaction. For validation purposes, we also record the Ether and token balances of each address before and after the transaction, to ensure that the correct amounts are transferred. Each experiment was run 1000 times, selecting a new set of addresses for each run.

### 9.1 EET vs. GSN vs. Standard ERC-20 token

For our first comparison, we ran the following three experiments:

- **Self-funded token transactions:** These experiments test the gas usage of typical, non-delegated use of our EET contract, i.e. by a user that does not want to interact with the delegation server (denoted  $M$  in our formalization) and has his own sufficiently funded Ethereum address. The sending address transfers some amount of tokens to the receiving address, submitting the transaction themselves and using their own Ether to pay the transaction fee. In this case, the relevant addresses are the sender and the receiver.
- **Delegated token transactions:** These experiments test the gas usage of our EET delegation mechanism. The sending address transfers some amount of tokens to the receiving address, but a third delegate address ( $M$ , which may be another user, another address owned by the same user, or a dedicated server,) submits the transaction and pays the ether fee, automatically (by contract conversion and execution) receiving an equivalent amount of tokens

from the sending address in the process. In this case, the relevant addresses are the sender, the receiver, and the delegate.

- GSN token transactions** These experiments test the gas usage of delegation through the GSN. The sending address transfers some amount of tokens to the receiving address, but delegates to the locally-running relay server, which submits the transaction and pays the ether fee, receiving a repayment of ether from the token contract (indirectly, from the token contract’s deposit with the `RelayHub` contract). The `EETPaymaster` contract then extracts an equivalent token fee from the sender. In this case, the relevant addresses are the sender, the receiver, and the relay server address. However, we only have control over the sender and receiver addresses; the relay server’s address is determined by the GSN testing infrastructure and cannot be easily changed.

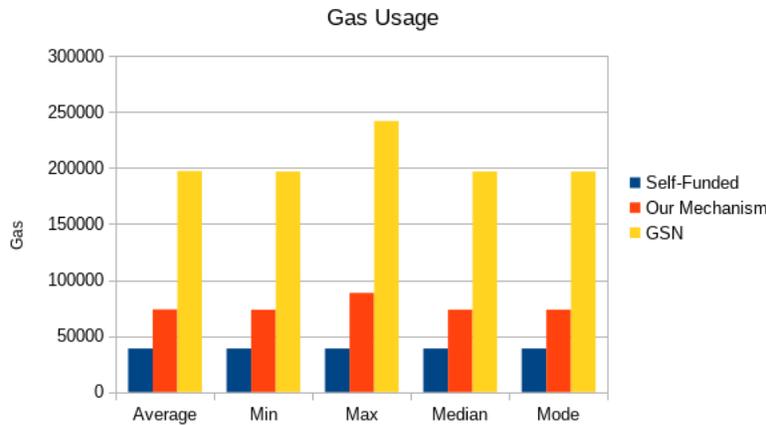


Figure 7: EETs vs GSNs

The results of our experiments are summarized in Figure 7. As one can observe, using the etherless functionality (delegation mechanism) of our contract consumes less than twice the gas of a standard self-funded token transaction, which we believe is a reasonable compromise for the added user experience. In contrast, using the GSN incurs a 4-5x increase in gas usage as compared to a self-funded transaction. This is the cost of the complexity of the GSN, a cost that is very unattractive for projects that do not require the extreme decentralization and genericity of the GSN.

## References

- [AB20] Ahmed Al-Balaghi. The state of meta transactions - 2020. <https://medium.com/biconomy/the-state-of-meta-transactions-2020-506840e37e75>, 2020.
- [BGK<sup>+</sup>18] Christian Badertscher, Peter Gazi, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 913–930. ACM Press, October 2018.

- [BMTZ17] Christian Badertscher, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. Bitcoin as a transaction ledger: A composable treatment. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 324–356. Springer, Heidelberg, August 2017.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.
- [Can03] Ran Canetti. Universally composable signatures, certification and authentication. Cryptology ePrint Archive, Report 2003/239, 2003. <https://eprint.iacr.org/2003/239>.
- [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.
- [DPS19] Phil Daian, Rafael Pass, and Elaine Shi. Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake. In Ian Goldberg and Tyler Moore, editors, *FC 2019*, volume 11598 of *LNCS*, pages 23–41. Springer, Heidelberg, February 2019.
- [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.
- [GKL17] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol with chains of variable difficulty. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 291–323. Springer, Heidelberg, August 2017.
- [Gri18] Austin Griffith. Ethereum meta transactions, 2018. [https://medium.com/@austin\\_48503/ethereum-meta-transactions-90ccf0859e84](https://medium.com/@austin_48503/ethereum-meta-transactions-90ccf0859e84).
- [gsn] Ethereum gas station network (gsn) documentation. <https://docs.opengsn.org/>.
- [Kia] Aggelos Kiayias. Babel fees-denominating transaction costs in native tokens. <https://iohk.io/en/blog/posts/2021/02/25/babel-fees/>.
- [KKKZ19] Thomas Kerber, Aggelos Kiayias, Markulf Kohlweiss, and Vassilis Zikas. Ouroboros cryptsinous: Privacy-preserving proof-of-stake. In *2019 IEEE Symposium on Security and Privacy*, pages 157–174. IEEE Computer Society Press, May 2019.
- [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.
- [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.

- [PS17] Rafael Pass and Elaine Shi. The sleepy model of consensus. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part II*, volume 10625 of *LNCS*, pages 380–409. Springer, Heidelberg, December 2017.
- [PSs17] Rafael Pass, Lior Seeman, and abhi shelat. Analysis of the blockchain protocol in asynchronous networks. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 643–673. Springer, Heidelberg, April / May 2017.
- [VB15] Fabian Vogelsteller and Vitalik Buterin. Eip 20: Erc-20 token standard. 2015. <https://eips.ethereum.org/EIPS/eip-20>.

## A Native vs. Contract-based Tokens

The experiments discussed below, use the same software and infrastructure setup as the first set (but different contracts).

- **Overhead of Native Tokens:** Adding native support for tokens on a cryptocurrency blockchain following the PPU principle means that every (even non-token) transaction processing will be slightly more (gas-)expensive than a transaction that does not support tokens. The reason is that miners/minters will at the very least need to check whether a transaction is native cryptocurrency (in which case it is added to a block as is) or a token transaction (in which case they will need to calculate if they are willing to fund its fees and compute the modified transaction to send to the network. As discussed, estimating this overhead in currently infeasible in lack of a relevant platform. Instead, here we attempt a lower-bound of this overhead, if it would be implemented in Ethereum. To this direction we implemented a simple contract `IfNoop` which performs a conditional branch on the value of an input byte—corresponding to the check of whether it is a token or native cryptocurrency transaction—and then exits in either case of the branch. This approximates the overhead of a single ‘if’ statement, followed by native execution of either an ether or token transfer. We also implemented an equivalent contract, `IfNoopYul`, in Yul, which omits the overhead of setting up the Solidity runtime and performing method dispatch, and is therefore potentially a closer approximation of the true overhead.
- **Overhead of Smart-Contract Tokens:** Our second experiment considers a contract `IfFull` which performs a conditional branch on the value of an input byte, and in one case of the branch transfers the ether value of the calling transaction to an address specified in the remainder of the input data. This approximates the overhead of a contract implementing a rough equivalent of native tokens, i.e. handling both ether and token transfers. (The case of transferring tokens is already simulated by the self-funded transaction experiments above; the leading ‘if’ statement can be assumed to be simulated by the Solidity method dispatch at the beginning of the contract execution.) As above, to get a closer estimate we also implemented a contract `IfFullYul` in the lower level EVM language Yul with the same functionality as `IfFull`.

For each of the above contracts, we submitted identical input (a 1 byte, indicating an ether transfer for the contracts that perform it, and a fixed address to transfer to) 100 times over,

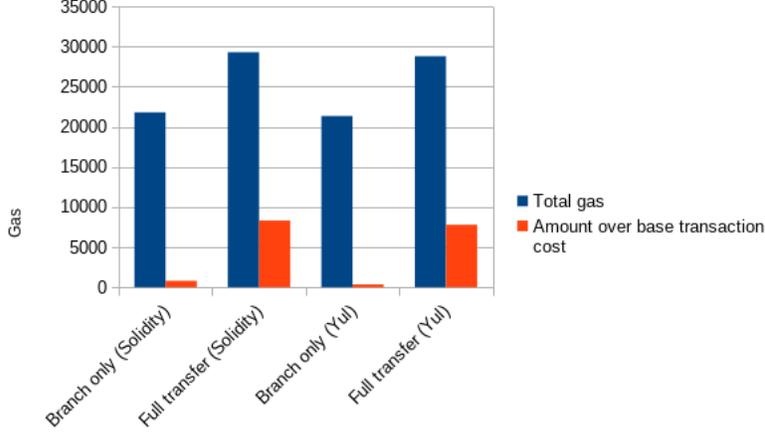


Figure 8: Native vs Smart-Contract-Simulated Tokens

measuring the gas usage for each. Unsurprisingly, since the contracts are deterministic and do not store or modify any state, each has a constant gas usage. Our experiments summarized, in Figure 8, demonstrate that even when only charging for the if-brach in a native implementation—which is clearly a favorable lower-bound on the overhead of every transaction on a native-tokens-enabled blockchain—the gas overhead of emulating tokens via a smart contract is  $\approx 33\%$ . We believe that this overhead is acceptable given the functionality and adaptability offered by smart contracts as opposed to natively-hardwired validation.

## B Functionalities with Dynamic Party Sets

UC provides support for functionalities in which the set of parties that might interact with the functionality is dynamic. We make this explicit by means of the following mechanism, (which we describe almost verbatim from [BMTZ17, Sec. 3.1]): 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:

- Upon receiving (REGISTER, 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 (REGISTER, sid,  $p_i$ ) to the caller.
- Upon receiving (DE-REGISTER, sid) from some party  $P_i \in \mathcal{P}$ , the functionality updates  $\mathcal{P} := \mathcal{P} \setminus \{P_i\}$  and returns (DE-REGISTER, sid,  $P_i$ ) to  $P_i$ .
- Upon receiving (IS-REGISTERED, sid) from some party  $P_i$ , return (REGISTER, sid,  $b$ ) to the caller, where the bit  $b$  is 1 if and only if  $P_i \in \mathcal{P}$ .
- Upon receiving (GET-REGISTERED, sid) from  $\mathcal{A}$ , the functionality returns the response (GET-REGISTERED, sid,  $\mathcal{P}$ ) to  $\mathcal{A}$ .

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 UC functionalities to register with them. Concretely, global setups include, in addition to the above party registration instructions, two registration/de-registration instructions for functionalities:

- Upon receiving  $(\text{REGISTER}, \text{sid}_G)$  from a functionality  $F$  with session-id  $\text{sid}$ , update  $F := F \cup \{(F, \text{sid})\}$ .
- Upon receiving  $(\text{DE-REGISTER}, \text{sid}_G)$  from a functionality  $F$  with session-id  $\text{sid}$ , update  $F := F \setminus \{(F, \text{sid})\}$ .
- Upon receiving  $(\text{GET-REGISTERED}_F, \text{sid}_G)$  from  $\mathcal{A}$ , return  $(\text{GET-REGISTERED}_F, \text{sid}_G, F)$  to  $\mathcal{A}$ .

We use the expression  $\text{sid}_G$  to refer to the encoding of the session identifier of global setups. By default (and if not otherwise stated), the above four (or, in the case of global setups, seven) instructions will be part of the code of all ideal functionalities considered in this work. However, to keep the description simple, we will omit these instructions from the formal descriptions unless deviations are defined.

## C Modeling Time and Clock-dependent Protocol Execution

Katz et al. [KMTZ13] proposed a methodology for casting synchronous protocols in UC by assuming they have access to an ideal functionality  $\mathcal{F}_{\text{CLOCK}}$ , *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  $(\text{CLOCK-READ}, \text{sid}_C)$  to  $\mathcal{F}_{\text{CLOCK}}$ . This value is updated only once all honest parties send the clock a  $(\text{CLOCK-UPDATE}, \text{sid}_C)$  command. We lift this idea to a shared setup: the global clock functionality  $\mathcal{F}_{\text{CLOCK}}$  is a shared clock that may interact with more than one protocol session. The global clock provides a means for parties to synchronize each of their sessions.<sup>14</sup> The clock can also be used as a local (not shared) hybrid functionality, in which case the number of sessions it will synchronize is simply one. The description is given in Figure 9.

Given a clock, the authors of [KMTZ13] describe how synchronous protocols can maintain their necessary round structure in UC: for every round  $\rho$ , each party first executes all of its round- $\rho$  instructions, and then sends the clock a  $\text{CLOCK-UPDATE}$  command. Subsequently, whenever activated, it sends the clock a  $\text{CLOCK-READ}$  command and does not advance to round  $\rho + 1$  until it sees that the clock’s variable has been updated. This ensures that no honest party will start round  $\rho + 1$  before every honest party has completed round  $\rho$ . In [KZZ16], this idea was transferred 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 this is the model we will also use here.

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

---

<sup>14</sup>The functionality presented here is different from shared clock functionalities used in prior work. We believe that the version here is closer to the spirit of the GUC/EUC version of UC.

### Functionality $\mathcal{F}_{\text{clock}}$

The functionality manages the set  $\mathcal{P}$  of registered identities, i.e. parties  $P = (\text{pid}, \text{sid})$ . It also manages the set  $F$  of functionalities (together with their session identifiers). Initially,  $\mathcal{P} := \emptyset$  and  $F := \emptyset$ .

For each session  $\text{sid}$ , the clock maintains a variable  $\tau_{\text{sid}}$ . For each identity  $P := (\text{pid}, \text{sid}) \in \mathcal{P}$ , it maintains a variable  $d_P$ . For each pair  $(\mathcal{F}, \text{sid}) \in F$ , it maintains a variable  $d_{(\mathcal{F}, \text{sid})}$ . All integer variables are initially 0.

*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 a session  $\text{sid}$  such that  $(\mathcal{F}, \text{sid}) \in F$ , set  $d_{(\mathcal{F}, \text{sid})} := 1$ , execute *Round-Update*, and return  $(\text{CLOCK-UPDATE}, \text{sid}_C, \mathcal{F})$  to the sending instance of  $\mathcal{F}$ .
- Upon receiving  $(\text{CLOCK-READ}, \text{sid}_C)$  from any participant (including the environment on behalf of a party, the adversary, or any ideal — shared or local — functionality), return  $(\text{CLOCK-READ}, \text{sid}_C, \tau_{\text{sid}})$  to the requestor, where  $\text{sid}$  is the sid of the calling instance.

*Procedure Round-Update:* For each session  $\text{sid}$  do: If  $d_{(\mathcal{F}, \text{sid})} := 1$  for all  $\mathcal{F} \in F$  and  $d_P = 1$  for all honest parties  $P = (\cdot, \text{sid}) \in \mathcal{P}$ , then set  $\tau_{\text{sid}} := \tau_{\text{sid}} + 1$ , and reset  $d_{(\mathcal{F}, \text{sid})} := 0$  and  $d_P := 0$  for all parties  $P = (\cdot, \text{sid}) \in \mathcal{P}$ .

Figure 9: The shared/global clock functionality. We assume lazy creation of variables, i.e. a variable is only created once it is needed.

sends the clock an update command. We capture this property in a generic manner in [Definition C.1](#) (the content of this section and of [Appendix D](#) are taken almost verbatim from [\[BMTZ17\]](#)).

To follow the definition, 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 the party either sending a message to some of its hybrid functionalities, sending an output to the environment, or sending no message at all. In any of these cases, the party loses the activation.<sup>15</sup>

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 in which 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 session  $\text{sid}$  in total,  $\vec{\mathcal{I}}_H$  is a vector of the form  $((x_1, id_1), \dots, (x_m, id_m))$ , where  $x_i$  is the  $i$ -th input that was given in this execution, and  $id_i$  is the corresponding identity (i.e.  $id_i = (\text{pid}_i, \text{sid})$  for some bitstring  $\text{pid}$ ) of the party that received this input in this session. 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 at which each input was given. If the timed honest-input sequence of an execution is  $\vec{\mathcal{I}}_H^T = ((x_1, id_1, \tau_1), \dots, (x_m, id_m, \tau_m))$ , this means that  $((x_1, id_1), \dots, (x_m, id_m))$  is the honest-input sequence corresponding to this execution, and for each  $i \in [m]$ ,  $\tau_i$  is the time of the global clock when input  $x_i$  was handed to  $id_i$ .

**Definition C.1.** A  $\mathcal{F}_{\text{clock}}$ -hybrid protocol  $\Pi$  has a *predictable synchronization pattern* iff there exists an algorithm  $\text{predict-time}_\Pi(\cdot)$  such that, for any possible execution of  $\Pi$  in a session  $\text{sid}$  (i.e. for any adversary and environment and any choice of random coins), the following holds: if  $\vec{\mathcal{I}}_H^T = ((x_1, id_1, \tau_1), \dots, (x_m, id_m, \tau_m))$  is the corresponding timed honest-input sequence for this

<sup>15</sup>In the latter case the activation goes to the environment by default.

Ledger Element	Description
$\mathcal{P}, \mathcal{H}, \mathcal{P}_{DS}$	The party sets and categories: Registered, honest, and honest-but-desynchronized, respectively.
$\vec{I}_H^T$	The timed honest-input sequence.
predict-time	The function to predict the real-world time advancement.
state	The ledger state, i.e. a sequence of blocks containing the content.
buffer	The buffer of submitted input values.
$\text{pt}_i, \text{state}_i$	The pointer of party $P_i$ into state <b>state</b> . This prefix is denoted <b>state</b> <sub><math>i</math></sub> for brevity.
$\vec{\tau}_{\text{state}}$	A vector containing for each state block the time when the block added to the ledger state.
$\tau_L$	The current time as reported by the clock.
NxtBC	Stores the current adversarial suggestion for extending the ledger state.
Validate	Decides on the validity of a transaction with respect to the current state. Used to clean the buffer of transactions.
ExtendPolicy	The function that specifies the ledger's guarantees in extending the ledger state (e.g., speed, content etc.).
Blockify	The function to format the ledger state output.
windowSize	The window size (number of blocks) of the sliding window.
Delay	A general delay parameter for the time it takes for a newly joining (after the onset of the computation) miner to become synchronized.

Figure 10: Overview of main ledger elements such as parameters and state variables.

session, then for any  $i \in [m - 1]$  :

$$\text{predict-time}_{\Pi}((x_1, id_1, \tau_1), \dots, (x_i, id_i, \tau_i)) = \tau_{i+1},$$

where  $\tau_{i+1}$  is the clock time for this session (cf. Figure 9).

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

## D The Basic Transaction-Ledger Functionality

## Functionality $\mathcal{F}_{\text{ledger}}$

**General:** The functionality is parametrized by four algorithms `Validate`, `ExtendPolicy`, `Blockify`, and `predict-time`, along with two parameters  $\text{windowSize}, \text{Delay} \in \mathbb{N}$ . The functionality manages the variables `state`, `NxtBC`, `buffer`,  $\tau_L$ , and  $\bar{\tau}_{\text{state}}$ , as described above. Initially,  $\text{state} := \bar{\tau}_{\text{state}} := \text{NxtBC} := \varepsilon$ ,  $\text{buffer} := \emptyset$ ,  $\tau_L = 0$ .

For each party  $P_i \in \mathcal{P}$  the functionality maintains a pointer  $\text{pt}_i$  (initially set to 1) and a current-state view  $\text{state}_i := \varepsilon$  (initially set to empty). The functionality keeps track of the timed honest-input sequence  $\bar{\mathcal{I}}_H^T$  (initially  $\bar{\mathcal{I}}_H^T := \varepsilon$ ).

**Party management:** 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  $\emptyset$ . If a new honest party is already registered with the clock at the time it is registered with the ledger, it is added to the party sets  $\mathcal{H}$  and  $\mathcal{P}$ , and the time of registration is recorded. If the current time is  $\tau_L > 0$ , the new party is also added to  $\mathcal{P}_{DS}$ . Similarly, when a party is deregistered, it is removed from  $\mathcal{P}$ , and therefore also from  $\mathcal{P}_{DS}$  and  $\mathcal{H}$ . The ledger maintains the invariant that it is registered (as a functionality) with the clock whenever  $\mathcal{H} \neq \emptyset$ . A party is considered fully registered if it is registered with both the ledger and the clock.

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

1. Let  $\hat{\mathcal{P}} \subseteq \mathcal{P}_{DS}$  denote the set of desynchronized honest parties that have been registered (continuously, with both ledger and clock) since time  $\tau' < \tau_L - \text{Delay}$ . Set  $\mathcal{P}_{DS} := \mathcal{P}_{DS} \setminus \hat{\mathcal{P}}$ . On the other hand, for any synchronized party  $P \in \mathcal{H} \setminus \mathcal{P}_{DS}$ , if  $P$  is not registered to the clock, then  $\mathcal{P}_{DS} \cup \{P\}$ .
2. If  $I$  was received from an honest party  $P_i \in \mathcal{P}$ :
  - (a) Set  $\bar{\mathcal{I}}_H^T := \bar{\mathcal{I}}_H^T \parallel (I, P_i, \tau_L)$ .
  - (b) Compute  $\bar{N} = (\bar{N}_1, \dots, \bar{N}_\ell) := \text{ExtendPolicy}(\bar{\mathcal{I}}_H^T, \text{state}, \text{NxtBC}, \text{buffer}, \bar{\tau}_{\text{state}})$  and, if  $\bar{N} \neq \varepsilon$ , set  $\text{state} := \text{state} \parallel \text{Blockify}(\bar{N}_1) \parallel \dots \parallel \text{Blockify}(\bar{N}_\ell)$  and  $\bar{\tau}_{\text{state}} := \bar{\tau}_{\text{state}} \parallel \tau_L^\ell$ , where  $\tau_L^\ell = \tau_L \parallel \dots \parallel \tau_L$ .
  - (c) For each  $\text{BTX} \in \text{buffer}$ : if  $\text{Validate}(\text{BTX}, \text{state}, \text{buffer}) = 0$ , then delete  $\text{BTX}$  from  $\text{buffer}$ .
  - (d) If there exists  $P_j \in \mathcal{H} \setminus \mathcal{P}_{DS}$  such that  $|\text{state}| - \text{pt}_j > \text{windowSize}$  or  $\text{pt}_j < |\text{state}_j|$ , then set  $\text{pt}_k := |\text{state}|$  for all  $P_k \in \mathcal{H} \setminus \mathcal{P}_{DS}$ .
3. Depending on the input  $I$  and the ID of the sender, execute the respective code:
  - *Submitting a transaction:*  
If  $I = (\text{SUBMIT}, \text{sid}, \text{tx})$  and  $I$  was 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  $\text{txid}$  and set  $\text{BTX} := (\text{tx}, \text{txid}, \tau_L, P_i)$ .
    - (b) If  $\text{Validate}(\text{BTX}, \text{state}, \text{buffer}) = 1$ , then  $\text{buffer} := \text{buffer} \cup \{\text{BTX}\}$ .
    - (c) Send  $(\text{SUBMIT}, \text{BTX})$  to  $\mathcal{A}$ .
  - *Reading the state:*  
If  $I = (\text{READ}, \text{sid})$  is received from a fully registered party  $P_i \in \mathcal{P}$ , then set  $\text{state}_i := \text{state}_{\min\{\text{pt}_i, |\text{state}|\}}$  and return  $(\text{READ}, \text{sid}, \text{state}_i)$  to the requestor. If the requestor is  $\mathcal{A}$ , then send  $(\text{state}, \text{buffer}, \bar{\mathcal{I}}_H^T)$  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  $\bar{\mathcal{I}}_H^T$  as above)  $\text{predict-time}(\bar{\mathcal{I}}_H^T) = \hat{\tau} > \tau_L$ , then send  $(\text{CLOCK-UPDATE}, \text{sid}_C)$  to  $\mathcal{F}_{\text{clock}}$ . Otherwise, send  $I$  to  $\mathcal{A}$ .
  - *The adversary proposing the next block:*  
If  $I = (\text{NEXT-BLOCK}, \text{hFlag}, (\text{txid}_1, \dots, \text{txid}_\ell))$  is sent from the adversary, update `NxtBC` as follows:
    - (a) Set  $\text{listOfTxdid} \leftarrow \varepsilon$ .
    - (b) For  $i = 1, \dots, \ell$ , if there exists  $\text{BTX} := (x, \text{txid}, \text{minerID}, \tau_L, P_i) \in \text{buffer}$  with ID  $\text{txid} = \text{txid}_i$ , then set  $\text{listOfTxdid} := \text{listOfTxdid} \parallel \text{txid}_i$ .
    - (c) Finally, set  $\text{NxtBC} := \text{NxtBC} \parallel (\text{hFlag}, \text{listOfTxdid})$  and output  $(\text{NEXT-BLOCK}, ok)$  to  $\mathcal{A}$ .
  - *The adversary setting state-slackness:*  
If  $I = (\text{SET-SLACK}, (P_{i_1}, \widehat{\text{pt}}_{i_1}), \dots, (P_{i_\ell}, \widehat{\text{pt}}_{i_\ell}))$ , with  $\{P_{i_1}, \dots, P_{i_\ell}\} \subseteq \mathcal{H} \setminus \mathcal{P}_{DS}$  is received from the adversary  $\mathcal{A}$  do the following:
    - (a) If for all  $j \in [\ell] : |\text{state}| - \widehat{\text{pt}}_{i_j} \leq \text{windowSize}$  and  $\widehat{\text{pt}}_{i_j} \geq |\text{state}_{i_j}|$ , set  $\text{pt}_{i_1} := \widehat{\text{pt}}_{i_1}$  for every  $j \in [\ell]$  and return  $(\text{SET-SLACK}, ok)$  to  $\mathcal{A}$ .
    - (b) Otherwise, set  $\text{pt}_{i_j} := |\text{state}|$  for all  $j \in [\ell]$ .
  - *The adversary setting the state for desynchronized parties:*  
If  $I = (\text{DESYNC-STATE}, (P_{i_1}, \text{state}'_{i_1}), \dots, (P_{i_\ell}, \text{state}'_{i_\ell}))$ , with  $\{P_{i_1}, \dots, P_{i_\ell}\} \subseteq \mathcal{P}_{DS}$  is received from the adversary  $\mathcal{A}$ , set  $\text{state}_{i_j} := \text{state}'_{i_j}$  for each  $j \in [\ell]$  and return  $(\text{DESYNC-STATE}, ok)$  to  $\mathcal{A}$ .

Figure 11: The ledger functionality. We write  $[n]$  to denote the set  $\{1, \dots, n\}$ .