



A Formal Model of Algorand Smart Contracts

Bartoletti, Massimo; Bracciali, Andrea; Lepore, Cristian; Scalas, Alceste; Zunino, Roberto

Published in:
Financial Cryptography and Data Security

Link to article, DOI:
[10.1007/978-3-662-64322-8_5](https://doi.org/10.1007/978-3-662-64322-8_5)

Publication date:
2021

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Bartoletti, M., Bracciali, A., Lepore, C., Scalas, A., & Zunino, R. (2021). A Formal Model of Algorand Smart Contracts. In *Financial Cryptography and Data Security* (pp. 93-114). Springer. https://doi.org/10.1007/978-3-662-64322-8_5

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A formal model of Algorand smart contracts

Massimo Bartoletti¹, Andrea Bracciali², Cristian Lepore²,
Alceste Scalas³, Roberto Zunino⁴

¹ Università degli Studi di Cagliari, Cagliari, Italy

² Stirling University, Stirling, UK

³ Technical University of Denmark, Lyngby, Denmark

⁴ Università degli Studi di Trento, Trento, Italy

Abstract. We develop a formal model of Algorand stateless smart contracts (stateless ASC1). We exploit our model to prove fundamental properties of the Algorand blockchain, and to establish the security of some archetypal smart contracts. While doing this, we highlight various design patterns supported by Algorand. We perform experiments to validate the coherence of our formal model w.r.t. the actual implementation.

1 Introduction

Smart contracts are agreements between two or more parties that are automatically enforced without trusted intermediaries. Blockchain technologies reinvented the idea of smart contracts, providing trustless environments where they are incarnated as computer programs. However, writing secure smart contracts is difficult, as witnessed by the multitude of attacks on smart contracts platforms (notably, Ethereum) — and since smart contracts control assets, their bugs may directly lead to financial losses.

Algorand [22] is a late-generation blockchain that features a set of interesting features, including high-scalability and a no-forking consensus protocol based on Proof-of-Stake [9]. Its smart contract layer (ASC1) aims to mitigate smart contract risks, and adopts a non-Turing-complete programming model, natively supporting atomic sets of transactions and user-defined assets. These features make it an intriguing smart contract platform to study.

The official specification and documentation of ASC1 consists of English prose and a set of templates to assist programmers in designing their contracts [1, 5]. This conforms to standard industry practices, but there are two drawbacks:

1. Algorand lacks a mathematical model of contracts and transactions suitable for formal reasoning on their behaviour, and for the verification of their properties. Such a model is needed to develop techniques and tools to ensure that contracts are correct and secure;
2. furthermore, even preliminary informal reasoning on non-trivial smart contracts can be challenging, as it may require, in some corner cases, to resort to experiments, or direct inspection of the platform source code.

Given these drawbacks, we aim at developing a formal model that:

- o1. is high-level enough to simplify the design of Algorand smart contracts and enable formal reasoning about their security properties;
- o2. expresses Algorand contracts in a simple declarative language, similar to PyTeal (the official Python binding for Algorand smart contracts) [7];
- o3. provides a basis for the automatic verification of Algorand smart contracts.

Contributions. This paper presents:

- a *formal model of stateless ASC1* providing a solid theoretical foundation to Algorand smart contracts (§2). Such a model formalises both Algorand accounts and transactions (§2.1–§2.4, §2.6), and smart contracts (§2.5);
- a validation of our model through experiments [6] on the Algorand platform;
- the formalisation and proof of some *fundamental properties of the Algorand state machine*: no double spending, determinism, value preservation (§2.7);
- an *analysis of Algorand contract design patterns* (§3.2), based on several non-trivial contracts (covering both standard use cases, and novel ones). Quite surprisingly, we show that stateless contracts are expressive enough to encode arbitrary finite state machines;
- the proof of relevant *security properties of smart contracts* in our model;
- a *prototype tool* that compiles smart contracts (written in our formal declarative language) into executable TEAL code (§4).

Our formal model is faithful to the actual ASC1 implementation; by objectives o1–o3, it strives at being high-level and simple to understand, while covering the most commonly used primitives and mechanisms of Algorand, and supporting the specification and verification of non-trivial smart contracts (§3.2, §4). To achieve these objectives, we introduce minor high-level abstractions over low-level details: e.g., since TEAL code has the purpose of accepting or rejecting transactions, we model it using expressions that evaluate to *true* or *false* (similarly to PyTEAL); we also formalise different transaction types by focusing on their function, rather than their implementation. Our objectives imply that we do *not* aim at covering all the possible TEAL contracts with bytecode-level accuracy, and our Algorand model is *not* designed as a full low-level formalisation of the behavior of the Algorand blockchain. We discuss the differences between our model and the actual Algorand platform in §5.

2 The Algorand state machine

We present our formal model of the Algorand blockchain, including its smart contracts (stateless ASC1), incrementally. We first define the basic transactions that generate and transfer assets (§2.1–§2.3), and then add atomic groups of transactions (§2.4), smart contracts (§2.5), and authorizations (§2.6). We discuss the main differences between our model and Algorand in §5.

$\mathbf{a}, \mathbf{b}, \dots$	Users (key pairs)	$T_{iv} \subseteq \mathbb{T}$	Transactions in last Δ_{max} rounds
$\mathbf{x}, \mathbf{y}, \dots \in \mathbb{X}$	Addresses	$f_{asst} \in \mathbb{A} \rightarrow \mathbb{X}$	Asset manager
$\tau, \tau', \dots \in \mathbb{A}$	Assets	$f_{lx} \in (\mathbb{X} \times \mathbb{N}) \rightarrow \mathbb{N}$	Lease map
$v, w, \dots \in 0..2^{64} - 1$	Values	$f_{frz} \in \mathbb{X} \rightarrow 2^{\mathbb{A}}$	Freeze map
$\sigma, \sigma' \in \mathbb{A} \rightarrow \mathbb{N}$	Balances	Γ, Γ', \dots	Blockchain states
$\mathbf{x}[\sigma]$	Accounts	$\models \sigma$	Valid balance
$t, t', \dots \in \mathbb{T}$	Transactions	$f_{lx}, r \models t$	Valid time constraint
e, e', \dots	Scripts	$\mathcal{W} \models \mathcal{J}, i$	Authorized transaction in group
$r, r', \dots \in \mathbb{N}$	Rounds	$\llbracket e \rrbracket_{\mathcal{J}, i}^{\mathcal{W}}$	Script evaluation

Table 1: Summary of notation.

2.1 Accounts and transactions

We use $\mathbf{a}, \mathbf{b}, \dots$ to denote public/private key pairs (k_a^p, k_a^s) . Users interact with Algorand through pseudonymous identities, obtained as a function of their public keys. Hereafter, we freely use \mathbf{a} to refer to the public or the private key of \mathbf{a} , or to the user associated with them, relying on the context to resolve the ambiguity. The purpose of Algorand is to allow users to exchange *assets* τ, τ', \dots . Besides the Algorand native cryptocurrency *Algo*, users can create custom assets.

We adopt the following notational convention:

- lowercase letters for single entities (e.g., a user \mathbf{a});
- uppercase letters for *sets* of entities (e.g., a set of users \mathbf{A});
- calligraphic uppercase letters for *sequences* of entities (e.g., list of users \mathcal{A}).

Given a sequence \mathcal{L} , we write $|\mathcal{L}|$ for its length, $set(\mathcal{L})$ for the set of its elements, and $\mathcal{L}.i$ for its i^{th} element ($i \in 1..|\mathcal{L}|$); ε denotes the empty sequence. We write:

- $\{x \mapsto v\}$ for the function mapping x to v , and having domain equal to $\{x\}$;
- $f\{x \mapsto v\}$ for the function mapping x to v , and y to $f(y)$ if $y \neq x$;
- $f\{x \mapsto \perp\}$ for the function undefined at x , and mapping y to $f(y)$ if $y \neq x$.

Accounts. An *account* is a deposit of one or more crypto-assets. We model accounts as terms $\mathbf{x}[\sigma]$, where \mathbf{x} is an *address* uniquely identifying the account, and σ is a *balance*, i.e., a finite map from assets to non-negative 64-bit integers. In the concrete Algorand, an address is a 58-characters word; for mathematical elegance, in our model we represent an address as either:

- a *single user* \mathbf{a} . Performing transactions on $\mathbf{a}[\sigma]$ requires \mathbf{a} 's authorization;
- a pair (\mathcal{A}, n) , where \mathcal{A} is a sequence of users, and $1 \leq n \leq |\mathcal{A}|$, are *multisig* (*multi-signature*) addresses. Performing transactions on $(\mathcal{A}, n)[\sigma]$ requires that at least n users out of those in \mathcal{A} grant their authorization;⁵
- a *script*⁶ e . Performing transactions on $e[\sigma]$ requires e to evaluate to *true*.

⁵ W.l.o.g., we consider a single-user address \mathbf{a} equivalent to (\mathcal{A}, n) with $\mathcal{A} = \langle \mathbf{a} \rangle$, $n = 1$.

⁶ We formalize scripts (i.e., smart contracts) later on, in §2.5.

<i>pay</i>	<i>snd</i> , <i>rcv</i> , <i>val</i> , <i>asst</i>	<i>snd</i> transfers <i>val</i> units of <i>asst</i> to <i>rcv</i> (possibly creating <i>rcv</i>)
<i>close</i>	<i>snd</i> , <i>rcv</i> , <i>asst</i>	<i>snd</i> gives <i>asst</i> to <i>rcv</i> and removes it (if <i>Algo</i> , closes <i>snd</i>)
<i>gen</i>	<i>snd</i> , <i>rcv</i> , <i>val</i>	<i>snd</i> mints <i>val</i> units of a new asset, managed by <i>rcv</i>
<i>optin</i>	<i>snd</i> , <i>asst</i>	<i>snd</i> opts in to receive units of asset <i>asst</i>
<i>burn</i>	<i>asst</i>	<i>asst</i> is removed from the creator (if sole owner)
<i>rvk</i>	<i>snd</i> , <i>rcv</i> , <i>val</i> , <i>asst</i>	<i>asst</i> 's manager transfers <i>val</i> units of <i>asst</i> from <i>snd</i> to <i>rcv</i>
<i>frz</i>	<i>snd</i> , <i>asst</i>	<i>asst</i> 's manager freezes <i>snd</i> 's use of asset <i>asst</i>
<i>unfrz</i>	<i>snd</i> , <i>asst</i>	<i>asst</i> 's manager unfreezes <i>snd</i> 's use of asset <i>asst</i>
<i>delegate</i>	<i>snd</i> , <i>asst</i> , <i>rcv</i>	<i>asst</i> 's manager delegates <i>asst</i> to new manager <i>rcv</i>

Fig. 1: Transaction types. Fields *type*, *fv*, *lv*, *lx* are common to all types.

Each balance is required to own *Algos*, have at least 100000 micro-*Algos* for each owned asset, and cannot control more than 1000 assets. Formally, we say that σ is a *valid balance* (in symbols, $\models \sigma$) when:⁷

$$\text{Algo} \in \text{dom}(\sigma) \wedge \sigma(\text{Algo}) \geq 100000 \cdot |\text{dom}(\sigma)| \wedge |\text{dom}(\sigma)| \leq 1001$$

Transactions. Accounts can append various kinds of *transactions* to the blockchain, in order to, e.g., alter their balance or set their usage policies. We model transactions as records with the structure in Fig. 1. Each transaction has a *type*, which determines which of the other fields are relevant.⁸ The field *snd* usually refers to the subject of the transaction (e.g., the *sender* in an assets transfer), while *rcv* refers to the *receiver* in an assets transfer. The fields *asst* and *val* refer, respectively, to the affected asset, and to its amount. The fields *fv* (“first valid”), *lv* (“last valid”) and *lx* (“lease”) are used to impose time constraints.

Algorand groups transactions into *rounds* $r = 1, 2, \dots$. To establish *when* a transaction t is valid, we must consider both the current round r , and a *lease map* f_{lx} binding pairs (address, lease identifier) to rounds: this is used to enforce mutual exclusion between two or more transactions (see e.g. the *periodic payment* contract in §3). Formally, we define the *temporal validity of a transaction* t by the predicate $f_{lx}, r \models t$, which holds whenever:

$$\begin{aligned} & t.fv \leq r \leq t.lv \quad \text{and} \quad t.lv - t.fv \leq \Delta_{max} \quad \text{and} \\ & (t.lx = 0 \quad \text{or} \quad (t.snd, t.lx) \notin \text{dom}(f_{lx}) \quad \text{or} \quad r > f_{lx}(t.snd, t.lx)) \end{aligned}$$

First, the current round must lie between $t.fv$ and $t.lv$, whose distance cannot exceed Δ_{max} rounds⁹. Second, t must have a null lease identifier, or the identifier has not been seen before (i.e., $f_{lx}(t.snd, t.lx)$ is undefined), or the lease has expired (i.e., $r > f_{lx}(t.snd, t.lx)$). When performed, a transaction with non-null lease identifier acquires the lease on $(t.snd, t.lx)$, which is set to $t.lv$.

⁷ Since the codomain of σ is \mathbb{N} , the balance entry $\sigma(\text{Algo})$ represents micro-*Algos*.

⁸ In Algorand, the actual behaviour of a transaction may depend on both its type and other conditions, e.g., which optional fields are set. For instance, *pay* transactions may also close accounts if the *CloseRemainderTo* field is set. For the sake of clarity, in our model we prefer to use a richer set of types; see §5 for other differences.

⁹ Δ_{max} is a global consensus parameter, set to 1000 at time of writing.

2.2 Blockchain states

We model the evolution of the Algorand blockchain as a labelled transition system. A *blockchain state* Γ has the form:

$$\mathbf{x}_1[\sigma_1] \mid \cdots \mid \mathbf{x}_n[\sigma_n] \mid r \mid T_{\mathbf{lv}} \mid f_{asst} \mid f_{lx} \mid f_{frz} \quad (1)$$

where all addresses \mathbf{x}_i are distinct, \mid is commutative and associative, and:

- r is the current round;
- $T_{\mathbf{lv}}$ is the set of performed transactions whose “last valid” time \mathbf{lv} has not expired. This set is used to avoid double spending (see Theorem 1);
- f_{asst} maps each asset to the addresses of its *manager* and *creator*;
- f_{lx} is the *lease map* (from pairs (address, integer) to integers), used to ensure mutual exclusion between transactions;
- f_{frz} is a map from addresses to sets of assets, used to *freeze assets*.

We define the *initial state* Γ_0 as $\mathbf{a}_0[\{\text{Algo} \mapsto v_0\}] \mid 0 \mid \emptyset \mid f_{asst} \mid f_{lx} \mid f_{frz}$, where $\text{dom}(f_{asst}) = \text{dom}(f_{lx}) = \text{dom}(f_{frz}) = \emptyset$, \mathbf{a}_0 is the initial user address, and $v_0 = 10^{16}$ (which is the total supply of 10 billions *Algos*).

We now formalize the ASC1 state machine, by defining how it evolves by single transactions (§2.3), and then including atomic groups of transactions (§2.4), smart contracts (§2.5), and the authorization of transactions (§2.6).

2.3 Executing single transactions

We write $\Gamma \xrightarrow{t}_1 \Gamma'$ to mean: *if* the transaction t is performed in blockchain state Γ , then the blockchain evolves to state Γ' .¹⁰ We specify the transition relation \rightarrow_1 through a set of inference rules (see Fig. 5 in the Appendix for the full definition): each rule describes the effect of a transaction t in the state Γ of eq. (1). We now illustrate all cases, depending on the transaction type ($t.\text{type}$).

When $\tau \in \text{dom}(\sigma)$, we use the shorthand $\sigma + v:\tau$ to update balance σ by adding v units to token τ ; similarly, we write $\sigma - v:\tau$ to decrease τ by v units:

$$\sigma + v:\tau \equiv \sigma\{\tau \mapsto \sigma(\tau) + v\} \quad \sigma - v:\tau \equiv \sigma\{\tau \mapsto \sigma(\tau) - v\}$$

Pay to a new account. Let $t.\text{snd} = \mathbf{x}_i$ for some $i \in 1..n$, let $t.\text{rcv} = \mathbf{y} \notin \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ (i.e., the sender account \mathbf{x} is already in the state, while the receiver \mathbf{y} is not), and let $t.\text{val} = v$. The rule has the following preconditions:

- c1. t does not cause double-spending ($t \notin T_{\mathbf{lv}}$);
- c2. the time interval of the transaction, and its lease, are respected ($f_{lx}, r \models t$);
- c3. the updated balance of \mathbf{x}_i is valid ($\models \sigma_i - v : \text{Algo}$);
- c4. the balance of the new account at address \mathbf{y} is valid ($\models \{\text{Algo} \mapsto v\}$).

¹⁰ Note that $\Gamma \xrightarrow{t}_1 \Gamma'$ does not imply that transaction t *can* be performed in Γ : in fact, t might require an authorization. We specify the required conditions in §2.6.

If these conditions are satisfied, the new state Γ' is the following:

$$\mathbf{x}_i[\sigma_i - v:\text{Algo}] \mid \mathbf{y}[\{\text{Algo} \mapsto v\}] \mid \dots \mid r \mid T_{\text{iv}} \cup \{t\} \mid f_{\text{asst}} \mid \text{upd}(f_{\text{lx}}, t) \mid f_{\text{frz}}$$

In the new state, the **Algo** balance of \mathbf{x}_i is decreased by v units, and a new account at \mathbf{y} is created, containing exactly the v units taken from \mathbf{x}_i . The balances of the other accounts are unchanged. The updated lease mapping is:

$$\text{upd}(f_{\text{lx}}, t) = \begin{cases} f_{\text{lx}}\{(t.\text{snd}, t.\text{lx}) \mapsto t.\text{lv}\} & \text{if } t.\text{lx} \neq 0 \\ f_{\text{lx}} & \text{otherwise} \end{cases}$$

Note that all transaction types check conditions **c1** and **c2** above; further, all transactions check that updated account balances are valid (as in **c3** and **c4**).

Pay to an existing account. Let $t.\text{snd} = \mathbf{x}_i$, $t.\text{rcv} = \mathbf{x}_j$, $t.\text{val} = v$, and $t.\text{asst} = \tau$. Besides the common checks, performing t requires that \mathbf{x}_j has “opted in” τ (formally, $\tau \in \text{dom}(\sigma_j)$), and τ must not be frozen in accounts \mathbf{x}_i and \mathbf{x}_j (formally, $\tau \notin f_{\text{frz}}(\mathbf{x}_i) \cup f_{\text{frz}}(\mathbf{x}_j)$). If $\mathbf{x}_i \neq \mathbf{x}_j$, then in the new state the balance of τ in \mathbf{x}_i is decreased by v units, and that of τ in \mathbf{x}_j is increased by v units:

$$\mathbf{x}_i[\sigma_i - v:\tau] \mid \mathbf{x}_j[\sigma_j + v:\tau] \mid \dots \mid r \mid T_{\text{iv}} \cup \{t\} \mid f_{\text{asst}} \mid \text{upd}(f_{\text{lx}}, t) \mid f_{\text{frz}}$$

where all accounts but \mathbf{x}_i and \mathbf{x}_j are unchanged. Otherwise, if $\mathbf{x}_i = \mathbf{x}_j$, then the balance of \mathbf{x}_i is unchanged, and the other parts of the state are as above.

Close. Let $t.\text{snd} = \mathbf{x}_i$, $t.\text{rcv} = \mathbf{x}_j \neq \mathbf{x}_i$, and $t.\text{asst} = \tau$. Performing t has two possible outcomes, depending on whether τ is **Algo** or a user-defined asset. If $\tau = \text{Algo}$, we must check that σ_i contains *only* **Algos**. If so, the new state is:

$$\mathbf{x}_j[\sigma_j + \sigma_i(\text{Algo}):\text{Algo}] \mid \dots \mid r \mid T_{\text{iv}} \cup \{t\} \mid f_{\text{asst}} \mid \text{upd}(f_{\text{lx}}, t) \mid f_{\text{frz}}$$

where the new state no longer contains the account \mathbf{x}_i , and all the **Algos** in \mathbf{x}_i are transferred to \mathbf{x}_j . Instead, if $\tau \neq \text{Algo}$, performing t requires to check only that \mathbf{x}_i actually contains τ , and that \mathbf{x}_j has “opted in” τ . Further, τ must not be frozen for addresses \mathbf{x}_i and \mathbf{x}_j , i.e. $\tau \notin f_{\text{frz}}(\mathbf{x}_i) \cup f_{\text{frz}}(\mathbf{x}_j)$. The new state is:

$$\mathbf{x}_i[\sigma_i\{\tau \mapsto \perp\}] \mid \mathbf{x}_j[\sigma_j + \sigma_i(\tau):\tau] \mid \dots \mid r \mid T_{\text{iv}} \cup \{t\} \mid f_{\text{asst}} \mid \text{upd}(f_{\text{lx}}, t) \mid f_{\text{frz}}$$

where τ is removed from \mathbf{x}_i , and all the units of τ in \mathbf{x}_i are transferred to \mathbf{x}_j .

Gen. Let $t.\text{snd} = \mathbf{x}_i$, $t.\text{rcv} = \mathbf{x}_j$, and $t.\text{val} = v$. Performing t requires that \mathbf{x}_i has enough **Algos** to own another asset, i.e. $\models \sigma_i\{\tau \mapsto v\}$, where τ is the (fresh) identifier of the new asset. In the new state, the balance of \mathbf{x}_i is extended with $\{\tau \mapsto v\}$, and f_{asst} is updated, making \mathbf{x}_j the manager of τ . The new state is:

$$\mathbf{x}_i[\sigma_i\{\tau \mapsto v\}] \mid \dots \mid r \mid T_{\text{iv}} \cup \{t\} \mid f_{\text{asst}}\{\tau \mapsto (\mathbf{x}_j, \mathbf{x}_i)\} \mid \text{upd}(f_{\text{lx}}, t) \mid f_{\text{frz}}$$

Opt in. Let $t.\text{snd} = \mathbf{x}_i$ and $t.\text{asst} = \tau$. Performing t requires that τ already occurs in Γ , and that \mathbf{x}_i has enough **Algos** to store it. If the balance σ_i does not have an entry for τ , in the new state σ_i is extended with a new entry for τ :

$$\mathbf{x}_i[\sigma_i\{\tau \mapsto 0\}] \mid \dots \mid r \mid T_{\text{iv}} \cup \{t\} \mid f_{\text{asst}} \mid \text{upd}(f_{\text{lx}}, t) \mid f_{\text{frz}}$$

Otherwise, if \mathbf{x}_i 's balance has already an entry for τ , then σ_i is unchanged.

Burn. Let $t.snd = \mathbf{x}_i$ and $t.asst = \tau$. Performing t requires that \mathbf{x}_i is the creator of τ , and that \mathbf{x}_i stores *all* the units of τ (i.e., there are no units of τ in other accounts). In the resulting state, the token τ no longer exists:

$$\mathbf{x}_i[\sigma_i\{\tau \mapsto \perp\}] \mid \cdots \mid r \mid T_{lv} \cup \{t\} \mid f_{asst}\{\tau \mapsto \perp\} \mid upd(f_{lx}, t) \mid f_{frz}$$

Note that this transaction requires an authorization by the asset manager of τ , which is recorded in f_{asst} . (We address this topic in §2.6.)

Revoke. Let $t.snd = \mathbf{x}_i$ and $t.rcv = \mathbf{x}_j$. Performing t requires that both \mathbf{x}_i and \mathbf{x}_j are already storing the asset τ , and that τ is not frozen for \mathbf{x}_i and \mathbf{x}_j . In the new state, the balance of \mathbf{x}_i is decreased by $v = t.val$ units of the asset $\tau = t.asst$, and the balance of \mathbf{x}_j is increased by the same amount:

$$\mathbf{x}_i[\sigma_i - v:\tau] \mid \mathbf{x}_j[\sigma_j + v:\tau] \mid \cdots \mid r \mid T_{lv} \cup \{t\} \mid f_{asst} \mid upd(f_{lx}, t) \mid f_{frz}$$

The effect of a *rvk* transaction is essentially the same as *pay*. The difference is that *rvk* must be authorized by the manager of the asset τ , while *pay* must be authorized by the sender \mathbf{x}_i (see §2.6).

Freeze and unfreeze. A *frz* transaction t with $t.snd = \mathbf{x}_i$ and $t.asst = \tau$ updates the mapping f_{frz} into f'_{frz} , such that $f'_{frz}(\mathbf{x}_i) = f_{frz}(\mathbf{x}_i) \cup \{\tau\}$, whenever the asset τ is owned by \mathbf{x}_i . This effectively prevents any transfers of the asset τ to/from the account \mathbf{x}_i . The dual transaction *unfrz* updates the mapping f_{frz} into f'_{frz} such that $f'_{frz}(\mathbf{x}_i) = f_{frz}(\mathbf{x}_i) \setminus \{\tau\}$.

Delegate. A *delegate* transaction t with $t.snd = \mathbf{x}_i$, $t.rcv = \mathbf{x}_j$ and $t.asst = \tau$ updates the manager of τ , provided that $f_{asst}(\tau) = (\mathbf{x}_i, \mathbf{x}_i)$, for some \mathbf{x}_k . In the updated mapping $f_{asst}\{\tau \mapsto (\mathbf{x}_j, \mathbf{x}_k)\}$, the manager of τ is \mathbf{x}_j .

Initiating a new round. We model the advancement to the next round of the blockchain as a state transition $\Gamma \xrightarrow{\checkmark}_1 \Gamma'$. In the new state Γ' , the round is increased, and the set T_{lv} is updated as $T'_{lv} = \{t \in T_{lv} \mid t.lv > r\}$. The other components of the state are unchanged.

2.4 Executing atomic groups of transactions

Atomic transfers allow state transitions to atomically perform *sequences* of transactions. To atomically perform a sequence $\mathcal{T} = t_1 \cdots t_n$ from a state Γ , we must check that all the transactions t_i can be performed *in sequence*, i.e. the following precondition must hold (for some $\Gamma_1, \dots, \Gamma_n$):

$$\Gamma \xrightarrow{t_1}_1 \Gamma_1 \quad \cdots \quad \Gamma_{n-1} \xrightarrow{t_n}_1 \Gamma_n$$

If so, the state Γ can take a *single-step* transition labelled \mathcal{T} . Denoting the new transition relation with \rightarrow , we write the atomic execution of \mathcal{T} in Γ as follows:

$$\Gamma \xrightarrow{\mathcal{T}} \Gamma_n$$

$e ::= v$	constant
$e \circ e$	arithmetic ($\circ \in \{+, -, <, \leq, =, \geq, >, *, /, \%, \text{and}, \text{or}\}$)
$\text{not } e$	negation
txlen	number of transactions in the atomic group
txpos	index of current transaction in the atomic group
$\text{txid}(n)$	identifier of n -th transaction in the atomic group
$\text{tx}(n).\mathbf{f}$	value of field \mathbf{f} of n -th transaction in the atomic group
$\text{arg}(n)$	n -th argument of the current transaction
$\text{H}(e)$	hash
$\text{versig}(e, e, e)$	signature verification

Syntactic sugar: $\text{false} ::= 1 = 0$ $\text{true} ::= 1 = 1$ $\text{tx.f} ::= \text{tx}(\text{txpos}).\mathbf{f}$ $\text{txid} ::= \text{txid}(\text{txpos})$
 $\text{if } e_0 \text{ then } e_1 \text{ else } e_2 ::= (e_0 \text{ and } e_1) \text{ or } ((\text{not } e_0) \text{ and } e_2)$

Fig. 2: Smart contract scripts (inspired by PyTeal [7]).

2.5 Executing smart contracts

In Algorand, custom authorization policies can be defined with a smart contract language called TEAL [8]. TEAL is a bytecode-based stack language, with an official programming interface for Python (called PyTeal): in our formal model, we take inspiration from the latter to abstract TEAL bytecode scripts as terms, with the syntax in Fig. 2. Besides standard arithmetic-logical operators, TEAL includes operators to count and index all transactions in the current atomic group, and to access their id and fields. When firing transaction involving scripts, users can specify a sequence of *arguments*; accordingly, the script language includes operators to know the number of arguments, and access them. Further, scripts include cryptographic operators to compute hashes and verify signatures.

The *script evaluation function* $\llbracket e \rrbracket_{\mathcal{J}, i}^{\mathcal{W}}$ (Fig. 3) evaluates e using 3 parameters: a sequence of arguments \mathcal{W} , a sequence of transactions \mathcal{J} forming an atomic group, and the index $i < |\mathcal{J}|$ of the transaction containing e . The script $\text{tx}(n).\mathbf{f}$ evaluates to the field \mathbf{f} of the n^{th} transaction in group \mathcal{J} . The size of \mathcal{J} is given by txlen , while txpos returns the index i of the transaction containing the script being evaluated. The script $\text{arg}(n)$ returns the n^{th} argument in \mathcal{W} . The script $\text{H}(e)$ applies a public hash function H to the evaluation of e . The script $\text{versig}(e_1, e_2, e_3)$ verifies a signature e_2 on the message obtained by concatenating the enclosing script and e_1 , using public key e_3 . All operators in Fig. 3 are *strict*: they fail if the evaluation of any operand fails.

2.6 Authorizing transactions, and user-blockchain interaction

As noted before, the mere existence of a step $\Gamma \xrightarrow{t} \Gamma'$ does not imply that t can actually be issued. For this to be possible, users must provide a sequence \mathcal{W} of *witnesses*, satisfying the *authorization predicate* associated with t ; such

$$\begin{aligned}
 \llbracket v \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= v & \llbracket e \circ e' \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= \llbracket e \rrbracket_{\mathcal{J},i}^{\mathcal{W}} \circ_{\perp} \llbracket e' \rrbracket_{\mathcal{J},i}^{\mathcal{W}} & \llbracket \text{not } e \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= \neg_{\perp} \llbracket e \rrbracket_{\mathcal{J},i}^{\mathcal{W}} \\
 \llbracket \text{tx}(n).\mathbf{f} \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= (\mathcal{J}.n).\mathbf{f} \quad (0 \leq n < |\mathcal{J}|) & \llbracket \text{txid}(n) \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= \mathcal{J}.n \quad (0 \leq n < |\mathcal{J}|) \\
 \llbracket \text{txlen} \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= |\mathcal{J}| & \llbracket \text{txpos} \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= i & \llbracket \text{arg}(n) \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= \mathcal{W}.n \quad (0 \leq n < |\mathcal{W}|) \\
 \llbracket \mathbf{H}(e) \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= H(\llbracket e \rrbracket_{\mathcal{J},i}^{\mathcal{W}}) & \llbracket \text{versig}(e_1, e_2, e_3) \rrbracket_{\mathcal{J},i}^{\mathcal{W}} &= \text{ver}_k(m, s) & \left(\begin{array}{l} m = (\mathcal{J}.i.\text{snd}, \llbracket e_1 \rrbracket_{\mathcal{J},i}^{\mathcal{W}}) \\ s = \llbracket e_2 \rrbracket_{\mathcal{J},i}^{\mathcal{W}} \quad k = \llbracket e_3 \rrbracket_{\mathcal{J},i}^{\mathcal{W}} \end{array} \right)
 \end{aligned}$$

Fig. 3: Evaluation of scripts in Fig. 2.

a predicate is uniquely determined by the **authorizer address** of t , written $\text{auth}(t, f_{\text{asst}})$. For transaction types *close*, *pay*, *gen*, *optin* the authorizer address is $t.\text{snd}$; for *burn*, *rvk*, *frz* and *unfrz* on an asset τ it is the asset manager $f_{\text{asst}}(\tau)$. Intuitively, if $\text{auth}(t, f_{\text{asst}}) = \mathbf{x}$, then \mathcal{W} authorizes t iff:

1. if \mathbf{x} is a multisig address (\mathcal{A}, n) , then \mathcal{W} contains at least n signatures of t , made by users in \mathcal{A} ; (if \mathbf{x} is a single-user address \mathbf{a} : see footnote 5)
2. if \mathbf{x} is a script e , then e evaluates to *true* under the arguments \mathcal{W} .

We now formalize the intuition above. Since the evaluation of scripts depends on a whole group of transactions \mathcal{J} , and on the index i of the current transaction within \mathcal{J} , we define the authorization predicate as $\mathcal{W} \models \mathcal{J}, i$ (read: “ \mathcal{W} authorizes the i^{th} transaction in \mathcal{J} ”). Let $\text{sig}_{\mathcal{A}}(m)$ stand for the set of signatures containing $\text{sig}_{\mathbf{a}}(m)$ for all $\mathbf{a} \in \mathcal{A}$; then, $\mathcal{W} \models \mathcal{J}, i$ holds whenever:

1. if $\text{auth}(\mathcal{J}.i, f_{\text{asst}}) = (\mathcal{A}, n)$, then $|\text{set}(\mathcal{W}) \cap \text{sig}_{\text{set}(\mathcal{A})}(\mathcal{J}, i)| \geq n$
2. if $\text{auth}(\mathcal{J}.i, f_{\text{asst}}) = e$, then $\llbracket e \rrbracket_{\mathcal{J},i}^{\mathcal{W}} = \text{true}$

Note that, in general, the sequence of witnesses \mathcal{W} is not unique, i.e., it may happen that $\mathcal{W} \models \mathcal{J}, i$ and $\mathcal{W}' \models \mathcal{J}, i$ for $\mathcal{W} \neq \mathcal{W}'$. For instance, the *Oracle* contract in §3 accepts transactions with witnesses of the form $0s$ or $1s'$, where the first element of the sequence represents the oracle’s choice, and the second element is the oracle’s signature.

Given a *sequence of sequences* of witnesses $\mathbf{W} = \mathcal{W}_0 \cdots \mathcal{W}_{n-1}$ with $n = |\mathcal{J}|$, the **group authorization predicate** $\mathbf{W} \models \mathcal{J}$ holds iff $\mathcal{W}_i \models \mathcal{J}, i$ for all $i \in 0..n-1$.

User-blockchain interaction. We model the interaction of users with the blockchain as a transition system. Its states are pairs (Γ, K) , where Γ is a blockchain state, while K is the set of authorization bitstrings currently known by users. The transition relation $\xRightarrow{\ell}$ (with $\ell \in \{w, \checkmark, \mathbf{W}:\mathcal{J}\}$) is given by the rules:

$$\frac{}{(\Gamma, K) \xRightarrow{w} (\Gamma, K \cup \{w\})} \quad \frac{\Gamma \xrightarrow{\checkmark} \Gamma'}{(\Gamma, K) \xRightarrow{\checkmark} (\Gamma', K)} \quad \frac{\Gamma \xrightarrow{\mathcal{J}} \Gamma' \quad \text{set}(\mathbf{W}) \subseteq K \quad \mathbf{W} \models \mathcal{J}}{(\Gamma, K) \xRightarrow{\mathbf{W}:\mathcal{J}} (\Gamma', K)}$$

With the first two rules, users can broadcast a witness w , or advance to the next round. The last rule gathers from K a sequence of witnesses \mathbf{W} , and lets the blockchain perform an atomic group of transactions \mathcal{J} if authorized by \mathbf{W} .

2.7 Fundamental properties of ASC1

We now exploit our formal model to establish some fundamental properties of ASC1. Theorem 1 states that the same transaction t cannot be issued more than once, i.e., there is no *double-spending*. In the statement, we use $\rightarrow^* \overset{\mathcal{T}}{\rightarrow} \rightarrow^*$ to denote an arbitrarily long series of steps including a group of transactions \mathcal{T} .

Theorem 1 (No double-spending). *Let $\Gamma_0 \rightarrow^* \overset{\mathcal{T}}{\rightarrow} \rightarrow^* \Gamma \xrightarrow{\mathcal{T}'} \Gamma'$. Then, no transaction occurs more than once in $\mathcal{T}\mathcal{T}'$.*

Define the *value* of an asset τ in a state $\Gamma = \mathbf{x}_1[\sigma_1] \mid \dots \mid \mathbf{x}_n[\sigma_n] \mid r \mid \dots$ as the sum of the balances of τ in all accounts in Γ :

$$val_{\tau}(\Gamma) = \sum_{i=1}^n val_{\tau}(\sigma_i) \quad \text{where } val_{\tau}(\sigma) = \begin{cases} \sigma(\tau) & \text{if } \tau \in \text{dom}(\sigma) \\ 0 & \text{otherwise} \end{cases}$$

Theorem 2 states that, once an asset is minted, its value remains constant, until the asset is eventually burnt. In particular, since `Algos` cannot be burnt (nor minted, unlike in Bitcoin and Ethereum), their amount remains constant.

Theorem 2 (Value preservation). *Let $\Gamma_0 \rightarrow^* \Gamma \rightarrow^* \Gamma'$. Then:*

$$val_{\tau}(\Gamma') = \begin{cases} val_{\tau}(\Gamma) & \text{if } \tau \text{ occurs in } \Gamma \text{ and it is not burnt in } \Gamma \rightarrow^* \Gamma' \\ 0 & \text{otherwise} \end{cases}$$

Theorem 3 establishes that the transition systems \rightarrow and \Rightarrow are deterministic: crucially, this allows reconstructing the blockchain state from the transition log. Notably, by item 3 of Theorem 3, witnesses only determine whether a state transition happens or not, but they do not affect the new state. This is unlike Ethereum, where arguments of function calls in transactions may affect the state.

Theorem 3 (Determinism). *For all $\lambda \in \{\checkmark, \mathcal{T}\}$ and $\ell \in \{\checkmark, w\}$:*

1. *if $\Gamma \xrightarrow{\lambda} \Gamma'$ and $\Gamma \xrightarrow{\lambda} \Gamma''$, then $\Gamma' = \Gamma''$;*
2. *if $(\Gamma, K) \xRightarrow{\ell} (\Gamma', K')$ and $(\Gamma, K) \xRightarrow{\ell} (\Gamma'', K'')$, then $(\Gamma', K') = (\Gamma'', K'')$;*
3. *if $(\Gamma, K) \xRightarrow{\mathbf{W}:\mathcal{T}} (\Gamma', K')$ and $(\Gamma, K) \xRightarrow{\mathbf{W}:\mathcal{T}} (\Gamma'', K'')$, then $\Gamma' = \Gamma''$ and $K' = K'' = K$.*

3 Designing secure smart contracts in Algorand

We now exploit our formal model to design some archetypal smart contracts, and establish their security (§3.2). First, we introduce an attacker model.

3.1 Attacker model

We assume that cryptographic primitives are secure, i.e., hashes are collision resistant and signatures cannot be forged (except with negligible probability). A **run** \mathcal{R} is a (possibly infinite) sequence of labels $\ell_1\ell_2\cdots$ such that $(\Gamma_0, K_0) \xrightarrow{\ell_1} (\Gamma_1, K_1) \xrightarrow{\ell_2} \cdots$, where Γ_0 is the initial state, and $K_0 = \emptyset$ is the initial (empty) knowledge; hence, as illustrated in §2.6, each label ℓ_i in a run \mathcal{R} can be either w (broadcast of a witness bitstring w), $\mathbf{W}:\mathcal{T}$ (atomic group of transactions \mathcal{T} authorized by \mathbf{W}), or \checkmark (advance to next round). We consider a setting where:

- each user \mathbf{a} has a **strategy** Σ , i.e. a PPTIME algorithm to select which label to perform among those permitted by the ASC1 transition system. A strategy takes as input a finite run \mathcal{R} (the past history) and outputs a single enabled label ℓ . Strategies are *stateful*: users can read and write a private unbounded tape to maintain their own state throughout the run. The initial state of \mathbf{a} 's tape contains \mathbf{a} 's private key, and the public keys of all users;¹¹
- an **adversary** \mathbf{Adv} who controls the scheduling with her stateful **adversarial strategy** $\Sigma_{\mathbf{Adv}}$: a PPTIME algorithm taking as input the current run \mathcal{R} and the labels output by the strategies of users (i.e., the steps that users are trying to make). The output of $\Sigma_{\mathbf{Adv}}$ is a single label ℓ , that is appended to the current run. We assume the adversarial strategy $\Sigma_{\mathbf{Adv}}$ can delay users' transactions by at most $\delta_{\mathbf{Adv}}$ rounds, where $\delta_{\mathbf{Adv}}$ is a given natural number.¹²

A set Σ of strategies of users and \mathbf{Adv} induces a distribution of runs; we say that run \mathcal{R} is **conformant** to Σ if \mathcal{R} is sampled from such a distribution. We assume that infinite runs contain infinitely many \checkmark : this *non-Zeno condition* ensures that neither users nor \mathbf{Adv} can perform infinitely many transactions in a round.

3.2 Smart contracts

We now exploit our model to specify some archetypal ASC1 contracts, and reason about their security. To simplify the presentation, we assume $\delta_{\mathbf{Adv}} = 0$, i.e., the adversary \mathbf{Adv} can start a new round (performing \checkmark) only if all users agree.¹³ The table below summarises our selection of smart contracts, highlighting the design patterns they implement.

¹¹ Notice that new public/private key pairs can be generated during the run, and their public parts can be communicated as labels w .

¹² Without this assumption, \mathbf{Adv} could arbitrarily disrupt deadlines: e.g., $\Sigma_{\mathbf{Adv}}$ could make \mathbf{a} *always* lose lottery games (like the ones below) by delaying \mathbf{a} 's transactions.

¹³ All results can be easily adjusted for $\delta_{\mathbf{Adv}} > 0$, but this would require more verbose statements to account for possible delays introduced by \mathbf{Adv} .

Use case / Pattern	Signed witness	Timeouts	Commit/reveal	State machine	Atomic transfer	Time windows
Oracle	✓	✓				
HTLC		✓	✓			
Mutual HTLC (§B.1)		✓	✓		✓	
$O(n^2)$ -collateral lottery		✓	✓		✓	
0-collateral lottery (§B.2)		✓	✓	✓	✓	
Periodic payment						✓
Escrow (§B.3)	✓			✓		
Two-phase authorization		✓		✓		✓
Limit order (§B.4)		✓			✓	
Split (§B.5)		✓			✓	

Oracle. We start by designing a contract which allows either \mathbf{a} or \mathbf{b} to withdraw all the \mathbf{Algo} s in the contract, depending on the outcome of a certain boolean event. Let \mathbf{o} be an oracle who certifies such an outcome, by signing the value 1 or 0. We model the contract as the following script:

$$\begin{aligned} \mathit{Oracle} \triangleq & \text{tx.type} = \mathit{close} \text{ and } \text{tx.asst} = \mathbf{Algo} \text{ and } ((\text{tx.fv} > r_{max} \text{ and } \text{tx.rcv} = \mathbf{a}) \\ & \text{or } (\mathbf{arg}(0) = 0 \text{ and } \mathbf{versig}(\mathbf{arg}(0), \mathbf{arg}(1), \mathbf{o}) \text{ and } \text{tx.rcv} = \mathbf{a}) \\ & \text{or } (\mathbf{arg}(0) = 1 \text{ and } \mathbf{versig}(\mathbf{arg}(0), \mathbf{arg}(1), \mathbf{o}) \text{ and } \text{tx.rcv} = \mathbf{b})) \end{aligned}$$

Once created, the contract accepts only close transactions, using two arguments as witnesses. The argument $\mathbf{arg}(0)$ contains the outcome, while $\mathbf{arg}(1)$ is \mathbf{o} 's signature on $(\mathit{Oracle}, \mathbf{arg}(0))$, i.e., the concatenation between the script and the first argument. The user \mathbf{b} can collect the funds in Oracle if \mathbf{o} certifies the outcome 1, while \mathbf{a} can collect the funds if the outcome is 0, or after round r_{max} .

Theorem 4 below proves that Oracle works as intended. To state it, we define $T_{\mathbf{p}}$ as the set of transactions allowing a user \mathbf{p} to withdraw the contract funds:

$$T_{\mathbf{p}} = \{t \mid t.\text{type} = \mathit{close}, t.\text{snd} = \mathit{Oracle}, t.\text{rcv} = \mathbf{p}, t.\text{asst} = \mathbf{Algo}\}$$

The theorem considers the following strategies for \mathbf{a} , \mathbf{b} , and \mathbf{o} :

- $\Sigma_{\mathbf{a}}$: wait for $s = \mathbf{sig}_{\mathbf{o}}(\mathit{Oracle}, 0)$; if s arrives at round $r \leq r_{max}$, then immediately send a transaction $t \in T_{\mathbf{a}}$ with $t.\text{fv} = r$ and witness $0s$; otherwise, at round $r_{max} + 1$, send a transaction $t \in T_{\mathbf{a}}$ with $t.\text{fv} = r_{max} + 1$;
- $\Sigma_{\mathbf{b}}$: wait for $s' = \mathbf{sig}_{\mathbf{o}}(\mathit{Oracle}, 1)$; if s' arrives at round r , immediately send a transaction $t \in T_{\mathbf{b}}$ with $t.\text{fv} = r$ and witness $1s'$;
- $\Sigma_{\mathbf{o}}$: do one of the following: (a) send \mathbf{o} 's signature on $(\mathit{Oracle}, 0)$ at any time, or (b) send \mathbf{o} 's signature on $(\mathit{Oracle}, 1)$ at any time, or (c) do nothing.

Theorem 4. Let \mathcal{R} be a run conforming to some set of strategies Σ , such that: (i) $\Sigma_{\mathbf{o}} \in \Sigma$; (ii) \mathcal{R} reaches, at some round before r_{max} , a state $\mathit{Oracle}[\sigma] \mid \dots$; (iii) \mathcal{R} reaches the round $r_{max} + 2$. Then, with overwhelming probability:

- (1) if $\Sigma_{\mathbf{a}} \in \Sigma$ and \mathbf{o} has not sent a signature on $(\mathit{Oracle}, 1)$, then \mathcal{R} contains a transaction in $T_{\mathbf{a}}$, transferring at least $\sigma(\mathbf{Algo})$ to \mathbf{a} ;
- (2) if $\Sigma_{\mathbf{b}} \in \Sigma$ and \mathbf{o} has sent a signature on $(\mathit{Oracle}, 1)$ at round $r \leq r_{max}$, then \mathcal{R} contains a transaction in $T_{\mathbf{b}}$, transferring at least $\sigma(\mathbf{Algo})$ to \mathbf{b} .

Notice that in item (1) we are only assuming that \mathbf{a} and \mathbf{o} use the strategies $\Sigma_{\mathbf{a}}$ and $\Sigma_{\mathbf{o}}$, while \mathbf{b} and \mathbf{Adv} can use any strategy (and possibly collude). Similarly, in item (2) we are only assuming \mathbf{b} 's and \mathbf{o} 's strategies.

Hash Time Lock Contract (HTLC). A user \mathbf{a} promises that she will either reveal a secret s_a by round r_{max} , or pay a penalty to \mathbf{b} . More sophisticated contracts, e.g. gambling games, use this mechanism to let players generate random numbers in a fair way. We define the HTLC as the following contract, parameterised on the two users \mathbf{a}, \mathbf{b} and the hash $h_a = H(s_a)$ of the secret:

$$HTLC(\mathbf{a}, \mathbf{b}, h_a) \triangleq \text{tx.type} = \text{close} \text{ and } \text{tx.asst} = \text{Algo} \text{ and} \\ ((\text{tx.rcv} = \mathbf{a} \text{ and } H(\text{arg}(0)) = h_a) \text{ or } (\text{tx.rcv} = \mathbf{b} \text{ and } \text{tx.fv} \geq r_{max}))$$

The contract accepts only *close* transactions with receiver \mathbf{a} or \mathbf{b} . User \mathbf{a} can collect the funds in the contract only by providing the secret s_a in $\text{arg}(0)$, effectively making s_a public.¹⁴ Instead, if \mathbf{a} does not reveal s_a , then \mathbf{b} can collect the funds after round r_{max} . We state the correctness of *HTLC* in Theorem 5; first, let T_p be the set of transactions allowing user \mathbf{p} to withdraw the contract funds:

$$T_p = \{t \mid t.\text{type} = \text{close}, t.\text{snd} = HTLC(\mathbf{a}, \mathbf{b}, h_a), t.\text{rcv} = \mathbf{p}, t.\text{asst} = \text{Algo}\}$$

We consider the following strategies for \mathbf{a} and \mathbf{b} :

- Σ_a : at a round $r < r_{max}$, send a $t \in T_a$ with $t.\text{fv} = r$ and witness s_a ;
- Σ_b : at round r_{max} , check whether any transaction in T_a occurs in \mathcal{R} . If not, then immediately send a transaction $t \in T_b$ with $t.\text{fv} = r_{max}$.

Theorem 5. *Let \mathcal{R} be a run conforming to some set of strategies Σ , such that: (i) \mathcal{R} reaches, at some round before r_{max} , a state $HTLC(\mathbf{a}, \mathbf{b}, h_a)[\sigma] \mid \dots$; (ii) \mathcal{R} reaches the round $r_{max} + 1$. Then, with overwhelming probability:*

- (1) if $\Sigma_a \in \Sigma$, then \mathcal{R} contains a transaction in T_a , transferring at least $\sigma(\text{Algo})$ to \mathbf{a} ;
- (2) if $\Sigma_b \in \Sigma$ and \mathcal{R} does not contain the secret s_a before round $r_{max} + 1$, then \mathcal{R} contains a transaction in T_b , transferring at least $\sigma(\text{Algo})$ to \mathbf{b} .

Lotteries. Consider a gambling game where n players bet 1Algo each, and the winner, chosen uniformly at random among them, can redeem $n\text{Algos}$. A simple implementation, inspired by [11–13] for Bitcoin, requires each player to deposit $n(n-1)\text{Algos}$ as collateral in an HTLC contract.¹⁵ For $n = 2$ players \mathbf{a} and \mathbf{b} , such deposits are transferred by the following transactions:

$$t_{Ha} = \{\text{type} : \text{pay}, \text{snd} : \mathbf{a}, \text{rcv} : HTLC(\mathbf{a}, \mathbf{b}, h_a), \text{val} : 2, \text{asst} : \text{Algo}, \dots\} \\ t_{Hb} = \{\text{type} : \text{pay}, \text{snd} : \mathbf{b}, \text{rcv} : HTLC(\mathbf{b}, \mathbf{a}, h_b), \text{val} : 2, \text{asst} : \text{Algo}, \dots\}$$

The bets are stored in the following contract, which determines the winner as a function of the secrets, and allows her to withdraw the whole pot:

$$\text{Lottery} \triangleq \text{tx.type} = \text{close} \text{ and } \text{tx.asst} = \text{Algo} \text{ and } H(\text{arg}(0)) = h_a \text{ and } H(\text{arg}(1)) = h_b \\ \text{and if } (\text{arg}(0) + \text{arg}(1)) \% 2 = 0 \text{ then tx.rcv} = \mathbf{a} \text{ else tx.rcv} = \mathbf{b}$$

with $h_a \neq h_b$.¹⁶ Players \mathbf{a} and \mathbf{b} start the game with the atomic transactions:

¹⁴ If s_a is a sufficiently long bitstring generated uniformly at random, collision resistance of the hash function ensures that only \mathbf{a} (who knows s_a) can provide such an $\text{arg}(0)$.

¹⁵ A zero-collateral lottery is presented in (§B.2).

¹⁶ This check prevents a replay attack: if \mathbf{a} chooses $h_a = h_b$, then \mathbf{b} cannot win.

$$\begin{aligned} \mathcal{T}_{a,b} &= t_{Ha} t_{Hb} t_{La} t_{Lb} \text{ where:} \\ t_{La} &= \{\text{type} : \text{pay}, \text{snd} : a, \text{rcv} : \text{Lottery}, \text{val} : 1, \text{asst} : \text{Algo}, \dots\} \\ t_{Lb} &= \{\text{type} : \text{pay}, \text{snd} : a, \text{rcv} : \text{Lottery}, \text{val} : 1, \text{asst} : \text{Algo}, \dots\} \end{aligned}$$

The transaction t_{La} creates the contract with a 's bet, and t_{Lb} completes it with b 's bet. At this point, there are two possible outcomes:

- (a) both players reveal their secret. Then, the winner can withdraw the pot, by performing a *close* action on the *Lottery* contract, providing as arguments the two secrets, and setting her identity in the *rcv* field;
- (b) one of the players does not reveal the secret. Then, the other player can withdraw the collateral in the other player's HTLC (and redeem her own).

To formalise the correctness of the lottery, consider the sets of transactions:

$$\begin{aligned} T_{p,q}^{secr} &= \{t \mid t.\text{type} = \text{close}, t.\text{snd} = \text{HTLC}(p, q, h_p), t.\text{rcv} = p, t.\text{asst} = \text{Algo}\} \\ T_{p,q}^{tout} &= \{t \mid t.\text{type} = \text{close}, t.\text{snd} = \text{HTLC}(p, q, h_p), t.\text{rcv} = q, t.\text{asst} = \text{Algo}\} \\ T_p^{lott} &= \{t \mid t.\text{type} = \text{close}, t.\text{snd} = \text{Lottery}, t.\text{rcv} = p, t.\text{asst} = \text{Algo}\} \end{aligned}$$

and consider the following strategy Σ_a for a (the one for b is analogous):

1. at some $r < r_{max}$, send a transaction $t \in T_{a,b}^{secr}$ with $t.\text{fv} = r$ and witness s_a ;
2. if some transaction in $T_{b,a}^{secr}$ occurs in \mathcal{R} at round $r' < r_{max}$, then extract its witness s_b and compute the winner; if a is the winner, immediately send a transaction $t \in T_a^{lott}$ with $t.\text{fv} = r'$ and witness $s_a s_b$;
3. if at round r_{max} no transaction in $T_{b,a}$ occurs in \mathcal{R} , immediately send a transaction $t \in T_{b,a}^{tout}$ with $t.\text{fv} = r_{max}$.

Theorem 6 below establishes that the lottery is fair, implying that the expected payoff of player a following strategy Σ_a is at least negligible; instead, if a does not follow Σ_a (e.g., by not revealing her secret), the expected payoff may be negative; analogous results hold for player b . This result can be generalised for $n > 2$ players, with a collateral of $n(n-1)$ *Algos*. As in the HTLC, we assume that s_a and s_b are sufficiently long bitstrings generated uniformly at random.

Theorem 6. *Let \mathcal{R} be a run conforming to a set of strategies Σ , such that: (i) \mathcal{R} contains, before r_{max} , the label $\mathcal{T}_{a,b}$; (ii) \mathcal{R} reaches round $r_{max} + 1$. For $p \neq q \in \{a, b\}$, if $\Sigma_p \in \Sigma$, then: (1) \mathcal{R} contains a transaction in $T_{p,q}^{secr}$, transferring at least 2 *Algo* to p ; (2) the probability that \mathcal{R} contains $T_{q,p}^{tout}$ or T_p^{lott} , which transfer at least 1 *Algo* to p , is $\geq \frac{1}{2}$ (up-to a negligible quantity).*

Periodic payment. We want to ensure that a can withdraw a fixed amount of v *Algos* at fixed time windows of p rounds. We can implement this behaviour through the following contract, which can be refilled when needed:

$$\begin{aligned} PP(p, d, n) &\triangleq \text{tx.type} = \text{pay} \text{ and } \text{tx.val} = v \text{ and } \text{tx.asst} = \text{Algo} \text{ and} \\ &\quad \text{tx.rcv} = a \text{ and } \text{tx.fv} \% p = 0 \text{ and } \text{tx.lv} = \text{tx.fv} + d \text{ and } \text{tx.lx} = n \end{aligned}$$

The contract accepts only *pay* transactions of v *Algos* to receiver a . The conditions $\text{tx.fv} \% p = 0$ and $\text{tx.lv} = \text{tx.fv} + d$ ensure that the contract only accepts transactions with validity interval $[k p, k p + d]$, for $k \in \mathbb{N}$. The condition $\text{tx.lx} = n$ ensures that *at most* one such transactions is accepted for each time window.

Finite-state machines. Consider a set of users A who want to stipulate a contract whose behaviour is given by a finite-state machine with states q_0, \dots, q_n . We can implement such a contract by representing each state q_i as a script e_i ; the current state/script holds the assets, and each state transition $q_i \rightarrow q_j$ is a clause in e_i which enables a *close* transaction to transfer the assets to e_j . This clause requires $\text{tx.rcv} = e_j$ — except in case of loops, which cannot be encoded directly:¹⁷ in this case, we identify the next state as $\text{tx.rcv} = \text{arg}(0)$, also requiring all users in A to sign $\text{arg}(0)$ to confirm its correctness. To ensure that any user in A can trigger a state transition (by firing the corresponding transaction), their signatures must be exchanged before the contract starts. We show an instance of this pattern as the two-phase authorization contract below.

An alternative technique is based on quines. As before, a state transition $q_i \rightarrow q_j$ is rendered as a transaction which closes e_i and transfers the balance to e_j . Here, all such scripts e_k have the same code, except for a single state constant k which occurs at a specific offset, and which represents the current state. To verify that tx.rcv represents a legit next state, e_i requires a witness w such that: (i) tx.rcv is equal to the hash of w , and the state constant j within w is indeed a next state for i ; (ii) tx.snd is equal to the hash of w' , where w' is obtained from w by replacing the state constant j with the current state i . Performing these checks could be possible by using concatenation and substring operators.¹⁸

Two-phase authorization. We want a contract to allow user c to withdraw some funds, but only if authorized by a and b . We want a to give her authorization first; if b 's authorization is not given within $p \geq \Delta_{max}$ rounds, then anyone can fire a transaction to reset the contract to its initial state. We model this contract with two scripts: $P1$ represents the state where no authorization has been given yet, while $P2$ represents the state where a 's authorization has been given. Conceptually, the contract implements a finite-state machine, looping between two states until the contract funds are withdrawn by c .

$$\begin{aligned}
 P1 &\triangleq \text{tx.type} = \textit{close} \text{ and } \text{tx.asst} = \textit{Algo} \text{ and } \text{versig}(\text{txid}, \text{arg}(0), a) \text{ and} \\
 &\quad \text{tx.rcv} = P2 \text{ and } \text{tx.fv} \% (4 * p) = 0 \text{ and } \text{tx.lv} = \text{tx.fv} + \Delta_{max} \\
 P2 &\triangleq \text{tx.type} = \textit{close} \text{ and } \text{tx.asst} = \textit{Algo} \text{ and} \\
 &\quad ((\text{versig}(\text{txid}, \text{arg}(0), b) \text{ and } \text{tx.rcv} = c) \text{ or} \\
 &\quad (\text{versig}(\text{arg}(0), \text{arg}(1), a_1) \text{ and } \text{versig}(\text{arg}(0), \text{arg}(2), b_1) \text{ and} \\
 &\quad \text{tx.rcv} = \text{arg}(0) \text{ and } \text{tx.fv} \% (4 * p) = 2 * p \text{ and } \text{tx.lv} = \text{tx.fv} + \Delta_{max}))
 \end{aligned}$$

The scripts $P1$ and $P2$ use a time window with 4 frames, each lasting p rounds. Script $P1$ only accepts *close* transactions which transfer the balance to $P2$; the time constraint ensures that such transactions are sent in the first time frame. The script $P2$ accepts two kinds of transactions: (a) transfer the balance to c , using an authorization by b ; (b) transfer the balance to $P1$, in the 4th time frame. Note that in $P2$ we cannot use the (intuitively correct) condition $\text{tx.rcv} = P1$, as it would introduce a circularity. Instead, we apply the state machines technique

¹⁷ This is because Algorand contracts cannot have circular references: contract accounts are referenced by script hashes, and no script can depend on its own hash.

¹⁸ In Algorand, these operators are available only for `LogicSigVersion` ≥ 2 .


```

{
  "type": "pay",
  "snd": x,
  "rcv": 0,
  "close": y,
  "amt": 0
}
{
  "type": "axfer",
  "snd": x,
  "asnd": x,
  "arcv": 0,
  "aclose": y,
  "xaid": tau,
  "aamt": 0
}

```

Fig. 4: Translation of a *close* transaction (left: $\tau = \text{Algo}$, right: $\tau \neq \text{Algo}$).

described above: we require $\text{tx.rcv} = \text{arg}(0)$, with $\text{arg}(0)$ signed by both **a** and **b**,¹⁹ and assume that these signatures are exchanged before the contract starts.

4 From the formal model to concrete Algorand

We now discuss how to translate transactions and scripts in our model to concrete Algorand. We first sketch how to compile our scripts into TEAL. The compilation of most constructs is straightforward. For instance, a script $e + e'$ is compiled by using the opcode `+`, and similarly for the other arithmetic and comparison operators, and for the cryptographic primitives. The logic operators `and`, or are compiled via the opcode `bnz`, to obtain the short-circuit semantics. The `not` operator is compiled via the opcode `!`. The operator $\text{txid}(n)$ is compiled as `gtxn n TxID`, txlen is compiled as `global GroupSize`, txpos is compiled as `txn GroupIndex`, and $\text{arg}(n)$ as `arg n`.

Finally, compiling the script $\text{tx}(n).\mathbf{f}$ depends on the field **f**. If **f** is `fv`, `lv`, or `lx`, then the compilation is `gtxn n i`, where *i* is, respectively, `FirstValid`, `LastValid`, or `Lease`. For the other cases of **f**, the compilation of $\text{tx}(n).\mathbf{f}$ generates a TEAL script which computes **f** by decoding the concrete Algorand transaction fields, and making them available in the scratch space. This decoding is detailed in Table 2 in §C. From the same table we can also infer how to translate transactions in the model to concrete Algorand transactions. For instance, translating a transaction of the form:

$$\{\text{type} : \text{close}, \text{snd} : \mathbf{x}, \text{rcv} : \mathbf{y}, \text{asst} : \tau\}$$

results in the concrete transaction in Fig. 4 (where we omit the irrelevant fields).

Our modelling approach is supported by a prototype tool, called `secteal` (*secure TEAL*), and accessible via a web interface at:

<http://secteal.cs.stir.ac.uk/>

The core of the tool is a compiler that translates smart contracts written as expressions, based on the script language (§2.5), into executable TEAL bytecode. In its current form, `secteal` supports experimentation with our model, and is provided with a series of examples from §3.2. Users can also compile their own `secteal` contracts, paving the way to a declarative approach to contract design

¹⁹ We use other key pairs \mathbf{a}_1 and \mathbf{b}_1 to avoid confusion with the signatures on `txid`.

and development. `secteal` is a first building block toward a comprehensive IDE for the design, verification, and deployment of contracts on Algorand.

5 Conclusions

This work is part of a wider research line on formal modelling of blockchain-based contracts, ranging from Bitcoin [14, 30, 35] to Ethereum [20, 25–28, 33], Cardano [21], Tezos [19], and Zilliqa [36]. These formal models are a necessary prerequisite to rigorously reason on the security of smart contracts, and they are the basis for automatic verification. Besides modelling the behaviour of transactions, in §3.1 we have proposed a model of attackers: this enables us to prove properties of smart contracts in the presence of adversaries, in the spirit of long-standing research in the cryptography area [10, 11, 15, 18, 24, 31, 32].

Differences between our model and Algorand Besides not modelling the consensus protocol, to keep the formalization simple, we chose to abstract from some aspects of ASC1, which do not appear to be relevant to the development of (the majority of) smart contracts. First, we are not modelling some transaction fields: among them, we have omitted the `fee` field, used to specify an amount of `Algos` to be paid to nodes, and the `note` field, used to embed arbitrary bitstrings into transactions. We associate a single manager to assets, while Algorand uses different managers for different operations (e.g., the freeze manager for `frz/unfrz` and the clawback manager for `rvk`). We use two different transactions types, `pay` and `close`, to perform asset transfers and account closures: in Algorand, a single `pay` transaction can perform both. Note that we can achieve the same effect by performing the `pay` and `close` transactions within the same atomic group. Although Algorand relies on 7 transaction types, the behaviour of some transactions needs to be further qualified by the combination of other fields (e.g., freeze and unfreeze are obtained by transactions with the same type `afrz`, but with a different value of the `AssetFrozen` field). While this is useful as an implementation detail, our model simplifies reasoning about different behaviours by explicitly exposing them in the transaction type. In the same spirit, while Algorand uses different transaction types to represent actions with similar functionality (e.g., transferring `Algos` and user-defined assets are rendered with different transaction types, `pay` and `axfer`), we use the same transaction type (e.g., `pay`) for such actions. Our model does not encompass some advanced features of Algorand, e.g.: rekeying accounts, key registration transactions (`keyreg`), some kinds of asset configuration transaction (e.g., decimals, default frozen, different managers), and application call transactions.²⁰ Our script language substantially covers TEAL opcodes with `LogicSigVersion=1`, but for a few exceptions, e.g. bitwise operations, different hash functions, jumps.

²⁰ Application call transactions are used to implement *stateful* contracts, and therefore are outside the scope of this paper.

Related work Besides featuring an original consensus protocol based on proof-of-stake [22], Algorand has also introduced a novel paradigm of (stateless) smart contracts, which differs from the paradigms of other mainstream blockchains. On the one hand, Algorand follows the *account-based* model, similarly to Ethereum (and differently from Bitcoin and Cardano, which follow the UTXO model). On the other hand, Algorand’s paradigm of stateless contracts diverges from Ethereum’s stateful contracts: rather, it resembles Bitcoin’s, where contracts are based upon custom transaction redeem conditions. Besides these differences, Algorand natively features user-defined assets, while other platforms render them as smart contracts (e.g., by implementing ERC20 and ERC721 interfaces in Ethereum). Overall, these differences demand for a formal model that is substantially different from the models devised for the other blockchain platforms.

Our formalization of the Algorand’s script language is close, with respect to the level of abstraction, to the model of Bitcoin script developed in [14]. Indeed, both works formalise scripts in an expression language, abstracting from the bytecode. A main difference between Algorand and Ethereum is that Ethereum contracts are stateful: their state can be updated by specific bytecode instructions; instead, (stateless) TEAL scripts merely authorize transactions. Consequently, a difference between our model and formal models of Ethereum contracts is that the semantics of our scripts has no side effects. In this way, our work departs from most literature on the formalization of Ethereum contracts, where the target of the formalization is either the bytecode language EVM [26, 27, 33], or the high-level language Solidity [16, 23, 29].

Future work Our formal model of Algorand smart contracts can be expanded depending on the evolution of the Algorand framework. In mid August 2020, Algorand has introduced *stateful* ASC1 contracts [4], enriching contract accounts with a persistent key-value store, accessible and modifiable through a new kind of transaction (which can use an extended set of TEAL opcodes). To accommodate stateful contracts in our model, we would need to embed the key-value store in contract accounts, and extend the script language with key-value store updates. The rest of our model (in particular, the semantics of transactions and the attacker model) is mostly unaffected by this extension. Future work could also investigate declarative languages for stateful Algorand smart contracts, and associated verification techniques. Another research direction is the mechanization of our formal model, using a proof assistant: this would allow machine-checking the proofs developed by pencil-and-paper in §D. Similar work has been done e.g. for Bitcoin [35] and for Tezos [19].

Acknowledgements The authors thank the anonymous reviewers of Financial Cryptography 2021 for their insightful comments on a preliminary version of this paper. Massimo Bartoletti is partially supported by Convenzione tra Fondazione di Sardegna e Atenei Sardi project F74I19000900007 *ADAM*. Cristian Lepore is partially supported by The Data Lab, Innovation Center. Alceste Scalas is partially supported by EU Horizon 2020 project 830929 *CyberSec4Europe*, and Industriens Fonds Cyberprogram 2020-0489 *Security-by-Design in Digital Den-*

mark (*Sb3D*). Roberto Zunino is partially supported by MIUR PON *Distributed Ledgers for Secure Open Communities*.

References

1. Algorand developer docs (2020), <https://developer.algorand.org/docs/>
2. Algorand developer docs: Limit order (contract owner has Algos) (2020), https://developer.algorand.org/docs/reference/teal/templates/limit_order/
3. Algorand developer docs: Split (2020), <https://developer.algorand.org/docs/reference/teal/templates/split/>
4. Algorand developer docs: stateful smart contracts (2020), <https://developer.algorand.org/docs/features/asc1/stateful/>
5. Algorand developer docs: Transaction Execution Approval Language (TEAL) (2020), <https://developer.algorand.org/docs/reference/teal>
6. ASC1 coherence-checking experiments (2020), <https://github.com/blockchain-unica/asc1-experiments>
7. PyTeal: Algorand smart contracts in Python (2020), <https://github.com/algorand/pyteal>
8. Transaction execution approval language (TEAL) specification (2020), <https://developer.algorand.org/docs/reference/teal/specification/>
9. Alturki, M.A., Chen, J., Luchangco, V., Moore, B.M., Palmskog, K., Peña, L., Rosu, G.: Towards a verified model of the Algorand consensus protocol in Coq. In: Formal Methods Workshops. Lecture Notes in Computer Science, vol. 12232, pp. 362–367. Springer (2019). https://doi.org/10.1007/978-3-030-54994-7_27
10. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Fair two-party computations via Bitcoin deposits. In: Financial Cryptography Workshops. LNCS, vol. 8438, pp. 105–121. Springer (2014). https://doi.org/10.1007/978-3-662-44774-1_8
11. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multiparty computations on Bitcoin. In: IEEE S & P. pp. 443–458 (2014). <https://doi.org/10.1109/SP.2014.35>, first appeared on Cryptology ePrint Archive, <http://eprint.iacr.org/2013/784>
12. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multiparty computations on Bitcoin. *Commun. ACM* **59**(4), 76–84 (2016). <https://doi.org/10.1145/2896386>
13. Atzei, N., Bartoletti, M., Cimoli, T., Lande, S., Zunino, R.: SoK: unraveling Bitcoin smart contracts. In: POST. LNCS, vol. 10804, pp. 217–242. Springer (2018). <https://doi.org/10.1007/978-3-319-89722-6>
14. Atzei, N., Bartoletti, M., Lande, S., Zunino, R.: A formal model of Bitcoin transactions. In: Financial Cryptography and Data Security. LNCS, vol. 10957. Springer (2018). <https://doi.org/10.1007/978-3-662-58387-6>
15. Banasik, W., Dziembowski, S., Malinowski, D.: Efficient zero-knowledge contingent payments in cryptocurrencies without scripts. In: ESORICS. LNCS, vol. 9879, pp. 261–280. Springer (2016). https://doi.org/10.1007/978-3-319-45741-3_14
16. Bartoletti, M., Galletta, L., Murgia, M.: A minimal core calculus for Solidity contracts. In: Cryptocurrencies and Blockchain Technology. LNCS, vol. 11737, pp. 233–243. Springer (2019). https://doi.org/10.1007/978-3-030-31500-9_15
17. Bartoletti, M., Zunino, R.: Constant-deposit multiparty lotteries on Bitcoin. In: Financial Cryptography Workshops. LNCS, vol. 10323. Springer (2017). <https://doi.org/10.1007/978-3-319-70278-0>

18. Bentov, I., Kumaresan, R.: How to use Bitcoin to design fair protocols. In: CRYPTO. LNCS, vol. 8617, pp. 421–439. Springer (2014). https://doi.org/10.1007/978-3-662-44381-1_24
19. Bernardo, B., Cauderlier, R., Hu, Z., Pesin, B., Tesson, J.: Mi-Cho-Coq, a framework for certifying Tezos smart contracts. In: Sekerinski, E., Moreira, N., Oliveira, J.N., Ratiu, D., Guidotti, R., Farrell, M., Luckcuck, M., Marmosoler, D., Campos, J., Astarte, T., Gonnord, L., Cerone, A., Couto, L., Dongol, B., Kutrib, M., Monteiro, P., Delmas, D. (eds.) Workshop on Formal Methods for Blockchains. LNCS, vol. 12232, pp. 368–379. Springer (2019). https://doi.org/10.1007/978-3-030-54994-7_28
20. Bhargavan, K., Delignat-Lavaud, A., Fournet, C., Gollamudi, A., Gonthier, G., Kobeissi, N., Rastogi, A., Sibut-Pinote, T., Swamy, N., Zanella-Beguelin, S.: Formal verification of smart contracts. In: PLAS (2016). <https://doi.org/10.1145/2993600.2993611>
21. Chakravarty, M.M., Chapman, J., MacKenzie, K., Melkonian, O., Jones, M.P., Wadler, P.: The extended UTXO model. In: Financial Cryptography Workshops. LNCS, vol. 12063, pp. 525–539. Springer (2020). https://doi.org/10.1007/978-3-030-54455-3_37
22. Chen, J., Micali, S.: Algorand: A secure and efficient distributed ledger. *Theoretical Computer Science* **777**, 155–183 (2019). <https://doi.org/10.1016/j.tcs.2019.02.001>
23. Crafa, S., Pirro, M.D., Zucca, E.: Is Solidity solid enough? In: Financial Cryptography Workshops. LNCS, vol. 11599, pp. 138–153. Springer (2019). https://doi.org/10.1007/978-3-030-43725-1_11
24. Delgado-Segura, S., Pérez-Solà, C., Navarro-Arribas, G., Herrera-Joancomartí, J.: A fair protocol for data trading based on bitcoin transactions. *Future Generation Computer Systems* (2017). <https://doi.org/10.1016/j.future.2017.08.021>
25. Grishchenko, I., Maffei, M., Schneidewind, C.: Foundations and tools for the static analysis of Ethereum smart contracts. In: CAV. LNCS, vol. 10981, pp. 51–78. Springer (2018). https://doi.org/10.1007/978-3-319-96145-3_4
26. Grishchenko, I., Maffei, M., Schneidewind, C.: A semantic framework for the security analysis of Ethereum smart contracts. In: Principles of Security and Trust (POST). LNCS, vol. 10804, pp. 243–269. Springer (2018). https://doi.org/10.1007/978-3-319-89722-6_10
27. Hildenbrandt, E., Saxena, M., Rodrigues, N., Zhu, X., Daian, P., Guth, D., Moore, B.M., Park, D., Zhang, Y., Stefanescu, A., Rosu, G.: KEVM: A complete formal semantics of the Ethereum Virtual Machine. In: IEEE Computer Security Foundations Symposium (CSF). pp. 204–217. IEEE Computer Society (2018). <https://doi.org/10.1109/CSF.2018.00022>
28. Hirai, Y.: Defining the Ethereum Virtual Machine for interactive theorem provers. In: Financial Cryptography Workshops. LNCS, vol. 10323, pp. 520–535. Springer (2017). https://doi.org/10.1007/978-3-319-70278-0_33
29. Jiao, J., Kan, S., Lin, S., Sanán, D., Liu, Y., Sun, J.: Semantic understanding of smart contracts: Executable operational semantics of Solidity. In: IEEE Symposium on Security and Privacy. pp. 1695–1712. IEEE (2020). <https://doi.org/10.1109/SP40000.2020.00066>
30. Klomp, R., Bracciali, A.: On symbolic verification of Bitcoin’s script language. In: Workshop on Cryptocurrencies and Blockchain Technology (CBT). LNCS, vol. 11025, pp. 38–56. Springer (2018). https://doi.org/10.1007/978-3-030-00305-0_3
31. Kumaresan, R., Bentov, I.: How to use Bitcoin to incentivize correct computations. In: ACM CCS. pp. 30–41 (2014). <https://doi.org/10.1145/2660267.2660380>

32. Kumaresan, R., Moran, T., Bentov, I.: How to use Bitcoin to play decentralized poker. In: ACM CCS. pp. 195–206 (2015). <https://doi.org/10.1145/2810103.2813712>
33. Luu, L., Chu, D.H., Olickel, H., Saxena, P., Hobor, A.: Making smart contracts smarter. In: ACM CCS. pp. 254–269 (2016). <https://doi.org/10.1145/2976749.2978309>
34. Miller, A., Bentov, I.: Zero-collateral lotteries in Bitcoin and Ethereum. In: EuroS&P Workshops. pp. 4–13 (2017). <https://doi.org/10.1109/EuroSPW.2017.44>
35. Rupic, K., Rozic, L., Derek, A.: Mechanized formal model of Bitcoin’s blockchain validation procedures. In: Workshop on Formal Methods for Blockchains (FMBC@CAV). OASICs, vol. 84, pp. 7:1–7:14. Schloss Dagstuhl - Leibniz-Zentrum fur Informatik (2020). <https://doi.org/10.4230/OASICs.FMBC.2020.7>
36. Sergey, I., Nagaraj, V., Johannsen, J., Kumar, A., Trunov, A., Hao, K.C.G.: Safer smart contract programming with Scilla. Proc. ACM Program. Lang. **3**(OOPSLA), 185:1–185:30 (2019). <https://doi.org/10.1145/3360611>

A The stateless ASC1 state machine

We assume the following sets:

- \mathbb{K} , the set of all users;
- \mathbb{S} , the set of all scripts;
- $\mathbb{X} = (\mathbb{K}^* \times \mathbb{N}) \cup \mathbb{S}$, the set of all addresses;
- \mathbb{A} , the set of all assets;
- \mathbb{T} , the set of all transactions;
- $\mathbb{U}_{64} = \{v \in \mathbb{Z} \mid 0 \leq v < 2^{64}\}$, the set of unsigned 64-bit integers.

We define the partial operator \circ on accounts as follows:

$$\sigma \circ v : \tau = \sigma\{\tau \mapsto \sigma(\tau) \circ v\} \quad \text{if } \tau \in \text{dom}(\sigma) \text{ and } \circ \in \{+, -\}$$

Note that the use of the \circ operator in an instance of the rules in Fig. 5 on a reachable state $\mathbf{x}[\sigma] \mid \dots$ never leads to overflows or underflows, since these rules correctly keep track of the number of assets. In particular, overflow never happens since the number of each asset is bounded (10^{16} for *Algo*, and the value *val* in *gen* transactions for the other assets).

We define a partial function $\text{auth}(t, f_{asst})$ which gives the manager of a transaction t :

$$\text{auth}(t, f_{asst}) = \begin{cases} t.\text{snd} & \text{if } t.\text{type} \in \{\text{close}, \text{pay}, \text{gen}, \text{optin}\} \\ f_{st}(f_{asst}(t.\text{asst})) & \text{if } t.\text{type} \in \{\text{rvk}, \text{frz}, \text{unfrz}, \text{burn}, \text{delegate}\} \end{cases}$$

Let $\nu_0, \nu_1 \in \mathbb{U}_{64} \cup \{\perp\}$. We define:

$$\nu_0 \circ_{\perp} \nu_1 \equiv \begin{cases} \nu_0 \circ \nu_1 & \text{if } \circ \in \{+, -, *\} \text{ and } \nu_0, \nu_1, \nu_0 \circ \nu_1 \in \mathbb{U}_{64} \\ \lfloor \nu_0 / \nu_1 \rfloor & \text{if } \circ = / \text{ and } \nu_0 \in \mathbb{U}_{64} \text{ and } \nu_1 \in \mathbb{U}_{64} \setminus \{0\} \\ \nu_0 \bmod \nu_1 & \text{if } \circ = \% \text{ and } \nu_0 \in \mathbb{U}_{64} \text{ and } \nu_1 \in \mathbb{U}_{64} \setminus \{0\} \\ 1 & \text{if } \circ \in \{<, \leq, =, \geq, >\} \text{ and } \nu_0, \nu_1 \in \mathbb{U}_{64} \text{ and } \nu_0 \circ \nu_1 \\ 0 & \text{if } \circ \in \{<, \leq, =, \geq, >\} \text{ and } \nu_0, \nu_1 \in \mathbb{U}_{64} \text{ and } \neg(\nu_0 \circ \nu_1) \\ 0 & \text{if } \circ = \text{and} \text{ and } \nu_0 = 0 \\ \nu_1 & \text{if } \circ = \text{and} \text{ and } \nu_0 \in \mathbb{U}_{64} \setminus \{0\} \\ 1 & \text{if } \circ = \text{or} \text{ and } \nu_0 \in \mathbb{U}_{64} \setminus \{0\} \\ \nu_1 & \text{if } \circ = \text{or} \text{ and } \nu_0 = 0 \\ \perp & \text{otherwise} \end{cases}$$

$$\neg_{\perp} \nu_0 = \begin{cases} 1 & \text{if } \nu_0 = 0 \\ 0 & \text{if } \nu_0 \in \mathbb{U}_{64} \setminus \{0\} \\ \perp & \text{otherwise} \end{cases}$$

A *blockchain state* Γ is a term with the following syntax:

$$\Gamma ::= \mathbf{x}[\sigma] \mid r \mid T_v \mid f_{asst} \mid f_{tx} \mid f_{frz} \mid \Gamma \mid \Gamma'$$

and subject to the following conditions:

- all the terms except $\mathbf{x}[\sigma]$ occur exactly once in a configuration;
- $r \in \mathbb{N}$ is the *current round*;
- $T_{\mathbf{lv}} \subseteq \mathbb{T}$ is the set of performed transactions whose “last valid” time \mathbf{lv} has not expired;
- $f_{\text{asst}} \in \mathbb{A} \rightarrow \mathbb{X} \times \mathbb{X}$ is the *asset map*;
- $f_{\text{lx}} \in \mathbb{X} \times \mathbb{N} \rightarrow \mathbb{N}$ is the *lease map*;
- $f_{\text{frz}} \in \mathbb{X} \rightarrow 2^{\mathbb{A}}$ is the *freeze map*;
- if $\mathbf{x}[\sigma]$ and $\mathbf{y}[\sigma']$ occur in a configuration, then $\mathbf{x} \neq \mathbf{y}$.
- configurations form a commutative monoid with respect to the composition operator $|$ (with identity 0).

We define in Fig. 5 a labelled transition relation \rightarrow between blockchain states, where labels are the following:

- \mathcal{T} performs a sequence of transactions;
- \checkmark advances one round.

B Additional smart contracts

B.1 Mutual HTLC

The *mutual HTLC* [11] is a variant of HTLC presented in §3.2: two users \mathbf{a} , \mathbf{b} choose their own secrets, pay a deposit, and the contract ensures that either (a) *both* users reveal their secret and get their deposits back, or (b) whoever does not reveal the secret loses the deposit (in favour of the other user).

Assuming that $h_{\mathbf{a}}$ and $h_{\mathbf{b}}$ are, respectively, the hashes of the secrets $s_{\mathbf{a}}$ and $s_{\mathbf{b}}$, we can implement the mutual HTLC by creating two HTLC contracts within an atomic group of transactions:

$$\{\text{snd} : \mathbf{a}, \text{rcv} : \text{HTLC}(\mathbf{a}, \mathbf{b}, h_{\mathbf{a}}), \dots\} \quad \{\text{snd} : \mathbf{b}, \text{rcv} : \text{HTLC}(\mathbf{b}, \mathbf{a}, h_{\mathbf{b}}), \dots\}$$

The mutual HTLC contract guarantees that once \mathbf{a} has stipulated it, then she will either learn \mathbf{b} ’s secret, or receive the compensation. Instead, if we create two instances of HTLC in a non-atomic way, this property is not guaranteed, since \mathbf{b} could refuse to stipulate its part of the contract.

Theorem 7 states the correctness of the mutual HTLC. For $\mathbf{p}, \mathbf{q} \in \{\mathbf{a}, \mathbf{b}\}$, let:

$$T_{\mathbf{p}, \mathbf{q}} = \left\{ t \mid \begin{array}{l} t.\text{type} = \text{close}, t.\text{snd} = \text{HTLC}(\mathbf{p}, \mathbf{q}, h_{\mathbf{p}}), t.\text{rcv} = \mathbf{p}, t.\text{asst} = \text{Algo} \end{array} \right\}$$

$$T'_{\mathbf{p}, \mathbf{q}} = \left\{ t \mid \begin{array}{l} t.\text{type} = \text{close}, t.\text{snd} = \text{HTLC}(\mathbf{p}, \mathbf{q}, h_{\mathbf{p}}), t.\text{rcv} = \mathbf{q}, t.\text{asst} = \text{Algo} \end{array} \right\}$$

and consider the following strategy for distinct users $\mathbf{p} \neq \mathbf{q} \in \{\mathbf{a}, \mathbf{b}\}$:

- $\Sigma_{\mathbf{p}}$: at a round $r < r_{\text{max}}$, send a transaction $t \in T_{\mathbf{p}, \mathbf{q}}$ with $t.\text{fv} = r$ and witness $s_{\mathbf{p}}$. Then, at round r_{max} check whether any transaction in $T_{\mathbf{q}, \mathbf{p}}$ occurs in \mathcal{R} : if not, immediately send a transaction $t \in T'_{\mathbf{q}, \mathbf{p}}$ with $t.\text{fv} = r_{\text{max}}$.

Theorem 7. *Let \mathcal{R} be a run conforming to some set of strategies Σ , such that: (i) \mathcal{R} reaches, before r_{max} , a state $\text{HTLC}(\mathbf{a}, \mathbf{b}, h_{\mathbf{a}})[\sigma] \mid \text{HTLC}(\mathbf{b}, \mathbf{a}, h_{\mathbf{b}})[\sigma'] \mid \dots$; (ii) \mathcal{R} reaches $r_{\text{max}} + 1$. Let $\mathbf{p} \neq \mathbf{q} \in \{\mathbf{a}, \mathbf{b}\}$. If $\Sigma_{\mathbf{p}} \in \Sigma$, then with overwhelming probability: (1) \mathcal{R} contains a transaction in $T_{\mathbf{p}, \mathbf{q}}$; (2) if \mathcal{R} does not contain the secret $s_{\mathbf{q}}$ before round $r_{\text{max}} + 1$, then \mathcal{R} contains a transaction in $T'_{\mathbf{q}, \mathbf{p}}$.*

B.2 Zero-collateral lottery

We show a variant of the two-players lottery in §3.2 which requires no collateral, similarly to [17,34]. The preconditions just require the 1Algo bets and the secrets, while the contract is the following:

$$\begin{aligned}
ZDL &\triangleq \text{tx.type} = \textit{close} \text{ and } \text{tx.asst} = \textit{Algo} \text{ and} \\
&\quad ((\text{tx.rcv} = ZDL2 \text{ and } H(\text{arg}(0)) = h_a) \text{ or} \\
&\quad (\text{tx.rcv} = \mathbf{b} \text{ and } \text{tx.fv} \geq r_0)) \\
ZDL2 &\triangleq \text{tx.type} = \textit{close} \text{ and } \text{tx.asst} = \textit{Algo} \text{ and } H(\text{arg}(0)) = h_a \text{ and} \\
&\quad (H(\text{arg}(1)) = h_b \text{ and} \\
&\quad \text{if } (\text{arg}(0) + \text{arg}(1))\%2 = 0 \text{ then } \text{tx.rcv} = \mathbf{a} \text{ else } \text{tx.rcv} = \mathbf{b}) \text{ or} \\
&\quad (\text{tx.rcv} = \mathbf{a} \text{ and } \text{tx.fv} \geq r_0 + r_1))
\end{aligned}$$

Here, \mathbf{b} must reveal first. If \mathbf{b} does not reveal his secret by the deadline r_0 , then \mathbf{a} can redeem the 2Algos stored in the contract. Otherwise, \mathbf{a} in turn must reveal by the deadline $r_0 + r_1$, or let \mathbf{b} redeem 2Algos. If both \mathbf{a} and \mathbf{b} reveal, then the winner is determined as a function of their secrets. As before, the rational strategy for each player is to reveal the secret. This makes the lottery fair, even in the absence of a collateral.

B.3 Escrow

User \mathbf{a} wants to buy from seller \mathbf{b} an item that costs v Algos. We want to guarantee that (1) \mathbf{b} will get paid if \mathbf{a} authorizes the payment; (2) \mathbf{a} will be refunded if \mathbf{b} authorizes it; (3) if neither \mathbf{a} nor \mathbf{b} give their authorization, an escrow service \mathbf{c} will resolve the dispute, by either fully refunding \mathbf{a} , or partially refunding \mathbf{a} with v_a Algos and giving the remaining $(v - v_a)$ Algos to \mathbf{b} .

$$\begin{aligned}
Escrow &\triangleq \text{tx.type} = \textit{close} \text{ and } \text{tx.asst} = \textit{Algo} \text{ and } (\\
&\quad \text{versig}(\text{txid}, \text{arg}(1), \mathbf{a}) \text{ and } (\text{tx.rcv} = \mathbf{b} \text{ or } \text{tx.rcv} = \textit{Resolve}) \\
&\quad \text{or } \text{versig}(\text{txid}, \text{arg}(1), \mathbf{b}) \text{ and } (\text{tx.rcv} = \mathbf{a} \text{ or } \text{tx.rcv} = \textit{Resolve})) \\
Resolve &\triangleq \text{tx.type} = \textit{pay} \text{ and } \text{tx.asst} = \textit{Algo} \text{ and } \text{versig}(\text{arg}(0), \text{arg}(1), \mathbf{c}) \text{ and} \\
&\quad ((\text{tx.rcv} = \mathbf{a} \text{ and } \text{tx.val} = \text{arg}(0)) \text{ or} \\
&\quad (\text{tx.rcv} = \mathbf{b} \text{ and } \text{tx.val} = v - \text{arg}(0)))
\end{aligned}$$

B.4 Limit order

The *limit order* contract [2] allows a user \mathbf{a} to exchange her Algos for units of a certain asset τ , provided by any user. The contract imposes a lower bound ρ_{min} to the exchange rate τ/Algo , and guarantees its operation as long as it has enough funds, and a deadline r_{max} is not reached (after then, \mathbf{a} can close it). To this purpose, the contract accepts two kinds of actions:

- an atomic group of two *pay* transactions: the first one transfers v_0 Algos from the contract to the sender of the second one; the second transaction transfers v_1 units of τ to user \mathbf{a} . The contract ensures that: (i) the ratio between v_1 and v_0 is greater than a given constant ρ_{min} ; (ii) v_0 is greater than a given

- constant v_{min} . Note that such a group can be issued by *any* user owning enough units of the asset τ , without requiring any interaction from \mathbf{a} .
- a single transaction where \mathbf{a} closes the contract after the deadline r_{max} .

The following script implements this specification:

```
(txlen = 2 and txpos = 0 and
 tx(0).type = pay and tx(0).asst = Algo and tx(0).rcv = tx(1).snd and
 tx(1).type = pay and tx(1).asst =  $\tau$  and tx(1).rcv =  $\mathbf{a}$  and
 tx(1).val/tx(0).val  $\geq$   $\rho_{min}$  and tx(0).val  $\geq$   $v_{min}$ )
or
(txlen = 1 and tx.fv >  $r_{max}$  and
 tx.type = close and tx.asst = Algo and tx(0).rcv =  $\mathbf{a}$ )
```

B.5 Split

The *split* contract [3] is created by a user \mathbf{a} , who want to transfer its funds to users \mathbf{b}_0 and \mathbf{b}_1 in a fixed ratio ρ . The contract is initially funded with some Algos from \mathbf{a} , and once started it accepts two kinds of actions:

- an atomic group of two *pay* transactions whose sender is the *split* contract: the first transaction transfers v_0 Algos to \mathbf{b}_0 , while the second one transfers v_1 Algos to \mathbf{b}_1 . The contract ensures that: (i) the ratio between v_1 and v_0 is equal to a given constant ρ ; (ii) v_0 is greater than a given constant v_{min} .
- a single transaction where \mathbf{a} closes the contract after a deadline r_{max} .

The following script implements this specification:

```
(txlen = 2 and tx(0).snd = tx(1).snd and
 tx(0).type = pay and tx(0).asst = Algo and tx(0).rcv =  $\mathbf{b}_0$  and
 tx(1).type = pay and tx(1).asst = Algo and tx(1).rcv =  $\mathbf{b}_1$  and
 tx(1).val =  $\rho * tx(0).val$  and tx(0).val  $\geq$   $v_{min}$ )
or
(txlen = 1 and tx.fv >  $r_{max}$  and
 tx.type = close and tx.asst = Algo and tx(0).rcv =  $\mathbf{a}$ )
```

C Compiling scripts from our model to concrete ASC1

Table 2 summarizes the main checks performed by the script, and the values stored in the scratch space. For brevity, we assume a sugared syntax of TEAL (e.g., we use comparison operators in infix notation), and we do not detail the checks needed to ensure that certain fields are not set. Notice that some of the TEAL operators used in Table 2, like e.g., `ConfigAsset`, require `LogicSigVer \geq 2` (see <https://developer.algorand.org/docs/reference/teal/opcodes/>). When none of the conditions in the left column of Table 2 is satisfied, we make the script generate an error, which corresponds to requiring $\llbracket tx(n).f \rrbracket = \perp$.

TEAL checks	Values stored in scratch space
<pre> gtxn n TypeEnum == pay gtxn n Amount == 0 gtxn n CloseRemainderTo != 0 ... </pre>	<pre> tx(n).type := close tx(n).asst := Algo tx(n).snd := gtxn n Sender tx(n).rcv := gtxn n CloseRemainderTo </pre>
<pre> gtxn n TypeEnum == pay gtxn n CloseRemainderTo == 0 ... </pre>	<pre> tx(n).type := pay tx(n).asst := Algo tx(n).snd := gtxn n Sender tx(n).rcv := gtxn n Receiver tx(n).val := gtxn n Amount </pre>
<pre> gtxn n TypeEnum == axfer gtxn n AssetSender != 0 gtxn n AssetAmount == 0 gtxn n AssetCloseTo != 0 ... </pre>	<pre> tx(n).type := close tx(n).asst := gtxn n XferAsset tx(n).snd := gtxn n AssetSender tx(n).rcv := gtxn n AssetCloseTo </pre>
<pre> gtxn n TypeEnum == axfer gtxn n Sender != 0 gtxn n AssetSender == 0 gtxn n AssetAmount != 0 gtxn n AssetCloseTo == 0 ... </pre>	<pre> tx(n).type := pay tx(n).asst := gtxn n XferAsset tx(n).snd := gtxn n Sender tx(n).rcv := gtxn n AssetReceiver tx(n).val := gtxn n AssetAmount </pre>
<pre> gtxn n TypeEnum == acfg gtxn n ConfigAsset == 0 (gtxn n ConfigAssetManager == gtxn n ConfigAssetFreeze == gtxn n ConfigAssetReserve == gtxn n ConfigAssetClawback) ... </pre>	<pre> tx(n).type := gen tx(n).snd := gtxn n Sender tx(n).rcv := gtxn n ConfigAssetManager tx(n).val := gtxn n AssetAmount </pre>
<pre> gtxn n TypeEnum == axfer gtxn n Sender == gtxn n AssetReceiver gtxn n AssetSender == 0 gtxn n AssetAmount == 0 ... </pre>	<pre> tx(n).type := optin tx(n).asst := gtxn n XferAsset tx(n).snd := gtxn n Sender </pre>
<pre> gtxn n TypeEnum == acfg len (gtxn n ConfigAssetManager) == 0 len (gtxn n ConfigAssetFreeze) == 0 len (gtxn n ConfigAssetReserve) == 0 len (gtxn n ConfigAssetClawback) == 0 ... </pre>	<pre> tx(n).type := burn tx(n).asst := gtxn n ConfigAsset </pre>
<pre> gtxn n TypeEnum == axfer gtxn n Sender != 0 gtxn n AssetSender != 0 ... </pre>	<pre> tx(n).type := rvk tx(n).asst := gtxn n XferAsset tx(n).snd := gtxn n AssetSender tx(n).rcv := gtxn n AssetReceiver tx(n).val := gtxn n AssetAmount </pre>
<pre> gtxn n TypeEnum == afrz gtxn n AssetFrozen == true ... </pre>	<pre> tx(n).type := frz tx(n).asst := gtxn n FreezeAsset tx(n).snd := gtxn n FreezeAccount </pre>
<pre> gtxn n TypeEnum == afrz gtxn n AssetFrozen == false ... </pre>	<pre> tx(n).type := unfrz tx(n).asst := gtxn n FreezeAsset tx(n).snd := gtxn n FreezeAccount </pre>
<pre> gtxn n TypeEnum == acfg gtxn n ConfigAsset != 0 len (gtxn n ConfigAssetManager) != 0 (gtxn n ConfigAssetManager == gtxn n ConfigAssetFreeze == gtxn n ConfigAssetReserve == gtxn n ConfigAssetClawback) ... </pre>	<pre> tx(n).type := delegate tx(n).snd := gtxn n Sender tx(n).rcv := gtxn n ConfigAssetManager tx(n).asst := gtxn n ConfigAsset </pre>

Table 2: Compilation of the script $\text{tx}(n).f$ in TEAL (sketch).

D Proofs

Proof of Theorem 1

By induction on the length of the run.

Conditions **c1** and **c2** are required at each step.

Condition **c1** prevents the same transaction t to appear again for the next Δ_{max} rounds ($t \notin T_{lv}$). Condition **c2** requires t to be valid only in a time window of at most Δ_{max} rounds ($t.lv - t.fv \leq \Delta_{max}$). Therefore, it is impossible for t to appear more than once. \square

Proof of Theorem 2

By induction on the length of the run, and by cases on the rules generating each step.

By inspecting the rules, we can see that they all preserve the value for all assets, except for $[\text{BURN}]$ which completely destroys the asset, and $[\text{GEN}]$ which can create a fresh asset. The rules prevent a burnt asset to be re-created later on.

The two special cases are taken into account in the definition of $val_\tau(I)$, so the theorem holds. \square

Proof of Theorem 3

By induction on the length of the run, and by cases on the rules generating each step.

Inspecting each pair of (different) rules, we can observe that they either have distinct labels or they have conflicting side conditions, so that at most one of them can apply in each case. Further, we observe that in each rule the final state is a function of the label and of the initial state.

For part 3 of the theorem involving different labels $\mathbf{W}:\mathcal{J}$ and $\mathbf{W}':\mathcal{J}$, we further add that $(\Gamma, K) \xrightarrow{\mathbf{W}:\mathcal{J}} (\Gamma', K)$ is possible only when $\Gamma \xrightarrow{\mathcal{J}} \Gamma'$ holds, and the same holds for $\mathbf{W}':\mathcal{J}$. We conclude that the difference between \mathbf{W} and \mathbf{W}' is immaterial, since the determinism of \Rightarrow is derived from the determinism of \rightarrow which does not involve witnesses. \square

Lemma 1. *Let $e = (\text{tx.type} = \text{close}$ and e'), and let \mathcal{R} be a run passing through a state (Γ, K) with $\Gamma = e[\sigma] \mid \dots$ and leading to (Γ', K') with $\Gamma' = e[\sigma'] \mid T'_{lv} \mid r' \mid f'_{lx} \dots$, such that all intermediate states contain the account $e[\dots]$. Then:*

1. $\text{dom } \sigma' = \{\text{Algo}\};$
2. $\sigma'(\text{Algo}) \geq \sigma(\text{Algo});$
3. *if t is a transaction such that $t.\text{snd} = e$, $t.\text{asst} = \text{Algo}$, $t \notin T'_{lv}$, and $f'_{lx}, r' \models t$, and $\llbracket e \rrbracket_t^{\mathbf{W}} = \text{true}$ with \mathbf{W} in K' , then $(\Gamma', K') \xrightarrow{\mathbf{W}:t}$.*

Proof. By construction, the script e only allows *close* transactions, and in particular it does not accept *optin* or *gen* transactions. Therefore, it must be $\text{dom } \sigma' = \{\text{Algo}\}$, proving item 1. Further, the value $\sigma'(\text{Algo})$ can only increase until e is closed, proving item 2. For item 3, since by hypothesis e evaluates to *true*, then it must be $t.\text{type} = \text{close}$. Then, by definition of *auth*, it must be $\text{auth}(t, \dots) = e$. Therefore, $\mathcal{W} \models t$ holds, since $\llbracket e \rrbracket_t^{\mathcal{W}} = \text{true}$.

Note that, if the script e does not constrain the *lx* field, the hypothesis $t \notin T'_V$ in item 3 can always be satisfied by setting a value of $t.\text{lx}$ not used in T'_V . Further, if e does not check signatures on t , doing this does not invalidate the possibility of appending t . Therefore, in the following proofs we will omit checking that the hypothesis $t \notin T'_V$ holds.

Proof of Theorem 4

For item (1), assume that $\Sigma_a, \Sigma_o \in \Sigma$. By hypothesis, o has not sent a signature $\text{sig}_o(\text{Oracle}, 1)$. By Lemma 1, in any reachable state $\text{Oracle}[\sigma'] \mid \dots$ it must be $\text{dom } \sigma' = \{\text{Algo}\}$, and the value $\sigma'(\text{Algo})$ can only increase until *Oracle* is closed. Consequently, to prove item (1) it enough to show that there exists some transaction t in \mathcal{R} such that $t \in T_a$, and t is the first transaction in \mathcal{R} whose sender is *Oracle*. By the monotonicity of the balance of *Oracle* discussed before, this implies the thesis, i.e. that t transfers at least $\sigma(\text{Algo})$ to a . We have the following two cases:

- (a) o has sent a signature s on $(\text{Oracle}, 0)$ at round $r \leq r_{max}$. By contradiction, assume that *Oracle* has not been closed in \mathcal{R} within round r . Since \mathcal{R} conforms to Σ_a , at round r , a has sent a transaction $t' \in T_a$ with $t'.\text{fv} = r$ and witness $0s$. Since no transactions have closed *Oracle*, the contract is still open, so let σ' be its balance. As observed above, $\text{dom } \sigma' = \{\text{Algo}\}$, and so t' satisfies the premises of either rule [CLOSE-OPEN] or [CLOSE-PAY]. Further, t' satisfies the script *Oracle*, so by item 3 of Lemma 1 it closes *Oracle* at round r — contradiction. Therefore, \mathcal{R} contains at least a transaction which closes *Oracle* within round r : let t be the first one of these transactions. Since t must satisfy the script *Oracle*, one of the following cases must apply:
 - (i) $t.\text{fv} > r_{max}$: impossible, because t has closed *Oracle* within $r \leq r_{max}$;
 - (ii) t is validated by o 's signature on $(\text{Oracle}, 1)$: impossible, since it would contradict the hypothesis of item (1);
 - (iii) t is validated by o 's signature on $(\text{Oracle}, 0)$: this implies that $t \in T_a$ — which proves the thesis.
- (b) o has *not* sent a signature on $(\text{Oracle}, 0)$ at any round $r \leq r_{max}$. By contradiction, assume that *Oracle* has not been closed in \mathcal{R} within round $r_{max} + 1$. Since \mathcal{R} conforms to Σ_a , at round $r_{max} + 1$, a has sent a transaction $t' \in T_a$ with $t'.\text{fv} = r_{max} + 1$. Since no transactions have closed *Oracle*, the contract is still open, so let σ' be its balance. As observed above, $\text{dom } \sigma' = \{\text{Algo}\}$, and so t' satisfies the premises of either rule [CLOSE-OPEN] or [CLOSE-PAY]. Further, t' satisfies the script *Oracle*, so by item 3 of Lemma 1 it closes *Oracle* at

- round $r_{max} + 1$ — contradiction. Therefore, \mathcal{R} contains at least a transaction which closes *Oracle* within round $r_{max} + 1$: let t be the first one of these transactions. Since t must satisfy the script *Oracle*, one of the following cases must apply:
- (i) $t.fv > r_{max}$ and $t.rcv = a$: this implies that $t \in T_a$ — which proves the thesis.
 - (ii) t is validated by o 's signature on $(Oracle, 1)$: impossible, since it would contradict the hypothesis of item (1);
 - (iii) t is validated by o 's signature on $(Oracle, 0)$: impossible, since it would contradict the hypothesis of case (b) of item (1).

For item (2), assume that $\Sigma_b, \Sigma_o \in \Sigma$. By Lemma 1, in any reachable state $Oracle[\sigma'] \mid \dots$ it must be $\text{dom } \sigma' = \{Algo\}$, and the value $\sigma'(Algo)$ can only increase until *Oracle* is closed. Consequently, to prove item (2) it enough to show that there exists some transaction t in \mathcal{R} such that $t \in T_b$, and t is the first transaction in \mathcal{R} whose sender is *Oracle*. By the monotonicity of the balance of *Oracle* discussed before, this implies the thesis, i.e. that t transfers at least $\sigma(Algo)$ to b .

By hypothesis, o has sent a signature s' on $(Oracle, 1)$ at round $r \leq r_{max}$. By contradiction, assume that *Oracle* has not been closed in \mathcal{R} within round r . Since \mathcal{R} conforms to Σ_b , at round r , b has sent a transaction $t' \in T_b$ with $t'.fv = r$ and witness $1 s'$. Since no transactions have closed *Oracle*, the contract is still open, so let σ' be its balance. As observed above, $\text{dom } \sigma' = \{Algo\}$, and so t' satisfies the premises of either rule [CLOSE-OPEN] or [CLOSE-PAY]. Further, t' satisfies the script *Oracle*, so by item 3 of Lemma 1 it closes *Oracle* at round r — contradiction. Therefore, \mathcal{R} contains at least a transaction which closes *Oracle* within round r : let t be the first one of these transactions. Since t must satisfy the script *Oracle*, one of the following cases must apply:

- (i) $t.fv > r_{max}$: impossible, because t has closed *Oracle* within round $r \leq r_{max}$;
- (ii) t is validated by o 's signature on $(Oracle, 1)$: this implies that $t \in T_b$ — which proves the thesis.
- (iii) t is validated by o 's signature on $(Oracle, 0)$: impossible, since \mathcal{R} would not conform to Σ_o , which requires that the oracle never signs both 0 and 1. \square

Proof of Theorem 5

For item (i), assume that $\Sigma_a \in \Sigma$. By contradiction, assume that \mathcal{R} does not contain a transaction in T_a . Since \mathcal{R} conforms to Σ_a , at round r , user a has sent a transaction $t \in T_a$ with $t.fv = r$ and witness s_a . By construction, t satisfies the script *HTLC*. Hence, by Lemma 1, to prove the thesis we show that *HTLC* was not closed earlier in \mathcal{R} . By inspecting the script *HTLC*, there are only two conditions under which a previous transaction t' can close the contract:

1. $t'.fv > r_{max}$: impossible, because $t.fv = r \leq r_{max}$;
2. t' is validated using a preimage of h_a : this implies that $t' \in T_a$ — contradiction.

For item (ii), assume that $\Sigma_b \in \Sigma$, and that s_a was not revealed before round $r_{max} + 1$. By contradiction, assume that \mathcal{R} does not contain a transaction in T_b . Since \mathcal{R} conforms to Σ_b , at round r_{max} , user b has sent a transaction $t \in T_b$ with $t.fv = r_{max}$. By construction, t satisfies the script *HTLC*. Hence, by Lemma 1, to prove the thesis we show that *HTLC* was not closed earlier in \mathcal{R} . By inspecting the script *HTLC*, there are only two conditions under which a previous transaction t' can close the contract:

1. $t'.fv > r_{max}$: this implies that $t' \in T_b$ — contradiction.
2. t' is validated using a preimage of h_a : this implies, with overwhelming probability, that s_a was used as the preimage. In such case, it was revealed and occurs in \mathcal{R} at round r_{max} — contradiction.

Proof of Theorem 6

The proof for the fairness of the lottery protocol when implemented on Bitcoin appeared in [11, 12]. This protocol relies on running, simultaneously, two timed commitments (for which we use *HTLC*) and a contract that transfers the bets to the winner, which can be computed after the secrets have been revealed (for which we use *Lottery*).

We already proved the security of *HTLC* in Theorem 5, which implies part 1 of the thesis: a honest participant p will reveal her secret in time with a transaction in $T_{p,q}^{secr}$.

For part 2, the argument is similar to the one for the original protocol. Briefly put, if the other participant q does not to reveal s_q in time, p wins the lottery by timeout using a transaction in $T_{q,p}^{tout}$. If instead s_q is revealed in time, since s_p was chosen in a uniformly random way, independently from s_q (which has a different hash), we have that the parity of $s_p + s_q$ is a random bit, uniformly distributed. Hence, if p won, her strategy makes her send a transaction in T_p^{lott} and claim the pot.

In the argument above we neglected the case where q reveals another preimage than s_q as his secret, since that can only happen with negligible probability. Summing up, an honest p wins with at least 1/2 probability (up-to a negligible quantity).

Proof of Theorem 7

The proof is analogous to the one for HTLC in Theorem 5.

$$\begin{array}{c}
 t.type = \text{pay} \quad t.snd = \mathbf{x} \quad t.asst = \text{Algo} \quad t.rcv = \mathbf{y} \quad f_{lx}, r \models t \quad t \notin T_v \\
 \hline
 \models \sigma - t.val : \text{Algo} \quad \sigma' = \{\text{Algo} \mapsto t.val\} \quad \models \sigma' \quad \mathbf{y}[\cdot\cdot\cdot] \text{ not in } \Gamma \\
 \hline
 \mathbf{x}[\sigma] \mid \Gamma \mid r \mid T_v \mid f_{asst} \mid f_{lx} \mid f_{frz} \xrightarrow{t}_1 \\
 \mathbf{x}[\sigma - t.val : \text{Algo}] \mid \mathbf{y}[\sigma'] \mid \Gamma \mid r \mid T_v \cup \{t\} \mid f_{asst} \mid upd(f_{lx}, t) \mid f_{frz} \\
 \text{[PAY-OPEN]}
 \end{array}$$

$$\begin{array}{c}
 t.type = \text{close} \quad t.snd = \mathbf{x} \quad t.rcv = \mathbf{y} \quad t.asst = \text{Algo} \quad f_{lx}, r \models t \quad t \notin T_v \\
 \hline
 \text{dom}(\sigma) = \{\text{Algo}\} \quad \mathbf{y}[\cdot\cdot\cdot] \text{ not in } \mathbf{x}[\sigma] \mid \Gamma \\
 \hline
 \mathbf{x}[\sigma] \mid \Gamma \mid r \mid T_v \mid f_{asst} \mid f_{lx} \mid f_{frz} \xrightarrow{t}_1 \\
 \mathbf{y}[\sigma] \mid \Gamma \mid r \mid T_v \cup \{t\} \mid f_{asst} \mid upd(f_{lx}, t) \mid f_{frz} \\
 \text{[CLOSE-OPEN]}
 \end{array}$$

$$\begin{array}{c}
 t.type = \text{close} \quad t.snd = \mathbf{x} \quad t.rcv = \mathbf{y} \quad t.asst = \text{Algo} \quad f_{lx}, r \models t \quad t \notin T_v \\
 \hline
 \text{dom}(\sigma) = \{\text{Algo}\} \\
 \hline
 \mathbf{x}[\sigma] \mid \mathbf{y}[\sigma'] \mid \Gamma \mid r \mid T_v \mid f_{asst} \mid f_{lx} \mid f_{frz} \xrightarrow{t}_1 \\
 \mathbf{y}[\sigma' + \sigma(\text{Algo}) : \text{Algo}] \mid \Gamma \mid r \mid T_v \cup \{t\} \mid f_{asst} \mid upd(f_{lx}, t) \mid f_{frz} \\
 \text{[CLOSE-PAY]}
 \end{array}$$

$$\begin{array}{c}
 t.type = \text{close} \quad t.snd = \mathbf{x} \quad t.rcv = \mathbf{y} \quad t.asst = \tau \quad f_{lx}, r \models t \quad t \notin T_v \\
 \hline
 \tau \neq \text{Algo} \quad \sigma(\tau) = v \quad \tau \in \text{dom}(\sigma') \quad \models \sigma' + v : \tau \quad \tau \notin f_{frz}(\mathbf{x}) \cup f_{frz}(\mathbf{y}) \\
 \hline
 \mathbf{x}[\sigma] \mid \mathbf{y}[\sigma'] \mid \Gamma \mid r \mid T_v \mid f_{asst} \mid f_{lx} \mid f_{frz} \xrightarrow{t}_1 \\
 \mathbf{x}[\sigma\{\tau \mapsto \perp\}] \mid \mathbf{y}[\sigma' + v : \tau] \mid \Gamma \mid r \mid T_v \cup \{t\} \mid f_{asst} \mid upd(f_{lx}, t) \mid f_{frz} \\
 \text{[CLOSE-ASST]}
 \end{array}$$

$$\begin{array}{c}
 t.type = \text{pay} \quad t.snd = \mathbf{x} \quad t.rcv = \mathbf{y} \quad t.val = v > 0 \quad t.asst = \tau \quad f_{lx}, r \models t \quad t \notin T_v \\
 \hline
 \models \sigma - v : \tau \quad \tau \in \text{dom}(\sigma) \cap \text{dom}(\sigma') \quad \tau \notin f_{frz}(\mathbf{x}) \cup f_{frz}(\mathbf{y}) \\
 \hline
 \mathbf{x}[\sigma] \mid \mathbf{y}[\sigma'] \mid \Gamma \mid r \mid T_v \mid f_{asst} \mid f_{lx} \mid f_{frz} \xrightarrow{t}_1 \\
 \mathbf{x}[\sigma - v : \tau] \mid \mathbf{y}[\sigma' + v : \tau] \mid \Gamma \mid r \mid T_v \cup \{t\} \mid f_{asst} \mid upd(f_{lx}, t) \mid f_{frz} \\
 \text{[PAY]}
 \end{array}$$

$$\begin{array}{c}
 t.type = \text{pay} \quad t.snd = \mathbf{x} \quad t.rcv = \mathbf{y} \neq \mathbf{x} \quad t.val = 0 \quad f_{lx}, r \models t \quad t \notin T_v \\
 \hline
 \mathbf{x}[\sigma] \mid \mathbf{y}[\sigma'] \mid \Gamma \mid r \mid T_v \mid f_{asst} \mid f_{lx} \mid f_{frz} \xrightarrow{t}_1 \\
 \mathbf{x}[\sigma] \mid \mathbf{y}[\sigma'] \mid \Gamma \mid r \mid T_v \cup \{t\} \mid f_{asst} \mid upd(f_{lx}, t) \mid f_{frz} \\
 \text{[PAY-ZERO]}
 \end{array}$$

$$\begin{array}{c}
 t.type = \text{gen} \quad t.snd = \mathbf{x} \quad t.rcv = \mathbf{y} \quad val = v \quad f_{lx}, r \models t \quad t \notin T_v \\
 \hline
 \models \sigma\{\tau \mapsto v\} \quad \tau \text{ next fresh asset identifier} \\
 \hline
 \mathbf{x}[\sigma] \mid \Gamma \mid r \mid T_v \mid f_{asst} \mid f_{lx} \mid f_{frz} \xrightarrow{t}_1 \\
 \mathbf{x}[\sigma\{\tau \mapsto v\}] \mid \Gamma \mid r \mid T_v \cup \{t\} \mid f_{asst}\{\tau \mapsto (\mathbf{y}, \mathbf{x})\} \mid upd(f_{lx}, t) \mid f_{frz} \\
 \text{[GEN]}
 \end{array}$$

$$\begin{array}{c}
 t.type = \text{optin} \quad t.snd = \mathbf{x} \quad t.asst = \tau \quad f_{lx}, r \models t \quad t \notin T_v \\
 \hline
 \tau \text{ occurs in } \mathbf{x}[\sigma] \mid \Gamma \quad \sigma' = \begin{cases} \sigma\{\tau \mapsto 0\} & \text{if } \tau \notin \text{dom}(\sigma) \\ \sigma & \text{otherwise} \end{cases} \quad \models \sigma' \\
 \hline
 \mathbf{x}[\sigma] \mid \Gamma \mid r \mid T_v \mid f_{asst} \mid f_{lx} \mid f_{frz} \xrightarrow{t}_1 \\
 \mathbf{x}[\sigma'] \mid \Gamma \mid r \mid T_v \cup \{t\} \mid f_{asst} \mid upd(f_{lx}, t) \mid f_{frz} \\
 \text{[OPTIN]}
 \end{array}$$

$$\begin{array}{c}
\frac{t.\text{type} = \text{burn} \quad t.\text{asst} = \tau \quad \text{snd}(f_{\text{asst}}(\tau)) = \mathbf{x} \quad f_{lx}, r \models t \quad t \notin T_{lv} \\
\tau \in \text{dom}(\sigma) \quad \text{for all } \mathbf{y}[\sigma'] \text{ in } \Gamma : \tau \notin \text{dom}(\sigma')}{\mathbf{x}[\sigma] \mid \Gamma \mid r \mid T_{lv} \mid f_{\text{asst}} \mid f_{lx} \mid f_{frz} \xrightarrow{t} \mathbf{1}} \text{ [BURN]} \\
\mathbf{x}[\sigma\{\tau \mapsto \perp\}] \mid \Gamma \mid r \mid T_{lv} \cup \{t\} \mid f_{\text{asst}}\{\tau \mapsto \perp\} \mid \text{upd}(f_{lx}, t) \mid f_{frz}
\end{array}$$

$$\begin{array}{c}
t.\text{type} = \text{rvk} \quad t.\text{snd} = \mathbf{x} \quad t.\text{rcv} = \mathbf{y} \quad t.\text{val} = v \quad t.\text{asst} = \tau \quad f_{lx}, r \models t \quad t \notin T_{lv} \\
\tau \in \text{dom}(\sigma) \cap \text{dom}(\sigma') \quad \models \sigma - v : \tau \quad \tau \notin f_{frz}(\mathbf{x}) \cup f_{frz}(\mathbf{y})}{\mathbf{x}[\sigma] \mid \mathbf{y}[\sigma'] \mid \Gamma \mid r \mid T_{lv} \mid f_{\text{asst}} \mid f_{lx} \mid f_{frz} \xrightarrow{t} \mathbf{1}} \text{ [REVOKE]} \\
\mathbf{x}[\sigma - v : \tau] \mid \mathbf{y}[\sigma' + v : \tau] \mid \Gamma \mid r \mid T_{lv} \cup \{t\} \mid f_{\text{asst}} \mid \text{upd}(f_{lx}, t) \mid f_{frz}
\end{array}$$

$$\begin{array}{c}
t.\text{type} = \text{frz} \quad t.\text{snd} = \mathbf{x} \quad t.\text{asst} = \tau \quad f_{lx}, r \models t \quad t \notin T_{lv} \\
\tau \in \text{dom}(\sigma) \quad f'_{frz} = f_{frz}\{\mathbf{x} \mapsto f_{frz}(\mathbf{x}) \cup \{\tau\}\}}{\mathbf{x}[\sigma] \mid \Gamma \mid r \mid T_{lv} \mid f_{\text{asst}} \mid f_{lx} \mid f_{frz} \xrightarrow{t} \mathbf{1}} \text{ [FREEZE]} \\
\mathbf{x}[\sigma] \mid \Gamma \mid r \mid T_{lv} \cup \{t\} \mid f_{\text{asst}} \mid \text{upd}(f_{lx}, t) \mid f'_{frz}
\end{array}$$

$$\begin{array}{c}
t.\text{type} = \text{unfrz} \quad t.\text{snd} = \mathbf{x} \quad t.\text{asst} = \tau \quad f_{lx}, r \models t \quad t \notin T_{lv} \\
\tau \in \text{dom}(\sigma) \quad f'_{frz} = f_{frz}\{\mathbf{x} \mapsto f_{frz}(\mathbf{x}) \setminus \{\tau\}\}}{\mathbf{x}[\sigma] \mid \Gamma \mid r \mid T_{lv} \mid f_{\text{asst}} \mid f_{lx} \mid f_{frz} \xrightarrow{t} \mathbf{1}} \text{ [UNFREEZE]} \\
\mathbf{x}[\sigma] \mid \Gamma \mid r \mid T_{lv} \cup \{t\} \mid f_{\text{asst}} \mid \text{upd}(f_{lx}, t) \mid f'_{frz}
\end{array}$$

$$\begin{array}{c}
t.\text{type} = \text{delegate} \quad t.\text{snd} = \mathbf{x} \quad t.\text{rcv} = \mathbf{y} \quad t.\text{asst} = \tau \quad f_{lx}, r \models t \quad t \notin T_{lv} \\
\text{fst}(f_{\text{asst}}(\tau)) = \mathbf{x}}{\Gamma \mid r \mid T_{lv} \mid f_{\text{asst}} \mid f_{lx} \mid f_{frz} \xrightarrow{t} \mathbf{1}} \text{ [DELEGATE]} \\
\Gamma \mid r \mid T_{lv} \cup \{t\} \mid f_{\text{asst}}\{\tau \mapsto (\mathbf{y}, \text{snd}(f_{\text{asst}}(\tau)))\} \mid \text{upd}(f_{lx}, t) \mid f_{frz}
\end{array}$$

$$\frac{\Gamma \mid T_{lv} \mid r \xrightarrow{\checkmark} \mathbf{1} \quad \Gamma \mid \{t \in T_{lv} \mid t.lv > r\} \mid r + 1}{\Gamma \mid T_{lv} \mid r \xrightarrow{\checkmark} \mathbf{1}} \text{ [ROUND]}$$

$$\frac{\mathcal{J} = t_1 \cdots t_n \quad \Gamma \xrightarrow{t_1} \Gamma_1 \quad \cdots \quad \Gamma_{n-1} \xrightarrow{t_n} \Gamma_n}{\Gamma \xrightarrow{\mathcal{J}} \Gamma_n} \text{ [TXG]}$$

$$\frac{}{(\Gamma, K) \xrightarrow{w} (\Gamma, K \cup \{w\})} \text{ [NET-WIT]}$$

$$\frac{\Gamma \xrightarrow{\checkmark} \Gamma'}{(\Gamma, K) \xrightarrow{\checkmark} (\Gamma', K)} \text{ [NET-ROUND]}$$

$$\frac{\Gamma \xrightarrow{\mathcal{J}} \Gamma' \quad \text{set}(\mathbf{W}) \subseteq K \quad \mathbf{W} \models \mathcal{J}}{(\Gamma, K) \xrightarrow{\mathbf{W}:\mathcal{J}} (\Gamma', K)} \text{ [NET-TXG]}$$

Fig. 5: The stateless ASC1 state machine.