

A formal model of Bitcoin transactions

Nicola Atzei¹, Massimo Bartoletti¹, Stefano Lande¹, Roberto Zunino²

¹ Università degli Studi di Cagliari, Cagliari, Italy

² Università degli Studi di Trento, Trento, Italy

Abstract. We propose a formal model of Bitcoin transactions, which is sufficiently abstract to enable formal reasoning, and at the same time is concrete enough to serve as an alternative documentation to Bitcoin. We use our model to formally prove some well-formedness properties of the Bitcoin blockchain, for instance that each transaction can only be spent once. We release an open-source tool through which programmers can write transactions in our abstract model, and compile them into standard Bitcoin transactions.

1 Introduction

In recent years we have observed a growing interest around *cryptocurrencies*. Bitcoin [10], the first decentralized cryptocurrency, was introduced in 2009, and through the years it has consolidated its position as the most popular one. Bitcoin and other cryptocurrencies have pushed forward the concept of decentralization, providing means for reliable interactions between mutually distrusting parties on an open network.

The nodes of the Bitcoin network maintain a public and immutable data structure, called *blockchain*. The blockchain stores the historical record of all transfers of bitcoins, which are referred to as *transactions*. When a node updates the blockchain, the other nodes verify if the appended transactions are valid, e.g. by checking if the conditions specified in *scripts* are satisfied. Scripts are programmable boolean functions: in their standard (and mostly used) form they verify a digital signature against a public key. Since the blockchain is immutable, tampering with a stored transaction would result in the invalidation of all the subsequent ones. Updating the state of the blockchain, i.e. appending new transactions, requires solving a moderately difficult cryptographic puzzle. In case of conflicting updates, the chain that required the largest computational effort is considered the valid one. Hence, the immutability and the consistency of the blockchain is bounded by the total computational power of honest nodes. An adversary with enough resources can append invalid transactions, e.g. with incorrect digital signatures, or rewrite a part of the blockchain, e.g. to perform a *double-spending attack*. The attack consists in paying someone by publishing a transaction on the blockchain, and then removing it (making the funds unspent).

Besides the intended monetary application, the Bitcoin blockchain can be seen as a way to consistently maintain the state of a system over a peer-to-peer network, without the need of a trusted authority. If the system is a currency, its

state is the amount of funds in each account. This concept can be generalised to the case where the system is a *smart contract* [12], namely an executable computer protocol which can also handle transfers of currency. The idea of exploiting the Bitcoin blockchain to build smart contracts has recently been explored by several works. Lotteries [2, 4, 5, 8], gambling games [7], contingent payments [3], covenants [9, 11], and other kinds of fair computations [1, 6] are some examples of the capabilities of Bitcoin as a smart contracts platform.

Smart contracts often rely on features of Bitcoin that go beyond the standard transfers of currency. For instance, while the vast majority of Bitcoin transactions uses scripts only to verify signatures, smart contracts like the ones above exploit more complex scripts, e.g. to determine the winner of a lottery, or to check if a secret has been revealed. Smart contracts may also exploit other (infrequently used) features of Bitcoin, e.g. various signature modifiers, and temporal constraints on transactions.

As a matter of fact, using these advanced features to design a new smart contract is not a trivial matter, for two reasons. First, while the overall behaviour of Bitcoin is clear, the details of many of its crucial aspects is poorly documented. To understand the details of how a mechanism actually works, one has to explore various web pages (often inaccurate, or inconsistent, or overly technical), and eventually resort to the source code of the Bitcoin client³ to have the correct answer. Second, the description of advanced features is often too concrete to be effectively used in the design and analysis of a smart contract (indeed, in many cases the only available description coincides with the implementation).

Contributions. We propose a formal model of Bitcoin transactions. This model is abstract enough to allow for formal reasoning on the behaviour of Bitcoin transactions. For instance, we use our model to formally prove some properties of the Bitcoin blockchain, e.g. that transactions cannot be spent twice (Theorem 1), and that the overall value contained in the blockchains (excluding the coinbase transactions) is decreasing (Theorem 2).

Our model formally specifies some poorly documented features of Bitcoin, e.g. transaction signatures and signature modifiers (Definition 4), output scripts (Definitions 1 and 7), multi-signature verification (Definition 6), Segregated Witnesses (Definitions 2 and 9), paving the way towards automatic verification.

We make available an open-source tool⁴ which translates transactions specified in our model to *standard* Bitcoin transactions.

Structure of the paper. Section 2 briefly recaps Bitcoin transactions, which we formalise in Section 3. Besides transactions, we also provide an high-level model of the blockchain, and we study its basic properties. In Section 4 we illustrate, through a basic case study, the impact of the Segregated Witness feature on the expressiveness of Bitcoin smart contracts. In Section 5 we show how to translate transactions from our model to standard Bitcoin transactions. We discuss the differences between our model and the actual Bitcoin in Section 6.

³ <https://github.com/bitcoin/bitcoin>

⁴ <https://github.com/bitcoin-transaction-model/bitcoin-transaction-model>

T_0
in: ...
wit: ...
out: ($\text{versig}_k(x), v_0\text{฿}$)

T_1
in: T_0
wit: σ
out: ($e, v_1\text{฿}$)

Fig. 1: Two Bitcoin transactions.

2 Bitcoin transactions in a nutshell

We now give a minimalistic introduction to the behaviour of Bitcoin transactions.

Users interact with Bitcoin through *addresses*, which they can freely generate. *Transactions* describe transfers of bitcoins (฿) between addresses. The log of all transactions is recorded on a public, immutable and decentralised data structure called *blockchain*. To explain how the blockchain works, consider the transactions T_0 and T_1 displayed in Figure 1. The transaction T_0 contains $v_0\text{฿}$, which can be *redeemed* by putting on the blockchain a transaction (e.g., T_1), whose in field is a reference to T_0 . To redeem T_0 , the *witness* of the redeeming transaction (the value in its wit field) must make the *output script* of T_0 (the first element of the pair in the out field) evaluate to true. When this happens, the value of T_0 is transferred to the new transaction, and T_0 is no longer redeemable.

In the example displayed before, the output script of T_0 evaluates to true when receiving a digital signature on the redeeming transaction T_1 , with a given key pair k . We denote with $\text{versig}_k(x)$ the verification of the signature x on the redeeming transaction: of course, since the signature must be included in the witness of the redeeming transaction, it will consider all the parts of that transaction *except* its wit field. We assume that σ is the signature of T_1 .

Now, assume that the blockchain contains T_0 , not yet redeemed, and someone tries to append T_1 . To validate this operation, the nodes of the Bitcoin network check that $v_1 \leq v_0$, and then they evaluate the output script of T_0 , by instantiating its formal parameter x to the signature σ in the witness of T_1 . The function $\text{versig}_k(\sigma)$ verifies that σ is actually the signature of T_1 : therefore, the output script succeeds, and T_1 redeems T_0 . Subsequently, a new transaction can redeem T_1 by satisfying its output script e (not specified in the figure).

Bitcoin transactions may be more general than the ones illustrated by the previous example. First, there can be multiple inputs and outputs. Each output has an associated output script and value, and can be redeemed independently from others. Consequently, in fields must specify which output they are redeeming. Similarly, a transaction with multiple inputs associates a witness to each of them. To be valid, the sum of the values of all the inputs must be greater or equal to the sum of the values of all the outputs. In its general form, the output script is a program in a (not Turing-complete) scripting language, featuring a limited set of logic, arithmetic, and cryptographic operators. Finally, a transaction can specify time constraints (absolute, or relative to its input transactions) about when it can appear on the blockchain.

$\mathbf{A}, \mathbf{B}, \dots \in \text{Part}$	Participants	$e, e', \dots \in \text{Exp}$	Script expressions
$x, y, \dots \in \text{Var}$	Variables	$\mathbf{T}, \mathbf{T}', \dots \in \text{Tx}$	Transactions
$\nu, \nu', \dots \in \text{Den}$	Denotations, i.e.:	μ, μ'	Signature modifier
$k, k' \dots \in \mathbb{Z}$	Constants	$\text{sig}_k^\mu(\mathbf{T})$	Transaction signature
$t, t' \dots \in \mathbb{N}$	Time	$\text{versig}_k((\nu, \mu), \mathbf{T})$	Signature verification
$v, v' \dots \in \mathbb{N}$	Currency values	$\mathbf{T}, i \models \lambda \mathbf{x}. e$	Script verification
$\sigma, \sigma', \dots \in \mathbb{Z}$	Signatures	$(\mathbf{T}, i, t) \overset{v}{\rightsquigarrow} (\mathbf{T}', j, t')$	Transaction redeem
$\text{true}, \text{false}$	Boolean values	$\mathbf{B} = (\mathbf{T}_1, t_1) \dots$	Blockchains
\perp	Undefined	$\mathbf{B} \triangleright (\mathbf{T}, t)$	Consistent update

Table 1: Summary of notation.

3 A formal model of Bitcoin transactions

In this section we introduce a formal model of Bitcoin transactions. We start in Section 3.1 by defining the scripts that can be used in transaction outputs. Then, in Section 3.2 we formalise transactions, and in Section 3.3 we define a signature scheme for them. Sections 3.4 and 3.5 give semantics, respectively, to scripts and transactions. In Section 3.6 we model the Bitcoin blockchain, and in particular we define the crucial notion of *consistency*, which corresponds to the one enforced by the Bitcoin consensus protocol. We then state a few results about consistent blockchains (their proofs are in Appendix A).

We start by introducing some auxiliary notation. We assume several sets, ranged over by meta-variables as shown in the left column of Table 1. We use the bold notation to denote finite sequences of elements, e.g. \mathbf{x} stands for a sequence of variables. We denote with \mathbf{x}_i the i -th element of the sequence, i.e. $\mathbf{x}_i = x_i$ if $\mathbf{x} = x_1 \dots x_n$, and with $\mathbf{x}_{i..j}$ the subsequence of \mathbf{x} starting from the i -th element and ending to the j -th element. We denote with $|\mathbf{x}|$ the number of elements of \mathbf{x} , and with \square the empty sequence. We denote with $f : A \multimap B$ a *partial* function f from A to B , with $\text{dom } f$ the *domain* of f , i.e. the subset of A where f is defined, and with $\text{ran } f$ the *range* of f , i.e. $\text{ran } f = \{f(x) \mid x \in \text{dom } f\}$. We use \perp to represent an “undefined” element; in particular, when the element is a partial function, \perp denotes the function with empty domain.

3.1 Scripts

Each output in a Bitcoin transaction contains a script, which is used to establish when the output can be redeemed by another transaction. Intuitively, a script is a first-order function (written in a non Turing-equivalent language), which is applied to the witness provided by the redeeming transaction. The output can be redeemed only if such function application evaluates to true.

In our model, we abstract from the actual stack-based scripting language implemented in Bitcoin⁵, by using instead a minimalistic language of expressions.

⁵ <https://en.bitcoin.it/wiki/Script>

Definition 1 (Scripts). We define the set Exp of script expressions (ranged over by e, e', \dots) as follows:

$$e ::= x \mid k \mid e + e \mid e - e \mid e = e \mid e < e \mid \text{if } e \text{ then } e \text{ else } e \mid |e| \mid \\ \text{H}(e) \mid \text{versig}_k(e) \mid \text{absAfter } t : e \mid \text{relAfter } t : e$$

We denote with Script the set of terms of the form $\lambda z.e$ such that all the variables in e occur in z .

Besides some basic arithmetic and logical operators, script expressions include a few operators inspired from the actual Bitcoin scripting language. The expression $|e|$ denotes the size, in bytes, of the evaluation of e . The expression $\text{H}(e)$ evaluates to the hash of e . The expression $\text{versig}_k(e)$ takes as arguments a sequence of m script expressions, representing signatures of the enclosing transactions, and a sequence of n public keys. Intuitively, it evaluates to true whenever the provided signatures are verified by using m out of the n provided keys. The expressions $\text{absAfter } t : e$ and $\text{relAfter } t : e$ define temporal constraints (see Section 3.4). They evaluate as e if the constraints are satisfied, otherwise they fail.

Notation 1. We use the following syntactic sugar for expressions: (i) *false* to denote $1 = 0$ (ii) *true* to denote $1 = 1$ (iii) $e \wedge e'$ to denote *if e then e' else false* (iv) $e \vee e'$ to denote *if e then true else e'* (v) *not e* to denote *if e then false else true*.

3.2 Transactions

The following definition formalises Bitcoin transactions.

Definition 2 (Transactions). We inductively define the set Tx of transactions as follows. A transaction \mathbb{T} is a tuple $(\text{in}, \text{wit}, \text{out}, \text{absLock}, \text{relLock})$, where:

- $\text{in} : \mathbb{N} \rightarrow \text{Tx} \times \mathbb{N}$
- $\text{wit} : \mathbb{N} \rightarrow \mathbb{Z}^*$, where $\text{dom wit} = \text{dom in}$
- $\text{out} : \mathbb{N} \rightarrow \text{Script} \times \mathbb{N}$
- $\text{absLock} : \mathbb{N}$
- $\text{relLock} : \mathbb{N} \rightarrow \mathbb{N}$, where $\text{dom relLock} = \text{dom in}$

where: $\forall i, j \in \text{dom in} : \text{fst}(\text{in}(i)).\text{wit} = \perp$ and $i \neq j \implies \text{in}(i) \neq \text{in}(j)$.

We denote with $\mathbb{T}.f$ the value of field f of \mathbb{T} , for $f \in \{\text{in}, \text{wit}, \text{out}, \text{absLock}, \text{relLock}\}$. We say that \mathbb{T} is *initial* when $\mathbb{T}.\text{in} = \mathbb{T}.\text{relLock} = \perp$ and $\mathbb{T}.\text{absLock} = 0$.

The fields in and out represent, respectively, the inputs and the outputs of a transaction. There is an input for each $i \in \text{dom in}$, and an output for each $j \in \text{dom out}$. When $\mathbb{T}.\text{in}(i) = (\mathbb{T}', j)$, it means that the i -th input of \mathbb{T} wants to redeem the j -th output of \mathbb{T}' . The side condition $i \neq j \implies \text{in}(i) \neq \text{in}(j)$ ensures that inputs are pairwise distinct. The side condition $\text{fst}(\text{in}(i)).\text{wit} = \perp$ is related to the *Segregated Witness (SegWit)* feature⁶, and it requires that the witness of

⁶ This feature, specified in the BIP 141 and activated on August 24th 2017, implies that witnesses are not used in the computation of transaction hashes.

the input transaction is left unspecified. The output $\mathsf{T}'.\text{out}(j)$ is a pair $(\lambda\mathbf{z}.e, v)$, meaning that v Satoshis ($1\text{B} = 10^8$ Satoshis) can be redeemed by whoever can provide a witness which satisfies $\lambda\mathbf{z}.e$. Such witness is defined by $\mathsf{T}.\text{wit}(i)$. The fields $\mathsf{T}.\text{absLock}$ and $\mathsf{T}.\text{relLock}(i)$ specify a constraint on when T can be put on the blockchain: the first in absolute terms, whereas the second is relative to the transaction in the input $\mathsf{T}.\text{in}(i)$. More specifically, $\mathsf{T}.\text{absLock} = t$ means that T can appear on the blockchain only after time t . If $\mathsf{T}.\text{relLock}(i) = t$, then T can appear only after time t since the transaction in $\mathsf{T}.\text{in}(i)$ appeared.

To improve readability, we use the following conventions: (i) if T has exactly one input, we denote it by $\mathsf{T}.\text{in}$ (omitting the index, which we assume to be 1); We act similarly for $\mathsf{T}.\text{wit}$, $\mathsf{T}.\text{out}$, and $\mathsf{T}.\text{relLock}$; (ii) if $\mathsf{T}.\text{absLock} = 0$, we omit it (similarly for $\mathsf{T}.\text{relLock}$ when it is \perp); (iii) we denote with $\text{script}(\mathsf{T}.\text{out}(i))$ and $\text{val}(\mathsf{T}.\text{out}(i))$, respectively, the first and the second element of the pair $\mathsf{T}.\text{out}(i)$.

3.3 Transaction signatures

We extend to transactions the signing and verification functions of the signature schemes, denoted respectively as $\text{sig}_k(\cdot)$ and $\text{ver}_k(\cdot, \cdot)$. For simplicity, although we will always use $k = (k_p, k_s)$ for key pairs, we implicitly assume that $\text{sig}_k(\cdot)$ only uses the private part k_s , while $\text{ver}_k(\cdot, \cdot)$ only uses the public part k_p .

In Bitcoin, transaction signatures never apply to the whole transaction: users can specify which parts of a transaction are signed (with the exception of the wit field, which is never signed). However, not all possible combinations of transaction parts are possible; the legit ones are listed in Definition 4. In order to specify which parts of a transaction are signed, we first introduce the auxiliary notion of *transaction substitution*.

Definition 3 (Transaction substitutions). A transaction substitution Σ is a function from $\mathsf{T}\mathbf{x}$ to $\mathsf{T}\mathbf{x}$. For a transaction field \mathbf{f} , we denote with $\{\mathbf{f} \mapsto d\}$ the substitution which replaces the value of \mathbf{f} with d . For $\mathbf{f} \neq \text{absLock}$ and $i \in \mathbb{N}$, we denote with $\{\mathbf{f}(i) \mapsto d\}$ the substitution which replaces $\mathbf{f}(i)$ with d . Further, for $\circ \in \{<, >, \neq\}$, we denote with $\{\mathbf{f}(\circ i) \mapsto d\}$ the substitution which replaces $\mathbf{f}(j)$ with d , for all $j \circ i \in \text{dom } \mathbf{f}$.

Definition 4 (Signature modifiers). We define signature modifiers μ_i (with $i \in \mathbb{N}$) in Figure 2. We associate to each modifier a substitution, and we denote with $\mu_i(\mathsf{T})$ the result of applying it to the transaction T .

Each modifier is represented by a pair of symbols, describing, respectively, the subset of inputs and of outputs being signed ($a = \text{all}$, $s = \text{single}$, $n = \text{none}$). In case of s , the parameter i determines the index of the signed input/output.

Definition 5 (Transaction signatures). We define the transaction signature (under modifier μ and index i) and verification as follows:

$$\text{sig}_k^{\mu, i}(\mathsf{T}) = (\text{sig}_k(\mu_i(\mathsf{T})), \mu) \quad \text{ver}_k((w, \mu), \mathsf{T}, i) = \text{ver}_k(w, \mu_i(\mathsf{T}))$$

Hereafter, we use σ, σ', \dots to range over transaction signatures.

$$\begin{aligned}
aa_i(\mathbb{T}) &= \mathbb{T}\{\text{wit} \mapsto \perp\} \\
an_i(\mathbb{T}) &= \mathbb{T}\{\text{wit} \mapsto \perp\}\{\text{out} \mapsto \perp\} \\
as_i(\mathbb{T}) &= \mathbb{T}\{\text{wit} \mapsto \perp\}\{\text{out}(< i) \mapsto (\text{false}, 0)\}\{\text{out}(> i) \mapsto \perp\} \\
sa_i(\mathbb{T}) &= \mathbb{T}\{\text{wit} \mapsto \perp\}\{\text{in}(1) \mapsto \mathbb{T}.\text{in}(i)\}\{\text{in}(\neq 1) \mapsto \perp\} \\
&\quad \{\text{relLock}(1) \mapsto \mathbb{T}.\text{relLock}(i)\}\{\text{relLock}(\neq 1) \mapsto \perp\} \\
sn_i(\mathbb{T}) &= sa_i(an_i(\mathbb{T})) \\
ss_i(\mathbb{T}) &= sa_i(as_i(\mathbb{T}))
\end{aligned}$$

Fig. 2: Signature modifiers.

Notation 2. Note that $\text{sig}_k^{\mu,i}(\mathbb{T})$ can meaningfully appear within $\mathbb{T}.\text{wit}(i)$, since the signature always neglects the wit field of transactions (as all signature modifiers set wit to \perp). In this case, as a shorthand we denote it with sig_k^μ , or just sig_k when $\mu = aa$.

We now extend the signature verification $\text{ver}_k(\sigma, \mathbb{T}, i)$ to the case where, instead of providing a single key k and a single signature σ , one has many keys and signatures, i.e. $\text{ver}_{\mathbf{k}}(\sigma, \mathbb{T}, i)$. Intuitively, if $|\sigma| = m$ and $|\mathbf{k}| = n$, the function $\text{ver}_{\mathbf{k}}(\sigma, \mathbb{T}, i)$ implements a m -of- n multi-signature scheme, i.e. it evaluates to true if all the m signatures match with (some of) the keys in \mathbf{k} . The actual definition is a bit more complex, to be coherent with the one implemented in Bitcoin.

Definition 6 (Multi-signature verification). Let \mathbf{k} and σ be sequences of (public) keys and signatures such that $|\mathbf{k}| \geq |\sigma|$. We define the function:

$$\text{ver}_{\mathbf{k}}^{n,m}(\sigma, \mathbb{T}, i) \equiv \begin{cases} \text{true} & \text{if } m = 0 \\ \text{false} & \text{if } m \neq 0 \text{ and } n = 0 \\ \text{ver}_{\mathbf{k}}^{n-1,m-1}(\sigma, \mathbb{T}, i) & \text{if } m, n \neq 0 \text{ and } \text{ver}_{\mathbf{k}_n}(\sigma_m, \mathbb{T}, i) \\ \text{ver}_{\mathbf{k}}^{n-1,m}(\sigma, \mathbb{T}, i) & \text{otherwise} \end{cases}$$

Then, we define $\text{ver}_{\mathbf{k}}(\sigma, \mathbb{T}, i) = \text{ver}_{\mathbf{k}}^{|\mathbf{k}|,|\sigma|}(\sigma, \mathbb{T}, i)$.

Our formalisation of multi-signature verification (Definition 6) follows closely the implementation of Bitcoin, whose stack-based scripting language imposes that the sequence σ is read in reverse order. Accordingly, the function ver tries to verify the last signature in σ with the last key in \mathbf{k} . If they match, the function ver proceeds to verify the previous signature in the sequence, otherwise it tries to verify the signature with the previous key.

Example 1 (2-of-3 multi-signature). Let $\mathbf{k} = k_a k_b k_c$, and let σ_p, σ_q be such that $\text{ver}_{k_a}(\sigma_p, \mathbb{T}, i) = \text{ver}_{k_b}(\sigma_q, \mathbb{T}, i) = \text{true}$, *false* otherwise. We can check that $\text{ver}_{\mathbf{k}}(\sigma_p \sigma_q, \mathbb{T}, i) = \text{true}$, while $\text{ver}_{\mathbf{k}}(\sigma_q \sigma_p, \mathbb{T}, i) = \text{false}$. \square

3.4 Semantics of scripts

Definition 7 gives the semantics to script expressions. This semantics will be then used in Section 3.5 to define when a transaction can redeem another one. We

$$\begin{aligned}
\llbracket x \rrbracket_{\mathbb{T}, i, \rho} &= \rho(x) \\
\llbracket k \rrbracket_{\mathbb{T}, i, \rho} &= k \\
\llbracket \text{versig}_k(e) \rrbracket_{\mathbb{T}, i, \rho} &= \text{ver}_k(\llbracket e \rrbracket_{\mathbb{T}, i, \rho}, \mathbb{T}, i) \\
\llbracket H(e) \rrbracket_{\mathbb{T}, i, \rho} &= H(\llbracket e \rrbracket_{\mathbb{T}, i, \rho}) \quad (H \text{ is a public hash function}) \\
\llbracket \text{absAfter } t : e \rrbracket_{\mathbb{T}, i, \rho} &= \text{if } \mathbb{T}.\text{absLock} \geq t \text{ then } \llbracket e \rrbracket_{\mathbb{T}, i, \rho} \text{ else } \perp \\
\llbracket \text{relAfter } t : e \rrbracket_{\mathbb{T}, i, \rho} &= \text{if } \mathbb{T}.\text{relLock}(i) \geq t \text{ then } \llbracket e \rrbracket_{\mathbb{T}, i, \rho} \text{ else } \perp \\
\llbracket e \circ e' \rrbracket_{\mathbb{T}, i, \rho} &= \llbracket e \rrbracket_{\mathbb{T}, i, \rho} \circ_{\perp} \llbracket e' \rrbracket_{\mathbb{T}, i, \rho} \quad (\circ \in \{+, -, =, <\}) \\
\llbracket e \rrbracket_{\mathbb{T}, i, \rho} &= \text{size}(\llbracket e \rrbracket_{\mathbb{T}, i, \rho}) \\
\llbracket \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \rrbracket_{\mathbb{T}, i, \rho} &= \text{if } \llbracket e_0 \rrbracket_{\mathbb{T}, i, \rho} \text{ then } \llbracket e_1 \rrbracket_{\mathbb{T}, i, \rho} \text{ else } \llbracket e_2 \rrbracket_{\mathbb{T}, i, \rho}
\end{aligned}$$

Fig. 3: Semantics of script expressions.

use an environment $\rho : \text{Var} \rightarrow \mathbb{Z}$ which associates a denotation to each variable occurring in it. Further, we use a transaction $\mathbb{T} \in \text{Tx}$ and an index $i \in \mathbb{N}$ to indicate the witness redeeming the script, both used to evaluate the timelock expressions. We use the denotation \perp to represent “failure” of the evaluation. This is the case e.g. of timelock expressions, when the temporal constraint is not satisfied. All the semantic operators used in Definition 7 are *strict*, i.e. they evaluate to \perp if some of their operands is \perp .

Definition 7 (Expression evaluation). Let $\rho : \text{Var} \rightarrow \mathbb{Z}$, let $\mathbb{T} \in \text{Tx}$ and $i \in \mathbb{N}$. We define the function $\llbracket \cdot \rrbracket_{\mathbb{T}, i, \rho} : \text{Exp} \rightarrow \text{Den}$ in Figure 3, where we use the following operators on denotations:

$$\begin{aligned}
\text{if } \nu_0 \text{ then } \nu_1 \text{ else } \nu_2 &\equiv \begin{cases} \nu_1 & \text{if } \nu_0 = \text{true} \\ \nu_2 & \text{if } \nu_0 = \text{false} \\ \perp & \text{otherwise} \end{cases} & \text{size}(\nu) &\equiv \begin{cases} \perp & \text{if } \nu \notin \mathbb{Z} \\ 0 & \text{if } \nu = 0 \\ \lceil \frac{\log_2 |\nu|}{7} \rceil & \text{otherwise} \end{cases} \\
\nu_0 \circ_{\perp} \nu_1 &\equiv \text{if } \nu_0, \nu_1 \in \mathbb{Z} \text{ then } \nu_0 \circ \nu_1 \text{ else } \perp \quad (\circ \in \{+, -, =, <\})
\end{aligned}$$

Definition 8 (Script verification). We say that the input i of \mathbb{T} verifies $\lambda \mathbf{x}.e$ (in symbols: $\mathbb{T}, i \models \lambda \mathbf{x}.e$) when $\mathbf{x} = x_1 \dots x_n$, $\mathbb{T}.\text{wit}(i) = k_1 \dots k_n$, and:

$$\llbracket e \rrbracket_{\mathbb{T}, i, \{x_j \mapsto k_j \mid j \in 1 \dots n\}} = \text{true}$$

Example 2. Let H be a hash function, let $s, h \in \mathbb{Z}$ be such that $h = H(s)$, and let \mathbb{T} be such that $\mathbb{T}.\text{wit}(1) = (\sigma, s)$, with $\sigma = \text{sig}_k^{aa}(\mathbb{T})$. We prove that:

$$\mathbb{T}, 1 \models \lambda(\varsigma, x).(\text{versig}_k(\varsigma) \text{ and } H(x) = h)$$

To do this, let $\rho = \{\varsigma \mapsto \sigma, x \mapsto s\}$. We have that:

$$\begin{aligned}
&\llbracket \text{versig}_k(\varsigma) \text{ and } H(x) = h \rrbracket_{\mathbb{T}, 1, \rho} = \llbracket \text{versig}_k(\varsigma) \rrbracket_{\mathbb{T}, 1, \rho} \text{ and } \llbracket H(x) = h \rrbracket_{\mathbb{T}, 1, \rho} \\
&= \text{ver}_k(\llbracket \varsigma \rrbracket_{\mathbb{T}, 1, \rho}, \mathbb{T}, 1) \text{ and } (\llbracket H(x) \rrbracket_{\mathbb{T}, 1, \rho} =_{\perp} \llbracket h \rrbracket_{\mathbb{T}, 1, \rho}) \\
&= \text{ver}_k(\rho(\varsigma), \mathbb{T}, 1) \text{ and } (H(\llbracket x \rrbracket_{\mathbb{T}, 1, \rho}) =_{\perp} h) = \text{ver}_k(\sigma, \mathbb{T}, 1) \text{ and } (H(\rho(x)) =_{\perp} h) \\
&= \text{true} \quad \square
\end{aligned}$$

T_0
in: ...
wit: ...
out: $(\lambda\varsigma.\text{versig}_k(\varsigma), v_0)$

T_1
in: $(T_0, 1)$
wit: sig_k
out: $(\lambda\varsigma.\text{versig}_{k'}(\varsigma), v_1)$

T'_1
in: $(T_0, 1)$
wit: sig_k
out: $(\lambda\varsigma.\text{versig}_{k'}(\varsigma), v_1)$
absLock: 5.1.2017
relLock: 2 days

Fig. 4: Three transactions. For notational conciseness, when displaying transactions we omit the substitution $\{\text{wit} \mapsto \perp\}$ for the transaction within the in field (e.g., we just write T_0 within $T_1.\text{in}$). Also, we use dates in time constraints.

3.5 Semantics of transactions

Definition 9 describes when the j -th input of a transaction T' (put on the blockchain at time t') can redeem v Satoshis from the i -th output of the transaction T (put on the blockchain at time t). We denote this by $(T, i, t) \overset{v}{\rightsquigarrow} (T', j, t')$.

Definition 9 (Output redeeming). We write $(T, i, t) \overset{v}{\rightsquigarrow} (T', j, t')$ iff all the following conditions hold:

- (a) $T'.\text{in}(j) = (T\{\text{wit} \mapsto \perp\}, i)$
- (b) $T', j \models \text{script}(T.\text{out}(i))$
- (c) $v = \text{val}(T.\text{out}(i))$
- (d) $t' \geq T'.\text{absLock}$
- (e) $t' - t \geq T'.\text{relLock}(j)$ if $j \in \text{dom } T'.\text{relLock}$

We write $(T, i, t) \not\rightsquigarrow (T', j, t')$ when for no v it holds that $(T, i, t) \overset{v}{\rightsquigarrow} (T', j, t')$.

Item (a) links the j -th input of T' to the i -th output of T . Note that, since we are modelling SegWit, the witness in the transaction $T'.\text{in}(j)$ is left unspecified: this is why we set to \perp also the witness of T . Item (b) requires that the j -th witness of T' verifies the i -th output script of T . Item (c) just defines v as the value in the i -th output of T . Items (d) and (e) check the absolute and relative timelocks, respectively. The first constraint states that T' cannot appear on the blockchain before $T'.\text{absLock}$; the second one states that T' cannot appear until at least $T'.\text{relLock}(j)$ time units have elapsed since T was put on the blockchain.

Example 3. With the transactions in Figure 4, we have $(T_0, 1, t_0) \overset{v_0}{\rightsquigarrow} (T_1, 1, t_1)$. Indeed, for item (a) we have that $T_1.\text{in}(1) = (T_0\{\text{wit} \mapsto \perp\}, 1)$; for item (b), $T_1, 1 \models \lambda\varsigma.\text{versig}_k(\varsigma)$; for item (c), $v_0 = \text{val}(T_0.\text{out}(1))$. The other two items trivially hold, as there are no time constraints. We also have $(T_0, 1, 2.1.2017) \overset{v_0}{\rightsquigarrow} (T'_1, 1, 6.1.2017)$. To show that, we have to check also items (d) and (e). For item (d), we have that $6.1.2017 \geq T'_1.\text{absLock} = 5.1.2017$. For item (e), we have that $6.1.2017 - 2.1.2017 \geq T'_1.\text{relLock}(1) = 2 \text{ days}$. \square

\mathbb{T}_1
in: \perp
wit: \perp
$1 \mapsto (\lambda\varsigma.\text{versig}_{k_1}(\varsigma), 3)$
$2 \mapsto (\lambda\varsigma.\text{versig}_{k_2}(\varsigma), 5)$
$3 \mapsto (\lambda\varsigma.\text{versig}_{k_3}(\varsigma), 7)$

\mathbb{T}_2
in: $1 \mapsto (\mathbb{T}_1, 2), 2 \mapsto (\mathbb{T}_1, 3)$
wit: $1 \mapsto \text{sig}_{k_2}, 2 \mapsto \text{sig}_{k_3}$
out: $(\lambda\varsigma.\text{versig}_{k_2}(\varsigma), 10)$

\mathbb{T}_3
in: $(\mathbb{T}_1, 2)$
wit: sig_{k_2}
out: $(\lambda\varsigma.\text{versig}_{k_2}(\varsigma), 5)$

Fig. 5: Three transactions for Examples 4 to 6.

3.6 Blockchain and consistency

In Definition 10 we model blockchains as sequences of *timed transactions* (\mathbb{T}, t) , where t represents the time when the transaction \mathbb{T} has been added. Note that our definition is very permissive: for instance, it allows a blockchain to contain transactions which do not redeem any transactions, or double-spent transactions. We will rule out such *inconsistent* blockchains later on in Definition 13.

Definition 10 (Blockchain). A blockchain \mathbf{B} is a sequence $(\mathbb{T}_1, t_1) \cdots (\mathbb{T}_n, t_n)$, where \mathbb{T}_1 is the only transaction with $\text{in} = \perp$, and $t_i \leq t_j$ for all $1 \leq i \leq j \leq n$.

We denote with $\text{trans}_{\mathbf{B}}$ the set of transactions occurring in \mathbf{B} , and with $\text{time}_{\mathbf{B}}(\mathbb{T}_i)$ the time t_i of transaction \mathbb{T}_i in \mathbf{B} . Given a transaction \mathbb{T} , we define $\text{match}_{\mathbf{B}}(\mathbb{T})$ as the set of transactions \mathbb{T}_i such that $\mathbb{T}\{\text{wit} \mapsto \perp\} = \mathbb{T}_i\{\text{wit} \mapsto \perp\}$.

Definition 11 (Unspent output). Let $\mathbf{B} = (\mathbb{T}_1, t_1) \cdots (\mathbb{T}_n, t_n)$ be a blockchain. We say that the output j of transaction \mathbb{T}_i is *unspent in \mathbf{B}* whenever:

$$\forall i' \leq n, j' \in \mathbb{N} : (\mathbb{T}_i, j, t_i) \not\rightsquigarrow (\mathbb{T}_{i'}, j', t_{i'})$$

Given a blockchain \mathbf{B} , we define:

- $\text{UTXO}_{\mathbf{B}}$, the Unspent Transaction Output of \mathbf{B} , as the set of pairs (\mathbb{T}_i, j) such that output j of \mathbb{T}_i is unspent in \mathbf{B} .
- $\text{val}(\mathbf{B})$, the value of \mathbf{B} , as the sum of the values of all outputs in its UTXO .

Example 4. Consider the transactions in Figure 5, and let $\mathbf{B} = (\mathbb{T}_1, 0)(\mathbb{T}_2, t_2)$. We have that $(\mathbb{T}_1, 2, 0) \overset{5}{\rightsquigarrow} (\mathbb{T}_2, 1, t_2)$ and $(\mathbb{T}_1, 3, 0) \overset{7}{\rightsquigarrow} (\mathbb{T}_2, 2, t_2)$, while the other outputs are unspent. Hence, the UTXO of \mathbf{B} is $\{(\mathbb{T}_1, 1), (\mathbb{T}_2, 1)\}$. \square

The following definition establishes when (\mathbb{T}, t) is a *consistent update* of \mathbf{B} .

Definition 12 (Consistent update). We write $\mathbf{B} \triangleright (\mathbb{T}, t)$ iff either $\mathbf{B} = []$, \mathbb{T} is initial and $t = 0$, or, given, for all $i \in \text{dom}(\mathbb{T}.\text{in})$:

$$\begin{aligned} \{\mathbb{T}'_i\} &= \text{match}_{\mathbf{B}}(\text{fst}(\mathbb{T}.\text{in}(i))) && \text{(redeemed transaction)} \\ o_i &= \text{snd}(\mathbb{T}.\text{in}(i)) && \text{(redeemed output index)} \\ t'_i &= \text{time}_{\mathbf{B}}(\mathbb{T}'_i) && \text{(time when } \mathbb{T}'_i \text{ was added to } \mathbf{B}) \\ v_i &= \text{val}(\mathbb{T}'_i.\text{out}(o_i)) && \text{(value of the redeemed output)} \end{aligned}$$

the following conditions hold:

- (1) $\forall i \in \text{dom } \mathbf{T}.\text{in} : (\mathbf{T}'_i, o_i) \in \text{UTXO}_{\mathbf{B}}$
- (2) $\forall i \in \text{dom } \mathbf{T}.\text{in} : (\mathbf{T}'_i, o_i, t'_i) \overset{v_i}{\rightsquigarrow} (\mathbf{T}, i, t)$
- (3) $\sum \{v_i \mid i \in \text{dom } \mathbf{T}.\text{in}\} \geq \sum \{\text{val}(\mathbf{T}.\text{out}(j)) \mid j \in \text{dom } \mathbf{T}.\text{out}\}$
- (4) $\mathbf{B} = \mathbf{B}'(\mathbf{T}', t') \implies t \geq t'$

Firstly, for each $\mathbf{T}.\text{in}(i)$ we obtain the singleton $\{\mathbf{T}'_i\}$ from the blockchain, using $\text{match}_{\mathbf{B}}$, such that $\text{fst}(\mathbf{T}.\text{in}(i))\{\text{wit} \mapsto \perp\} = \mathbf{T}'_i\{\text{wit} \mapsto \perp\}$. The update is inconsistent if $\text{match}_{\mathbf{B}}(\text{fst}(\mathbf{T}.\text{in}(i)))$ is not a singleton for some i . Condition (1) requires that the redeemed outputs are currently unspent in \mathbf{B} . Condition (2) asks that each input of \mathbf{T} redeems an output of a transaction in \mathbf{B} . Condition (3) requires that the sum of the values of the outputs of \mathbf{T} is not greater than the total value it redeems. Finally, (4) requires that the time of \mathbf{T} is greater than or equal to the time of the last transaction in \mathbf{B} .

Example 5. Consider again the transactions in Figure 5, and let $\mathbf{B} = (\mathbf{T}_1, 0)$. We prove that $\mathbf{B} \triangleright (\mathbf{T}_2, t_2)$. Let $o_1 = 2, o_2 = 3, t'_1 = t'_2 = 0, v_1 = 5, v_2 = 7$. We now prove that the conditions of Definition 12 are satisfied. For condition (1), note that both $(\mathbf{T}_1, 2)$ and $(\mathbf{T}_1, 3)$ are unspent, according to Definition 11. For condition (2), note that:

$$(\mathbf{T}_1, 2, 0) \overset{v_1}{\rightsquigarrow} (\mathbf{T}_2, 1, t_2) \quad (\mathbf{T}_1, 3, 0) \overset{v_2}{\rightsquigarrow} (\mathbf{T}_2, 2, t_2)$$

hold, according to Definition 9. Finally, for condition (3), we have that:

$$\sum \{v_i \mid i \in \{1, 2\}\} = 5 + 7 \geq \sum \{\text{val}(\mathbf{T}_2.\text{out}(j)) \mid j \in \text{dom } \mathbf{T}_2.\text{out}\} = 10$$

Therefore, (\mathbf{T}_2, t_2) is a consistent update of \mathbf{B} . \square

Example 6 (Double spending). Consider again the transactions in Figure 5, and let $\mathbf{B} = (\mathbf{T}_1, 0)(\mathbf{T}_2, t_2)$. We prove that (\mathbf{T}_3, t_3) is *not* a consistent update of \mathbf{B} . Although condition (2) of Definition 12 holds:

$$(\mathbf{T}_1, 2, 0) \overset{5}{\rightsquigarrow} (\mathbf{T}_3, 1, t_3)$$

we have that condition (1) is *not* satisfied. In fact, according to Definition 11, $(\mathbf{T}_1, 2)$ is already spent in \mathbf{B} because

$$(\mathbf{T}_1, 2, 0) \overset{5}{\rightsquigarrow} (\mathbf{T}_2, 1, t_2)$$

holds and both \mathbf{T}_1 and \mathbf{T}_2 are in \mathbf{B} . Since \mathbf{T}_3 is trying to spend an output already spent, this transaction should not be appended to \mathbf{B} . \square

We now define when a blockchain is consistent. Intuitively, consistency holds when the blockchain has been constructed, started from the empty one, by appending consistent updates, only. The actual definition is given by induction.

Definition 13 (Consistency). We say that a blockchain \mathbf{B} is consistent if either $\mathbf{B} = []$, or $\mathbf{B} = \mathbf{B}'(\mathbb{T}, t)$ with \mathbf{B}' consistent and $\mathbf{B}' \triangleright (\mathbb{T}, t)$.

Note that the empty blockchain is consistent; the blockchain with a single transaction (\mathbb{T}_1, t_1) is consistent iff \mathbb{T}_1 is initial and $t_1 = 0$. The transaction \mathbb{T}_1 models the first transaction in the *genesis block* (as discussed in Section 6, we are abstracting away the *coinbase* transactions, which forge new bitcoins).

We now establish some basic properties of consistent blockchains. Lemma 1 states that, in a consistent blockchain, the inputs of a transaction point backwards to the output of some transaction in the blockchain.

Lemma 1. *If $(\mathbb{T}_1, t_1) \cdots (\mathbb{T}_n, t_n)$ is consistent, then:*

$$\forall i \in 2 \dots n : \forall (\mathbb{T}, h) \in \text{ran}(\mathbb{T}_i.\text{in}) : \exists j < i : \mathbb{T}_j\{\text{wit} \mapsto \perp\} = \mathbb{T} \wedge h \in \text{dom}(\mathbb{T}_j.\text{out})$$

The following theorem establishes that a transaction output cannot be redeemed twice in a consistent blockchain.

Theorem 1 (No double spending). *If $(\mathbb{T}_1, t_1) \cdots (\mathbb{T}_n, t_n)$ is consistent, then:*

$$\forall i \neq j \in 1 \dots n : \text{ran}(\mathbb{T}_i.\text{in}) \cap \text{ran}(\mathbb{T}_j.\text{in}) = \emptyset$$

The following lemma states that there can be at most a single match of an arbitrary transaction within a consistent blockchain. This implies that the in field of an arbitrary transaction points at most to one transaction output within the blockchain.

Lemma 2. *If \mathbf{B} is consistent, then for all transactions \mathbb{T} , $\text{match}_{\mathbf{B}}(\mathbb{T})$ contains at most one element.*

Lemma 3 ensures that all the transactions on a consistent blockchain are pairwise distinct, even when neglecting their witnesses.

Lemma 3. *If $(\mathbb{T}_1, t_1) \cdots (\mathbb{T}_n, t_n)$ is consistent, then:*

$$\forall i \neq j \in 1 \dots n : \mathbb{T}_i\{\text{wit} \mapsto \perp\} \neq \mathbb{T}_j\{\text{wit} \mapsto \perp\}$$

The following theorem states that the overall value of a blockchain decreases as the blockchain grows. This is because our model does not keep track of the *coinbase transactions*, which in Bitcoin allow miners to collect transaction fees (the difference between inputs and outputs of a transaction), and block rewards.

Theorem 2 (Decreasing value). *Let \mathbf{B} be a consistent blockchain, and let \mathbf{B}' be a non-empty prefix of \mathbf{B} . Then, $\text{val}(\mathbf{B}') \geq \text{val}(\mathbf{B})$.*

Note that the scripting language and its semantics are immaterial in all the statements above. Actually, proving these results never involves checking condition (b) of Definition 9. Of course, the choice of the scripting language affects the expressiveness of the smart contracts built upon Bitcoin.

T_{AB}	T_{BC}
in: $(T_A, 1)$ wit: \perp out: $(\lambda_{\varsigma_A \varsigma_B}. \text{versig}_{k_A k_B}(\varsigma_A \varsigma_B), 1.1\text{B})$	in: $(T_{AB}, 1)$ wit: \perp out: $1 \mapsto (\lambda_{\varsigma_B}. \text{versig}_{k_B}(\varsigma_B), 0.1\text{B})$ $2 \mapsto (\lambda_{\varsigma_C}. \text{versig}_{k_C}(\varsigma_C), 1\text{B})$

Fig. 6: Transactions of the chain contract.

4 Example: static chains of transactions

We now formally specify in our model a simple smart contract⁷, which illustrates the impact of SegWit on the expressiveness of Bitcoin contracts.

A participant **A** wants to send an indirect payment of 1B to **C**, routing it through **B**. To authorize the payment, **B** wants to keep a fee of 0.1B . However, **A** is afraid that **B** will keep all the money for himself, so she exploits the following contract. She creates a *chain* of transactions, as shown in Figure 6. The transaction T_{AB} transfers 1.1B from **A** to **B** (but it is not signed by **A**, yet), while T_{BC} transfers 1B from **B** to **C**. We assume that $(T_A, 1)$ is a transaction output redeemable by **A** through her key k_A , and that k_B is the key of **B**.

The protocol of **A** is the following: **A** starts by asking **B** for his signature on T_{BC} , ensuring that **C** will be paid. After receiving and verifying the signature, **A** puts T_{AB} on the blockchain, adding her signature on the wit field. Then, she also appends T_{BC} , replacing the wit field with her signature and **B**'s one. Since **A** takes care of publishing the transactions, the behaviour of **B** consists just in sending his signature on T_{BC} .

Remarkably, this contract relies on the SegWit feature: indeed, without SegWit it no longer works. We can disable SegWit by changing our model as follows:

- in Definition 2, we no longer require that $\forall i \in \text{dom in} : \text{fst}(\text{in}(i)).\text{wit} = \perp$
- in Definition 9, we replace item (a) with the condition: $T'.\text{in}(j) = (T, i)$
- in Definition 10, we let $\text{match}_B(T) = \{T\}$ if T occurs in **B**, empty otherwise.

To see why disabling SegWit breaks the contract, assume that the transaction $T = T_{AB}\{\text{wit} \mapsto \text{sig}_{k_A}^{aa}(T_{AB})\}$ is unspent on the blockchain, when participant **A** attempts to append also $T' = T_{BC}\{\text{wit} \mapsto \text{sig}_{k_A}^{aa}(T_{BC}) \text{sig}_{k_B}^{aa}(T_{BC})\}$. To be a consistent update, by item (2) of Definition 12 we must have (for some $t_1 \leq t_2$):

$$(T, 1, t_1) \stackrel{1\text{B}}{\rightsquigarrow} (T', 1, t_2) \quad (1)$$

For this, all the conditions in Definition 9 must hold. However, since we have disabled SegWit, for item (a) we no longer check that:

$$T'.\text{in}(1) = (T\{\text{wit} \mapsto \perp\}, 1)$$

⁷ https://www.bitcoinhk.org/media/presentations/2016-03-16/2016-03-16-Segregated_Witness.pdf

but instead we need to check the condition:

$$\tilde{T}'.\text{in}(1) = (\tilde{T}, 1) \tag{2}$$

where the transactions \tilde{T}, \tilde{T}' correspond to the non-SegWit versions of T, T' , i.e. their in fields point to their actual parents, according to the new Definition 2.

Hence, condition (2) checks the equality between \tilde{T}_{AB} (the transaction in the input of \tilde{T}') and $\tilde{T}_{AB}\{\text{wit} \mapsto \text{sig}_{k_A}^{aa}(\tilde{T}_{AB})\}$ (the transaction \tilde{T}). Note that all the fields of the second transaction — but the wit field — are equal to those of the first transaction. Instead, the witness of \tilde{T}_{AB} is \perp , while the one of \tilde{T} contains the signature of A . This difference in the wit field is ignored with the SegWit semantics, while it is discriminating for the older version of Bitcoin.

A naïve attempt to amend the contract would be to set the input field of \tilde{T}' to \tilde{T} . However, this would invalidate the signature of A on \tilde{T}' .

5 Compiling to standard Bitcoin transactions

We now sketch how to compile the transactions of our abstract model into concrete Bitcoin transactions. In particular, we aim at producing *standard* Bitcoin transactions, which respect further constraints on their fields⁸. This is crucial, because non-standard transactions are mostly discarded by the Bitcoin network.

Our compiler produces output scripts of the following kinds, which are all allowed in standard transactions:

Pay to Public Key Hash (P2PKH) takes as parameters a public key and a signature, and checks that (i) the hash of the public key matches the hash hardcoded in the script; (ii) the signature is verified against the public key.

Pay to Script Hash (P2SH) contains only a hash (say, h). The actual script $\lambda x.e$ — which is *not* required to be standard — is contained instead in the wit field of the redeeming transaction, alongside with the actual parameters k . The evaluation succeeds if $H(\lambda x.e) = h$ and $(\lambda x.e)k$ evaluates to *true*.

The only constraint imposed by P2SH is on the size of the script, which is limited to the size of a stack element (520 bytes).

OP_RETURN allows to put up to 80 bytes of data in an output script, making the output unredemable.

We compile the scripts of the form $\lambda \varsigma.\text{versig}_k(\varsigma)$ to P2PKH, and those of the form λk to OP_RETURN. All other scripts are compiled to P2SH when they comply with the size constraint, otherwise compilation fails. In this way, our compiler always produces standard transactions.

Our compiler exploits the *alternative stack* as temporary storage of the variable values. In this way we cope with the stack-based nature of the Bitcoin scripting language. For instance, for the script $\lambda x.H(x) = H(x + 1)$, the variable x is pushed on the alternative stack beforehand, then duplicated and copied in the main stack before each operation involving x .

⁸ <https://bitcoin.org/en/developer-guide#standard-transactions>

6 Conclusions

We have proposed a formal model for Bitcoin transactions. Our model abstractly describes their essential aspects, at the same time enabling formal reasoning, and providing a formal specification to some of Bitcoin’s less documented features. This work provides the theoretical foundations to model Bitcoin smart contracts, reducing the gap between cryptography and programming languages communities. A formal description of smart contracts would enable their automated verification and analysis, which are of crucial importance in a context where bugs in design or implementation may result in loss of money.

There are some differences between our model and the actual Bitcoin, which we outline below.

In Definition 2, we stipulate that the `in` field of a transaction points to another transaction. Instead, in Bitcoin the `in` field contains the identifier of the input transaction. More specifically, this identifier is defined as $H(\mu(\mathbb{T}))$, where:

- $\mu = \{\text{wit} \mapsto \perp\}$ since the activation of the SegWit feature;
- $\mu = \perp$, beforehand.

Consequently, the condition $(\mathbb{T}, i, t) \overset{v}{\rightsquigarrow} (\mathbb{T}', j, t')$ item (a) of Definition 9 would be translated in Bitcoin as: $\mathbb{T}'.\text{in}(j) = (H(\mu(\mathbb{T}')), i)$, where $H(\mu(\mathbb{T}')) = H(\mu(\mathbb{T}))$. Intuitively, the `in` field specifies the transaction (and the output index) to redeem. Since the activation of SegWit, the computation of the transaction identifier does not take in account the `wit` field.

The scripting language in Definition 1 is a bit more expressive than Bitcoin’s. For instance, the script $\lambda x. H(x) < k$ is admissible in our model, while it is not in Bitcoin. Indeed, the Bitcoin scripting language only admits the comparison (via the `OP_LESSTHANOREQUAL` opcode) on 32-bit integers, while two arbitrary values can only be tested for equality (via the `OP_EQUAL` opcode). Similar restrictions apply to arithmetic operations. It is straightforward to adapt our model to apply the same restrictions on Bitcoin scripts. Indeed, our compiler already implements a simple type system which rules away scripts not admissible in Bitcoin.

Definition 10 represents blockchains as sequences of transactions. Instead, in Bitcoin they are sequences of *blocks* of transactions. In this way, we are abstracting both from the cryptographic puzzle that miners have to solve to append a new block to the blockchain, and from the *coinbase transactions*, which (like our initial transaction) do not redeem other transactions, and mint new bitcoins (the block rewards). Coinbase transactions are also used in Bitcoin to collect transaction fees, which are just discarded in our model. Note that extending our model with coinbase transactions would falsify Theorem 2, since the overall value in the blockchain would no longer be decreasing.

In Definitions 2 and 9, the `absLock` and `relLock` fields specify the time when a transaction can be appended to the blockchain. In Bitcoin transactions, besides the time we can also use the *block height*, i.e. the distance between any given block and the genesis block. Setting the block height to h implies that the transaction can be mined from the block h onward.

Acknowledgments. This work is partially supported by Aut. Reg. of Sardinia project P.I.A. 2013 “NOMAD”. Stefano Lande gratefully acknowledges Sardinia Regional Government for the financial support of his PhD scholarship (P.O.R. Sardegna F.S.E. Operational Programme of the Autonomous Region of Sardinia, European Social Fund 2014-2020).

References

1. 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)
2. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multiparty computations on Bitcoin. In: IEEE Symposium on Security and Privacy. pp. 443–458 (2014)
3. 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)
4. Bartoletti, M., Zunino, R.: Constant-deposit multiparty lotteries on Bitcoin. In: Financial Cryptography Workshops. LNCS, vol. 10323. Springer (2017)
5. Bentov, I., Kumaresan, R.: How to use Bitcoin to design fair protocols. In: CRYPTO. LNCS, vol. 8617, pp. 421–439. Springer (2014)
6. Kumaresan, R., Bentov, I.: How to use Bitcoin to incentivize correct computations. In: ACM CCS. pp. 30–41 (2014)
7. Kumaresan, R., Moran, T., Bentov, I.: How to use Bitcoin to play decentralized poker. In: ACM CCS. pp. 195–206 (2015)
8. Miller, A., Bentov, I.: Zero-collateral lotteries in Bitcoin and Ethereum. In: EuroS&P Workshops. pp. 4–13 (2017)
9. Möser, M., Eyal, I., Sirer, E.G.: Bitcoin covenants. In: Financial Cryptography Workshops. pp. 126–141. Springer (2016)
10. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf> (2008)
11. O’Connor, R., Piekarska, M.: Enhancing Bitcoin transactions with covenants. In: Financial Cryptography Workshops (2017)
12. Szabo, N.: Formalizing and securing relationships on public networks. First Monday 2(9) (1997), <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/548>

A Proofs

Proof of Lemma 1

By Definition 13, (T_i, t_i) is a consistent update of $(T_1, t_1) \cdots (T_{i-1}, t_{i-1})$. The thesis follows from condition (2) of Definition 12. \square

Proof of Theorem 1

Let $\mathbf{B} = (\mathbb{T}_1, t_1) \cdots (\mathbb{T}_n, t_n)$ be consistent. By contradiction, assume that there exist $i < j$ and i', j' such that $\mathbb{T}_i.\text{in}(i') = \mathbb{T}_j.\text{in}(j')$. By consistency, there exist h, h' such that $(\mathbb{T}_h\{\text{wit} \mapsto \perp\}, h') = \mathbb{T}_i.\text{in}(i')$. Since $\mathbf{B}_{1..i-1} \triangleright (\mathbb{T}_i, t_i)$, then by item (2) of Definition 12 it must be $(\mathbb{T}_h, h', t_h) \rightsquigarrow (\mathbb{T}_i, i', t_i)$. Hence, by Definition 11 it follows that (\mathbb{T}_h, h') is already *spent* in \mathbf{B} . Since $\mathbf{B}_{1..j-1} \triangleright (\mathbb{T}_j, t_j)$, by item (1) of Definition 12, (\mathbb{T}_h, h') must be *unspent* — contradiction. \square

Proof of Lemma 2

Let $\mathbf{B} = (\mathbb{T}_1, t_1) \cdots (\mathbb{T}_n, t_n)$ be consistent. By contradiction, assume that $\mathbb{T}_i, \mathbb{T}_j \in \text{match}_{\mathbf{B}}(\mathbb{T})$, with $\mathbb{T}_i \neq \mathbb{T}_j$ (and so, $i \neq j$). By Definition 10 it must be $\mathbb{T}_i\{\text{wit} \mapsto \perp\} = \mathbb{T}\{\text{wit} \mapsto \perp\} = \mathbb{T}_j\{\text{wit} \mapsto \perp\}$, hence in particular $\mathbb{T}_i.\text{in} = \mathbb{T}_j.\text{in}$. There are two cases. If $\mathbb{T}_i.\text{in} = \mathbb{T}_j.\text{in} = \perp$, then by Definition 10 \mathbf{B} is not a blockchain, since $i \neq j$. Hence, $\text{ran}(\mathbb{T}_i.\text{in}) \cap \text{ran}(\mathbb{T}_j.\text{in}) = \text{ran}(\mathbb{T}_i.\text{in}) \neq \emptyset$. By Theorem 1, this cannot happen because \mathbf{B} is consistent — contradiction. \square

Proof of Lemma 3

Straightforward from Lemma 2, taking $\mathbb{T} = \mathbb{T}_j$. \square

Proof of Theorem 2

Let $\mathbf{B} = (\mathbb{T}_1, t_1) \cdots (\mathbb{T}_n, t_n)$. By contradiction, there exists some $i < n$ such that, given $\mathbf{B}_i = (\mathbb{T}_1, t_1) \cdots (\mathbb{T}_i, t_i)$:

$$\text{val}(\mathbf{B}_i) < \text{val}(\mathbf{B}_i(\mathbb{T}_{i+1}, t_{i+1}))$$

Let U_i and U_{i+1} be the UTXOs of \mathbf{B}_i and of $\mathbf{B}_i(\mathbb{T}_{i+1}, t_{i+1})$, respectively, and let $U = U_i \cap U_{i+1}$. Since $\text{val}(U_i) < \text{val}(U_{i+1})$, then it must be $\text{val}(U_i \setminus U) < \text{val}(U_{i+1} \setminus U)$. The set $U_i \setminus U$ contains the outputs redeemed by \mathbb{T}_{i+1} , while the set $U_{i+1} \setminus U$ contains exactly the outputs in \mathbb{T}_{i+1} . Since \mathbf{B} is consistent, then $\mathbf{B}_i \triangleright (\mathbb{T}_{i+1}, t_{i+1})$. Then, by Definition 12, for each $k \in \text{dom } \mathbb{T}_{i+1}.\text{in}$, there exists a unique $j \leq i$ such that, given $o_k = \text{snd}(\mathbb{T}_{i+1}.\text{in}(k))$ and $v_k = \text{val}(\mathbb{T}_j.\text{out}(o_k))$:

$$(\mathbb{T}_j, o_k, t_j) \xrightarrow{v_k} (\mathbb{T}_{i+1}, k, t_{i+1})$$

Then, by item (3) of Definition 12:

$$\begin{aligned} \text{val}(U_i \setminus U) &= \sum \{v_k \mid k \in \text{dom } \mathbb{T}_{i+1}.\text{in}\} \\ &\geq \sum \{\text{val}(\mathbb{T}_{i+1}.\text{out}(h)) \mid h \in \text{dom } \mathbb{T}_{i+1}.\text{out}\} = \text{val}(U_{i+1} \setminus U) \end{aligned}$$

while we assumed $\text{val}(U_i \setminus U) < \text{val}(U_{i+1} \setminus U)$ — contradiction. \square