



UNIVERSITÀ DEGLI STUDI DI CAGLIARI

DIPARTIMENTO DI MATEMATICA E INFORMATICA  
DOTTORATO DI RICERCA IN MATEMATICA E INFORMATICA  
CICLO XXXIII

PH.D. THESIS

# Formal Methods for Secure Bitcoin Smart Contracts

S.S.D. INF/01

CANDIDATE  
Stefano Lande

SUPERVISOR  
Prof. Massimo Bartoletti

PHD COORDINATOR  
Prof. Michele Marchesi  
Prof. Roberto Tonelli

Academic year 2019 - 2020  
Thesis defended in April 2021



# Acknowledgments

Stefano Lande gratefully acknowledges Sardinian 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 - Axis III Education and training, Thematic goal 10, Investment Priority 10ii), Specific goal 10.5.





# Abstract

The notion of *smart contracts* was introduced in 1997 by Nick Szabo, to describe agreements among mutually distrusting parties that can be automatically enforced without resorting to a trusted intermediary. Then, the idea was mostly forgotten due to the technical impossibility to implement it. The advent of distributed ledger technologies, pioneered by Bitcoin, provided a technical foundation to reshape and develop smart contracts.

Since smart contracts handle the ownership of valuable assets, attackers may be tempted to exploit vulnerabilities in their implementation to steal or tamper with these assets. For instance, a series of vulnerabilities in Ethereum contracts have been exploited, causing money losses in the order of hundreds of millions of dollars.

Over the last years, a variety of smart contracts for Bitcoin have been proposed, both by the academic community and by that of developers. However, the heterogeneity in their treatment, the informal (often incomplete or imprecise) descriptions, and the use of poorly documented Bitcoin features, poses obstacles to the development of secure smart contracts. Using formal models and domain-specific languages to describe the behaviour of the underlying platform, and to model contracts, could help to overcome these security issues, by reducing the distance between the intended behaviour of a contract and the implementation.

In this thesis, we propose a formal model of Bitcoin transactions, which is the foundation for a new process algebra for defining Bitcoin smart contracts. Furthermore, we present a toolchain for developing smart contracts in BitML, a domain-specific language based on the contributions of this thesis. Moreover, we propose a new extension to Bitcoin, called neighbourhood covenants, which extends its expressiveness as a smart contract platform. We then exploit neighbourhood covenants to implement fungible tokens on Bitcoin.



# Contents

<b>Introduction</b>	<b>1</b>
<b>I Background</b>	<b>7</b>
<b>1 Blockchain and Smart Contracts</b>	<b>9</b>
1.1 Bitcoin . . . . .	10
1.1.1 Transactions . . . . .	11
1.1.2 Smart contracts . . . . .	12
<b>2 BitML: a calculus for Bitcoin smart contracts</b>	<b>13</b>
2.1 BitML in a nutshell . . . . .	13
2.1.1 A basic example . . . . .	15
2.2 BitML semantics . . . . .	15
2.2.1 Runs and strategies. . . . .	18
2.3 Liquidity . . . . .	18
2.3.1 Verifying liquidity . . . . .	20
<b>II Contributions</b>	<b>21</b>
<b>3 A Formal Model of Bitcoin Transactions</b>	<b>23</b>
3.1 The model . . . . .	24
3.1.1 Scripts . . . . .	25
3.1.2 Transactions . . . . .	26
3.1.3 Transaction signatures . . . . .	27
3.1.4 Semantics of scripts . . . . .	30
3.1.5 Semantics of transactions . . . . .	32
3.1.6 Blockchain and consistency . . . . .	34
3.2 Example: static chains of transactions . . . . .	39

3.3	Compiling to standard Bitcoin transactions . . . . .	41
<b>4</b>	<b>Extending Bitcoin with Neighborhood Covenants</b>	<b>43</b>
4.1	Neighbourhood covenants . . . . .	44
4.2	Use cases . . . . .	48
4.2.1	Crowdfunding . . . . .	49
4.2.2	Non-fungible tokens . . . . .	50
4.2.3	Vaults . . . . .	51
4.2.4	A pyramid scheme . . . . .	53
4.2.5	King of the Ether Throne . . . . .	54
4.3	Using covenants in high-level contract languages . . . . .	55
4.4	Implementing neighbourhood covenants . . . . .	58
<b>5</b>	<b>Bitcoin Smart Contracts as Endpoint Protocols</b>	<b>61</b>
5.1	Modelling Bitcoin contracts . . . . .	61
5.2	Smart Contracts . . . . .	66
5.2.1	Oracle . . . . .	66
5.2.2	Crowdfunding . . . . .	67
5.2.3	Escrow . . . . .	68
5.2.4	Mixing bitcoins . . . . .	70
5.2.5	Intermediated payment . . . . .	71
5.2.6	Timed commitment . . . . .	72
5.2.7	Micropayment channels . . . . .	74
5.2.8	Fair lotteries . . . . .	76
5.2.9	Contingent payments . . . . .	78
<b>6</b>	<b>A Toolchain for Developing BitML Contracts</b>	<b>81</b>
6.1	Designing BitML contracts . . . . .	83
6.2	Verifying BitML contracts . . . . .	85
6.3	Compiling BitML to Bitcoin . . . . .	87
6.4	Evaluation . . . . .	87
6.5	BitML toolchain tutorial . . . . .	90
6.5.1	Developing BitML contracts in a nutshell . . . . .	90
6.5.2	Verifying contracts with the BitML toolchain . . . . .	95
<b>7</b>	<b>On-chain Fungible Tokens on Bitcoin</b>	<b>103</b>
7.1	Overview of the approach . . . . .	104
7.2	A symbolic model of tokens . . . . .	111
7.3	Implementing tokens in Bitcoin . . . . .	114
	<b>Conclusions</b>	<b>121</b>



<b>List of Figures</b>	<b>131</b>
<b>List of Tables</b>	<b>133</b>
<b>Bibliography</b>	<b>135</b>
<b>Appendices</b>	<b>147</b>
7.4 Supplementary material for Chapter 7 . . . . .	149



# Introduction

In recent years we have observed a growing interest around *cryptocurrencies*. Bitcoin [116], 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.

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* [72], 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 [4, 26, 62, 23], gambling games [57], contingent payments [12], covenants [63, 66], and other kinds of fair computations [1, 56] are some examples of the capabilities of Bitcoin as a platform for smart contracts.

The term “smart contract” was conceived by Nick Szabo [72] to describe agreements between two or more parties, that can be automatically enforced without a trusted intermediary. Fallen into oblivion for several years, the idea of smart contract has been resurrected with the recent surge of distributed ledger technologies, led by Ethereum [98] and Hyperledger [109]. In such incarnations, smart contracts are rendered as computer programs. Users can request the execution of contracts by sending suitable *transactions* to the nodes of a peer-to-peer network. These nodes collectively maintain the history of all transactions in a public, append-only data structure, called *blockchain*. The sequence of transactions on the blockchain determines the state of each contract, and, accordingly, the assets of each user.

A crucial feature of smart contracts is that their correct execution does

not rely on a trusted authority: rather, the nodes which process transactions are assumed to be mutually untrusted. Potential conflicts in the execution of contracts are resolved through a *consensus* protocol, whose nature depends on the specific platform (e.g., it is based on “proof-of-work” in Ethereum). Ideally, contracts execute correctly whenever the adversary does not control the majority of some resource (e.g., computational power for “proof-of-work” consensus).

Since smart contracts handle the ownership of valuable assets, attackers may be tempted to exploit vulnerabilities in their implementation to steal or tamper with these assets. Although analysis tools [60, 27, 52] may improve the security of contracts, so far they have not been able to completely prevent attacks. For instance, a series of vulnerabilities in Ethereum contracts [8]) have been exploited, causing money losses in the order of hundreds of millions of dollars [124, 120, 77].

Unlike Ethereum, where contracts can be expressed as computer programs with a well-defined semantics [51, 125], Bitcoin contracts are usually realised as cryptographic protocols relying on features of Bitcoin that go beyond the standard transfers of currency. Roughly, participants of the protocols send/receive messages, verify signatures, and put/search custom-designed transactions on the blockchain. While the vast majority of Bitcoin transactions uses scripts only to verify signatures, the transactions of smart contracts 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 are 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 client [83] 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).

The informal (often incomplete or imprecise) narration of these protocols, together with the use of poorly documented features of Bitcoin (e.g., segregated witnesses, scripts, signature modifiers, temporal constraints), and the overall heterogeneity in their treatment, pose serious obstacles to the research on smart contracts in Bitcoin.

Using formal models and domain-specific languages (possibly, not

Turing-complete) to describe the behaviour of the underlying platform and to model contracts could help to overcome these security issues. by reducing the distance between contract specification and implementation.

## Contributions

**Bitcoin formalization** 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 3.1.2), and that the overall value contained in the blockchains (excluding the coinbase transactions) is decreasing (Theorem 3.1.5).

Our model formally specifies some poorly documented features of Bitcoin, e.g. transaction signatures and signature modifiers (Definition 3.5), output scripts (Definitions 3.1 and 3.9), multi-signature verification (Definition 3.8), Segregated Witnesses (Definitions 3.3 and 3.11), paving the way towards automatic verification.

Then, we provide the first systematic survey of smart contracts on Bitcoin. In order to obtain a uniform and precise treatment, we propose a new formal model of contracts, which exploits the previous one about Bitcoin transactions. This model is based on a process calculus with primitives to construct Bitcoin transactions, to put them on the blockchain, and to search the blockchain for transactions matching given patterns. Our calculus allows us to give smart contracts a precise operational semantics, which describes the interactions of the (possibly dishonest) participants involved in the execution of a contract.

We systematically formalise a large portion of the contracts proposed so far both by researchers and Bitcoin developers. In many cases, we find that specifying a contract with the intended security properties is significantly more complex than expected after reading the informal descriptions of the contract. Usually, such informal descriptions focus on the case where all participants are honest, neglecting the cases where one needs to compensate for some unexpected behaviour of the dishonest environment.

**New Bitcoin extensions** We propose a variant of covenants, named *neighbourhood covenants*, which can inspect not only the redeeming transaction, but also the siblings and the parent of the spent one. This extension preserves the basic UTXO design of Bitcoin, adding only a few opcodes to its script language, which is kept efficient, loop-free, and *non*

Turing-complete. Still, neighbourhood covenants significantly increase the expressiveness of Bitcoin as a smart contracts platform, allowing to execute arbitrary smart contracts by appending a *chain* of transactions to the blockchain. Technically, we prove that neighbourhood covenants make Bitcoin Turing-complete.

We summarise the contribution as follows:

- we propose neighbourhood covenants as a Bitcoin extension (Section 4.1), and we show that they make Bitcoin Turing-powerful (Theorem 4.1.1). We then discuss how to efficiently implement them on Bitcoin (Section 4.4);
- we use our formal model to specify complex Bitcoin contracts, which largely extend the set of use cases expressible in pure Bitcoin.
- we show how this form of covenants can be exploited in a high-level language like BitML.

**A toolchain for Developing BitML contract** We consider BitML, a high-level language for smart contracts, featuring a computationally sound embedding into Bitcoin [22], and a sound and complete verification technique of relevant trace properties [24]. BitML can express many of the smart contracts appeared in the literature [15, 9], and execute them by appending suitable transactions to the Bitcoin blockchain.

We develop a toolchain for writing and verifying BitML contracts, and for deploying them on Bitcoin. More specifically, our contribution can be summarised as follows:

- A BitML embedding in Racket [42], which allows for programming BitML contracts within the DrRacket IDE.
- A security analyzer which can check arbitrary LTL properties of BitML contracts. In particular, the analysis can decide *liquidity*, a landmark property of smart contracts. requiring that the funds within a contract do not remain frozen forever.
- A compiler from BitML contracts to standard Bitcoin transactions.
- A collection of BitML contracts, which we use as a benchmark to evaluate our toolchain. This collection contains some of the most complex contracts ever developed for Bitcoin, e.g. financial services, auctions, timed commitments, lotteries, and a variety of other gambling games. We use our benchmarks to discuss the expressiveness and the limitations of Bitcoin as a smart contracts platform.

The toolchain is open-source and its component are available at [95].

**On-chain Fungible Tokens on Bitcoin** Bitcoin, due to the limitation of its simple scripting language, does not support fungible tokens directly. Nevertheless, the first implementations of tokens were developed before Ethereum, on top of Bitcoin. All these implementations have a common drawback: the correctness of the token actions is *not* guaranteed by the consensus protocol of the blockchain. In fact, the blockchain is used just to notarize the actions that manipulate tokens, but not to check that these actions are actually permitted.

By contrast, modern blockchain platforms support *on-chain* tokens, whose correctness is guaranteed by the consensus protocol of the blockchain. Since adding native tokens to Bitcoin appears to be out of reach, we exploit neighborhood covenants, which increases the expressiveness of Bitcoin enough to support tokens. Summarising:

- we introduce a symbolic model of fungible tokens, which formalises their archetypal features: their minting and burning, the split and join operations, and the exchange of tokens with other tokens or with bitcoins (Section 7.2);
- we exploit neighbourhood covenants to implement tokens on Bitcoin (Section 7.3);

## Structure of the thesis

We briefly describe the overall structure this thesis below.

**Chapter 1: Blockchain and Smart Contracts** overviews blockchain technologies and Bitcoin in particular. Part of this material borrows from [11, 7].

**Chapter 2: BitML: a calculus for Bitcoin smart contracts** briefly overviews BitML, an high-level language for specifying smart contracts which can be executed on Bitcoin. Then, it gives an intuition of the notion of liquidity of a contract, and its verification technique. Part of this material borrows from [24].

**Chapter 3: A Formal Model of Bitcoin Transactions** proposes 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. This chapter is based on [11].

**Chapter 4: Extending Bitcoin with Neighborhood Covenants** proposes a formal model of covenants, i.e. linguistic primitives that

allow transactions to constrain the scripts of the redeeming ones. Our proposal increases the expressivity of traditional covenants by allowing the spending conditions in a transaction to depend on the neighbour transactions. We use our model to specify some complex Bitcoin contracts, and we discuss how to exploit covenants to design high-level language primitives for Bitcoin contracts. This chapter is based on [18].

### **Chapter 5: Bitcoin Smart Contracts as Endpoint Protocols**

presents a comprehensive survey of smart contracts on Bitcoin, in a uniform framework. Our treatment is based on a new formal specification language for smart contracts, which also helps us to highlight some subtleties in existing informal descriptions, making a step towards automatic verification. We discuss some obstacles to the diffusion of smart contracts on Bitcoin, and we identify the most promising open research challenges. This chapter is based on [9].

### **Chapter 6: A Toolchain for Developing BitML Contracts**

present a BitML toolchain for developing and verifying smart contracts that can be executed on Bitcoin. Our toolchain translates BitML contracts into sets of standard Bitcoin transaction, and automatically verifies relevant properties of contracts. This chapter is based on [10].

### **Chapter 7: On-chain Fungible Tokens on Bitcoin**

proposes a secure and efficient implementation of fungible tokens on Bitcoin, based on neighbourhood covenants. We propose a symbolic model for fungible token which can be applied to UTXO-based blockchains, and prove some relevant security properties. This chapter is based on [19].

**Conclusions** summarises our work and shows a comparison with the related ones. Finally, it outlines some future perspectives and extensions.



Part I

Background



# Chapter 1

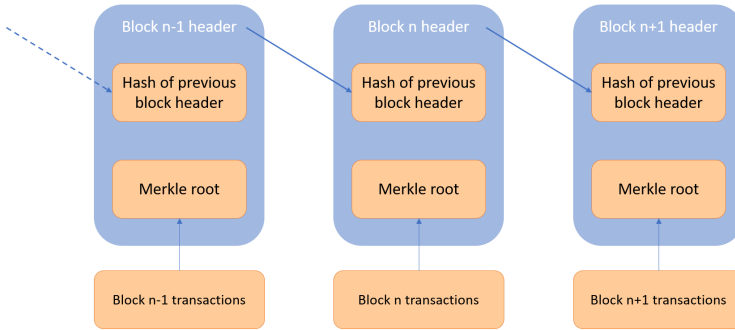
# Blockchain and Smart Contracts

A blockchain is an immutable data structure that maintains an ordered set of *transactions*. For implementation reasons, transactions are grouped into blocks. Each block is chained together with the previous one by including its hash. Changing the past, i.e. altering the data contained in a block, would change its hash, consequently invalidating the chain. The blockchain is distributed and maintained by a network of nodes: each node maintain a local copy of the blockchain and use a *consensus mechanism* that ensures that each node agrees on the new block to append.

The process of append a new block is called *mining*. A subset of nodes, called *miners*, try to build a new block to append to the blockchain. To avoid conflicts and regulate this process, miners reach consensus on the blockchain by voting. In Bitcoin and Ethereum, the vote is done though *Proof of Work*, i.e. a miner must solve a cryptographic puzzle that requires time-consuming computations. Conceptually, the more computing power a miner has, the more voting power she has on deciding the next block. In other blockchain platforms, like Algorand, miners voting power is proportional to the amount of currency they own. This mechanism is called *Proof of Stake*.

The purpose of a transaction depends on the blockchain platform, although the main goal is usually to exchange a currency. In Bitcoin, transactions register a transfer of bitcoins between users, albeit alternative uses have been studied to store arbitrary data on the blockchain [21, 17].

This chapter provides an overview of Bitcoin. Section 1.1 explains its main aspects, detailed enough to understand the formal model of Bitcoin



**Figure 1.1:** *Simplified Blockchain.*

transactions and its extension proposed in Chapters 3 and 4, and the model of Bitcoin contracts as protocols in Chapter 5.

## 1.1 Bitcoin

Bitcoin [116], is the first decentralized cryptocurrency. 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).

### 1.1.1 Transactions

Users interact with Bitcoin through *addresses*, which they can freely generate. *Transactions* describe transfers of bitcoins ( $\text{฿}$ ) 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 1.2.

$T_0$	$T_1$
in: $\dots$ wit: $\dots$ out: $(\lambda x.\text{versig}(k, x), v_0\text{฿})$	in: $T_0$ wit: $\sigma$ out: $(\lambda y.e, v_1\text{฿})$

**Figure 1.2:** *Two Bitcoin transactions.*

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  $\sigma$  is the signature of  $T_1$ , computed with the key pair  $k$ .

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  $\sigma$  in the witness of  $T_1$ . The function  $\text{versig}(k, \sigma)$  verifies that  $\sigma$  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  $\lambda y.e$  (not specified in the figure). The formalism used to represent  $T_0$  and  $T_1$  is fully presented in Chapter 3.

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 the others. Consequently, **in** fields must specify which output they are redeeming. A transaction with multiple in-

puts associates a witness to each of them. The sum of the values of all the inputs must be greater or equal to the sum of the values of all the outputs, otherwise the transaction is considered invalid. In its general form, the output script is a program in a (non 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.

### 1.1.2 Smart contracts

Albeit the primary usage of Bitcoin is to exchange currency, its blockchain and consensus mechanism can also be exploited to securely execute some forms of *smart contracts*. These are agreements among mutually distrusting parties, which can be automatically enforced without resorting to a trusted intermediary.

Bitcoin scripting language permits the definition of several smart contracts. Differently from Ethereum, where smart contracts are long-lived programs stored and invoked on the blockchain through transactions, in Bitcoin they are modelled as cryptographic protocols that may spread multiple transactions, they have no state, and cannot be reused once terminated.

A new process algebra to express smart contracts in Bitcoin is presented in Chapter 5. This formalism is expressive enough to model real use-cases in the Bitcoin community, providing a clear semantics and enabling formal reasoning. Briefly, the participants involved in a smart contract create and publish new transactions on the blockchain and interact with other participants to exchange signatures.

## Chapter 2

# BitML: a calculus for Bitcoin smart contracts

In this section we briefly overview BitML, a domain-specific language for Bitcoin smart contracts. BitML is a process calculus, with primitives to stipulate contracts and to exchange currency according to the contract terms. In this respect, BitML departs from the current practice of representing Bitcoin contracts as cryptographic protocols: rather, BitML pioneers the “contracts-as-program” paradigm for Bitcoin, by completely abstracting from Bitcoin transactions and cryptographic details. We then give some intuition about liquidity and our verification technique.

We assume a set of *participants*, ranged over by  $A, B, \dots$ , and a set of names, of two kinds:  $x, y, \dots$  denote *deposits* of  $\mathfrak{B}$ , while  $a, b, \dots$  denote *secrets*. We write  $\mathbf{x}$  (resp.  $\mathbf{a}$ ) for a finite sequence of deposit (resp. secrets) names.

### 2.1 BitML in a nutshell

BitML allows participants to exchange cryptocurrency according to pre-agreed contract rules. In BitML, any participant can broadcast a *contract advertisement*  $\{G\}C$ , where  $C$  is the actual contract, specifying the rules to transfer bitcoins ( $\mathfrak{B}$ ), while  $G$  is a set of *preconditions* to its execution.

Preconditions (Figure 2.1) may require participants to deposit some  $\mathfrak{B}$  in the contract (either upfront or at runtime), or to commit to some secret. More in detail,  $A : ! v @ x$  requires  $A$  to own  $v\mathfrak{B}$  in a deposit  $x$ , and to spend it for stipulating a contract  $C$ . Instead,  $A : ? v @ x$  only requires  $A$  to pre-authorize the spending of  $x$ , which can be gathered by the contract

$G ::= A : ? v @ x$	volatile deposit of $v\text{฿}$ , expected from $A$
$A : ! v @ x$	deposit of $v\text{฿}$ put by $A$
$A : \text{secret } a$	secret committed by $A$
$G \mid G$	composition

**Figure 2.1:** *Syntax of BitML contract preconditions.*

$C ::= \sum_{i \in I} D_i$	contract		
$D ::=$	guarded contract		
put $x$ & reveal $a$ if $p$ . $C$	put deposits and reveal secrets (if $p$ is true)		
withdraw $A$	transfer the balance to $A$		
split $v \rightarrow C$	split the balance		
$A : D$	wait for $A$ 's authorization		
after $t : D$	wait until time $t$		
$p ::=$	predicate		
true	truth	$E ::=$	contract expression
$p \wedge p$	conjunction	$a$	secret
$\neg p$	negation	$E + E$	addition
$E = E$	equality	$E - E$	subtraction
$E < E$	less than		

**Figure 2.2:** *Syntax of BitML contracts.*

at run-time. The precondition  $A : \text{secret } a$  requires  $A$  to commit to a secret  $a$  before  $C$  starts.

After  $\{G\}C$  has been advertised, each participant can choose whether to accept it, or not. When all the preconditions  $G$  have been satisfied, and all the involved participants have accepted, the contract  $C$  becomes *stipulated*. The contract starts its execution with a balance, initially set to the sum of the  $!$ -deposits required by its preconditions. Running  $C$  will affect this balance, when participants deposit/withdraw funds to/from the contract.

A contract  $C$  (Figure 2.2) is a *choice* among zero or more branches. Each branch is a *guarded contract* which enables an action, and possibly proceeds with a continuation  $C'$ . The guarded contract **withdraw**  $A$  transfers the whole balance to  $A$ , while **split**  $v_1 \rightarrow C_1 \mid \dots \mid v_n \rightarrow C_n$  decomposes the contract into  $n$  parallel components  $C_i$ , each one with



balance  $v_i$ . The guarded contract **put**  $x$  & **reveal**  $a$  **if**  $p$  atomically performs the following: (i) spend all the  $\text{?}$ -deposits  $x$ , adding their values to the contract balance; (ii) check that all the secrets  $a$  have been revealed and satisfy the predicate  $p$ . When enabled, the above-mentioned actions can be fired by anyone, at anytime. To restrict *who* can execute actions and *when*, one can use the decoration  $\mathbf{A} : D$ , which requires the authorization of  $\mathbf{A}$ , and the decoration **after**  $t : D$ , which requires to wait until time  $t$ .

### 2.1.1 A basic example

As a first example, we express in BitML the *timed commitment* [4], a basic protocol to construct more complex contracts, like e.g. lotteries and other games [5]. In the timed commitment, a participant  $\mathbf{A}$  wants to choose a secret, and promises to reveal it before some time  $t$ . The contract ensures that if  $\mathbf{A}$  does not reveal the secret in time, then she will pay a penalty of  $1\text{B}$  to  $\mathbf{B}$  (e.g., the opponent player in a game). In BitML, this is modelled as follows:

$$\{\mathbf{A} : !1\text{A}x \mid \mathbf{A} : \text{secret } a\} (\text{reveal } a. \text{withdraw } \mathbf{A} + \text{after } t : \text{withdraw } \mathbf{B})$$

The precondition requires  $\mathbf{A}$  to pay upfront  $1\text{B}$ , and to commit to a secret  $a$ . The contract is a non-deterministic choice between two branches. Only  $\mathbf{A}$  can choose the first branch, by performing **reveal**  $a$  (syntactic sugar for **put**  $\varepsilon$  & **reveal**  $a$  **if** *true*, where  $\varepsilon$  is the empty sequence). Subsequently, anyone can transfer  $1\text{B}$  to  $\mathbf{A}$ . Only after  $t$ , if the **reveal** has not been fired, any participant can fire **withdraw**  $\mathbf{B}$  in the second branch, moving  $1\text{B}$  to  $\mathbf{B}$ . So, before  $t$ ,  $\mathbf{A}$  has the option to reveal  $a$  (avoiding the penalty), or to keep it secret (paying the penalty). If no branch is taken by  $t$ , the first one who fires its **withdraw** gets  $1\text{B}$ .

## 2.2 BitML semantics

We briefly recall from [22] the semantics of BitML. The semantics is a labelled transition system between configurations of the following form:

- $\{G\}C$ , representing the advertisement of contract  $C$  with preconditions  $G$ ;
- $\langle C, v \rangle_x$ , representing a stipulated contract, holding a current balance of  $v\text{B}$ . The name  $x$  uniquely identifies the contract in a configuration;

- $\langle \mathbf{A}, v \rangle_x$  representing a fund of  $v\mathbb{B}$  owned by  $\mathbf{A}$ , and with unique name  $x$ ;
- $\mathbf{A} \triangleright \chi$ , representing  $\mathbf{A}$ 's *authorizations* to perform some operation  $\chi$ . We refer to [22] for the syntax of authorizations (some of them are exemplified below);
- $\{\mathbf{A} : a\#N\}$ , representing that  $\mathbf{A}$  has committed to a random secret  $a$  with (secret) length  $N$ ;
- $\mathbf{A} : a\#N$ , representing that  $\mathbf{A}$  has revealed her secret  $a$  (with its length  $N$ ).
- $\Gamma \mid \Delta$  is the parallel composition of two configurations (with identity  $\mathbf{0}$ );
- $\Gamma \mid t$  is a *timed* configuration, where  $t \in \mathbb{N}$  is a global time.

We now illustrate the BitML semantics by examples; when time is immaterial, we only show the steps of the untimed semantics. We omit labels on transitions.

**Deposits.** When  $\mathbf{A}$  owns a deposit  $\langle \mathbf{A}, v \rangle_x$ , she can use it in various ways: she can divide the deposit into two smaller deposits, or join it with another deposit of hers to form a larger one; the deposit can also be transferred to another participant, or destroyed. For instance, to donate a deposit  $x$  to  $\mathbf{B}$ ,  $\mathbf{A}$  must first issue the authorization  $\mathbf{A} \triangleright x \triangleright \mathbf{B}$ ; then, anyone can transfer the money to  $\mathbf{B}$ :

$$\langle \mathbf{A}, v \rangle_x \mid \cdots \rightarrow \langle \mathbf{A}, v \rangle_x \mid \mathbf{A} \triangleright x \triangleright \mathbf{B} \mid \cdots \rightarrow \langle \mathbf{B}, v \rangle_y \mid \cdots \quad (y \text{ fresh})$$

**Advertisement.** Any participant can advertise a new contract  $\mathbf{C}$  (with preconditions  $\mathbf{G}$ ). This is obtained by performing the step  $\Gamma \rightarrow \Gamma \mid \{\mathbf{G}\}\mathbf{C}$ .

**Stipulation.** Stipulation turns a contract advertisement into an active contract. For instance, let  $\mathbf{G} = \mathbf{A} : !1\textcircled{x} \mid \mathbf{A} : ?1\textcircled{y} \mid \mathbf{A} : \text{secret } a$ . Given a contract  $\mathbf{C}$ , the stipulation of  $\{\mathbf{G}\}\mathbf{C}$  is done in a few steps:

$$\langle \mathbf{A}, 1 \rangle_x \mid \langle \mathbf{A}, 1 \rangle_y \mid \{\mathbf{G}\}\mathbf{C} \rightarrow^* \langle \mathbf{A}, 1 \rangle_y \mid \langle \mathbf{C}, 1 \rangle_z \mid \{\mathbf{A} : a\#N\}$$

Above, the funds in the deposit  $x$  are transferred to the newly created contract, to fulfill the precondition  $\mathbf{A} : !1\textcircled{x}$ . Instead, the deposit  $y$  remains in the configuration, to be possibly spent after some time. The component  $\{\mathbf{A} : a\#N\}$  represents the secret committed to by  $\mathbf{A}$ , with its length  $N$ .

**Withdraw.** Executing `withdraw A` terminates the contract, and transfers its whole balance to `A` by creating a fresh deposit owned by `A`:

$$\langle \text{withdraw } A + C', v \rangle_x \rightarrow \langle A, v \rangle_y \quad (y \text{ fresh})$$

Above, `withdraw A` is executed as a branch within a choice: as usual, taking a branch discards the other ones (denoted as `C'`).

**Split.** The `split` primitive can be used to spawn several new concurrent contracts, dividing the balance among them. For instance:

$$\langle (\text{split } v_1 \rightarrow C_1 \mid v_2 \rightarrow C_2), v_1 + v_2 \rangle_x \rightarrow \langle C_1, v_1 \rangle_y \mid \langle C_2, v_2 \rangle_z \\ (y, z \text{ fresh})$$

**Reveal.** A prefix `reveal a if p` can be fired when the previously committed secret `a` (satisfying the predicate `p`) has been revealed. For instance:

$$\langle \text{reveal } a \text{ if } a = N. C, v \rangle_x \mid \{A : a \# N\} \\ \rightarrow \langle \text{reveal } a \text{ if } a = N. C, v \rangle_x \mid A : a \# N \\ \rightarrow \langle C, v + v' \rangle_y \mid A : a \# N$$

In the first step, `A` reveals her secret `a`. In the second step, any participant fires the prefix.

**Authorizations.** When a branch is decorated by `A : ...` it can be taken only after `A` has provided her authorization. For instance:

$$\langle A : \text{withdraw } B + A : \text{withdraw } C, v \rangle_x \\ \rightarrow \langle A : \text{withdraw } B + A : \text{withdraw } C, v \rangle_x \mid A \triangleright x \triangleright A : \text{withdraw } B \\ \rightarrow \langle B, v \rangle_y$$

In the first step, `A` authorizes to take the branch `withdraw B`. After that, any participant can fire such branch.

**Time.** We always allow time `t` to advance by a delay  $\delta > 0$ , through a transition  $\Gamma \mid t \rightarrow \Gamma \mid t + \delta$ . Advancing time can enable branches decorated with `after t`. For instance, if  $t_0 + \delta \geq t$ , we have the following computation:

$$\langle (\text{after } t : \text{withdraw } B) + C', v \rangle_x \mid t_0 \\ \rightarrow \langle (\text{after } t : \text{withdraw } B) + C', v \rangle_x \mid t_0 + \delta \rightarrow \langle B, v \rangle_y \mid t_0 + \delta$$

### 2.2.1 Runs and strategies.

A *run*  $\mathcal{S}$  is a (possibly infinite) sequence:

$$\Gamma_0 \mid t_0 \xrightarrow{\alpha_0} \Gamma_1 \mid t_1 \xrightarrow{\alpha_1} \dots$$

where  $\alpha_i$  are the transition labels,  $\Gamma_0$  contains only deposits, and  $t_0 = 0$ . If  $\mathcal{S}$  is finite, we write  $\Gamma_{\mathcal{S}}$  for its last untimed configuration, and  $\delta_{\mathcal{S}}$  for its last time. A *strategy*  $\Sigma_{\mathbf{A}}^s$  is a PPTIME algorithm which allows  $\mathbf{A}$  to select which actions to perform (possibly, time delays), among those permitted by the BitML semantics. The choice among these actions is controlled by the adversary strategy  $\Sigma_{\mathbf{Adv}}^s$ , which acts on behalf of all the dishonest participants. Given the strategies of all participants (including  $\mathbf{Adv}$ ), there is a unique run *conforming* to all of them.

## 2.3 Liquidity

A desirable property of smart contracts is *liquidity*, which requires that the contract balance is always eventually transferred to some participant. In a non-liquid contract, funds can be frozen forever, unavailable to anyone, hence effectively destroyed. There are many possible flavours of liquidity, depending e.g. on which participants are assumed to be honest, and on which are their strategies. The simplest form of liquidity is to consider the case where everyone cooperates: i.e. a contract is liquid if there exists some strategy for each participant such that no funds are ever frozen. However, this notion does not capture the essence of smart contracts, i.e. to allow mutually untrusted participants to safely interact.

For instance, consider the following contract, where  $\mathbf{A}$  and  $\mathbf{B}$  contribute  $1\mathbb{B}$  each for a donation of  $2\mathbb{B}$  to either  $\mathbf{C}$  or  $\mathbf{D}$  (we omit the preconditions for brevity):

$$\mathbf{A} : \mathbf{B} : \text{withdraw } \mathbf{C} + \mathbf{A} : \mathbf{B} : \text{withdraw } \mathbf{D}$$

In order to unlock the funds,  $\mathbf{A}$  and  $\mathbf{B}$  must agree on the recipient of the donation, by giving their authorization on the same branch. This contract would be liquid only by assuming the cooperation between  $\mathbf{A}$  and  $\mathbf{B}$ : indeed,  $\mathbf{A}$  alone cannot guarantee that the  $2\mathbb{B}$  will eventually be donated, as  $\mathbf{B}$  can choose a different recipient, or even refuse to give any authorization. Consequently, unless  $\mathbf{A}$  trusts  $\mathbf{B}$ , it makes sense to consider this contract as non-liquid, from the point of view of  $\mathbf{A}$  (and for similar reasons, also from that of  $\mathbf{B}$ ).

Consider now the timed commitment contract discussed before:

$$\text{reveal } a. \text{withdraw } \mathbf{A} + \text{after } t : \text{withdraw } \mathbf{B}$$

This contract is liquid from **A**'s point of view (even if **B** is dishonest), because **A** can reveal the secret and then redeem the funds from the contract. The timed commitment is also liquid from **B**'s point of view: if **A** does not reveal the secret (making the first branch stuck), the funds in the contract can be redeemed through the second branch, after time  $t$ .

In a *mutual* timed commitment contract, where **A** and **B** have to exchange their secrets or pay a  $1\text{฿}$  penalty, achieving liquidity is a bit more challenging. We first consider a wrong attempt:

$$\begin{aligned} & \text{reveal } a. \text{reveal } b. \text{split } (1\text{฿} \rightarrow \text{withdraw } \mathbf{A} \mid 1\text{฿} \rightarrow \text{withdraw } \mathbf{B}) \\ & + \text{after } t : \text{withdraw } \mathbf{B} \end{aligned}$$

Intuitively, **A** has only the following strategies, according to when she decides to reveal her secret  $a$ : (i) **A** chooses to reveal  $a$  unconditionally, and to perform the `reveal  $a$`  action. This strategy is *not* liquid: indeed, if **B** does not reveal  $b$ , the contract is stuck. (ii) **A** chooses to reveal  $a$  only *after* **B** has revealed  $b$ . This strategy is *not* liquid: indeed, if **B** chooses not to reveal  $b$ , the contract will never advance. (iii) **A** chooses to wait until **B** reveals secret  $b$ , or until time  $t' \geq t$ , whichever comes first. If  $b$  was revealed, **A** reveals  $a$ , and splits the contract balance between **A** and **B**. Otherwise, if the deadline  $t'$  is expired, **A** transfers the whole balance to **B**. Note that, although this strategy is liquid, it is not satisfactory for **A**, since in the second case she will lose money.

This example highlights a crucial point: participants' strategies have to be taken into account when defining liquidity. Indeed, the mere fact that a liquid strategy exists does not imply that it is the ideal strategy for the honest participant. To fix this issue, we revise the mutual timed commitment as follows:

$$\begin{aligned} & \text{reveal } a. (\text{reveal } b. \text{split } (1\text{฿} \rightarrow \text{withdraw } \mathbf{A} \mid 1\text{฿} \rightarrow \text{withdraw } \mathbf{B}) \\ & \quad + \text{after } t' : \text{withdraw } \mathbf{A}) \\ & + \text{after } t : \text{withdraw } \mathbf{B} \end{aligned}$$

where  $t < t'$ . Now, **A** has a liquid strategy where she does not pay the penalty. First, **A** reveals  $a$  before time  $t$ . After that, if **B** reveals  $b$ , then **A** can execute the `split`, transferring  $1\text{฿}$  to herself and  $1\text{฿}$  to **B** (note that this does not require **B**'s cooperation); otherwise, after time  $t'$ , **A** can withdraw  $2\text{฿}$  by executing the `withdraw  $\mathbf{A}$`  in the `after  $t'$  : ...` branch.

These examples, albeit elementary, show that detecting if a strategy is liquid for a contract is not straightforward, in general. The problem of determining a liquid strategy for a given contract seems even more demanding. Automatic techniques for the verification and inference of liquid strategies can be useful tools for the developers of smart contracts.

### 2.3.1 Verifying liquidity

One of the main features of BitML is a verification technique for the liquidity of BitML contracts. The technique is based on a more general result, i.e. a strict correspondence between the concrete semantics of BitML and a new abstract semantics, which is finite-state. The abstraction is a correct and complete approximation of the concrete semantics with respect to a given set of contracts [24]. To obtain a finite-state abstraction, we need to cope with three sources of infiniteness of the concrete semantics of BitML: the unbounded passing of time, the advertisement/stipulation of new contracts, and the operations on deposits. The abstraction replaces the time  $t$  in concrete configurations with a finite number of time intervals  $T = [t_0, t_1)$ , and it disables the transitions to advertise new contracts. Further, the only operations on deposits allowed by the abstract semantics are the ones for transferring them to contracts and for destroying them. The latter is needed e.g. to properly model the situation where a participant spends a  $\text{?}$ -deposit.

The intended use of our abstraction is to start from a configuration containing an arbitrary (but finite) set of contracts, and then analyse their possible evolutions in the presence of an honest participant and an adversary. This produces a finite set of (finite) traces, which we can model-check for liquidity. Soundness and completeness of the abstraction are exploited to prove that liquidity is decidable. The computational soundness of the BitML compiler [22] guarantees that if a contract is verified to be liquid according to our analysis, this property is preserved when executing it on Bitcoin.

**Part II**

**Contributions**





## Chapter 3

# A Formal Model of Bitcoin Transactions

Bitcoin [116], 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.

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* [72], 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 [4, 26, 62, 23], gambling games [57], contingent payments [12], covenants [63, 66], and other kinds of fair computations [1, 56] are some examples of the capabilities of Bitcoin as a platform for smart contracts.

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 above-mentioned ones 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 are 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 client [83] 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).

This Chapter is structured as follows. In Section 3.1 we formalise Bitcoin transactions. Besides transactions, we also provide an high-level model of the blockchain, and we study its basic properties. In Section 3.2 we illustrate, through a basic case study, the impact of the Segregated Witness feature on the expressiveness of Bitcoin smart contracts. In Section 3.3 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 7.3.

## 3.1 The model

In this section 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.

In Section 3.1.1 we define the scripts that can be used in transaction outputs. Then, in Section 3.1.2 we formalise transactions, and in Section 3.1.3 we define a signature scheme for them. Sections 3.1.4 and 3.1.5 give semantics, respectively, to scripts and transactions. In Section 3.1.6 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.

We start by introducing some auxiliary notation. We assume several sets, ranged over by meta-variables as shown in the left column of Table 3.1. We use the bold notation to denote finite sequences of elements. We denote with  $\mathbf{x}_i$  the  $i$ -th element of a sequence  $\mathbf{x}$ , 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 num-

$A, B, \dots \in \text{Part}$	Participants	$e, e', \dots \in \text{Exp}$	Script expressions
$x, y, \dots \in \text{Var}$	Variables	$T, T', \dots \in \text{Tx}$	Transactions
$\nu, \nu', \dots \in \text{Den}$	Denotations, i.e.:	$\mu, \mu'$	Signature modifier
$k, k' \dots \in \mathbb{Z}$	Constants	$\text{sig}_k^{\mu, i}(T)$	Transaction signature
$t, t' \dots \in \mathbb{N}$	Time	$\text{ver}_k(\sigma, T, i)$	Signature verification
$v, v' \dots \in \mathbb{N}$	Currency values	$T, i \models \lambda \mathbf{x}. e$	Script verification
$\sigma, \sigma', \dots \in \mathbb{Z}$	Signatures	$(T, i, t) \overset{v}{\rightsquigarrow} (T', j, t')$	Transaction redeem
$\text{true}, \text{false}$	Boolean values	$\mathbf{B} = (T_1, t_1) \dots$	Blockchains
$\perp$	Undefined	$\mathbf{B} \triangleright (T, t)$	Consistent update

**Table 3.1:** *Summary of notation.*

ber of elements of  $\mathbf{x}$ , and with  $\varepsilon$  the empty sequence. We denote with  $f : A \rightarrow B$  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. For a pair  $(x, y)$ , we define  $\text{fst}(x, y) = x$  and  $\text{snd}(x, y) = y$ .

### 3.1.1 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 Bitcoin [89], by using instead a minimalistic language of expressions.

**Definition 3.1** (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  $\lambda \mathbf{z}. e$  such that all the variables in  $e$  occur in  $\mathbf{z}$ .

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  $H(e)$  evaluates to the hash of  $e$ . The expression  $\text{versig}(\mathbf{k}, e)$  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 3.1.4). They evaluate to  $e$  if the constraints are satisfied, otherwise they fail.

**Notation 3.2.** *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.*

### 3.1.2 Transactions

The following definition formalises Bitcoin transactions.

**Definition 3.3** (Transactions). *We inductively define the set  $\text{Tx}$  of transactions as follows. A transaction  $\mathbf{T}$  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, for all  $i, j \in \text{dom in}$ ,  $\text{fst}(\text{in}(i)).\text{wit} = \perp$  and  $i \neq j \implies \text{in}(i) \neq \text{in}(j)$ .

We denote with  $\mathbf{T.f}$  the value of field  $\mathbf{f}$  of  $\mathbf{T}$ , for  $\mathbf{f} \in \{\text{in}, \text{wit}, \text{out}, \text{absLock}, \text{relLock}\}$ .

We say that  $\mathbf{T}$  is initial when  $\mathbf{T.in} = \mathbf{T.relLock} = \perp$  and  $\mathbf{T.absLock} = 0$ .

The fields  $\text{in}$  and  $\text{out}$  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  $\mathbf{T.in}(i) = (\mathbf{T}', j)$ , it means that the  $i$ -th input

of  $\mathsf{T}$  wants to redeem the  $j$ -th output of  $\mathsf{T}'$ . The side condition  $i \neq j \Rightarrow \mathsf{in}(i) \neq \mathsf{in}(j)$  ensures that inputs are pairwise distinct. The side condition  $\mathit{fst}(\mathsf{in}(i)).\mathit{wit} = \perp$  is related to the *Segregated Witness* (SegWit) feature and it requires that the witness of the input transaction is left unspecified. This feature, specified in the BIP 141 [113] and activated on August 24th 2017, implies that witnesses are not used in the computation of transaction hashes. The output  $\mathsf{T}'.\mathit{out}(j)$  is a pair  $(\lambda z.e, v)$ , meaning that  $v$  Satoshis ( $1\text{B} = 10^8$  Satoshis) can be redeemed by whoever can provide a witness which makes the term  $\lambda z.e$  evaluate to *true*. Such witness is defined by  $\mathsf{T}.\mathit{wit}(i)$ . The fields  $\mathsf{T}.\mathit{absLock}$  and  $\mathsf{T}.\mathit{relLock}(i)$  specify a constraint on when  $\mathsf{T}$  can be put on the blockchain: the first in absolute terms, whereas the second is relative to the transaction in the input  $\mathsf{T}.\mathit{in}(i)$ . More specifically,  $\mathsf{T}.\mathit{absLock} = t$  means that  $\mathsf{T}$  can appear on the blockchain only after time  $t$ . If  $\mathsf{T}.\mathit{relLock}(i) = t$ , then  $\mathsf{T}$  can appear only after time  $t$  since the transaction in  $\mathsf{T}.\mathit{in}(i)$  appeared.

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

### 3.1.3 Transaction signatures

We extend to transactions the signing and verification functions of the public key signature schemes, denoted respectively as  $\mathit{sig}_k(\cdot)$  and  $\mathit{ver}_k(\cdot, \cdot)$ . For simplicity, although we will always use  $k = (k_p, k_s)$  for key pairs, we implicitly assume that  $\mathit{sig}_k(\cdot)$  only uses the private part  $k_s$ , while  $\mathit{ver}_k(\cdot, \cdot)$  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 3.5. In order to specify which parts of a transaction are signed, we first introduce the auxiliary notion of *transaction substitution*.

**Definition 3.4** (Transaction substitutions). *A transaction substitution  $\Sigma$  is a function from  $\mathsf{Tx}$  to  $\mathsf{Tx}$ . For a transaction field  $\mathsf{f}$ , we denote with  $\{\mathsf{f} \mapsto d\}$  the substitution which replaces the value of  $\mathsf{f}$  with  $d$ . For  $\mathsf{f} \neq \mathit{absLock}$  and  $i \in \mathbb{N}$ , we denote with  $\{\mathsf{f}(i) \mapsto d\}$  the substitution which replaces  $\mathsf{f}(i)$  with  $d$ . Further, for  $\circ \in \{<, >, \neq\}$ , we denote with  $\{\mathsf{f}(\circ i) \mapsto$*

$$\begin{aligned}
**_i(\mathbb{T}) &= \mathbb{T}\{\mathbf{wit}(1) \mapsto i\}\{\mathbf{wit}(\neq 1) \mapsto \perp\} \\
*0_i(\mathbb{T}) &= **_i(\mathbb{T}\{\mathbf{out} \mapsto \perp\}) \\
*1_i(\mathbb{T}) &= **_i(\mathbb{T}\{\mathbf{out}(< i) \mapsto (\mathit{false}, 0)\}\{\mathbf{out}(> i) \mapsto \perp\}) \\
1*_i(\mathbb{T}) &= **_1(\mathbb{T}\{\mathbf{in}(1) \mapsto \mathbb{T}.\mathbf{in}(i)\}\{\mathbf{in}(\neq 1) \mapsto \perp\} \\
&\quad \{\mathbf{relLock}(1) \mapsto \mathbb{T}.\mathbf{relLock}(i)\}\{\mathbf{relLock}(\neq 1) \mapsto \perp\}) \\
10_i(\mathbb{T}) &= 1*_i(*0_i(\mathbb{T})) \\
11_i(\mathbb{T}) &= 1*_i(*1_i(\mathbb{T}))
\end{aligned}$$

**Figure 3.1:** Signature modifiers.

$d\}$  the substitution which replaces  $\mathbf{f}(j)$  with  $d$ , for all  $j \circ i \in \text{dom } \mathbf{f}$ .

**Definition 3.5** (Signature modifiers). *We define signature modifiers  $\mu[i]$  (with  $i \in \mathbb{N}$ ) in Figure 3.1. We associate to each modifier a substitution, and we denote with  $\mu[i](\mathbb{T})$  the result of applying it to the transaction  $\mathbb{T}$ .*

Each modifier is represented by a pair of symbols, describing, respectively, the set of inputs and of outputs being signed ( $*$  = all,  $1$  = single,  $0$  = none), and an index  $i \in \mathbb{N}$ . The index has different meanings, depending on the modifier. Regarding the first symbol of the modifier, if it is  $*$ , then  $i$  is the index of the witness where the signature will be included, so to ensure that a signature computed for being included in the witness at index  $i$  can not be used in any witness with index  $j \neq i$  (see Example 3.1.4). If the first symbol of the modifier is  $1$ , then only the  $i$ -th input is signed, while all the other inputs are removed from the transaction. With respect to the second symbol of the modifier, if it is  $1$ , then  $i$  is the index of the signed output; otherwise,  $i$  has no effect on the outputs to be signed. Note that a single index is used for both inputs and outputs: in any case, the index refers to the witness where the signature will be included.

**Definition 3.6** (Transaction signatures). *We define the transaction signature (under modifier  $\mu$  and index  $i$ ) and verification functions as follows:*

$$\begin{aligned}
\mathbf{sig}_k^{\mu,i}(\mathbb{T}) &= (\mathit{sig}_k(\mu[i](\mathbb{T}), \mu), \mu) \\
\mathbf{ver}_k(\sigma, \mathbb{T}, i) &= \mathit{ver}_k(w, (\mu[i](\mathbb{T}), \mu)) \text{ if } \sigma = (w, \mu)
\end{aligned}$$

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

Note that a signature  $\sigma = (\text{sig}_k((\mu[i](\mathbb{T}), \mu)), \mu)$  does not contain the index  $i$ . Consequently, the verification function requires  $i$  to be passed as parameter, i.e. we write  $\text{ver}_k(\sigma, \mathbb{T}, i)$ . The parameter  $i$  will be instantiated by the script verification function (see Definition 3.10). Besides the modified transaction  $\mu[i](\mathbb{T})$ , the signature also applies to the modifier  $\mu$ . In this way, signing a single-input transaction  $\mathbb{T}$  with modifier  $**_1$  and with modifier  $1*_1$  results in two different signatures, even though  $**_1(\mathbb{T}) = 1*_1(\mathbb{T})$ .

**Notation 3.7.** Note that  $\text{sig}_k^{\mu,i}(\mathbb{T})$  can meaningfully appear within  $\mathbb{T}.\text{wit}(i)$ , since such signature does not depend on the  $\text{wit}$  field of transactions (as all signature modifiers overwrite all the witnesses). When a signature of  $\mathbb{T}$  appears within  $\mathbb{T}.\text{wit}(i)$ , as a shorthand we denote it with  $\text{sig}_k^{\mu}$  (so, neglecting the enclosing transaction  $\mathbb{T}$  and the index  $i$ ), or just  $\text{sig}_k$  when  $\mu = **$ .

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  $\sigma$ , one has many keys and signatures, i.e.  $\text{ver}_{\mathbf{k}}(\boldsymbol{\sigma}, \mathbb{T}, i)$ . Intuitively, if  $|\boldsymbol{\sigma}| = m$  and  $|\mathbf{k}| = n$ , the function  $\text{ver}_{\mathbf{k}}(\boldsymbol{\sigma}, \mathbb{T}, i)$  implements a  $m$ -of- $n$  multi-signature scheme, i.e. it evaluates to true if all the  $m$  signatures match (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 3.8** (Multi-signature verification). Let  $\mathbf{k}$  and  $\boldsymbol{\sigma}$  be sequences of (public) keys and signatures such that  $|\mathbf{k}| \geq |\boldsymbol{\sigma}|$ , and let  $i \in \mathbb{N}$ . For all  $m, n \in \mathbb{N}$ , we define the function:

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

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

Our formalisation of multi-signature verification (Definition 3.8) follows closely the implementation of Bitcoin, whose stack-based scripting language imposes that the sequence  $\boldsymbol{\sigma}$  is read in reverse order. Accordingly, the function  $\text{ver}$  tries to verify the last signature in  $\boldsymbol{\sigma}$  with the last key in  $\mathbf{k}$ . If they match, the function  $\text{ver}$  proceeds to verify the previous

signature in the sequence, otherwise it tries to verify the signature with the previous key.

**Example 3.1.1** (2-of-3 multi-signature). Let  $\mathbf{k} = k_a k_b k_c$ , and let  $\sigma = \sigma_p \sigma_q$  be such that  $\text{ver}_{k_a}(\sigma_p, \mathbb{T}, 1) = \text{ver}_{k_b}(\sigma_q, \mathbb{T}, 1) = \text{true}$ , and false otherwise. We have that:

$$\begin{aligned}
 \text{ver}_{\mathbf{k}}(\sigma, \mathbb{T}, 1) &= \text{ver}_{\mathbf{k}}^{3,2}(\sigma, \mathbb{T}, 1) && (\text{as } |\mathbf{k}| = 3 \text{ and } |\sigma| = 2) \\
 &= \text{ver}_{\mathbf{k}}^{2,2}(\sigma, \mathbb{T}, 1) && (\text{as } \text{ver}_{k_c}(\sigma_q, \mathbb{T}, 1) = \text{false}) \\
 &= \text{ver}_{\mathbf{k}}^{1,1}(\sigma, \mathbb{T}, 1) && (\text{as } \text{ver}_{k_b}(\sigma_q, \mathbb{T}, 1) = \text{true}) \\
 &= \text{ver}_{\mathbf{k}}^{0,0}(\sigma, \mathbb{T}, 1) && (\text{as } \text{ver}_{k_a}(\sigma_p, \mathbb{T}, 1) = \text{true}) \\
 &= \text{true} && (\text{as } m = 0)
 \end{aligned}$$

Note that, if we let  $\sigma' = \sigma_q \sigma_p$ , the resulting evaluation will be:

$$\begin{aligned}
 \text{ver}_{\mathbf{k}}(\sigma', \mathbb{T}, 1) &= \text{ver}_{\mathbf{k}}^{3,2}(\sigma', \mathbb{T}, 1) && (\text{as } |\mathbf{k}| = 3 \text{ and } |\sigma'| = 2) \\
 &= \text{ver}_{\mathbf{k}}^{2,2}(\sigma', \mathbb{T}, 1) && (\text{as } \text{ver}_{k_c}(\sigma_p, \mathbb{T}, 1) = \text{false}) \\
 &= \text{ver}_{\mathbf{k}}^{1,2}(\sigma', \mathbb{T}, 1) && (\text{as } \text{ver}_{k_b}(\sigma_p, \mathbb{T}, 1) = \text{false}) \\
 &= \text{ver}_{\mathbf{k}}^{0,1}(\sigma', \mathbb{T}, 1) && (\text{as } \text{ver}_{k_a}(\sigma_p, \mathbb{T}, 1) = \text{true}) \\
 &= \text{false} && (\text{as } m \neq 0 \text{ and } n = 0) \quad \square
 \end{aligned}$$

### 3.1.4 Semantics of scripts

Definition 3.9 gives the semantics of script expressions. This semantics will be used in Section 3.1.5 to define when a transaction can redeem another one. We 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 3.9 are *strict*, i.e. they evaluate to  $\perp$  if some of their operands is  $\perp$ .

**Definition 3.9** (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 3.2, where*



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

**Figure 3.2:** Semantics of script expressions.

$\text{Den} = \mathbb{Z} \cup \{\text{true}, \text{false}\}$ . We use the following operators on denotations:

$$\begin{aligned}
\text{if } \nu_0 \text{ then } \nu_1 \text{ else } \nu_2 &\equiv \begin{cases} \nu_1 & \text{if } \nu_0 = \text{true} \\ \nu_2 & \text{if } \nu_0 = \text{false} \\ \perp & \text{otherwise} \end{cases} \\
\text{size}(\nu) &\equiv \begin{cases} \perp & \text{if } \nu \notin \mathbb{Z} \\ 0 & \text{if } \nu = 0 \\ \left\lceil \frac{\log_2 |\nu| + 1}{7} \right\rceil & \text{otherwise} \end{cases} \\
\nu_0 \circ_{\perp} \nu_1 &\equiv \text{if } \nu_0, \nu_1 \in \mathbb{Z} \text{ then } \nu_0 \circ \nu_1 \text{ else } \perp \\
&\quad (\circ \in \{+, -, =, <\})
\end{aligned}$$

The function  $\text{size}(\nu)$  denotes the amount of bytes that are necessary to represent the value  $\nu$ , as defined in the Bitcoin scripting language [89].

**Definition 3.10** (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 3.1.2.** 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  $\sigma = \text{sig}_k^{**}(\mathbb{T})$ .

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:

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

### 3.1.5 Semantics of transactions

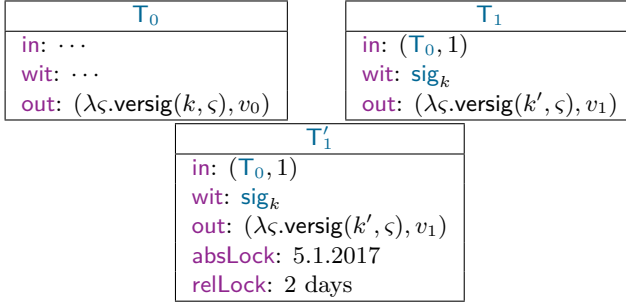
Definition 3.11 describes when the  $j$ -th input of a transaction  $\mathbb{T}'$  (put on the blockchain at time  $t'$ ) can redeem  $v$  Satoshis from the  $i$ -th output of the transaction  $\mathbb{T}$  (put on the blockchain at time  $t$ ). We denote this by  $(\mathbb{T}, i, t) \overset{v}{\rightsquigarrow} (\mathbb{T}', j, t')$ .

**Definition 3.11** (Output redeeming). *We write  $(\mathbb{T}, i, t) \overset{v}{\rightsquigarrow} (\mathbb{T}', j, t')$  iff all the following conditions hold:*


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

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

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



**Figure 3.3:** 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  $\overline{T_0}$  within  $\overline{T_1}.\text{in}$ ). Also, we use dates in time constraints.



**Figure 3.4:** Three transactions for Example 3.1.4. Note that, by Definition 3.7, the first witness of  $T'_3$  is  $\text{sig}_k^{**,1}(T'_3)$ , while the second is  $\text{sig}_k^{**,2}(T'_3)$ .

**Example 3.1.3.** With the transactions in Figure 3.3, we have  $(\overline{T_0}, 1, t_0) \overset{v_0}{\rightsquigarrow} (\overline{T_1}, 1, t_1)$ . Indeed, for item (a) we have that  $\overline{T_1}.\text{in}(1) = (\overline{T_0}\{\text{wit} \mapsto \perp\}, 1)$ ; for item (b),  $\overline{T_1}, 1 \models \lambda\varsigma.\text{versig}(k, \varsigma)$ ; for item (c),  $v_0 = \text{val}(\overline{T_0}.\text{out}(1))$ . The other two items trivially hold, as there are no time constraints. We also have  $(\overline{T_0}, 1, 2.1.2017) \overset{v_0}{\rightsquigarrow} (\overline{T'_1}, 1, 6.1.2017)$ . To show that, we have to check also items (d) and (e). For item (d), we have that  $6.1.2017 \geq \overline{T'_1}.\text{absLock} = 5.1.2017$ . For item (e), we have that  $6.1.2017 - 2.1.2017 \geq \overline{T'_1}.\text{relLock}(1) = 2$  days.  $\square$

**Example 3.1.4.** Consider the transactions in Figure 3.4. The signature in  $\overline{T'_3}.\text{wit}(1)$  is computed as follows:

$$\begin{aligned} \text{sig}_k^{**,1}(T'_3) &= (\text{sig}_k(**1(T'_3, **)), **) && \text{by Definition 3.6} \\ &= (\text{sig}_k(\overline{T'_3}\{\text{wit}(1) \mapsto 1\}\{\text{wit}(\neq 1) \mapsto \perp\}, **), **) && \text{by Definition 3.5} \end{aligned}$$

We prove that, when verifying  $(\mathbb{T}'_1, 1, t) \rightsquigarrow^1 (\mathbb{T}'_3, 1, t')$ , item (b) of Definition 3.11 holds, i.e.  $\mathbb{T}'_3, 1 \models \text{scr}(\mathbb{T}'_1.\text{out}(1))$ . To this purpose, let  $\rho = \{\varsigma \mapsto (w, **)\}$ , where  $w = \text{sig}_k(\mathbb{T}'_3\{\text{wit}(1) \mapsto 1\}\{\text{wit}(\neq 1) \mapsto \perp\}, **)$ . We have that:

$$\begin{aligned}
\llbracket \text{versig}(k, \varsigma) \rrbracket_{\mathbb{T}'_3, 1, \rho} &= \text{ver}_k(\llbracket \varsigma \rrbracket_{\mathbb{T}'_3, 1, \rho}, \mathbb{T}'_3, 1) && \text{by Def. 3.9} \\
&= \text{ver}_k((w, **), \mathbb{T}'_3, 1) && \rho(\varsigma) = (w, **) \\
&= \text{ver}_k(w, (**_1(\mathbb{T}'_3), **)) && \text{by Def. 3.6} \\
&= \text{ver}_k(w, (\mathbb{T}'_3\{\text{wit}(1) \mapsto 1\}\{\text{wit}(\neq 1) \mapsto \perp\}, **)) && \text{by Def. 3.5} \\
&= \text{true} && \text{by Def. of } w
\end{aligned}$$

We now show that  $w$  is *not* valid for the other witness, i.e.  $(\mathbb{T}'_2, 1, t) \rightsquigarrow^2 (\mathbb{T}''_3, 2, t')$ , where  $\mathbb{T}''_3 = \mathbb{T}'_3\{\text{wit}(2) \mapsto \text{sig}_k^{**,1}(\mathbb{T}'_3)\}$ . Let  $\rho = \{\varsigma \mapsto (w, **)\}$ . Item (b) of Definition 3.11 does not hold:

$$\begin{aligned}
\llbracket \text{versig}(k, \varsigma) \rrbracket_{\mathbb{T}''_3, 2, \rho} &= \text{ver}_k((w, **), \mathbb{T}''_3, 2) && \text{as above} \\
&= \text{ver}_k(w, (**_2(\mathbb{T}''_3), **)) && \text{by Def. 3.6} \\
&= \text{ver}_k(w, (\mathbb{T}'_3\{\text{wit}(1) \mapsto 2\}\{\text{wit}(\neq 1) \mapsto \perp\}, **)) && \text{by Def. 3.5} \\
&= \text{false}
\end{aligned}$$

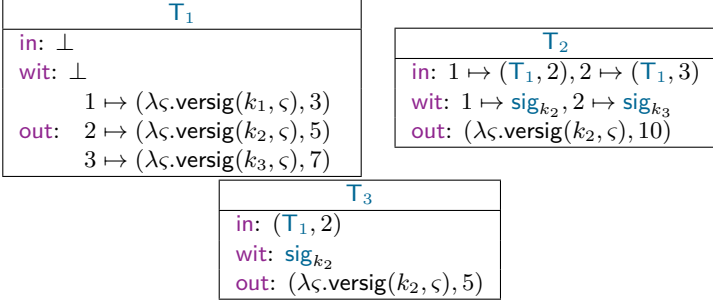
In the last equation,  $w$  is not a valid signature for  $\mathbb{T}''_3\{\text{wit}(1) \mapsto 2\}\{\text{wit}(\neq 1) \mapsto \perp\}$  because it is computed on  $\mathbb{T}'_3\{\text{wit}(1) \mapsto 1\}\{\text{wit}(\neq 1) \mapsto \perp\}$ , and the two transactions differ on  $\text{wit}(1)$ .  $\square$

### 3.1.6 Blockchain and consistency

In Definition 3.12 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 3.15.

**Definition 3.12** (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  $\mathbb{T}\{\text{wit} \mapsto \perp\} = \mathbb{T}_i\{\text{wit} \mapsto \perp\}$ .*



**Figure 3.5:** Three transactions for Examples 3.1.5 to 3.1.7.

**Definition 3.13** (Unspent output). Let  $\mathbf{B} = (\overline{T}_1, t_1) \cdots (\overline{T}_n, t_n)$  be a blockchain. We say that the output  $j$  of transaction  $\overline{T}_i$  is unspent in  $\mathbf{B}$  whenever:

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

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

- $\text{UTXO}_{\mathbf{B}}$ , the Unspent Transaction Output of  $\mathbf{B}$ , as the set of pairs  $(\overline{T}_i, j)$  such that output  $j$  of  $\overline{T}_i$  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  $\text{UTXO}$ .

**Example 3.1.5.** Consider the transactions in Figure 3.5, and let  $\mathbf{B} = (\overline{T}_1, 0)(\overline{T}_2, t_2)$ . We have that  $(\overline{T}_1, 2, 0) \xrightarrow{5} (\overline{T}_2, 1, t_2)$  and  $(\overline{T}_1, 3, 0) \xrightarrow{7} (\overline{T}_2, 2, t_2)$ , while the other outputs are unspent. Hence, the  $\text{UTXO}$  of  $\mathbf{B}$  is  $\{(\overline{T}_1, 1), (\overline{T}_2, 1)\}$ .  $\square$

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

**Definition 3.14** (Consistent update). We write  $\mathbf{B} \triangleright (\overline{T}, t)$  iff either  $\mathbf{B} = \varepsilon$ ,  $\overline{T}$  is initial and  $t = 0$ , or, for all  $i \in \text{dom}(\overline{T}.\text{in})$ :

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

the following conditions hold:

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

Firstly, for each  $\mathbf{T.in}(i)$  we obtain the singleton  $\{\mathbf{T}'_i\}$  from the blockchain, using  $\text{match}_{\mathbf{B}}$ , such that  $\text{fst}(\mathbf{T.in}(i))\{\text{wit} \mapsto \perp\} = \mathbf{T}'_i\{\text{wit} \mapsto \perp\}$ . The update is inconsistent if  $\text{match}_{\mathbf{B}}(\text{fst}(\mathbf{T.in}(i)))$  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 3.1.6.** Consider again the transactions in Figure 3.5, 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 3.14 are satisfied. For condition (1), note that both  $(\mathbf{T}_1, 2)$  and  $(\mathbf{T}_1, 3)$  are unspent, according to Definition 3.13. 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 3.11. 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 3.1.7** (Double spending). Consider again the transactions in Figure 3.5, 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 3.14 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 3.13,  $(\mathbf{T}_1, 2)$  is already spent in  $\mathbf{B}$  because

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

holds and both  $\mathsf{T}_1$  and  $\mathsf{T}_2$  are in  $\mathbf{B}$ . Since  $\mathsf{T}_3$  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, starting from the empty one, by appending consistent updates, only. The actual definition is given by induction.

**Definition 3.15** (Consistency). *We say that a blockchain  $\mathbf{B}$  is consistent if either  $\mathbf{B} = \varepsilon$ , or  $\mathbf{B} = \mathbf{B}'(\mathsf{T}, t)$  with  $\mathbf{B}'$  consistent and  $\mathbf{B}' \triangleright (\mathsf{T}, t)$ .*

Note that the empty blockchain is consistent; the blockchain with a single transaction  $(\mathsf{T}_1, t_1)$  is consistent iff  $\mathsf{T}_1$  is initial and  $t_1 = 0$ . The transaction  $\mathsf{T}_1$  models the first transaction in the *genesis block* (as discussed in Section 7.3, we are abstracting away the *coinbase* transactions, which forge new bitcoins).

We now establish some basic properties of consistent blockchains. Theorem 3.1.1 states that, in a consistent blockchain, the inputs of a transaction point backwards to the output of some transaction in the blockchain.

**Lemma 3.1.1.** If  $(\mathsf{T}_1, t_1) \cdots (\mathsf{T}_n, t_n)$  is consistent, then:

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

*Proof.* By Definition 3.15,  $(\mathsf{T}_i, t_i)$  is a consistent update of  $(\mathsf{T}_1, t_1) \cdots (\mathsf{T}_{i-1}, t_{i-1})$ . The thesis follows from condition (2) of Definition 3.14.  $\square$

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

**Theorem 3.1.2** (No double spending). If  $(\mathsf{T}_1, t_1) \cdots (\mathsf{T}_n, t_n)$  is consistent, then:

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

*Proof.* Let  $\mathbf{B} = (\mathsf{T}_1, t_1) \cdots (\mathsf{T}_n, t_n)$  be consistent. By contradiction, assume that there exist  $i < j$  and  $i', j'$  such that  $\mathsf{T}_i.\text{in}(i') = \mathsf{T}_j.\text{in}(j')$ . By consistency, there exist  $h, h'$  such that  $(\mathsf{T}_h\{\text{wit} \mapsto \perp\}, h') = \mathsf{T}_i.\text{in}(i')$ . Since  $\mathbf{B}_{1..i-1} \triangleright (\mathsf{T}_i, t_i)$ , then by item (2) of Definition 3.14 it must be  $(\mathsf{T}_h, h', t_h) \rightsquigarrow (\mathsf{T}_i, i', t_i)$ . Hence, by Definition 3.13 it follows that  $(\mathsf{T}_h, h')$  is already *spent* in  $\mathbf{B}$ . Since  $\mathbf{B}_{1..j-1} \triangleright (\mathsf{T}_j, t_j)$ , by item (1) of Definition 3.14,  $(\mathsf{T}_h, h')$  must be *unspent* — contradiction.  $\square$

The following theorem 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 3.1.3.** If  $\mathbf{B}$  is consistent, then for all transactions  $\mathbf{T}$ ,  $\text{match}_{\mathbf{B}}(\mathbf{T})$  contains at most one element.

*Proof.* Let  $\mathbf{B} = (\mathbf{T}_1, t_1) \cdots (\mathbf{T}_n, t_n)$  be consistent. By contradiction, assume that  $\mathbf{T}_i, \mathbf{T}_j \in \text{match}_{\mathbf{B}}(\mathbf{T})$ , with  $\mathbf{T}_i \neq \mathbf{T}_j$  (and so,  $i \neq j$ ). By Definition 3.12 it must be  $\mathbf{T}_i\{\text{wit} \mapsto \perp\} = \mathbf{T}\{\text{wit} \mapsto \perp\} = \mathbf{T}_j\{\text{wit} \mapsto \perp\}$ , hence in particular  $\mathbf{T}_i.\text{in} = \mathbf{T}_j.\text{in}$ . There are two cases. If  $\mathbf{T}_i.\text{in} = \mathbf{T}_j.\text{in} = \perp$ , then by Definition 3.12  $\mathbf{B}$  is not a blockchain, since  $i \neq j$ . Hence,  $\text{ran}(\mathbf{T}_i.\text{in}) \cap \text{ran}(\mathbf{T}_j.\text{in}) = \text{ran}(\mathbf{T}_i.\text{in}) \neq \emptyset$ . By Theorem 3.1.2, this cannot happen because  $\mathbf{B}$  is consistent — contradiction.  $\square$

Theorem 3.1.4 ensures that all the transactions on a consistent blockchain are pairwise distinct, even when neglecting their witnesses.

**Lemma 3.1.4.** If  $(\mathbf{T}_1, t_1) \cdots (\mathbf{T}_n, t_n)$  is consistent, then:

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

*Proof.* Straightforward from Theorem 3.1.3, taking  $\mathbf{T} = \mathbf{T}_j$ .  $\square$

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

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

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

Let  $U_i$  and  $U_{i+1}$  be the UTXOs of  $\mathbf{B}_i$  and of  $\mathbf{B}_i(\mathbf{T}_{i+1}, t_{i+1})$ , 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  $\mathbf{T}_{i+1}$ , while the set  $U_{i+1} \setminus U$  contains exactly the outputs in



$T_{AB}$	$T_{BC}$
in: $(T_A, 1)$ wit: $\perp$ out: $(\lambda_{\zeta_A \zeta_B}. \text{versig}(k_A k_B, \zeta_A \zeta_B), 1.1\text{฿})$	in: $(T_{AB}, 1)$ wit: $\perp$ out: $1 \mapsto (\lambda_{\zeta_B}. \text{versig}(k_B, \zeta_B), 0.1\text{฿})$ $2 \mapsto (\lambda_{\zeta_C}. \text{versig}(k_C, \zeta_C), 1\text{฿})$

**Figure 3.6:** Transactions of the chain contract.

$T_{i+1}$ . Since  $\mathbf{B}$  is consistent, then  $\mathbf{B}_i \triangleright (T_{i+1}, t_{i+1})$ . Then, by Definition 3.14, for each  $k \in \text{dom } T_{i+1}.\text{in}$ , there exists a unique  $j \leq i$  such that, given  $o_k = \text{snd}(T_{i+1}.\text{in}(k))$  and  $v_k = \text{val}(T_j.\text{out}(o_k))$ :

$$(T_j, o_k, t_j) \stackrel{v_k}{\rightsquigarrow} (T_{i+1}, k, t_{i+1})$$

Then, by item (3) of Definition 3.14:

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

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

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 3.11. Of course, the choice of the scripting language affects the expressiveness of the smart contracts built upon Bitcoin.

## 3.2 Example: static chains of transactions

We now formally specify in our model a simple smart contract, which illustrates the impact of Segregated Witness on the expressiveness of Bitcoin contracts [112].

A participant  $A$  wants to send an indirect payment of  $1\text{฿}$  to  $C$ , routing it through  $B$ . To authorize the payment,  $B$  wants to keep a fee of  $0.1\text{฿}$ . 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 3.6. The transaction  $T_{AB}$  transfers  $1.1\text{฿}$  from  $A$  to  $B$  (but it is not signed by  $A$ , yet), while  $T_{BC}$  transfers  $1\text{฿}$  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 3.3, we no longer require that  $\forall i \in \text{dom in} : \text{fst}(\text{in}(i)).\text{wit} = \perp$
- in Definition 3.11, we replace item (a) with the condition:  $T'.\text{in}(j) = (T, i)$
- in Definition 3.12, 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  $T = T_{AB}\{\text{wit} \mapsto \text{sig}_{k_A}^{**}(T_{AB})\}$  is unspent on the blockchain, when participant **A** attempts to append also  $T' = T_{BC}\{\text{wit} \mapsto \text{sig}_{k_A}^{**}(T_{BC}) \text{ sig}_{k_B}^{**}(T_{BC})\}$ . To be a consistent update, by item (2) of Definition 3.14 we must have (for some  $t_1 \leq t_2$ ):

$$(T, 1, t_1) \stackrel{1_B}{\rightsquigarrow} (T', 1, t_2) \quad (3.1)$$

For this, all the conditions in Definition 3.11 must hold. However, since we have disabled SegWit, for item (a) we no longer check that:

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

but instead we need to check the condition:

$$\tilde{T}'.\text{in}(1) = (\tilde{T}, 1) \quad (3.2)$$

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 3.3.

Hence, condition (3.2) checks the equality between  $\tilde{T}_{AB}$  (the transaction in the input of  $\tilde{T}'$ ) and  $\tilde{T}_{AB}\{\text{wit} \mapsto \text{sig}_{k_A}^{**}(\tilde{T}_{AB})\}$  (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  $\tilde{T}'$ .

### 3.3 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 fields [85]. This is crucial, because non-standard transactions are mostly discarded by the Bitcoin network.

The compiler, and a in-depth documentation are available at [81]. It 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 unredeemable.

We compile the scripts of the form  $\lambda \varsigma.\text{versig}(k, \varsigma)$  to P2PKH, and those of the form  $\lambda k$  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  $x$ .

There exist other standard scripts: P2PK and MULTISIG. P2PK is considered obsolete and replaced by P2PKH, while MULTISIG has some limitations (e.g. in the number of keys) that are overcome using P2SH to express the same semantics.

**Related works** Several works have proposed to use Bitcoin beyond the sole purpose of exchanging currency, by exploiting the flexibility of its scripting language. They propose to implement smart contracts, intended as sets of protocols of the participants involved in them.

Smart contracts requiring external state, namely oracles [87] and escrows [84], can be easily implemented using multi-signature transactions. Such implementations, however, rely on trusted third parties. The work [4] showed that Bitcoin can be used to implement timed commitments through deposit transactions. The commitments are then used to perform multiparty computations [1], such as calculating the winner of a lottery. The main drawback of this approach is indeed the deposit, which grows quadratically with the number of participants. More recently, [62] and [23] have proposed lottery smart contracts that require, respectively, zero and constant ( $\geq 0$ ) deposit. However, [62] requires the computation of an exponential number of signature w.r.t. the number of participants, while [23] only a quadratic one. The work [12] proposed a contingent payment protocol that can be implemented relying only on standard Bitcoin transaction. It allows to sell solutions for a class of NP problems (e.g. the factorization of a number), the use of zero-knowledge proofs ensure its correctness to the buyer.

## Chapter 4

# Extending Bitcoin with Neighborhood Covenants

Bitcoin is a decentralised infrastructure to transfer cryptocurrency between users. The log of all the currency transactions is recorded in a public, append-only, distributed data structure, called blockchain. Bitcoin implements a model of computation called *Unspent Transaction Output (UTXO)*: each transaction holds an amount of currency, and specifies conditions under which this amount can be redeemed by a subsequent transaction, which spends the old one. Compared to the *account-based* model, implemented e.g. by Ethereum, the UTXO model does not require a shared mutable state: the current state is given just by the set of unspent transaction outputs on the blockchain. While, on the one hand, this design choice fits well with the inherent concurrency of transactions, on the other hand the lack of a shared mutable state substantially complicates leveraging Bitcoin to implement *contracts*, i.e. protocols which transfer cryptocurrency according to programmable rules.

The literature has shown that Bitcoin contracts support a surprising variety of use cases, including e.g. crowdfunding [86, 9], lotteries and other gambling games [1, 23, 9, 26, 57, 62], contingent payments [12], micro-payment channels [122, 9], and other kinds of fair computations [4, 56]. Despite this apparent richness, the fact is that Bitcoin contracts cannot express most of the use cases that are mainstream in other blockchain platforms (e.g., decentralised finance). There are several factors that limit the expressiveness of Bitcoin contracts. Among them, the crucial one is the script language used to express the redeeming conditions within transactions. This language only features a limited set of logic, arithmetic,

and cryptographic operators, but it does not have loops, and it cannot access parts of the spent and of the redeeming transaction.

Several extensions of the Bitcoin script language have been proposed, with the aim to improve the expressiveness of Bitcoin contracts, while adhering to the UTXO model. Among these extensions, *covenants* are a class of script operators that allow a transaction to constrain how its funds can be used by the redeeming transactions. Covenants may also be recursive, by requiring the script of the redeeming transaction to contain the same covenant of the spent one. As noted by [66], recursive covenants would allow to implement Bitcoin contracts that execute state machines, by appending transactions to trigger state transitions.

Although the first proposals of covenants date back at least to 2013 [114], and that they are supported by Bitcoin fork “Bitcoin Cash” [110], their inclusion into Bitcoin is still uncertain, mainly because of the extremely cautious approach to implement changes to Bitcoin [102]. Still, the emerging of Bitcoin layer-2 protocols, like e.g. the Lightning Network [122], has revived the interest in covenants, as witnessed by a recent Bitcoin Improvement Proposal (BIP 119 [123, 70]), and by the incorporation of covenants in Liquid’s extensions to Bitcoin Script [117].

We propose a variant of covenants, named *neighbourhood covenants*, which can inspect not only the redeeming transaction, but also the siblings and the parent of the spent one. This extension preserves the basic UTXO design of Bitcoin, adding only a few opcodes to its script language, which is kept efficient, loop-free, and *non* Turing-complete. Still, neighbourhood covenants significantly increase the expressiveness of Bitcoin as a smart contracts platform, allowing to execute arbitrary smart contracts by appending a *chain* of transactions to the blockchain. Technically, we prove that neighbourhood covenants make Bitcoin Turing-complete.

Although this expressiveness result is of theoretical interest, in itself it does not enable an efficient implementation of tokens. To recover efficiency, we implement our new use cases using scripts which exploit covenants.

## 4.1 Neighbourhood covenants

To extend Bitcoin with neighbourhood covenants, we amend the model of pure Bitcoin in the previous Chapter as follows:

- in transactions, we add a field to outputs, making them records of the form  $\{\mathit{arg} : \mathbf{a}, \mathit{scr} : e, \mathit{val} : v\}$ , where  $\mathbf{a}$  is a sequence of values;

- in scripts, we add operators to access all the outputs of the redeeming transaction, and a relevant subset of those of the sibling and parent transactions (by contrast, pure Bitcoin scripts can only access the redeeming transaction, and only as a whole, to verify it against a signature);
- in scripts, we add operators for covenants.

We now formalize our Bitcoin extension. We use  $\circ$  to refer to the following transaction outputs:

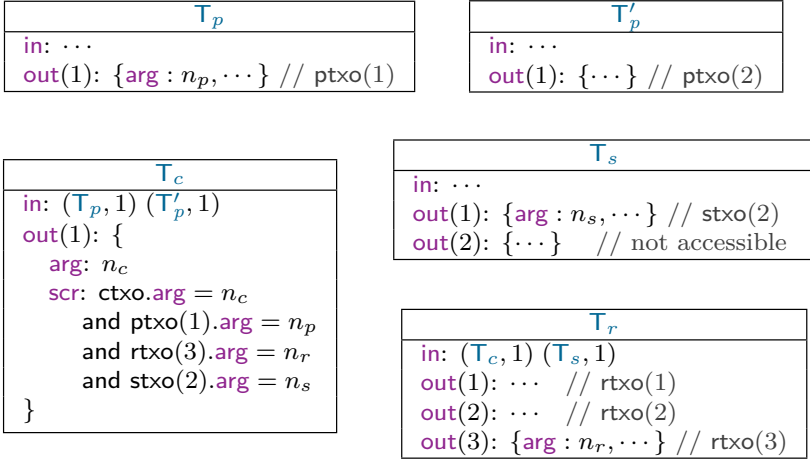
$\circ ::=$	$\text{rtxo}(e)$	output of the redeeming tx
	$  \text{stxo}(e)$	output of a sibling tx
	$  \text{ptxo}(e)$	output of a parent tx

More precisely, consider the case where a transaction output  $(T', j)$  is redeemed by a transaction  $T$ , through its  $i$ -th input. When used within the script of  $(T', j)$ :  $\text{rtxo}(n)$  refers to the  $n$ -th output of  $T$ ;  $\text{stxo}(n)$  refers to the output redeemed by the  $n$ -th input of  $T$ ;  $\text{ptxo}(n)$  refers to the output redeemed by the  $n$ -th input of  $T'$ . In Figure 4.1, we exemplify these outputs in relation to the transaction  $T_c$ . The semantics of  $\circ$  is defined in the first line of Figure 4.2; its result is a pair  $(T, n)$ .

We extend scripts as follows, where  $f \in \{\text{arg}, \text{val}\}$ :

$e ::=$	$\dots   \circ.f$	field of a tx output
	$  \text{verscr}(e, \circ)$	basic covenant
	$  \text{verrec}(\circ)$	recursive covenant
	$  \text{indx}$	index of redeeming tx input
	$  \text{outidx}$	index of redeemed tx output
	$  \text{inlen}(\circ)$	number of inputs
	$  \text{outlen}(\circ)$	number of outputs
	$  \text{txid}(\circ)$	hash of (tx,output)

The script  $\circ.f$  gives access to the field  $f$  of a transaction output  $\circ$  (where  $f$  is either  $\text{arg}$  or  $\text{val}$ ). The basic covenant  $\text{verscr}(e, \circ)$  checks that the script in the transaction output  $\circ$  is syntactically equal to  $e$  (note that  $e$  is not evaluated). The “recursive” covenant  $\text{verrec}(\circ)$  checks that the script in  $\circ$  is syntactically equal to the script which is currently being evaluated. We call it recursive because it is a covenant that enforce another one equal to itself. The operators  $\text{indx}$  and  $\text{outidx}$  evaluate, respectively, to the index of the redeeming input and redeemed output. We call *current transaction output* ( $\text{ctxo}$ ) the transaction output which includes the script which is



**Figure 4.1:** Accessing transaction outputs through a script.

currently being evaluated, i.e.:

$$\text{ctxo} \triangleq \text{stxo}(\text{indx})$$

The scripts  $\text{inlen}(\circ)$  and  $\text{outlen}(\circ)$  evaluate, respectively, to the number of inputs and to the number of outputs of the transaction containing  $\circ$ . Finally,  $\text{txid}(\circ)$  evaluates to a unique identifier of the transaction output  $\circ$ .

**Example 4.1.1.** Consider the transactions in Figure 4.1. The script in  $T_c.\text{out}(1)$  checks that (i) the **arg** field of  $\text{ctxo}$ , i.e. the same output which contains the script, equals to  $n_c$ ; (ii) the **arg** field of  $\text{ptxo}(1)$ , i.e. the parent transaction output redeemed by  $T_c.\text{in}(1)$ , equals to  $n_p$ ; (iii) the **arg** field of  $\text{rtxo}(3)$ , i.e. the third output of the redeeming transaction  $T_r$ , equals to  $n_r$ ; (iv) the **arg** field of  $\text{stxo}(2)$ , i.e. the sibling transaction output redeemed by  $T_r.\text{in}(2)$ , equals to  $n_s$ . Note that, in general, any script used in  $T_c$  can not access the parents of  $T_p$  and those of  $T_s$  — and in general all the transactions which are farther than those shown in the figure.  $\diamond$

Figure 4.2 defines the semantics of extended scripts. As in Chapter 3, the function  $\llbracket \cdot \rrbracket$  takes as parameters the redeeming transaction  $T$  and the index  $i$  of the redeeming input. We denote with  $\equiv$  syntactic equality between two scripts, i.e.  $e \equiv e'$  is 1 when  $e$  and  $e'$  are exactly the same, 0 otherwise.



$$\begin{aligned}
 \llbracket \text{rtxo}(e) \rrbracket_{\mathsf{T},i} &= (\mathsf{T}, \llbracket e \rrbracket_{\mathsf{T},i}) & \llbracket \text{stxo}(e) \rrbracket_{\mathsf{T},i} &= \mathsf{T}.\text{in}(\llbracket e \rrbracket_{\mathsf{T},i}) \\
 \llbracket \text{ptxo}(e) \rrbracket_{\mathsf{T},i} &= \mathsf{T}'.\text{in}(\llbracket e \rrbracket_{\mathsf{T},i}) \text{ if } \mathsf{T}.\text{in}(i) = (\mathsf{T}', j) & \llbracket \text{o.f} \rrbracket_{\mathsf{T},i} &= \llbracket \text{o} \rrbracket_{\mathsf{T},i}.\text{f} \\
 \llbracket \text{verscr}(e, \text{o}) \rrbracket_{\mathsf{T},i} &= e \equiv \llbracket \text{o} \rrbracket_{\mathsf{T},i}.\text{scr} & \llbracket \text{verrec}(\text{o}) \rrbracket_{\mathsf{T},i} &= \mathsf{T}.\text{in}(i).\text{scr} \equiv \llbracket \text{o} \rrbracket_{\mathsf{T},i}.\text{scr} \\
 \llbracket \text{txid}(\text{o}) \rrbracket_{\mathsf{T},i} &= H(\llbracket \text{o} \rrbracket_{\mathsf{T},i}) & \llbracket \text{outidx} \rrbracket_{\mathsf{T},i} &= j \text{ if } \mathsf{T}.\text{in}(i) = (\mathsf{T}, j) & \llbracket \text{inidx} \rrbracket_{\mathsf{T},i} &= i \\
 \llbracket \text{outlen}(\text{o}) \rrbracket_{\mathsf{T},i} &= |\mathsf{T}'.\text{out}| & \llbracket \text{inlen}(\text{o}) \rrbracket_{\mathsf{T},i} &= |\mathsf{T}'.\text{in}| \text{ if } \llbracket \text{o} \rrbracket_{\mathsf{T},i} = (\mathsf{T}', j)
 \end{aligned}$$

**Figure 4.2:** *Semantics of neighbourhood covenants (extending Figure 3.2).*

**Turing completeness** Our neighbourhood covenants make Bitcoin Turing-complete. To prove this, we describe how to simulate in the extended Bitcoin any counter machine [41], a well-known Turing-complete computational model. A *counter machine* is a pair  $(n, s)$ , where  $n \in \mathbb{N}$  is the number of integer registers of the machine, and  $s$  is a sequence of instructions. Instructions are the following: *inc*  $i$  increments register  $i$ , *dec*  $i$  decrements it, *zero*  $i$  sets it to zero, *if*  $i \neq 0$  *goto*  $j$  conditionally jumps to instruction  $j$  when register  $i$  is not zero, *halt* terminates the machine. The state of a counter machine  $(n, s)$  is a tuple  $(v_1, \dots, v_n, p)$  where each  $v_i$  represents the current value of register  $i$ , and  $p$  is the number of the next instruction to execute (i.e., the program counter). To exploit the currency transfer capabilities of Bitcoin, we slightly extend the counter machine model, by requiring that the machine has an initial  $\mathbb{B}$  balance which, upon termination, is transferred to the user **A** if the content of the first register is 0, or to **B** otherwise. This allows the machine to execute programs that transfer bitcoins. We call this extended model *UTXO-counter machine*.

**Theorem 4.1.1.** Neighbourhood covenants can simulate any UTXO-counter machine. Hence, they are Turing-complete.

*Proof. (sketch)* We represent the state of the counter machine as a single transaction having one output. We simulate an execution step by appending a new transaction, which redeems the output representing the old state, and transfers its balance to a new output representing the new state. This is done until the machine halts, at which point we transfer its balance to the user determined by the final registers state. We remark that this simulation is made possible by the use of unbounded integers in our model, while Bitcoin only supports 32-bit integers.

We represent the machine state in the **arg** field of the transaction output **o**, storing it as a sequence of integers. As a shorthand, we write


$\mathbf{o.r}_i$  for  $\mathbf{o.arg.i}$ , and  $\mathbf{o.p}$  for  $\mathbf{o.arg.(n+1)}$ . To simulate the execution steps of the machine, we define the script  $e_{CM}$ , which checks that the new transaction  $\mathbf{T}$  indeed represents the next state. More in detail, we first check whether the instruction in  $s$  at position  $\mathbf{ctxo.p}$  is `halt`, in which case we require that  $\mathbf{T}$  distributes the balance to users in the intended manner, depending on the current state. When  $\mathbf{ctxo.p}$  points to any other instruction, we start by requiring that  $\mathbf{T}$  has only one output, the same balance, and the same script. We use a recursive covenant `verrec` on the redeeming transaction to ensure the last part. Then, we check that the new state in  $\mathbf{T.out(1).arg}$  agrees with the counter machine semantics, by cases on the instruction pointed to by  $\mathbf{o.p}$ . If the instruction is `inc i`, then we require  $\mathbf{rtxo(1).r}_i = \mathbf{ctxo.r}_i + 1$ ,  $\mathbf{rtxo(1).r}_k = \mathbf{ctxo.r}_k$  for all  $k \neq i$ , and  $\mathbf{rtxo(1).p} = \mathbf{ctxo.p} + 1$ . The `dec i` and `zero i` cases are analogous. For the instruction `if i  $\neq$  0 goto j`, we check the value of  $\mathbf{ctxo.r}_i$ : if nonzero, we require that  $\mathbf{rtxo(1).p} = j$ , otherwise that  $\mathbf{rtxo(1).p} = \mathbf{ctxo.p} + 1$ . In both cases, we require that  $\mathbf{rtxo(1).r}_k = \mathbf{ctxo.r}_k$  for all  $k$ .

Users start the simulation by appending to the blockchain a transaction  $\mathbf{T}_0$  having one output with the desired value, the script  $e_{CM}$ , and a sequence of  $n + 1$  zeros as `arg`. After that, the balance is effectively locked inside the transaction, and the only way to transfer it back to the users is to simulate all the steps of the machine, until it halts. So, our simulated execution is a form of *secure multiparty computation* [5, 75, 46].  $\square$

Although, for simplicity, we use UTXO-counter machines just to transfer funds to  $\mathbf{A}$  or  $\mathbf{B}$  upon termination, it would be easy to generalise the computational model and simulation technique to encompass *interactive* computations, which at run-time can receive inputs and perform currency transfers. Doing so, we can execute arbitrary smart contracts, with the same expressiveness of Turing-complete smart contracts platforms. In principle, we could craft a smart contract which implements user-defined tokens in an account-based fashion: this contract would record the balance of the tokens of all users, and execute token actions. However, in practice this construction would be highly inefficient: performing a single token transfer would require to append a large number of transactions, and to pay the related fees.

## 4.2 Use cases

We illustrate the expressive power of our extension through a series of use cases, which, at the best of our knowledge, cannot be expressed in Bitcoin. We denote with  $\mathbf{U}_A^{v\mathbb{B}}$  an unspent transaction output  $\{\mathbf{arg} : \varepsilon, \mathbf{scr} :$



**Figure 4.3:** Transactions for the crowdfunding contract.

$\text{versig}(A, \text{rtx}.\text{wit}), \text{val} : v \mathbb{B}$ }, where  $\varepsilon$  denotes the empty sequence (we will usually omit  $\text{arg}$  when empty).

### 4.2.1 Crowdfunding

Assume that a start-up  $Z$  wants to raise funds through a crowdfunding campaign. The target of the campaign is to gather at least  $v \mathbb{B}$  by time  $t$ . The contributors want the guarantee that if this target is not reached, then they will get back their funds after the expiration date. The start-up wants to ensure that contributions cannot be retracted before time  $t$ , or once  $v \mathbb{B}$  have been gathered.

We implement this use case without covenants, but just constraining the  $\text{val}$  field of the redeeming transaction. To fund the campaign, a contributor  $A_i$  publishes the transaction  $T_i$  in Figure 4.3 (left), which uses the following script:

$$CF = (\text{versig}(Z, \text{rtxo}(1).\text{wit}) \text{ and } \text{rtxo}(1).\text{val} \geq v) \text{ or } \text{absAfter } t : \text{versig}(A_i, \text{rtxo}(1).\text{wit})$$

This script is a disjunction between two conditions. The first condition allows  $Z$  to redeem the bitcoins deposited in this output, provided that the output at index 1 of the redeeming transaction pays at least  $v \mathbb{B}$  (note that this constraint, rendered as  $\text{rtxo}(1).\text{val} \geq v$ , is not expressible in pure Bitcoin). The second condition allows  $A_i$  to get back her contribution after the expiration date  $t$ .

Once contributors have deposited enough funds (i.e., there are  $n$  transactions  $T_1, \dots, T_n$  with  $v' = v_1 + \dots + v_n \geq v$ ),  $Z$  can get  $v' \mathbb{B}$  by appending  $T_Z$  to the blockchain. Note that, compared to the assurance contract in the Bitcoin wiki [86], ours offers more protection to the start-up. Indeed, while in [86] any contributor can retract her funds at any time, this is not possible here until time  $t$ .

$T_0$	$T_1$
in: $U_A^{1\text{B}}$ wit: $\text{sig}_A(T_0)$ out: $\{\text{arg} : A, \text{scr} : \text{NFT}, \text{val} : 1\text{B}\}$	in: $(T_0, 1)$ wit: $\text{sig}_A(T_1)$ out: $\{\text{arg} : B, \text{scr} : \text{NFT}, \text{val} : 1\text{B}\}$

**Figure 4.4:** *A creates a token with  $T_0$ , and transfers it to B with  $T_1$ .*

$T_2$
in: $(T_A, 1) (T'_A, 1)$ wit: $\text{sig}_A(T_2) \text{sig}_A(T_2)$ out(1): $\{\text{arg} : A, \text{scr} : \text{NFT}, \text{val} : 1\text{B}\}$ out(2): $\{\text{arg} : \varepsilon, \text{scr} : \text{versig}(A, \text{rtxo}(1).\text{wit}), \text{val} : 1\text{B}\}$

**Figure 4.5:** *A exploits the flaw to destroy a token, redeeming its value.*

## 4.2.2 Non-fungible tokens

A non-fungible token represents the ownership of a physical or logical asset, which can be transferred between users. Unlike fungible tokens (e.g., ERC-20 tokens in Ethereum [106]), where each token unit is interchangeable with every other unit, non-fungible ones have unique identities. Further, they do not support split and join operations, unlike fungible tokens.

We start by implementing a subtly flawed version of the non-fungible token. Consider the transactions in Figure 4.4, which use the following script:

$$\text{NFT} = \text{versig}(\text{ctxo.arg}, \text{rtxo}(1).\text{wit}) \text{ and } \text{verrec}(1) \text{ and } \text{rtxo}(1).\text{val} = 1$$

User **A** mints a token by depositing  $1\text{B}$  in  $T_0$ : to declare her ownership over the token, she sets  $\text{out}(1).\text{arg}$  to her public key. To transfer the token to **B**, **A** appends the transaction  $T_1$ , setting its  $\text{out}(1).\text{arg}$  to **B**'s public key.

To spend  $T_0$ , the transaction  $T_1$  must satisfy the conditions specified by the script  $\text{NFT}$ : (i) the  $\text{wit}$  field must contain the signature of the current owner; (ii) the script at index 1 must be equal to that at the same index in  $T_0$ ; (iii) the output at index 1 must have  $1\text{B}$  value, to preserve the integrity of the token. Once  $T_1$  is on the blockchain, **B** can transfer the token to another user, by appending a transaction which redeems  $T_1$ .

The script  $\text{NFT}$  has a design flaw, already spotted in [63]: we show how **A** can exploit this flaw in Figure 4.5. Suppose we have two unspent transactions:  $T_A$  and  $T'_A$ , both representing a token owned by **A** (in their first and only output). The transaction  $T_2$  can spend both of them, since it complies with all the validity conditions: indeed,  $\text{NFT}$  only constrains the script in the first output of the redeeming transaction, while the other

outputs are only subject to the standard validity conditions (in particular, that the sum of their values does not exceed the value in input). Actually,  $T_2$  destroys one of the two tokens, and removes the covenant from the other one.

To solve this issue, we can amend the *NFT* script as follows:

$$\begin{aligned} \mathit{NFT}' &= \text{versig}(\text{ctxo.argoutidx}, \text{rtxo}(1).\text{wit}) \text{ and} \\ &\text{verrec}(\text{inidx}) \text{ and } \text{rtxo}(\text{inidx}).\text{val} = 1 \end{aligned}$$

The amended script correctly handles the case of a transaction which uses different outputs to store different tokens. *NFT'* uses `ctxo.argoutidx`, instead of `ctxo.arg1` in *NFT*, to ensure that, when redeeming a given output, the signature of the owner of the token at *that* output is checked. Further, *NFT'* uses `verrec(inidx)`, instead of `verrec(1)` in *NFT*, to ensure that the covenant is propagated exactly to the transaction output which is redeeming that token (i.e., the one at index `inidx`). Notice that the amendment would make  $T_2$  invalid: indeed, the script in  $T'_A.\text{out}(1)$  would evaluate to false:

$$\begin{aligned} \llbracket \mathit{NFT}' \rrbracket_{T_2,2} &= \llbracket \text{verrec}(\text{inidx}) \rrbracket_{T_2,2} \wedge \dots \\ &= (T_2, \llbracket \text{inidx} \rrbracket_{T_2,2}).\text{scr} \equiv T_2.\text{in}(2).\text{scr} \wedge \dots \\ &= (T_2, 2).\text{scr} \equiv (T'_A, 1).\text{scr} \wedge \dots \\ &= \text{versig}(A, \text{rtxo}(1).\text{wit}) \equiv \mathit{NFT}' \wedge \dots \\ &= \text{false} \end{aligned}$$

An alternative patch, originally proposed in [63], is to add a unique identifier *id* to each token, e.g. by amending the *NFT* script as follows:

$$\mathit{NFT} \text{ and } id = id$$

This allows to mint distinguishable tokens. For instance, if the tokens in  $T_A$  and  $T'_A$  are distinguishable,  $T_2$  cannot redeem both of them.

### 4.2.3 Vaults

Transaction outputs are usually secured by cryptographic keys (e.g. through the script `versig(pkA, rtx.wit)`). Whoever knows the corresponding private key (e.g., *sk<sub>A</sub>*) can redeem such an output: in case of key theft, the legitimate owner is left without defence. Vault transactions, introduced in [63], are a technique to mitigate this issue, by allowing the legitimate owner to abort the transfer.

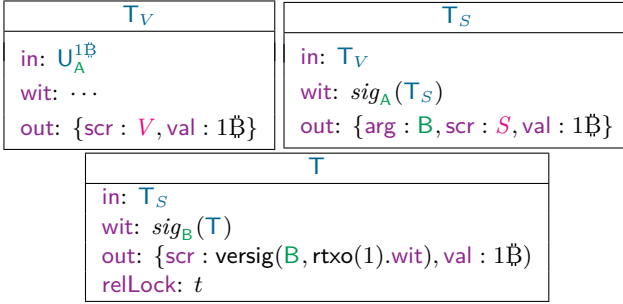


Figure 4.6: Transactions for the basic vault.

To create a vault,  $A$  deposits  $1\text{B}$  in a transaction  $T_V$  with the script  $V$ :

$V = \text{versig}(A, \text{rtxo}(1).\text{wit})$  and  $\text{verscr}(1, S)$

$S = (\text{relAfter } t : \text{versig}(\text{ctxo.arg}, \text{rtxo}(1).\text{wit}))$  or  $\text{versig}(Ar, \text{rtxo}(1).\text{wit})$

The transaction  $T_V$  can be redeemed with the signature of  $A$ , but only by a *de-vaulting* transaction like  $T_S$  in Figure 4.6, which uses the script  $S$ . The output of the de-vaulting transaction  $T_S$  can be spent by the user set in its  $\text{arg}$  field, but only after a certain time  $t$  (e.g., by the transaction  $T$  in Figure 4.6). Before time  $t$ ,  $A$  can cancel the transfer by spending  $T_S$  with her recovery key  $Ar$ .

**A recursive vault** The vault in Figure 4.6 has a potential issue, in that the recovery key may also be subject to theft. Although this issue is mitigated by hardware wallets (and by the infrequent need to interact with the recovery key), the vault modelled above does not discourage any attempt at stealing the key.


The issue can be solved by using a recursive covenant in the vault script  $R$ :

```

if ctxo.arg.1 = 0 // current state: vault
  then versig(A, rtxo(1).wit) and verrec(1) and
  rtxo(1).arg.1 = 1 // next state: de-vaulting
else // current state: de-vaulting
  (relAfter t : versig(ctxo.arg.2, rtxo(1).wit)) or // current state: de-vaulting
  versig(Ar, rtxo(1).wit) and verrec(1) and // next state: vault
  rtxo(1).arg.1 = 0

```

In this version of the contract, the vault and de-vaulting transactions (in Figure 4.7) have the same script. The first element of the  $\text{arg}$  sequence



**Figure 4.7:** Transactions for the recursive vault.

encodes the contract state (0 models the vault state, and 1 the de-vaulting state), while the second element is the user who can receive the bitcoin deposited in the vault. The recovery key  $Ar$  can only be used to append the re-vaulting transaction  $T_R$ , locking again the bitcoin into the vault.

Note that key theft becomes ineffective: indeed, even if both keys are stolen, the thief cannot take control of the bitcoin in the vault, as  $A$  can keep re-vaulting.


#### 4.2.4 A pyramid scheme

Ponzi schemes are financial frauds which lure users under the promise of high profits, but which actually repay them only with the investments of new users. A pyramid scheme is a Ponzi scheme where the scheme creator recruits other investors, who in turn recruit other ones, and so on. Unlike in Ethereum, where several Ponzi schemes have been implemented as smart contracts [14, 34], the limited expressive power of Bitcoin contract only allows for off-chain schemes [74].

We design the first “smart” pyramid scheme in Bitcoin using the transactions in Figure 4.8, where:

$$\begin{aligned}
 P &= \text{verscr}(1, X) \text{ and } \text{rtxo}(1).\text{arg} = \text{ctxo}.\text{arg} \text{ and } \text{rtxo}(1).\text{val} = 2 \\
 &\quad \text{and } \text{verrec}(2) \text{ and } \text{verrec}(3) \\
 X &= \text{versig}(\text{ctxo}.\text{arg}, \text{rtxo}(1).\text{wit})
 \end{aligned}$$

To start the scheme, a user  $A_0$  deposits  $1\mathbb{B}$  in the transaction  $T_0$  (we burn this bitcoin for uniformity, so that each user earns at most  $1\mathbb{B}$  from the scheme). To make a profit,  $A_0$  must convince other two users, say  $A_1$  and  $A_2$ , to join the scheme. This requires the cooperation of  $A_1$  and  $A_2$  to publish a transaction which redeems  $T_0$ . The script  $P$  ensures that this redeeming transaction has the form of  $T_1$  in Figure 4.8, i.e.  $\text{out}(1)$



**Figure 4.8:** Transactions for the pyramid scheme.

transfers  $2\text{฿}$  to  $A_0$ , while the scripts in **out(2)** and **out(3)** ensure that the same behaviour is recursively applied to  $A_1$  and  $A_2$ .

Overall, the contract ensures that, as long as new users join the scheme, each one earns  $1\text{฿}$ . Of course, as in any Ponzi scheme, at a certain point it will no longer be possible to find new users, so those at the leaves of the transaction tree will just lose their investment.

## 4.2.5 King of the Ether Throne

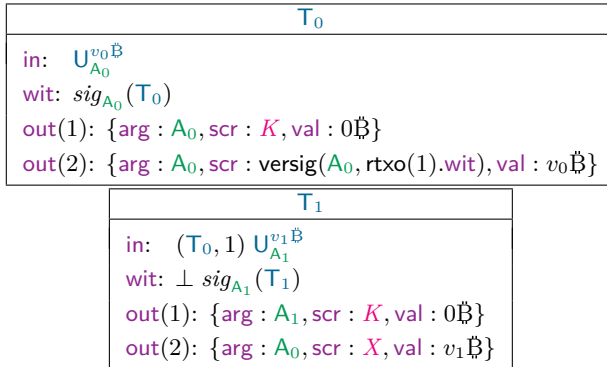
*King of the Ether Throne* [111] is an Ethereum contract, which has been popular for a while around 2016, until a bug caused its funds to be frozen. The contract is initiated by a user, who pays an entry fee  $v_0$  to become the “king”. Another user can usurp the throne by paying  $v_1 = 1.5v_0$  fee to the old king, and so on until new usurpers are available. Of course this leads to an exponential growth of the fee needed to become king, so subsequent versions of the contract introduced mechanisms to make the current king die if not ousted within a certain time. Although the logic to distribute money substantially differs from that in Section 4.2.4, this is still an instance of Ponzi scheme, since investors are only paid with the funds paid by later investors.

We implement the original version of the contract, fixing the multiplier to 2 instead of 1.5, since Bitcoin scripts do not support multiplication. The contract uses the transactions in Figure 4.9 for the first two kings,  $A_0$  and  $A_1$ , where:

$$\begin{aligned}
K &= \text{verrec}(1) \text{ and } \text{rtxo}(2).\text{arg} = \text{ctxo}.\text{arg} \text{ and} \\
&\quad \text{rtxo}(2).\text{val} \geq \text{ctxo}.\text{val} + \text{ctxo}.\text{val} \text{ and } \text{verscr}(2, X) \\
X &= \text{versig}(\text{ctxo}.\text{arg}, \text{rtxo}(1).\text{wit})
\end{aligned}$$

We use the **arg** field in **out(1)** to record the new king, and that in **out(2)** for the old one. The clause  $\text{rtxo}(2).\text{arg} = \text{ctxo}.\text{arg}$  in  $K$  preserves the old king in the redeeming transaction. The clause  $\text{rtxo}(2).\text{val} \geq$





**Figure 4.9:** Transactions for *King of the Ether Throne*.

$\text{ctxo.val} + \text{ctxo.val}$  ensures that his compensation is twice the value he paid. Finally,  $\text{verscr}$  guarantees that the old king can redeem his compensation via  $\text{out}(2)$ .

### 4.3 Using covenants in high-level contract languages

As witnessed by the use cases in Section 4.2, crafting a contract at the level of Bitcoin transactions can be complex and error-prone. To simplify this task, the work described in [22] has introduced a high-level contract language, called BitML, with a secure compiler to pure Bitcoin transactions. BitML has primitives to **withdraw** funds from a contract, to **split** a contract (and its funds) into subcontracts, to request the authorization from a participant  $A$  before proceeding with a subcontract  $C$  (written  $A : C$ ), to postpone the execution of  $C$  after a given time  $t$  (written **after**  $t : C$ ), to **reveal** committed secrets, and to branch between two contracts (written  $C + C'$ ). A recent paper [20] extends BitML with a new primitive that allows participants to (consensually) renegotiate a contract, still keeping the ability to compile to pure Bitcoin.

Despite the variety of use cases shown in [10, 15], BitML has known expressiveness limits, given by the requirement to have pure Bitcoin as its compilation target. For instance, BitML cannot specify recursive contracts (just as pure Bitcoin cannot), unless all participants agree to perform the recursive call [20]. In this section we discuss how to improve the expressiveness of BitML, assuming to use Bitcoin with covenants as compilation target. We illustrate our point by a couple of examples, post-

poning the formal treatment of this extended BitML and of its secure compilation to future work.

Covenants allow us to extend BitML with the construct:

$$?x \text{ if } b. X(x)$$

Intuitively, the prefix  $?x \text{ if } b$  can be fired whenever a participant provides a sequence of arguments  $x$  and makes the predicate  $b$  true. Once the prefix is fired, the contract proceeds as the continuation  $X(x)$ , which will reduce according to the equation defining  $X$ .

Using this construct, we can model the “King of the Ether Throne” contract of Section 4.2.5 (started by  $A$  with an investment of  $1\text{฿}$ ) as  $X(A, 1)$ , where:

$$\begin{aligned} X(a, v) &= ?b \text{ if } \text{val} \geq 2v. Y(a, b, \text{val}) \\ Y(a, b, v) &= \text{split } (0 \rightarrow X(b, v) \mid v \rightarrow \text{withdraw } a) \end{aligned}$$

The contract  $X(a, v)$  models a state where  $a$  is the current king, and  $v$  is his investment. The guard  $\text{val} \geq 2v$  becomes true when some participant injects funds into the contract, making its value ( $\text{val}$ ) greater than  $2v$ . This participant can choose the value for  $b$ , i.e. the new king. The contract proceeds as  $Y(a, b, \text{val})$ , which has two parallel branches. The first branch makes  $\text{val}\text{฿}$  available to the old king; the second branch has zero value, and it reboots the game, recording the new king  $b$  and his investment.

A possible computation of  $A$  starting the scheme with  $1\text{฿}$  is the following, where we represent a contract  $C$  storing  $v\text{฿}$  as a term  $\langle C, v\text{฿} \rangle$ :

$$\begin{aligned} &\langle X(A, -), 1\text{฿} \rangle \\ &\rightarrow \langle Y(A, B, 2), 2\text{฿} \rangle && (B \text{ pays } 2\text{฿} \text{ fee}) \\ &\rightarrow \langle X(B, 2), 0\text{฿} \rangle \mid \langle \text{withdraw } A, 2\text{฿} \rangle && (\text{contract splits}) \\ &\rightarrow \langle X(B, 2), 0\text{฿} \rangle \mid \langle A, 2\text{฿} \rangle && (A \text{ redeems } 2\text{฿}) \\ &\rightarrow \langle Y(B, C, 4), 4\text{฿} \rangle \mid \langle A, 2\text{฿} \rangle && (C \text{ pays } 4\text{฿} \text{ fee}) \\ &\rightarrow \langle X(C, 4), 0\text{฿} \rangle \mid \langle \text{withdraw } B, 4\text{฿} \rangle \mid \langle A, 2\text{฿} \rangle && (\text{contract splits}) \\ &\rightarrow \langle X(C, 4), 0\text{฿} \rangle \mid \langle B, 4\text{฿} \rangle \mid \langle A, 2\text{฿} \rangle && (B \text{ redeems } 4\text{฿}) \end{aligned}$$

Executing a step of the BitML contract corresponds, in Bitcoin, to appending a transaction containing in `out(1)` the script in Figure 4.10. The script implements a state machine, using `arg.1` to record the current state, and the other parts of `arg` for the old king, the new king, and  $v$ . The `verrec(1)` at line 8 preserves the script in `out(1)`. To pay the old king, we use the `verscr` at line 20, which constrains the script in `out(2)` of the transaction corresponding to the BitML state  $\langle X(b, v), 0\text{฿} \rangle \mid \langle \text{withdraw } a, v\text{฿} \rangle$ .

```

1 def arg.1 = q      // state    0 = <X(A,-),1>
2                   // state    1 = <Y(a,b,v),v>
3                   // state    2 = <X(b,v),0> | <withdraw a,v>
4 def arg.2 = oldK  // old King
5 def arg.3 = newK  // new king
6 def arg.4 = v     // paid fee
7
8 verrec(1) and     // out(1) preserves covenant
9 if ctxo(1).q = 0 then // state 0
10    rtxo(1).q = 1 // state transition 0 -> 1
11    and rtxo(1).oldK = ctxo(1).newK // usurp the throne
12    and rtxo(1).val >= ctxo(1).val
13        + ctxo(1).val // fee at least doubled
14    and rtxo(1).v = rtxo(1).val // instantiate v
15 else if ctxo(1).q = 1 then // state 1
16    rtxo(1).q = 2 // state transition 1 -> 2
17    and rtxo(1).newK = ctxo(1).newK // preserve new king
18    and rtxo(1).v = ctxo(1).v // preserve v
19    and rtxo(1).val = 0 // reset value in out(1)
20
21    and rtxo(2).oldK = ctxo(1).oldK // set old king
22    and verscr(2, // covenant to pay old king
23        versig(ctxo(2).oldK, rtx.wit))
24 else if ctxo(1).q = 2 then // state 2
25    rtxo(1).q = 1 // state transition 2 -> 1
26    and rtxo(1).oldK = ctxo(1).newK // usurp the throne
27    and rtxo(1).val >= ctxo(1).v
28        + ctxo(1).v // fee at least doubled
29    and rtxo(1).v = rtxo(1).val // update v
    
```

**Figure 4.10:** *Script for King of the Ether Throne, obtained by compiling BitML.*

We now apply our extended BitML to specify a more challenging use case, i.e. a recursive coin-flipping game where two players **A** and **B** repeatedly flip coins, and the one who wins two consecutive flips takes the pot. The precondition to stipulate the contract requires each player to deposit  $1\text{\$}$  as a bet. The game first makes each player commit to a secret, using a timed-commitment protocol [29]. The secrets are then revealed, and the winner of a flip is determined as a function of the two secrets. The game starts another flip if the current winner is different from that of the previous flip, otherwise the pot is transferred to the winner.

We model the recursive coin-flipping game as the (extended) BitML

contract  $X_A \langle C \rangle$ , where  $C \neq A, B$ , using the following defining equations:

$$\begin{aligned}
 X_A \langle w \rangle &= A : ?h_A . X_B \langle w, h_A \rangle + \text{afterRel } t : \text{withdraw } B \\
 X_B \langle w, h_A \rangle &= B : ?h_B . Y_A \langle w, h_A, h_B \rangle + \text{afterRel } t : \text{withdraw } A \\
 Y_A \langle w, h_A, h_B \rangle &= ?s_A \text{ if } H(s_A) = h_A . Y_B \langle w, s_A, h_B \rangle \\
 &\quad + \text{afterRel } t : \text{withdraw } B \\
 Y_B \langle w, s_A, h_B \rangle &= ?s_B \text{ if } H(s_B) = h_B \text{ and } 0 \leq s_B \leq 1 . W \langle w, s_A, s_B \rangle \\
 &\quad + \text{afterRel } t : \text{withdraw } A \\
 W \langle w, s_A, s_B \rangle &= \text{if } s_A = s_B \text{ and } w = A : \text{withdraw } A \quad // \text{ } A \text{ won twice} \\
 &\quad + \text{if } s_A = s_B \text{ and } w \neq A : X_A \langle A \rangle \quad // \text{ } A \text{ won last flip} \\
 &\quad + \text{if } s_A \neq s_B \text{ and } w = B : \text{withdraw } B \quad // \text{ } B \text{ won twice} \\
 &\quad + \text{if } s_A \neq s_B \text{ and } w \neq B : X_A \langle B \rangle \quad // \text{ } B \text{ won last flip}
 \end{aligned}$$

The contract  $X_A \langle w \rangle$  models a state where  $w$  is the last winner, and  $A$  must commit to her secret. To do that,  $A$  must authorize an input  $h_A$ , which represents the hash of her secret. If  $A$  does not commit within  $t$ , then the pot can be redeemed by  $B$  as a compensation (here, the primitive  $\text{afterRel } t : C$  models a relative timeout). Similarly,  $X_B \langle w \rangle$  models  $B$ 's turn to commit. In  $Y_A \langle w, h_A, h_B \rangle$ ,  $A$  must reveal her secret  $s_A$ , or otherwise lose her deposit. The contract  $Y_B \langle w, s_A, h_B \rangle$  is the same for  $B$ , except that here we additionally check that  $B$ 's secret is either 0 or 1 (this is needed to ensure fairness, as in the two-player lottery in [22]). The flip winner is  $A$  if the secrets of  $A$  and  $B$  are equal, otherwise it is  $B$ . If the winner is the same as the previous round, the winner can withdraw the pot, otherwise the game restarts, recording the last winner.

This coin flipping game is fair, i.e. the expected payoff of a *rational* player is always non-negative, notwithstanding the behaviour of the other player.

## 4.4 Implementing neighbourhood covenants

To extend Bitcoin with neighborhood covenants, only small changes to the script language are needed. Currently, scripts can only access the redeeming transaction, and only for signature verification. To enable covenants, scripts need to access the fields of the redeeming transaction, and those of the parent and the *sibling* transaction outputs. First, the UTXO data structure must be extended to record the parents of unspent outputs. Full nodes could simply access these transactions from their identifiers. Lightweight nodes, i.e. Bitcoin clients with limited resources that store the UTXO instead of the whole blockchain, must also store parents beside the

UTXO; when all the children of a transaction are spent, the parent can be deleted.

To implement the covenants `verrec` and `verscr`, the Bitcoin script language must be extended with new opcodes. A former covenant-enabling opcode is the `CheckOutputVerify` of [63], which uses placeholders to represent variable parts of the script (e.g., `versig(<pubKey>, rtx.wit)`). However, its implementation requires string substitutions at run-time to insert the wanted values in the script before checking script equality. Instead, our neighbourhood covenants can be implemented more efficiently. In our covenants, we can use exactly the same script along a chain of transactions, relying on the `arg` sequence to record state updates. The script can access the state through the operator `o.arg`, which require suitable opcodes. Since the script is fixed, nodes do not have to perform string substitutions, and checking script equality can be efficiently performed by comparing their hashes.

Adding the `arg` field to outputs does not require to alter the structure of pure Bitcoin transactions. Indeed, the `arg` values can be stored at the beginning of the script as push operations on the alternative stack, and copied to the main stack when the script refers to them, using the same technique used in [81]. When hashing the scripts for comparison, we discard these `arg` values: this just requires to skip the prefix of the script comprising all the push to the alternative stack.



## Chapter 5

# Bitcoin Smart Contracts as Endpoint Protocols

Albeit the primary usage of Bitcoin is to exchange currency, its blockchain and consensus mechanism can also be exploited to securely execute some forms of *smart contracts*. These are agreements among mutually distrusting parties, which can be automatically enforced without resorting to a trusted intermediary. Over the last few years a variety of smart contracts for Bitcoin have been proposed, both by the academic community and by that of developers. However, the heterogeneity in their treatment, the informal (often incomplete or imprecise) descriptions, and the use of poorly documented Bitcoin features, pose obstacles to the research. In this chapter we present a comprehensive survey of smart contracts on Bitcoin, in a uniform framework. Our treatment is based on a new formal specification language for smart contracts, which also helps us to highlight some subtleties in existing informal descriptions, making a step towards automatic verification. We discuss some obstacles to the diffusion of smart contracts on Bitcoin, and we identify the most promising open research challenges.

### 5.1 Modelling Bitcoin contracts

In this section we introduce a formal model of the behaviour of the participants in a contract, building upon the model of Bitcoin transactions in Chapter 3.

We start by formalising a simple language of expressions, which represent both the messages sent over the network, and the values used in internal computations made by the participants. Hereafter, we assume a

$$\begin{aligned}
\llbracket \nu \rrbracket &= \nu & \llbracket \text{sig}_k^{\mu,i}(T) \rrbracket &= \text{sig}_k^{\mu,i}(\llbracket T \rrbracket) \\
\llbracket \text{versig}_k(\mathbf{E}, T, i) \rrbracket &= \text{ver}_k(\llbracket \mathbf{E} \rrbracket, \llbracket T \rrbracket, i) \\
\llbracket T\{\mathbf{f}(i) \mapsto \mathbf{E}\} \rrbracket &= \llbracket T \rrbracket\{\mathbf{f}(i) \mapsto \llbracket \mathbf{E} \rrbracket\} & \llbracket (E, E') \rrbracket &= (\llbracket E \rrbracket, \llbracket E' \rrbracket) \\
\llbracket E \circ E' \rrbracket &= \llbracket E \rrbracket \circ \llbracket E' \rrbracket \quad \text{for } \circ \in \{ \text{and}, \text{or}, +, \dots \} & \llbracket \text{not } E \rrbracket &= \neg \llbracket E \rrbracket \\
\llbracket \mathbf{E} \rrbracket &= \llbracket E_1 \rrbracket \cdots \llbracket E_n \rrbracket \quad \text{if } \mathbf{E} = E_1 \cdots E_n
\end{aligned}$$

**Figure 5.1:** Semantics of contract expressions.

set  $\text{Var}$  of variables, and we define the set  $\text{Val}$  of values comprising constants  $k \in \mathbb{Z}$ , signatures  $\sigma$ , scripts  $\lambda z.e$ , transactions  $\mathbf{T}$ , and currency values  $v$ .

**Definition 5.1** (Contract expressions). *We define contract expressions through the following syntax:*

$E, T ::= \nu$	value ( $\nu \in \text{Val}$ )
$x$	variable ( $x \in \text{Var}$ )
$\text{sig}_k^{\mu,i}(T)$	signature ( $\mu$ signature modifier)
$\text{versig}_k(\mathbf{E}, T, i)$	(multi) signature verification
$T\{\mathbf{f}(i) \mapsto \mathbf{E}\}$	transaction field update
$(E, E)$	pair
$E \text{ and } E \mid E \text{ or } E \mid \text{not } E$	logical expressions
$E + E \mid \dots$	arithmetic expressions

where  $\mathbf{E}$  denotes a finite sequence of expressions (i.e.,  $\mathbf{E} = E_1 \cdots E_n$ ). We define the function  $\llbracket \cdot \rrbracket$  from (variable-free) contract expressions to values in Figure 5.1. As a notational shorthand, we omit the index  $i$  in  $\text{sig}$  (resp.  $\text{versig}$ ) when the signed (resp. verified) transactions have a single input.

Intuitively, when  $T$  evaluates to a transaction  $\mathbf{T}$ , the expression  $T\{\mathbf{f}(i) \mapsto \mathbf{E}\}$  represents the transaction obtained from  $\mathbf{T}$  by substituting the field  $\mathbf{f}(i)$  with the sequence of values obtained by evaluating  $\mathbf{E}$ . For instance,  $\mathbf{T}\{\text{wit}(1) \mapsto \sigma\}$  denotes the transaction obtained from  $\mathbf{T}$  by replacing the witness at index 1 with the signature  $\sigma$ . Further,  $\text{sig}_k^{\mu,i}(T)$  evaluates to the signature of the transaction represented by  $T$ , and  $\text{versig}_k(\mathbf{E}, T)$  represents the  $m$ -of- $n$  multi-signature verification of the



transaction represented by  $T$ . Both for the signing and verification, the parameter  $i$  represents the index where the signature will be used. We assume a simple type system that rules out ill-formed expressions, like e.g.  $k\{\text{wit}(1) \mapsto \mathbf{T}\}$ .

We formalise the behaviour of a participant as an *endpoint protocol*, i.e. a process where the participant can perform the following actions: (i) send/receive messages to/from other participants; (ii) put a transaction on the ledger; (iii) wait until some transactions appear on the blockchain; (iv) do some internal computation. Note that the last kind of operation allows a participant to craft a transaction before putting it on the blockchain, e.g. setting the `wit` field to her signature, and later on adding the signature received from another participant.

**Definition 5.2** (Endpoint protocols). *Assume a set of participants (named  $A, B, C, \dots$ ). We define prefixes  $\pi$ , and protocols  $P, Q, R, \dots$  as follows:*

$\pi ::= A!E$	<i>send messages to <math>A</math></i>
$A?x$	<i>receive messages from <math>A</math></i>
$\text{put } T$	<i>append transaction <math>T</math> to the blockchain</i>
$\text{ask } \mathbf{T} \text{ as } x$	<i>wait until all transactions in <math>\mathbf{T}</math> are on the blockchain</i>
$\text{check } E$	<i>test condition</i>
$P ::= \sum_{i \in I} \pi_i . P_i$	<i>guarded choice (<math>I</math> finite set)</i>
$P   P$	<i>parallel composition</i>
$X(E)$	<i>named process</i>

We assume that each name  $X$  has a unique defining equation  $X(x) = P$  where the free variables in  $P$  are included in  $x$ . We use the following syntactic sugar:

- $\tau \triangleq \text{check true}$ , the internal action;
- $\mathbf{0} \triangleq \sum_{\emptyset} P$ , the terminated protocol (as usual, we omit trailing  $\mathbf{0}$ s);
- $\text{if } E \text{ then } P \text{ else } Q \triangleq \text{check } E . P + \text{check not } E . Q$ ;
- $\pi_1.Q_1 + P \triangleq \sum_{i \in I \cup \{1\}} \pi_i.Q_i$ , provided that  $P = \sum_{i \in I} \pi_i.Q_i$  and  $1 \notin I$ ;
- $\text{let } x = E \text{ in } P \triangleq P\{E/x\}$ , i.e.  $P$  where  $x$  is replaced by  $E$ .

The behaviour of protocols is defined in terms of a LTS between *sys-*

$$\begin{array}{c}
\mathbf{A}[\mathbf{B}! \mathbf{E}. P + R \mid Q] \mid \mathbf{B}[\mathbf{A} ? \mathbf{x}. P' + R' \mid Q'] \mid S \rightarrow \mathbf{A}[P \mid Q] \mid \mathbf{B}[P'\{\llbracket \mathbf{E} \rrbracket / \mathbf{x}\} \mid Q'] \mid S \quad [\text{COM}] \\
\\
\frac{\llbracket \mathbf{E} \rrbracket = \text{true}}{\mathbf{A}[\text{check } \mathbf{E}. P + R \mid Q] \mid S \rightarrow \mathbf{A}[P \mid Q] \mid S} \quad [\text{CHECK}] \\
\\
\frac{\llbracket \mathbf{T} \rrbracket = \mathbf{T} \quad \mathbf{B} \triangleright (\mathbf{T}, t)}{\mathbf{A}[\text{put } \mathbf{T}. P + R \mid Q] \mid S \mid (\mathbf{B}, t) \rightarrow \mathbf{A}[P \mid Q] \mid S \mid (\mathbf{B}(\mathbf{T}, t), t)} \quad [\text{PUT}] \\
\\
\frac{\llbracket \mathbf{T} \rrbracket = \mathbf{T}_1 \cdots \mathbf{T}_n \quad \forall i \in 1..n : \text{match}_{\mathbf{B}}(\mathbf{T}_i) = \mathbf{T}'_i \neq \perp}{\mathbf{A}[\text{ask } \mathbf{T} \text{ as } \mathbf{x}. P + R \mid Q] \mid S \mid (\mathbf{B}, t) \rightarrow \mathbf{A}[P\{\mathbf{T}'_1 \cdots \mathbf{T}'_n / \mathbf{x}\} \mid Q] \mid S \mid (\mathbf{B}, t)} \quad [\text{ASK}] \\
\\
\frac{\mathbf{X}(\mathbf{x}) = P \quad \mathbf{A}[P\{\llbracket \mathbf{E} \rrbracket / \mathbf{x}\} \mid Q] \mid S \rightarrow S'}{\mathbf{A}[\mathbf{X}(\mathbf{E}) \mid Q] \mid S \rightarrow S'} \quad [\text{DEF}] \quad \frac{t' > 0}{S \mid (\mathbf{B}, t) \xrightarrow{t'} S' \mid (\mathbf{B}, t + t')} \quad [\text{DELAY}]
\end{array}$$

**Figure 5.2:** Semantics of endpoint protocols.

tems, i.e. the parallel composition of the protocols of all participants, and the blockchain.

**Definition 5.3** (Semantics of protocols). *A system  $S$  is a term of the form  $\mathbf{A}_1[P_1] \mid \cdots \mid \mathbf{A}_n[P_n] \mid (\mathbf{B}, t)$ , where (i) all the  $\mathbf{A}_i$  are distinct; (ii) there exists a single component  $(\mathbf{B}, t)$ , representing the current state of the blockchain  $\mathbf{B}$ , and the current time  $t$ ; (iii) systems are up-to commutativity and associativity of  $\mid$ . We define the relation  $\rightarrow$  between systems in Figure 5.2, where  $\text{match}_{\mathbf{B}}(\mathbf{T})$  is the set of all the transactions in  $\mathbf{B}$  that are equal to  $\mathbf{T}$ , except for the witnesses. When writing  $S \mid S'$  we intend that the conditions above are respected.*

Intuitively, a guarded choice  $\sum_i \pi_i.P_i$  can behave as one of the branches  $P_i$ . A parallel composition  $P \mid Q$  executes concurrently  $P$  and  $Q$ . All the rules (except the last two) specify how a protocol  $(\pi.P + Q) \mid R$  evolves within a system. Rule [COM] models a message exchange between  $\mathbf{A}$  and  $\mathbf{B}$ : participant  $\mathbf{A}$  sends messages  $\mathbf{E}$ , which are received by  $\mathbf{B}$  on variables  $\mathbf{x}$ . Communication is synchronous, i.e.  $\mathbf{A}$  is blocked until  $\mathbf{B}$  is ready to receive. Rule [CHECK] allows the branch  $P$  of a sum to proceed if the condition represented by  $\mathbf{E}$  is true. Rule [PUT] allows  $\mathbf{A}$  to append a transaction to the blockchain, provided that the update is consistent. Rule [ASK] allows the branch  $P$  of a sum to proceed only when the blockchain contains some transactions  $\mathbf{T}'_1 \cdots \mathbf{T}'_n$  obtained by instantiating some  $\perp$  fields in  $\mathbf{T}$  (see Section 1.1). This form of pattern matching is crucial because the value of some fields (e.g., `wit`), may not be known at the time the protocol is written. When the `ask` prefix unblocks, the variables  $\mathbf{x}$  in  $P$  are bound to  $\mathbf{T}'_1 \cdots \mathbf{T}'_n$ , so making it possible to inspect their actual

$T$	$T'_A$
$\text{in: } (T_A, 1)$ $\text{wit: } \perp$ $\text{out: } (\lambda\zeta\zeta'.\text{versig}(k_A k_B, \zeta\zeta'), 1\mathfrak{B})$	$\text{in: } (T, 1)$ $\text{wit: } \perp$ $\text{out: } (\lambda\zeta.\text{versig}(k_A, \zeta), 1\mathfrak{B})$
$T'_B$	
$\text{in: } (T, 1)$ $\text{wit: } \perp$ $\text{out: } (\lambda\zeta.\text{versig}(k_B, \zeta), 1\mathfrak{B})$	

**Figure 5.3:** Transactions of the naïve escrow contract.

fields. Rule [DEF] allows a named process  $X(\mathbf{E})$  to evolve as  $P$ , assuming a defining equation  $X(\mathbf{x}) = P$ . The variables  $\mathbf{x}$  in  $P$  are substituted with the results of the evaluation of  $\mathbf{E}$ . Such defining equations can be used to specify recursive behaviours. Finally, rule [DELAY] allows time to pass. To keep our presentation simple, we have not included time-constraining operators in endpoint protocols. In case one needs a finer-grained control of time, well-known techniques [64] exist to extend a process algebra like ours with these operators.

**Example 5.1.1** (Naïve escrow). A buyer  $A$  wants to buy an item from the seller  $B$ , but they do not trust each other. So, they would like to use a contract to ensure that  $B$  will get paid if and only if  $A$  gets her item. In a naïve attempt to realise this, they use the transactions in Figure 5.3, where we assume that  $(T_A, 1)$  used in  $T.\text{in}$ , is a transaction output redeemable by  $A$  through her key  $k_A$ . The transaction  $T$  makes  $A$  deposit  $1\mathfrak{B}$ , which can be redeemed by a transaction carrying the signatures of both  $A$  and  $B$ . The transactions  $T'_A$  and  $T'_B$  redeem  $T$ , transferring the money to  $A$  or  $B$ , respectively.

The protocols of  $A$  and  $B$  are, respectively,  $P_A$  and  $Q_B$ :

$$\begin{aligned}
 P_A &= \text{put } T \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(T) \}. P' \\
 P' &= \tau.B ! \text{sig}_{k_A}^{**}(T'_B) + \tau.B ? x. \text{put } T'_A \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(T'_A) x \} \\
 Q_B &= \text{ask } T. (\tau.A ? x. \text{put } T'_B \{ \text{wit} \mapsto x \text{sig}_{k_B}^{**}(T'_B) \} + \tau.A ! \text{sig}_{k_B}^{**}(T'_A))
 \end{aligned}$$

First,  $A$  adds her signature to  $T$ , and puts it on the blockchain. Then, she internally chooses whether to unblock the deposit for  $B$  or to request a refund. In the first case,  $A$  sends  $\text{sig}_{k_A}^{**}(T'_B)$  to  $B$ . In the second case, she waits to receive the signature  $\text{sig}_{k_B}^{**}(T'_A)$  from  $B$  (saving it in the variable  $x$ ); afterwards, she puts  $T'_A$  on the blockchain (after setting  $\text{wit}$ ) to

redeem the deposit. The seller **B** waits to see **T** on the blockchain. Then, he chooses either to receive the signature  $\text{sig}_{k_A}^{**}(\mathsf{T}'_B)$  from **A** (and then redeem the payment by putting  $\mathsf{T}'_B$  on the blockchain), or to refund **A**, by sending his signature  $\text{sig}_{k_B}^{**}(\mathsf{T}'_A)$ .

This contract is not secure if either **A** or **B** are dishonest. On the one hand, a dishonest **A** can prevent **B** from redeeming the deposit, even if she had already received the item (to do that, it suffices not to send her signature, taking the rightmost branch in  $P'$ ). On the other hand, a dishonest **B** can just avoid to send the item and the signature (taking the leftmost branch in  $Q_B$ ): in this way, the deposit gets frozen. For instance, let  $S = A[P_A] \mid B[Q_B] \mid (\mathbf{B}, t)$ , where **B** contains  $\mathsf{T}_A$  unredeemed. The scenario where **A** has never received the item, while **B** dishonestly attempts to receive the payment, is modelled as follows:

$$\begin{aligned} S &\rightarrow A[P'] \mid B[Q_B] \mid (\mathbf{B}(\mathsf{T}, t), t) \\ &\rightarrow A[P'] \mid B[\tau.A ? x. \text{put } \mathsf{T}'_B \{ \text{wit} \mapsto x \text{sig}_{k_B}^{**}(\mathsf{T}'_B) \} + \tau.A ! \text{sig}_{k_B}^{**}(\mathsf{T}'_A)] \mid \dots \\ &\rightarrow A[B ? x. \text{put } \mathsf{T}'_A \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\mathsf{T}'_A) x \}] \mid B[A ? x. \text{put } \mathsf{T}'_B \{ \text{wit} \mapsto x \text{sig}_{k_B}^{**}(\mathsf{T}'_B) \}] \mid \dots \end{aligned}$$

At this point the computation is stuck, because both **A** and **B** are waiting a message from the other participant. We will show in Section 5.2.3 how to design a secure escrow contract, with the intermediation of a trusted arbiter.

## 5.2 Smart Contracts

We now present a comprehensive survey of smart contracts on Bitcoin, comprising those published in the academic literature, and those found online. To this aim we exploit the model of computation introduced in Section 5.1. Remarkably, all the following contracts can be implemented by only using so-called *standard* transactions [85], e.g. via the compilation technique in [11]. This is crucial, because non-standard transactions are currently discarded by the Bitcoin network.

### 5.2.1 Oracle

In many concrete scenarios one would like to make the execution of a contract depend on some real-world events, e.g. results of football matches for a betting contract, or feeds of flight delays for an insurance contract. However, the evaluation of Bitcoin scripts can not depend on the environment, so in these scenarios one has to resort to a trusted third-party, or

$T$	$T'_B$
$\text{in: } (T_A, 1)$ $\text{wit: } \text{sig}_{k_A}^{**}(T)$ $\text{out: } (\lambda \varsigma \varsigma'. \text{versig}(k_B k_O, \varsigma \varsigma'), v_{\mathbb{B}})$	$\text{in: } (T, 1)$ $\text{wit: } \perp$ $\text{out: } (\lambda \varsigma. \text{versig}(k_B, \varsigma), v_{\mathbb{B}})$

**Figure 5.4:** Transactions of a contract relying on an oracle.

oracle [87, 90], who notifies real-world events by providing signatures on certain transactions.

For example, assume that  $A$  wants to transfer  $v_{\mathbb{B}}$  to  $B$  only if a certain event, notified by an oracle  $O$ , happens. To do that,  $A$  puts on the blockchain the transaction  $T$  in Figure 5.4, which can be redeemed by a transactions carrying the signatures of both  $B$  and  $O$ . Further,  $A$  instructs the oracle to provide his signature to  $B$  upon the occurrence of the expected event.

We model the behaviour of  $B$  as the following protocol:

$$P_B = O ? x. \text{put } T'_B \{ \text{wit} \mapsto \text{sig}_{k_B}^{**}(T'_B) x \}$$

Here,  $B$  waits to receive the signature  $\text{sig}_{k_O}^{**}(T'_B)$  from  $O$ , then he puts  $T'_B$  on the blockchain (after setting its  $\text{wit}$ ) to redeem  $T$ . In practice, oracles like the one needed in this contract are available as services in the Bitcoin ecosystem, eg [119].

Notice that, in case the event certified by the oracle never happens, the  $v_{\mathbb{B}}$  within  $T$  are frozen forever. To avoid this situation, one can add a time constraint to the output script of  $T$ , e.g. as in the transaction  $T_{\text{bond}}$  in Figure 5.10.

## 5.2.2 Crowdfunding

Assume that the curator  $C$  of a crowdfunding campaign wants to fund a venture  $V$  by collecting  $v_{\mathbb{B}}$  from a set  $\{A_i\}_{i \in I}$  of investors. The investors want to be guaranteed that either the required amount  $v_{\mathbb{B}}$  is reached, or they will be able to redeem their funds. To this purpose,  $C$  can employ the following contract. She starts with a canonical transaction  $U_V^{v_{\mathbb{B}}}$  (with empty  $\text{in}$  field) which has a single output of  $v_{\mathbb{B}}$  to be redeemed by  $V$ . Intuitively, each  $A_i$  can invest money in the campaign by “filling in” the  $\text{in}$  field of the  $U_V^{v_{\mathbb{B}}}$  with a transaction output under their control. To do this,  $A_i$  sends to  $C$  a transaction output  $(T_i, j_i)$ , together with the signature  $\sigma_i$  required to redeem it. We denote with  $\text{val}(T_i, j_i)$  the value of such output. Notice that, since the signature  $\sigma_i$  has been made on  $U_V^{v_{\mathbb{B}}}$ ,

the only valid output is the one of  $v\mathbb{B}$  to be redeemed by  $\mathbf{V}$ . Upon the reception of the message from  $\mathbf{A}_i$ ,  $\mathbf{C}$  updates  $\mathbf{U}_V^{v\mathbb{B}}$ : the provided output is appended to the **in** field, and the signature is added to the corresponding **wit** field. If all the outputs  $(\mathbf{T}_i, j_i)$  are distinct (and not redeemed) and the signatures are valid, when  $\sum_i \text{val}(\mathbf{T}_i, j_i) \geq v$  the filled transaction  $\mathbf{U}_V^{v\mathbb{B}}$  can be put on the blockchain. If  $\mathbf{C}$  collects  $v' > v\mathbb{B}$ , the difference  $v' - v$  goes to the miners as transaction fee.

The endpoint protocol of the curator is defined as  $\mathbf{X}(\mathbf{U}_V^{v\mathbb{B}}, 1, 0)$ , where:

$$\begin{aligned} \mathbf{X}(x, n, d) &= \text{if } d < v \text{ then } P \text{ else put } x \\ P &= \sum_i \mathbf{A}_i ? (y, j, \sigma). \mathbf{X}(x \{ \text{in}(n) \mapsto (y, j) \} \{ \text{wit}(n) \mapsto \sigma \}, n + 1, d + \text{val}(y, j)) \end{aligned}$$

while the protocol of each investor  $\mathbf{A}_i$  is the following:

$$P_{\mathbf{A}_i} = \mathbf{C} ! (\mathbf{T}_i, j_i, \text{sig}_{k_{\mathbf{A}_i}}^{1* \cdot 1} (\mathbf{U}_V^{v\mathbb{B}} \{ \text{in}(1) \mapsto (\mathbf{T}_i, j_i) \}))$$

Note that the transactions sent by investors are not known *a priori*, so they cannot just create the final transaction and sign it. Instead, to allow  $\mathbf{C}$  to complete the transaction  $\mathbf{U}_V^{v\mathbb{B}}$  without invalidating the signatures, they compute them using the modifier  $1* \cdot 1$ . In this way, only a single input is signed, and when verifying the corresponding signature, the others are neglected.

### 5.2.3 Escrow

In Example 5.1.1 we have discussed a naïve escrow contract, which is secure only if both the buyer  $\mathbf{A}$  and the seller  $\mathbf{B}$  are honest (so making the contract pointless). Rather, one would like to guarantee that, even if either  $\mathbf{A}$  or  $\mathbf{B}$  (or both) are dishonest, exactly one them will be able to redeem the money: in case they disagree, a trusted participant  $\mathbf{C}$ , who plays the role of arbiter, will decide who gets the money (possibly splitting the initial deposit in two parts) [84, 90].

The output script of the transaction  $\mathbf{T}$  in Figure 5.5 is a *2-of-3* multi-signature schema. This means that  $\mathbf{T}$  can be redeemed either with the signatures  $\mathbf{A}$  and  $\mathbf{B}$  (in case they agree), or with the signature of  $\mathbf{C}$  (with key  $k_C$ ) and the signature of  $\mathbf{A}$  or that of  $\mathbf{B}$  (in case they disagree). The transaction  $\mathbf{T}'_{\mathbf{AB}}(z)$  in Figure 5.5 allows the arbiter to issue a *partial* refund of  $z\mathbb{B}$  to  $\mathbf{A}$ , and of  $(1 - z)\mathbb{B}$  to  $\mathbf{B}$ . Instead, to issue a full refund to either  $\mathbf{A}$  or  $\mathbf{B}$ , the arbiter signs, respectively, the transactions  $\mathbf{T}'_{\mathbf{A}} = \mathbf{U}_{\mathbf{A}}^{1\mathbb{B}\mathbb{B}} \{ \text{in}(1) \mapsto (\mathbf{T}, 1) \}$  or  $\mathbf{T}'_{\mathbf{B}} = \mathbf{U}_{\mathbf{B}}^{1\mathbb{B}\mathbb{B}} \{ \text{in}(1) \mapsto (\mathbf{T}, 1) \}$  (not shown in the figure). The protocols of  $\mathbf{A}$  and  $\mathbf{B}$  are similar to those

$\overline{T}$
$\text{in: } (\overline{T}_A, 1)$ $\text{wit: } \perp$ $\text{out: } (\lambda \varsigma \varsigma'. \text{versig}(k_A k_B k_C, \varsigma \varsigma'), 1\text{B})$

$\overline{T}'_{AB}(z)$
$\text{in: } (\overline{T}, 1)$ $\text{wit: } \perp$ $\text{out: } 1 \mapsto (\lambda \varsigma. \text{versig}(k_A, \varsigma), z\text{B}), 2 \mapsto (\lambda \varsigma. \text{versig}(k_B, \varsigma), (1-z)\text{B})$

**Figure 5.5:** Transactions of the escrow contract.

in Example 5.1.1, except for the part where they ask  $C$  for an arbitration:

$$\begin{aligned}
P_A &= \text{put } \overline{T} \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\overline{T}) \}. (\tau. \text{B} ! \text{sig}_{k_B}^{**}(\overline{T}'_B) + \tau. P') \\
P' &= (\text{B} ? x. (\text{put } \overline{T}'_A \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\overline{T}'_A) x \} + P''')) + P'' \\
P'' &= \text{C} ? (z, x). (\text{check } z = 1 . \text{put } \overline{T}'_A \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\overline{T}'_A) x \} \\
&\quad + \text{check } 0 < z < 1 . (\text{put } \overline{T}'_{AB}(z) \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\overline{T}'_{AB}(z)) x \} + \tau. \mathbf{0}) \\
&\quad + \text{check } z = 0 . \mathbf{0})
\end{aligned}$$

In the summation within  $P_A$ , participant  $A$  internally chooses whether to send her signature to  $B$  (so allowing  $B$  to redeem  $1\text{B}$  via  $\overline{T}'_B$ ), or to proceed with  $P'$ . There,  $A$  waits to receive either  $B$ 's signature (which allows  $A$  to redeem  $1\text{B}$  by putting  $\overline{T}'_A$  on the blockchain), or a response from the arbiter, in the process  $P''$ . The three cases in the summation of  $\text{check}$  in  $P''$  correspond, respectively, to the case where  $A$  gets a full refund ( $z = 1$ ), a partial refund ( $0 < z < 1$ ), or no refund at all ( $z = 0$ ).

The protocol for  $B$  is dual to that of  $A$ :

$$\begin{aligned}
Q_B &= \text{ask } \overline{T}. (\tau. \text{A} ! \text{sig}_{k_B}^{**}(\overline{T}'_A) + \tau. Q') \\
Q' &= (\text{A} ? x. (\text{put } \overline{T}'_B \{ \text{wit} \mapsto x \text{sig}_{k_B}^{**}(\overline{T}'_B) \} + Q'')) + Q'' \\
Q'' &= \text{C} ? (z, x). (\text{check } z = 0 . \text{put } \overline{T}'_B \{ \text{wit} \mapsto \text{sig}_{k_B}^{**}(\overline{T}'_B) x \} \\
&\quad + \text{check } 0 < z < 1 . (\text{put } \overline{T}'_{AB}(z) \{ \text{wit} \mapsto \text{sig}_{k_B}^{**}(\overline{T}'_{AB}(z)) x \} + \tau. \mathbf{0}) \\
&\quad + \text{check } z = 1 . \mathbf{0})
\end{aligned}$$

If an arbitration is requested,  $C$  internally decides (through the  $\tau$  actions) who between  $A$  and  $B$  can redeem the deposit in  $\overline{T}$ , by sending its signature to one of the two participants, or decide for a partial refund of  $z$  and  $1 - z$  bitcoins, respectively, to  $A$  and  $B$ , by sending its signature on

T	T'
<b>in:</b> $1 \mapsto (\mathbf{T}_A, 1)$ <b>wit:</b> $\perp$ <b>out:</b> $1 \mapsto (\lambda\varsigma.\text{versig}(k_1, \varsigma), v_1\mathfrak{B})$ $2 \mapsto (\lambda\varsigma.\text{versig}(k_2, \varsigma), v_2\mathfrak{B})$	<b>in:</b> $1 \mapsto (\mathbf{T}_A, 1), 2 \mapsto (\mathbf{T}_B, 1)$ <b>wit:</b> $1 \mapsto \text{sig}_{k_A}^{1*1}, 2 \mapsto \text{sig}_{k_B}^{**}$ <b>out:</b> $1 \mapsto (\lambda\varsigma.\text{versig}(k_1, \varsigma), v_1\mathfrak{B})$ $2 \mapsto (\lambda\varsigma.\text{versig}(k_2, \varsigma), v_2\mathfrak{B})$

**Figure 5.6:** Transactions of the Bitcoin mixing contract.

$\mathbf{T}'_{AB}$  to both participants:

$$R_C = \text{ask } \mathbf{T}. (\tau.\mathbf{A}!(1, \text{sig}_{k_C}^{**}(\mathbf{T}'_A)) + \tau.\mathbf{B}!(1, \text{sig}_{k_C}^{**}(\mathbf{T}'_B)) + \tau.R_{AB})$$

$$R_{AB} = \sum_{0 < z < 1} \tau.(\mathbf{A}!(z, \text{sig}_{k_C}^{**}(\mathbf{T}'_{AB}(z))) | \mathbf{B}!(z, \text{sig}_{k_C}^{**}(\mathbf{T}'_{AB}(z))))$$

Note that, in the unlikely case where both **A** and **B** choose to send their signature to the other participant, the  $1\mathfrak{B}$  deposit becomes “frozen”. In a more concrete version of this contract, a participant could keep listening for the signature, and attempt to redeem the deposit when (unexpectedly) receiving it.

## 5.2.4 Mixing bitcoins

Assume two participants, **A** and **B** (with keys  $k_A$  and  $k_B$ , respectively), who want to pay, respectively,  $v_1\mathfrak{B}$  to  $\mathbf{C}_1$  and  $v_2\mathfrak{B}$  to  $\mathbf{C}_2$  (with keys  $k_1$  and  $k_2$ , respectively). Assume also that the blockchain **B** contains two unspent transactions,  $\mathbf{T}_A$  and  $\mathbf{T}_B$ , such that:

$$\mathbf{T}_A.\text{out} = (\lambda\varsigma.\text{versig}(k_A, \varsigma), v_1\mathfrak{B}) \quad \mathbf{T}_B.\text{out} = (\lambda\varsigma.\text{versig}(k_B, \varsigma), v_2\mathfrak{B})$$

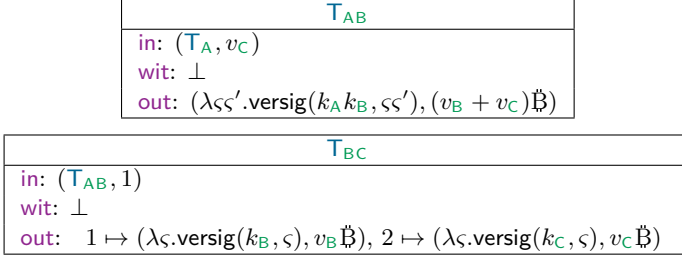
which can be redeemed by **A** and **B**, respectively.

To avoid de-anonymization, **A** mixes her bitcoins with those of **B**, using the following protocol:

$$P_A = \mathbf{B}! \text{sig}_{k_A}^{1*1}(\mathbf{T})$$

where the transaction  $\mathbf{T}$  is displayed in Figure 5.6 (left). The transaction  $\mathbf{T}$  has one input (the only output of  $\mathbf{T}_A$ ), and two outputs: one redeemable with  $k_1$ , and the other one with  $k_2$ . Participant **A** sends her signature (computed using the modifier  $1*1$ ) to participant **B**. With the modifier  $1*1$ , the verification of such signature will consider all outputs, but only the input with index 1. Therefore, **B** will be able to add inputs to  $\mathbf{T}$ , without invalidating **A**'s signature.





**Figure 5.7:** Transactions of the intermediated payment contract.

Note that  $T$  tries to redeem  $v_1\mathring{\text{B}}$  from  $T_A.\text{out}(1)$ , but it makes available  $v_1 + v_2 > v_1$  bitcoins through its outputs. Hence,  $T$  is *not* a consistent update of  $B$ , since it violates condition (3) of Definition 3.14. However, since  $A$  has signed  $T$  with the modifier  $1*_1$ ,  $B$  can add another input without breaking her signature (but he cannot steal her bitcoins, since all the outputs are signed).

The protocol of  $B$  is the following:

$$Q_B = A ? x. \text{put } T \{ \text{in}(2) \mapsto (T_B, 1) \} \{ \text{wit}(1) \mapsto x \} \{ \text{wit}(2) \mapsto \text{sig}_{k_B}^{**}(\{ \text{in}(2) \mapsto (T_B, 1) \}) \}$$

Here,  $B$  completes the transaction  $T$ , adding an input,  $A$ 's signature (in  $x$ ), and his signature on  $T$  (signing all inputs and all outputs, so that the transaction cannot be changed anymore). This modified transaction, displayed in Figure 5.6 (right), is a consistent update of  $B$ , hence  $B$  can put it on the blockchain.

This contract works as an anonymity schema because it falsifies the premise of the *multi-input* address clustering heuristic, firstly introduced by [61]. It assumes that all the inputs of a transaction are signed using keys controlled by the same user, which is not true in  $T'$ .

## 5.2.5 Intermediated payment

Assume that  $A$  wants to send an indirect payment of  $v_C\mathring{\text{B}}$  to  $C$ , routing it through an intermediary  $B$  who retains a fee of  $v_B < v_C$  bitcoins. Since  $A$  does not trust  $B$ , she wants to use a contract to guarantee that: (i) if  $B$  is honest, then  $v_C\mathring{\text{B}}$  are transferred to  $C$ ; (ii) if  $B$  is *not* honest, then  $A$  does not lose money. The contract uses the transactions in Figure 5.7:  $T_{AB}$  transfers  $(v_B + v_C)\mathring{\text{B}}$  from  $A$  to  $B$ , and  $T_{BC}$  splits the amount to  $B$  ( $v_B\mathring{\text{B}}$ ) and to  $C$  ( $v_C\mathring{\text{B}}$ ). We assume that  $(T_A, 1)$  is a transaction output

redeemable by **A**. The behaviour of **A** is as follows:

$$\begin{aligned}
 P_A &= (\mathbf{B} ? x. \text{if } \text{versig}_{k_B}(x, T_{BC}) \text{ then } P' \text{ else } \mathbf{0}) + \tau \\
 P' &= \text{put } T_{AB} \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(T_{AB}) \}. \text{put } T_{BC} \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(T_{BC}) x \}
 \end{aligned}$$

Here, **A** receives from **B** his signature on  $T_{BC}$ , which makes it possible to pay **C** later on. The  $\tau$  branch and the else branch ensure that **A** will correctly terminate also if **B** is dishonest (i.e., **B** does not send anything, or he sends an invalid signature). If **A** receives a valid signature, she puts  $T_{AB}$  on the blockchain, adding her signature to the **wit** field. Then, she also appends  $T_{BC}$ , adding to the **wit** field her signature and **B**'s one. Since **A** takes care of publishing both transactions, the behaviour of **B** consists just in sending his signature on  $T_{BC}$ . Therefore, **B**'s protocol can just be modelled as  $Q_B = A ! \text{sig}_{k_B}^{**}(T_{BC})$ .

This contract relies on SegWit. In Bitcoin without SegWit, the identifier of  $T_{AB}$  is affected by the instantiation of the **wit** field. So, when  $T_{AB}$  is put on the blockchain, the input in  $T_{BC}$  (which was computed before) does not point to it.

## 5.2.6 Timed commitment

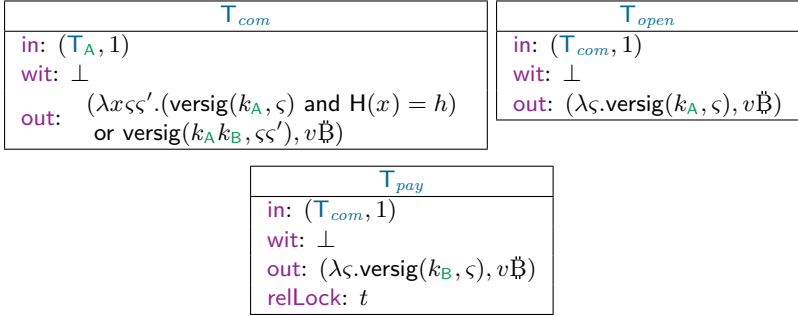
Assume that **A** wants to choose a secret  $s$ , and reveal it after some time — while guaranteeing that the revealed value corresponds to the chosen secret (or paying a penalty otherwise). This can be obtained through a *timed commitment* [29], a protocol with applications e.g. in gambling games [39, 71, 47], where the secret contains the player move, and the delay in the revelation of the secret is intended to prevent other players from altering the outcome of the game. Here we formalise the version of the timed commitment protocol presented in [4].

Intuitively, **A** starts by exposing the hash of the secret, i.e.  $h = H(s)$ , and at the same time depositing some amount  $v\mathfrak{B}$  in a transaction. The participant **B** has the guarantee that after  $t$  time units, he will either know the secret  $s$ , or he will be able to redeem  $v\mathfrak{B}$ .

The transactions of the protocol are shown in Figure 5.8, where we assume that  $(T_A, 1)$  is a transaction output redeemable by **A**. The behaviour of **A** is modelled as the following protocol:

$$\begin{aligned}
 P_A &= \text{put } T_{com} \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(T_{com}) \}. \mathbf{B} ! \text{sig}_{k_A}^{**}(T_{pay}) . P' \\
 P' &= \tau . \text{put } T_{open} \{ \text{wit} \mapsto s \text{ sig}_{k_A}^{**}(T_{open}) \perp \} + \tau
 \end{aligned}$$

Participant **A** starts by putting the transaction  $T_{com}$  on the blockchain. Note that within this transaction **A** is committing the hash of the chosen



**Figure 5.8:** Transactions of the timed commitment.

secret: indeed,  $h$  is encoded within the output script  $T_{com}.\text{out}$ . Then,  $A$  sends to  $B$  her signature on  $T_{pay}$ . Note that this transaction can be redeemed by  $B$  only when  $t$  time units have passed since  $T_{com}$  has been published on the blockchain, because of the relative timelock declared in  $T_{pay}.\text{relLock}$ . After sending her signature on  $T_{pay}$ ,  $A$  internally chooses whether to reveal the secret, or do nothing (via the  $\tau$  actions). In the first case,  $A$  must put the transaction  $T_{open}$  on the blockchain. Since it redeems  $T_{com}$ , she needs to write in  $T_{open}.\text{wit}$  both the secret  $s$  and her signature, so making the former public.

A possible behaviour of the receiver  $B$  is the following:

$$Q_B = (A ? x. \text{if } \text{versig}_{k_A}(x, T_{pay}) \text{ then } Q \text{ else } 0) + \tau$$

$$Q = \text{put } T_{pay} \{ \text{wit} \mapsto \perp x \text{ sig}_{k_B}^{**}(T_{pay}) \} + \text{ask } T_{open} \text{ as } o. Q'(\text{get}_{secret}(o))$$

In this protocol,  $B$  first receives from  $A$  (and saves in  $x$ ) her signature on the transaction  $T_{pay}$ . Then,  $B$  checks if the signature is valid: if not, he aborts the protocol. Even if the signature is valid,  $B$  cannot put  $T_{pay}$  on the blockchain and redeem the deposit immediately, since the transaction has a timelock  $t$ . Note that  $B$  cannot change the timelock: indeed, doing so would invalidate  $A$ 's signature on  $T_{pay}$ . If, after  $t$  time units,  $A$  has not published  $T_{open}$  yet,  $B$  can proceed to put  $T_{pay}$  on the blockchain, writing  $A$ 's and his own signatures in the witness. Otherwise,  $B$  retrieves  $T_{open}$  from the blockchain, from which he can obtain the secret, and use it in  $Q'$ .

A variant of this contract, which implements the timeout in  $T_{com}.\text{out}$ , and does not require the signature exchange, is used in Section 5.2.8.

**Timed commitment without signature exchange** The timed commitment protocol in Section 5.2.6 has a major drawback:  $A$  could put the

<div style="text-align: center; border-bottom: 1px solid black; margin-bottom: 5px;"><math>\overline{T}_{com}</math></div> <p>in: <math>(\overline{T}_A, 1)</math>  wit: <math>\perp</math>  out: <math>(\lambda x \varsigma. (\text{versig}(k_A, \varsigma) \text{ and } H(x) = h)</math>  or <math>\text{relAfter } t : \text{versig}(k_B, \varsigma), v\mathfrak{B})</math></p>	<div style="text-align: center; border-bottom: 1px solid black; margin-bottom: 5px;"><math>\overline{T}_{open}</math></div> <p>in: <math>(\overline{T}_{com}, 1)</math>  wit: <math>\perp</math>  out: <math>(\lambda \varsigma. \text{versig}(k_A, \varsigma), v\mathfrak{B})</math></p>
<div style="text-align: center; border-bottom: 1px solid black; margin-bottom: 5px;"><math>\overline{T}_{pay}</math></div> <p>in: <math>(\overline{T}_{com}, 1)</math>  wit: <math>\perp</math>  out: <math>(\lambda \varsigma. \text{versig}(k_B, \varsigma), v\mathfrak{B})</math>  relLock: <math>t</math></p>	

**Figure 5.9:** Transactions of the timed commitment without signature exchange.

transaction  $\overline{T}_{com}$  on the blockchain, and then avoid to send  $\text{sig}_{k_A}^{**}(\overline{T}_{pay})$ . In this way, even if she started the contract, she could maintain the secret hidden without losing the deposit. We now present a variant of the protocol which addresses this problem. The new transactions are in Figure 5.9.

The new version of  $\overline{T}_{com}$  can be redeemed by **B** using only his own signature. In order to enforce **A** to reveal the secret before  $t$  time units have passed, the output script of  $\overline{T}_{com}$  uses the expression  $\text{relAfter } t : \text{versig}(k_B, \varsigma)$ . This has the effect of verifying the signature in the witness of the redeeming transaction (i.e.,  $\overline{T}_{pay}$ ) only if such transaction has **relLock** greater or equal than  $t$ .

The participants' protocols are similar to those in Section 5.2.6, except for the signature exchange, which is no longer necessary.

$$P_A = \text{put } \overline{T}_{com} \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\overline{T}_{com}) \}. (\tau. (\text{put } \overline{T}_{open} \{ \text{wit} \mapsto s \text{ sig}_{k_A}^{**}(\overline{T}_{open}) \}) + \tau)$$

$$Q_B = \text{ask } \overline{T}_{com}. (\text{put } \overline{T}_{pay} \{ \text{wit} \mapsto \perp \text{ sig}_{k_B}^{**}(\overline{T}_{pay}) \} + \tau)$$

## 5.2.7 Micropayment channels

Assume that **A** wants to make a series of micropayments to **B**, e.g. a small fraction of  $\mathfrak{B}$  every few minutes. Doing so with one transaction per payment would result in conspicuous fees [82], so **A** and **B** use a micropayment channel contract [108]. **A** starts by depositing  $k\mathfrak{B}$ ; then, she signs a transaction that pays  $v\mathfrak{B}$  to **B** and  $(k - v)\mathfrak{B}$  back to herself, and she sends that transaction to **B**. Participant **B** can choose to publish that transaction immediately and redeem its payment, or to wait in case **A** sends another transaction with increased value. **A** can stop sending

$\overline{T}_{bond}$ <p>in: <math>(\overline{T}_A, 1)</math>  wit: <math>\perp</math>  out: <math>(\lambda\varsigma\varsigma'.\text{versig}(k_A k_B, \varsigma\varsigma')</math> or  <math>\text{relAfter } t : \text{versig}(k_A, \varsigma), k\mathfrak{B})</math></p>	$\overline{T}_{pay}(v)$ <p>in: <math>(\overline{T}_{bond}, 1)</math>  wit: <math>\perp</math>  out: <math>1 \mapsto (\lambda\varsigma.\text{versig}(k_A, \varsigma), (k - v)\mathfrak{B})</math>  <math>2 \mapsto (\lambda\varsigma.\text{versig}(k_C, \varsigma), v\mathfrak{B})</math></p>
$\overline{T}_{ref}$ <p>in: <math>(\overline{T}_{bond}, 1)</math>  wit: <math>\perp</math>  out: <math>(\lambda\varsigma.\text{versig}(k_A, \varsigma), v\mathfrak{B})</math>  relLock: <math>t</math></p>	

**Figure 5.10:** Transactions of the micropayment channel contract.

signatures at any time. If **B** redeems, then **A** can get back the remaining amount. If **B** does not cooperate, **A** can redeem all the amount after a timeout.

The protocol of **A** is the following (the transactions are in Figure 5.10). **A** publishes the transaction  $\overline{T}_{bond}$ , depositing  $k\mathfrak{B}$  that can be spent with her signature and that of **B**, or with her signature alone, after time  $t$ . **A** can redeem the deposit by publishing the transaction  $\overline{T}_{ref}$ . To pay for the service, **A** sends to **B** the amount  $v$  she is paying, and her signature on  $\overline{T}_{pay}(v)$ . Then, she can decide to increase  $v$  and recur, or to terminate.

$$P_A = \text{put } \overline{T}_{bond} \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\overline{T}_{bond}) \}. (P(1) \mid \text{put } \overline{T}_{ref} \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\overline{T}_{ref}) \})$$

$$P(v) = \text{B}!(v, \text{sig}_{k_A}^{**}(\overline{T}_{pay}(v))). (\tau + \tau.P(v+1))$$

The participant **B** waits for  $\overline{T}_{bond}$  to appear on the blockchain, then receives the first value  $v$  and **A**'s signature  $\sigma$ . Then, **B** checks if  $\sigma$  is valid, otherwise he aborts the protocol. At this point, **B** waits for another pair  $(v', \sigma')$ , or, after a timeout, he redeems  $v\mathfrak{B}$  using  $\overline{T}_{pay}(v)$ .

$$Q_B = \text{ask } \overline{T}_{bond}. \text{A}?(v, \sigma). \text{if } \text{versig}_{k_A}(\sigma, \overline{T}_{pay}(v)) \text{ then } P'(v, \sigma) \text{ else } \tau$$

$$P'(v, \sigma) = \tau.P_{pay}(v, \sigma) +$$

$$\text{A}?(v', \sigma'). \text{if } v' > v \text{ and } \text{versig}_{k_A}(\sigma', \overline{T}_{pay}(v')) \text{ then } P'(v', \sigma') \text{ else } P'(v, \sigma)$$

$$P_{pay}(v, \sigma) = \text{put } \overline{T}_{pay}(v) \{ \text{wit} \mapsto \sigma \text{sig}_{k_B}^{**}(\overline{T}_{pay}(v)) \}$$

Note that  $Q_B$  should redeem  $\overline{T}_{pay}$  before the timeout expires, which is not modelled in  $Q_B$ . This could be obtained by enriching the calculus with time-constraining operators.

### 5.2.8 Fair lotteries

A multiparty lottery is a protocol where  $N$  players put their bets in a pot, and a winner — uniformly chosen among the players — redeems the whole pot. Various contracts for multiparty lotteries on Bitcoin have been proposed in [4, 5, 80, 26, 23, 62]. These contracts enjoy a *fairness* property, which roughly guarantees that: (i) each honest player will have (on average) a non-negative payoff, even in the presence of adversaries; (ii) when all the players are honest, the protocol behaves as an ideal lottery: one player wins the whole pot (with probability  $1/N$ ), while all the others lose their bets (with probability  $N-1/N$ ).

Here we illustrate the lottery in [4], for  $N = 2$ . Consider two players **A** and **B** who want to bet 1 $\text{฿}$  each. Their protocol is composed of two phases. The first phase is a timed commitment (as in Section 5.2.6): each player chooses a secret ( $s_{\mathbf{A}}$  and  $s_{\mathbf{B}}$ ) and commits its hash ( $h_{\mathbf{A}} = H(s_{\mathbf{A}})$  and  $h_{\mathbf{B}} = H(s_{\mathbf{B}})$ ). In doing that, both players put a deposit of 2 $\text{฿}$  on the ledger, which is used to compensate the other player in case one chooses not to reveal the secret later on. In the second phase, the two bets are put on the ledger. After that, the players reveal their secrets, and redeem their deposits. Then, the secrets are used to compute the winner of the lottery in a fair manner. Finally, the winner redeems the bets.

The transactions needed for this lottery are displayed in Figure 5.11 (we only show **A**'s transactions, as those of **B** are similar). The transactions for the commitment phase ( $T_{com}, T_{open}, T_{pay}$ ) are similar to those in Section 5.2.6: they only differ in the script of  $T_{com.out}$ , which now also checks that the length of the secret is either 128 or 129. This check forces the players to choose their secret so that it has one of these lengths, and reveal it (using  $T_{open}$ ) before the `absLock` deadline, since otherwise they will lose their deposits (enabling  $T_{pay}$ ).

The bets are put using  $T_{lottery}$ , whose output script computes the winner using the secrets, which can then be revealed. For this, the secret lengths are compared: if equal, **A** wins, otherwise **B** wins. In this way, the lottery is equivalent to a coin toss. Note that, if a malicious player chooses a secret having another length than 128 or 129, the  $T_{lottery}$  transaction will become stuck, but its opponent will be compensated using the deposit.

The endpoint protocol  $P_{\mathbf{A}}$  of player **A** follows (the one for **B** is similar):

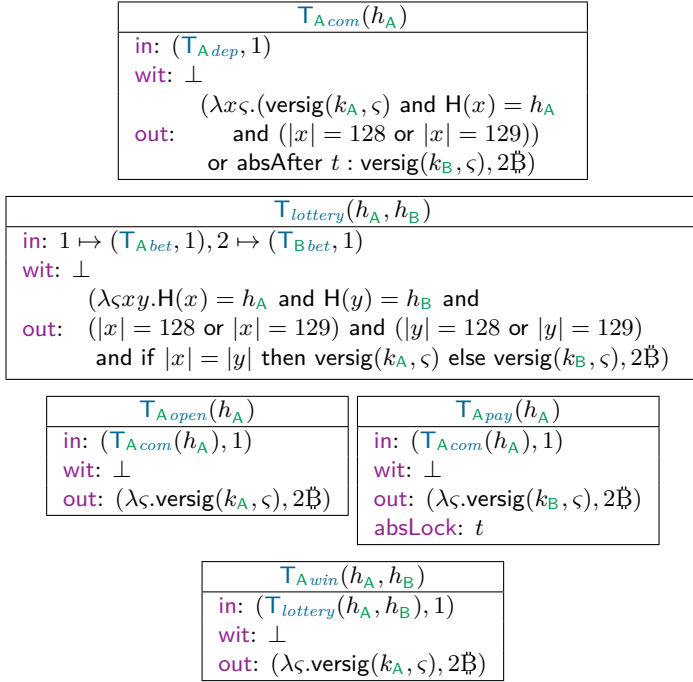


Figure 5.11: Transactions of the fair lottery with deposit.

$$\begin{aligned}
P_A &= \text{put } \overline{T}_{Acom} \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\overline{T}_{Acom}) \}. (\text{ask } \overline{T}_{Bcom} \text{ as } y. P' + \tau.P_{open}) \\
P' &= \text{let } h_B = \text{get}_{hash}(y) \text{ in if } h_B \neq h_A \text{ then } P_{pay} \mid P''' \text{ else } P_{pay} \mid P_{open} \\
P''' &= B ? x. P'' + \tau.P_{open} \\
P'' &= \text{let } \sigma = \text{sig}_{k_A}^{**,1}(\overline{T}_{lottery}(h_A, h_B)) \text{ in} \\
&\quad (\text{put } \overline{T}_{lottery}(h_A, h_B) \{ \text{wit}(1) \mapsto \sigma \} \{ \text{wit}(2) \mapsto x \}. (P_{open} \mid P_{win})) + \tau.P_{open} \\
P_{pay} &= \text{put } \overline{T}_{Bpay} \{ \text{wit} \mapsto \perp \text{ sig}_{k_A}^{**}(\overline{T}_{Bpay}) \} \\
P_{open} &= \text{put } \overline{T}_{Aopen} \{ \text{wit} \mapsto s_A \text{ sig}_{k_A}^{**}(\overline{T}_{Aopen}) \} \\
P_{win} &= \text{ask } \overline{T}_{Bopen} \text{ as } z. P'_{win} \\
P'_{win} &= \text{put } \overline{T}_{Awin}(h_A, h_B) \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\overline{T}_{Awin}(h_A, h_B)) s_A \text{ get}_{secret}(z) \}
\end{aligned}$$

Player **A** starts by putting  $\overline{T}_{Acom}$  on the blockchain, then she waits for **B** doing the same. If **B** does not cooperate, **A** can safely abort the protocol taking its  $\tau.P_{open}$  branch, so redeeming her deposit with  $\overline{T}_{Aopen}$  (as usual, here with  $\tau$  we are modelling a timeout). If **B** commits his

secret, **A** executes  $P'$ , extracting the hash  $h_B$  of **B**'s secret, and checking whether it is distinct from  $h_A$ . If the hashes are found to be equal, **A** aborts the protocol using  $P_{open}$ . Otherwise, **A** runs  $P''' \mid P_{pay}$ . The  $P_{pay}$  component attempts to redeem **B**'s deposit, as soon as the `absLock` deadline of  $T_{B_{pay}}$  expires, forcing **B** to timely reveal his secret. Instead,  $P'''$  proceeds with the lottery, asking **B** for his signature of  $T_{lottery}$ . If **B** does not sign, **A** aborts using  $P_{open}$ . Then, **A** runs  $P''$ , finally putting the bets ( $T_{bet}$ ) on the ledger. If this is not possible (e.g., because one of the  $T_{bet}$  is already spent), **A** aborts using  $P_{open}$ . After  $T_{lottery}$  is on the ledger, **A** reveals her secret and redeems her deposit with  $P_{open}$ . In parallel, with  $P_{win}$  she waits for the secret of **B** to be revealed, and then attempts to redeem the pot ( $T_{A_{win}}$ ).

The fairness of this lottery has been established in [4]. This protocol can be generalised to  $N > 2$  players [4, 5] but in this case the deposit grows quadratically with  $N$ . The works [62, 23] have proposed fair multiparty lotteries that require, respectively, zero and constant ( $\geq 0$ ) deposit. More precisely, [62] devises two variants of the protocol: the first one only relies on SegWit, but requires each player to statically sign  $O(2^N)$  transactions; the second variant reduces the number of signatures to  $O(N^2)$ , at the cost of introducing a custom opcode. Also the protocol in [23] assumes an extension of Bitcoin, i.e. the malleability of in fields, to obtain an ideal fair lottery with  $O(N)$  signatures per player (see Section 7.3).

## 5.2.9 Contingent payments

Assume a participant **A** who wants to pay  $v\mathfrak{B}$  to receive a value  $s$  which makes a public predicate  $p$  true, where  $p(s)$  can be verified efficiently. A seller **B** who knows such  $s$  is willing to reveal it to **A**, but only under the guarantee that he will be paid  $v\mathfrak{B}$ . Similarly, the buyer wants to pay only if guaranteed to obtain  $s$ .

A naïve attempt to implement this contract in Bitcoin is the following: **A** creates a transaction  $T$  such that  $T.out(\varsigma, x)$  evaluates to true if and only if  $p(x)$  holds and  $\varsigma$  is a signature of **B**. Hence, **B** can redeem  $v\mathfrak{B}$  from  $T$  by revealing  $s$ . In practice, though, this approach is arguably useful, since it requires coding  $p$  in the Bitcoin scripting language, whose expressiveness is quite limited.

More general contingent payment contracts can be obtained by exploiting zero-knowledge proofs [12, 115, 38]. In this setting, the seller generates a fresh key  $k$ , and sends to the buyer the encryption  $e_s = E_k(s)$ , together with the hash  $h_k = H(k)$ , and a zero-knowledge proof guaranteeing that such messages have the intended form. After verifying this proof, **A** is sure that **B** knows a preimage  $k'$  of  $h_k$  (by collision resistance,  $k' = k$ ) such



<div style="text-align: center; border-bottom: 1px solid black; margin-bottom: 5px;"><math>\mathbb{T}_{cp}(h)</math></div> <p>in: <math>(\mathbb{T}_A, 1)</math>  wit: <math>\perp</math>  out: <math>(\lambda x \varsigma. (\text{versig}(k_B, \varsigma)</math> and <math>H(x) = h)</math>  or <math>\text{relAfter } t : \text{versig}(k_A, \varsigma), v\mathbb{B})</math></p>	<div style="text-align: center; border-bottom: 1px solid black; margin-bottom: 5px;"><math>\mathbb{T}_{open}(h)</math></div> <p>in: <math>(\mathbb{T}_{cp}(h), 1)</math>  wit: <math>\perp</math>  out: <math>(\lambda \varsigma. \text{versig}(k_B, \varsigma), v\mathbb{B})</math></p>
<div style="text-align: center; border-bottom: 1px solid black; margin-bottom: 5px;"><math>\mathbb{T}_{refund}(h)</math></div> <p>in: <math>(\mathbb{T}_{cp}(h), 1)</math>  wit: <math>\perp</math>  out: <math>(\lambda \varsigma. \text{versig}(k_A, \varsigma), v\mathbb{B})</math>  relLock: <math>t</math></p>	

**Figure 5.12:** Transactions of the contingent payment.

that  $D_{k'}(e_s)$  satisfies the predicate  $p$ , and so she can buy the preimage  $k$  of  $h_k$  with the naïve protocol, so obtaining the solution  $s$  by decrypting  $e_s$  with  $k$ .

The transactions implementing this contract are displayed in Figure 5.12. The  $\text{relAfter } :$  clause in  $\mathbb{T}_{cp}$  allows  $A$  to redeem  $v\mathbb{B}$  if no solution is provided by the deadline  $t$ . The behaviour of the buyer  $A$  can be modelled as follows:

$$\begin{aligned}
P_A &= B?(e_s, h_k, z). P + \tau \\
P &= \text{if } \text{verify}(e_s, h_k, z) \text{ then put } \mathbb{T}_{cp}(h_k) \{ \text{wit} \mapsto \text{sig}_{k_A}^{**}(\mathbb{T}_{cp}(h_k)) \}. P' \text{ else } 0 \\
P' &= \text{ask } \mathbb{T}_{open}(h_k) \text{ as } x. P'''(D_{\text{get}_k(x)}(e_s)) + \\
&\quad \text{put } \mathbb{T}_{refund}(h_k) \{ \text{wit} \mapsto \perp \text{ sig}_{k_A}^{**}(\mathbb{T}_{refund}(h_k)) \}
\end{aligned}$$

For simplicity, here we model the zero-knowledge proof as a single message. More concretely, it should be modelled as a sub-protocol. Upon receiving  $e_s$ ,  $h_k$  and the proof  $z$  the buyer verifies  $z$ . If the verification succeeds,  $A$  puts  $\mathbb{T}_{cp}(h_k)$  on the blockchain. Then, she waits for  $\mathbb{T}_{open}$ , from which she can retrieve the key  $k$ , and so use the solution  $D_{\text{get}_k(x)}(e_s)$  in  $P'''$ . In this way,  $B$  can redeem  $v\mathbb{B}$ . If  $B$  does not put  $\mathbb{T}_{open}$ , after  $t$  time units  $A$  can get her deposit back through  $\mathbb{T}_{refund}$ . The protocol of  $B$  is simple, so it is omitted.



## Chapter 6

# A Toolchain for Developing BitML Contracts

In the last five years much outstanding research has been devoted to showing how to exploit Bitcoin to execute *smart contracts* — computer protocols which allow for exchanging cryptocurrency according to complex pre-agreed rules [4, 3, 1, 5, 12, 23, 26, 56, 57, 55, 58, 62]. Despite the wide variety of use cases witnessed by these works, no tool support has been provided yet to facilitate the development of multi-transaction Bitcoin contracts. Today, this task requires to devise complex protocols which, besides using the standard cryptographic primitives, (e.g., encryption, hash functions, signatures), can read and append transactions on the Bitcoin blockchain. Creating a new protocol requires a significant effort to establish its correctness and security: this is an error-prone task, usually performed manually, with the risk of overlooking some corner cases. Crafting the transactions used by these protocols is burdensome as well, since it requires to struggle with low-level, poorly documented features of Bitcoin, like e.g. its stack-based scripting language.

We consider BitML, which features a computationally sound embedding into Bitcoin [22], and a sound and complete verification technique of relevant trace properties [24]. BitML can express many of the smart contracts appeared in the literature [15, 9], and execute them by appending suitable transactions to the Bitcoin blockchain. The computational soundness of the embedding guarantees that security properties at the level of the BitML semantics are preserved at the level of Bitcoin transactions,

even in the presence of adversaries. Still, BitML lives in a theoretical limbo, as no tool support exists yet to develop contracts and deploy them on the Bitcoin blockchain.

We develop a toolchain for writing and verifying BitML contracts, and for deploying them on Bitcoin. More specifically, the toolchain is composed of:

1. A BitML embedding in Racket [42], which allows for programming BitML contracts within the DrRacket IDE.
2. A security analyser which can check arbitrary LTL properties of BitML contracts. In particular, the analysis can decide *liquidity*, a landmark property of smart contracts requiring that the funds within a contract do not remain frozen forever.
3. A compiler from BitML contracts to standard Bitcoin transactions. The computational soundness result in [22] ensures that attacks to compiled contracts are also observable at the BitML level. Therefore, the properties verified by our security analyzer also hold for compiled contracts.

The architecture of our toolchain is displayed in Figure 6.1. The development workflow is the following: (a) write the BitML contract, and specify the required properties. Optionally, specify some constraints on the participants' strategies, e.g. to partially define the behaviour of the honest participants; (b) verify that the contract satisfies the required properties through the security analyser; (c) compile the contract to Bitcoin transactions; (d) execute the contract, by appending these transactions to the Bitcoin blockchain according to the chosen strategy. The first two steps are performed within the DrRacket IDE, while the third one requires participants to use the Balzac tool to compute signatures and compile Balzac transactions into Bitcoin transactions. We remark that the last step can be performed on the Bitcoin main network, without requiring any extensions or customizations.

The toolchain also supports an extension of BitML that enables recursion and renegotiation of contracts [16]. In this variant, the security analyser implements a different semantics, and can only check the liquidity of a contract.

We provide a collection of BitML contracts [93], which we use as a benchmark to evaluate our toolchain. This collection contains some of the most complex contracts ever developed for Bitcoin, e.g. financial services, auctions, timed commitments, lotteries, and a variety of other gambling games. We use our benchmarks to discuss the expressiveness and the limitations of Bitcoin as a smart contracts platform.

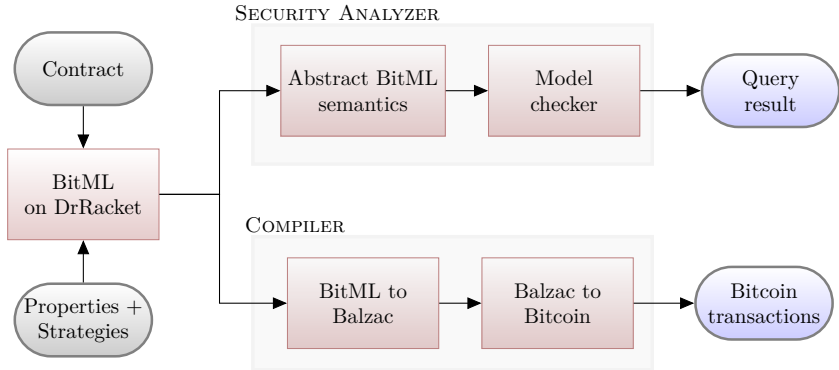


Figure 6.1: Toolchain architecture.

All the components of our toolchain are open-source [92, 96, 81], as well as the contracts in our benchmark. This Chapter is accompanied by a brief video that introduces the features of the toolchain (*Demo Video URL*: <https://bit.ly/2MxOBsT>).

## 6.1 Designing BitML contracts

BitML contracts allow two or more participants to exchange their bitcoins (₿) according to a given logic. A contract consists of two parts: a precondition, describing requirements that participants must fulfil to stipulate the contract, and a process, which specifies the execution logic of the contract. Here, rather than providing the syntax and semantics of BitML (see [22] for a formalization), we illustrate it through a simple but paradigmatic example, the *mutual timed commitment* contract [4]. This contract involves two participants (named below **A** and **B**) each one choosing a secret and depositing a certain amount of cryptocurrency (say, 1₿). The goal of the contract is to ensure that each participant will either learn the other participant’s secret, or otherwise receive the other participant’s deposit as a compensation. The contract gives some time to the participants to reveal their secrets. If a participant reveals her secret in time, then she can get her deposit back; otherwise, after the time is up, the other participant can withdraw that deposit.

In our tool, we can specify this contract as follows:

```

(participant "A" "029c...cced") ; A's public key
(participant "B" "022c...af30") ; B's public key

(contract
  (pre
    (deposit "A" 1 "1a34...6f3800") ; tx output id (1BTC)
    (secret "A" a "628f...de71") ; hash of A's secret
    (deposit "B" 1 "19e7...85ff02") ; tx output id (1BTC)
    (secret "B" b "9d48...bb35")) ; hash of B's secret

  (choice
    (reveal (a) (choice
      (reveal (b) (split
        (1 -> (withdraw "A"))
        (1 -> (withdraw "B")))))
      (after 100050 (withdraw "A"))))
    (after 100000 (withdraw "B"))))

```

The first two lines create aliases for the participant names, specifying their public keys. The contract preconditions are in the `pre` part: each participant must specify the identifier of a transaction output, and the hash of the chosen secret. The transaction output must be unspent, must contain the required 1 $\text{\textcircled{B}}$ , and must be redeemable using the participant's private key. The hash is used during the contract execution: when the participant provides a value, claiming that it is the chosen secret, the hash of this value is required to be equal to the one in the precondition.

The contract logic is specified after the preconditions. The top-level `choice` defines two alternative branches of the contract. The first branch can only be taken if `A` reveals her secret (named `a`); when this happens, the contract continues with the innermost `choice`. The second branch can only be taken after a timeout, specified as the block at height 100000, and it allows `B` to redeem all the funds deposited within the contract (i.e., 2 $\text{\textcircled{B}}$ ) by executing `withdraw "B"`. So, to avoid losing her deposit, `A` is incentivized to reveal her secret in time. Similarly, the innermost `choice` is used to incentivize `B` to reveal his secret before the block at height 100050. If `B` reveals, then the `split` subcontract is executed: this divides the balance of the contract in two parts of 1 $\text{\textcircled{B}}$  each, allowing the participants to withdraw their deposits back.

The language is defined exploiting the Racket macro system, which is used to rewrite BitML syntactic constructs to Racket code. This approach benefits from the Racket language ecosystem, and allows us to write BitML contracts in the DrRacket IDE. Indeed, our toolchain integrates within the DrRacket IDE the contract editor, the security analyzer and the BitML compiler. The implementation of BitML in Racket extends the idealized version of BitML in [22] to make the language usable in practice. For instance, it introduces special deposits of type `fee`,

which are automatically spread over all the transactions obtained by the compiler. We also implement static checks for a number of errors that could prevent the correct execution of contracts, e.g. committing secrets with the same hash, double spending a transaction output, etc.

## 6.2 Verifying BitML contracts

The tool verifies various forms of *liquidity*, requiring that no funds (or funds up-to a certain amount) are frozen forever within a contract. Further, the tool can verify arbitrary LTL formulae, where state predicates can specify, e.g., the funds owned by participants, the provided authorizations, and the revealed secrets. By default, the tool verifies the required property against *all* possible behaviours of each participant: for instance, if a contract contains `reveal a`, the verifier considers both the case where the secret is revealed and the one where it is not. Authorizations are handled similarly, by considering both cases. However, in most cases, a participant wishes to verify a contract with respect to a given behaviour for herself, making no assumptions on the other participants' behaviour (unless some other participants are considered trusted, in which case it would make sense to fix a behaviour also for them). For instance, a participant `A` may want to give her authorization to perform a given branch only after participant `B` has revealed his secret. The tool allows for constraining the behaviour of participants, specifying the conditions upon which secrets are revealed and authorizations are provided. Actions which can be performed by everyone, like `withdraw` and `split`, cannot be constrained.

For instance, we can verify that the mutual timed commitment contract is liquid whatever strategies are chosen by participants. The query `check-liquid` correctly answers true, since:

- if `A` does not reveal, then anyone (after the block at height 100000) can perform `withdraw "B"`, which transfers the whole contract balance to `B`;
- if `A` reveals but `B` does not reveal, then anyone (after the block at height 100050) can perform `withdraw "A"`, which transfers the whole contract balance to `A`;
- if both `A` and `B` reveal, then anyone can perform `split`, which transfers the balance in equal parts to `A` and `B`.

Note that if we remove the `after` branch at line 16, the contract is no longer liquid. However, it becomes liquid when `A`'s strat-

egy is to reveal the secret. We can verify that this holds through the query `check-liquid (strategy "A" (do-reveal a))`. Liquidity is lost again if `A` chooses to reveal only after `B` has revealed, i.e. when her strategy is `"A" (do-reveal a) if ("B" (do-reveal b))`.

Besides liquidity, we can check specific LTL properties of contracts through the command `check-query`. E.g., in the mutual timed commitment we can verify that, after `A` reveals, she will eventually get back at least her  $1\text{B}$  deposit. In LTL, this property is formalised as the following formula, where  $10^8$  satoshi =  $1\text{B}$ :

```
□ (a revealed => <>A has-deposit=> 100000000 satoshi)
```

We also verify that if `A` reveals the secret, then eventually either `B` reveals, or `A` will get `B`'s deposit, too. The LTL query is the following:

```
□ (a revealed =>
  <>(b revealed \/ A has-deposit=> 200000000 satoshi))
```

Our verification technique is based on model-checking the state space of BitML contracts. Since this state space is infinite, before running the model-checker we reduce it to a finite-state one, by exploiting the abstraction in [24]. This abstraction resolves the three sources of infiniteness of the concrete semantics of BitML: the passing of time, the advertisement/stipulation of contracts, and the off-contract bitcoin transfers. To obtain a finite-state system, the abstraction: (i) quotients time in a finite number of time intervals, (ii) disables the advertisement of new contracts, and (iii) limits the off-contract operations to those for transferring funds to contracts and for destroying them. This abstraction is shown in [24] to enjoy a strict correspondence with the concrete BitML semantics: namely, each concrete step of the contract under analysis is mimicked by an abstract step, and vice versa.

Our tool implements the abstract BitML semantics in Maude, a model-checking framework based on rewriting logic [36]. Maude is particularly convenient for this purpose: we use its equational logic to express structural equivalence between BitML terms, and its conditional rewriting rules to encode the abstract semantics of BitML. In this way, we naturally obtain an executable abstract semantics of BitML. Once a BitML contract is translated in Maude, we use the Maude LTL model-checker [40] to verify the required security properties, under the strategies specified by the user. The various forms of liquidity are also translated to corresponding LTL formulae. The computational soundness of the BitML compiler guarantees that the properties verified by the model checker are preserved when executing the contract on Bitcoin.



## 6.3 Compiling BitML to Bitcoin

Our compiler operates in two phases: first, it translates BitML contracts into Balzac [81], an abstraction layer over Bitcoin transactions based on the formal model of [11]; then, it translates Balzac transactions into standard Bitcoin transactions.

The compiler from BitML to Balzac implements the algorithm in [22], extending it with transaction fees. In particular, the compiler guarantees that each transaction contains enough fees to be publishable in the blockchain.

The compiler from Balzac to Bitcoin produces *standard* Bitcoin transactions [85]: this is crucial since non-standard ones are discarded by the Bitcoin network. To this purpose, Balzac produces standard output scripts of the form “Pay to Public Key Hash” (P2PKH) or “Pay to Script Hash” (P2SH). P2PKH is used for encoding signature verification (e.g., to redeem the deposit obtained by a `withdraw`), while P2SH is used for complex redeeming conditions (e.g., to check that the revealed secret matches the committed hash). Since Bitcoin requires that all the values pushed by standard scripts fit within 520 bytes, our compiler checks that this constraint is satisfied for each generated script. Balzac outputs serialized raw transactions, which can be directly broadcast to the Bitcoin network.

## 6.4 Evaluation

To evaluate our toolchain, we use a benchmark of representative use cases, including financial contracts [68, 28], auctions, lotteries [5, 62] and gambling games [91]. For each contract in the benchmark, we display in Table 6.1 the number  $N$  of involved participants, the number  $T$  of transactions obtained by the compiler, and the verification time  $V$  for checking liquidity. For uniformity, in the performance evaluation we focus on liquidity (other queries to verify the functional correctness of the contracts in Table 6.1 are on the repository). We carry out our experiments on a PC equipped with a hexa-core Intel Core i7-7800X CPU @ 3.50GHz, and 64GB of RAM. The participants’ strategies are constrained only as needed to ensure liquidity: in most cases, we do not put any constraints at all. For the contracts which involve predicates on secrets (e.g., all the lotteries), in principle one would need to check liquidity against all the possible choices of secrets. To make verification feasible, since each contract only checks a finite set of predicates, we partition the infinite choices of secrets into a finite set of regions, and sample one choice from each region. In this way,

Contract	$N$	$T$	$V$
Mutual timed commitment	2	15	83ms
Mutual timed commitment	3	34	103ms
Mutual timed commitment	4	75	454ms
Mutual timed commitment	5	164	13s
Escrow (early fees)	3	12	8s
Escrow (late fees)	3	11	3.4s
Zero Coupon Bond	3	8	86ms
Coupon Bond	3	18	1.3s
Future( $C$ )	3	$5 + T_C$	$80\text{ms} + V_C$
Option( $C, D$ )	3	$14 + T_C + T_D$	$90\text{ms} + V_C + V_D$
Lottery ( $O(N^2)$ collateral)	2	15	427ms
Lottery (0 collateral)	2	8	142ms
Lottery (0 collateral)	4	587	67h
Rock-Paper-Scissors	2	23	781ms
Morra game	2	40	674ms
Shell game	2	23	27s
Auction (2 turns)	2	42	3.3s

**Table 6.1:** Benchmarks for the BitML toolchain.

the liquidity check is performed a finite number of times, ensuring that the verifier explores every reachable state of the contract. For instance, in the 4-players lottery we explore  $3^4$  regions, which explains the 67 hours needed to verify its liquidity. Another feature which significantly affects the verification time is the fact that we are considering *all* the possible strategies of *all* the participants.

We also evaluate the same set of benchmark contract with the variant of the analyser that supports recursion. The verification time for all the benchmarks is in the order of milliseconds on a consumer-grade laptop. This is due to the fact that the semantics is more abstract (e.g. it does not consider secrets values) and does not support LTL queries or other variant of liquidity other than the plain one.

The only work against which we can compare the performance of our tool is [3], which models Bitcoin contracts in Uppaal, a model-checking framework based on Timed Automata. The most complex contract modelled in [3] is the mutual timed commitment with 2 participants: this

requires  $\sim 30$ s to be verified in Uppaal, while our tool verifies the same property in  $< 100$ ms. This speedup is due to the higher abstraction level of BitML over [3], which operates at the (lower) level of Bitcoin transactions.

One of the main difficulties that we have encountered in developing contracts is that some complex BitML specifications can not be compiled to Bitcoin, because Bitcoin has a 520-byte limit on the size of each value pushed to the evaluation stack [76]. In some cases, we managed to massage the BitML contract so to make its compilation respect the 520-byte constraint. For instance, a common pattern that easily violates the 520-byte constraint is the following:

```
(choice (revealif (b) (pred (p0)) (C0))
        (revealif (b) (pred (p1)) (C1))
        (after T      (C2)))
```

The `choice` is compiled into a transaction whose redeem script encodes the disjunction of *three* logical conditions, corresponding to the three branches of the `choice`. Depending on the predicates `p0` and `p1`, and on the number of participants in the contract, this script may violate the 520-byte constraint. A workaround is to rewrite the pattern above into the following one:

```
(choice (revealif (b) (pred (p0)) (C0))
        (after T (tau (choice
                      (revealif (b) (pred (p1)) (C1))
                      (after T1      (C2)))))))
```

In this case the compilation includes two transactions, corresponding to the two `choice`s. The scripts of these transactions encode the disjunction of *two* logical conditions, corresponding to the two branches of the `choice`s. Using this workaround we have managed to compile the 4-players lottery into standard transactions, at the price of increasing the number of transactions (587 for the standard version *vs.* 138 for the nonstandard one). Similar techniques (e.g. simplification of predicates) allowed us to compile all the contracts in Table 6.1 into standard Bitcoin transactions.

In general, the 520-byte constraint intrinsically limits the expressiveness of Bitcoin contracts: for instance, since public keys are 33 bytes long, a contract which needs to simultaneously verify 15 signatures can not be implemented using standard transactions.

## 6.5 BitML toolchain tutorial

The following Section describes how to exploit the toolchain to design, compile and verify BitML contracts. We proceed in an incremental fashion: each step introduces a new feature through a small, self-enclosed use-case. Some lines have been shortened due to space constraints, refer to [94] for the full examples.

### 6.5.1 Developing BitML contracts in a nutshell

BitML contracts allow two or more participants to exchange their bitcoins according to complex pre-agreed rules.

The first step in designing a BitML contract is to declare the involved participants. For instance, we can declare three participants "A", "B" and "C" as follows:

```
(participant "A" "029c5f6...095f547799a6289fbc90c70209c1cced")
(participant "B" "0316589...a876e4f7315fa20a07114d5fb8866553")
(participant "C" "03c7e15...928d986b26fe0dc2533f304c19268a2f")

(debug-mode)
```

Each participant is associated to a public key: for instance, "A" has the public key "029c...cced". The command `(debug-mode)` is needed to generate auxiliary keys which are used by the BitML compiler, instead of declaring them as you are supposed to when executing a contract in a real life scenario.

**Simple payments** Assume that "A" simply wants to donate 1 $\text{B}$  to "B". To this purpose, "A" must first declare that she owns a transaction output with 1 $\text{B}$ . We can define this transaction output as follows:

```
(define (txA) "tx:0200000000...c6bdb51600@1")
```

In the definition above, "0200000000...c6bdb51600" are the bytes of the serialized transaction, and the trailing "@1" is the index of the output.

The contract advertised by "A" is the following:

```
(contract
  (pre (deposit "A" 1 (ref (txA))))
  (withdraw "B"))
```

The contract precondition (`pre (deposit "A" 1 (ref (txA)))`) declares that "A" agrees to transfer the  $1\text{B}$  referenced by the transaction output `txA` under the control of the contract. The actual contract is (`withdraw "B"`): this just transfers the funds deposited into the contract to "B".

In the previous contract, the initial deposit has been provided by a transaction output; more in general, a contract can gather money from more than one transaction. For instance, assume that another participant "C" wants to contribute  $1\text{B}$  to the donation. The contract precondition is modified as follows:

```
(contract
  (pre (deposit "A" 1 (ref (txA)))
        (deposit "C" 1 "tx:0200000...3b6cdef8ac0000000000"))
  (withdraw "B"))
```

**Procrastinating payments** Assume now that "A" wants to donate  $1\text{B}$  to "B", but only after a certain time  $t$ . For instance, the  $1\text{B}$  could be a birthday present to be withdrawn only after the birthday date; or the amount of a rent to the landlord, to be paid only after the 1st of the month. We represent the time  $t$  as a block height. For instance, we set  $t$  to 500000 (note that the block at this height was actually mined on 2017-12-18).

To craft this contract we use the primitive `after height contract`, which locks the `contract` until the block at the given `height` is appended to the blockchain. We also reuse the transaction output `txA` from the previous example:

```
(define (t) 500000)

(contract
  (pre (deposit "A" 1 (ref (txA))))
  (after (ref (t)) (withdraw "B")))
```

This contract ensures that only after the block at height  $t$  has been appended to the blockchain, "B" will be able to redeem  $1\text{B}$  from the contract, by performing the action (`withdraw "B"`).

The following contract allows "A" to recover her deposit if "B" has not withdrawn within a given deadline  $t1 > t$ :

```
(define (t) 500000)
(define (t1) 510000)

(contract
```

```
(pre (deposit "A" 1 (ref (txA))))

(choice
  (after (ref (t)) (withdraw "B"))
  (after (ref (t1)) (withdraw "A"))))
```

The contract allows two (mutually exclusive) behaviours: either "A" or "B" can withdraw 1 $\text{\textcircled{B}}$ . Before the deadline  $t$  no one can withdraw; after  $t$  (but before  $t1$ ) only "B" can withdraw, while after the  $t1$  both withdraw actions are enabled, so the first one who performs their withdraw will get the money.

**Authorizing payments** Assume that "A" is willing to pay 1 $\text{\textcircled{B}}$  to "A", but only if an "Oracle" gives his authorization. We can use the authorization primitive `auth Participant Contract` as follows:

```
(contract
  (pre (deposit "A" 1 (ref (txA))))
  (auth "Oracle" (withdraw "B")))
```

This contract ensures that (`withdraw "B"`) is performed whenever "Oracle" authorizes it.

We can play with authorizations and summations to construct more complex contracts. For instance, assume we want to design an *escrow* contract, which allows "A" to buy an item from "B", authorizing the payment only after she gets the item. Further, "B" can authorize a full refund to "A", in case there is some problem with the item. A naïve attempt to model this contract is the following:

```
(define (Naive-escrow)
  (choice
    (auth "A" (withdraw "B"))
    (auth "B" (withdraw "A"))))
```

If both participants are honest, everything goes smoothly: when "A" receives the item, she authorizes the payment to "B", otherwise "B" authorizes the refund. The problem with this contract is that, if neither "A" nor "B" give the authorization, the money in the contract is frozen. To cope with this issue, we can refine the escrow contract, by introducing a trusted arbiter "O" which resolves the dispute:

```

(define (Oracle-escrow)
  (choice
    (auth "A" (withdraw "B")) ; same as
    (auth "B" (withdraw "A")) ; Naive-escrow
    (auth "0" (withdraw "A"))
    (auth "0" (withdraw "B"))))

(contract
  (pre (deposit "A" 1 (ref (txA))))
  (ref (Oracle-escrow)))

```

The last two branches are used if neither "A" nor "B" give their authorizations: in this case, the arbiter chooses whether to authorize "A" or "B" to redeem the deposit.

**Splitting deposits** In all the previous examples, the deposit within the contract is transferred to a single participant. More in general, deposits can be split in many parts, to be transferred to different participants. For instance, assume that "A" wants her 1 $\text{B}\text{\$}$  deposit to be transferred in equal parts to "B1" and to "B2". We can model this behaviour as follows:

```

(define (Pay-split)
  (split
    (0.5 -> (withdraw "B1"))
    (0.5 -> (withdraw "B2"))))

```

The split construct splits the contract in two or more parallel subcontracts, each with its own balance. Of course, the choice of their balances must be less than or equal to the deposit of the whole contract.

We can use split together with the other primitives presented so far to craft more complex contracts. For instance, assume that "A" wants to pay 0.9 $\text{B}\text{\$}$  to "B", routing the payment through an intermediary "I" who can choose whether to authorize it (in this case retaining a 0.1 $\text{B}\text{\$}$  fee), or not. Since "A" does not trust "I", she wants to use a contract to guarantee that: (i) if "I" authorizes the payment then 0.9 $\text{B}\text{\$}$  are transferred to "B"; (ii) otherwise, "A" does not lose money.

We can model this behaviour as follows:

```

(contract
  (pre (deposit "A" 1 (ref (txA))))
  (choice
    (auth "I" (split (0.1 -> (withdraw "I"))
                    (0.9 -> (withdraw "B"))))
    (after (ref (d)) (withdraw "A"))))

```

The first branch can only be taken if "I" authorizes the payment: in this case, "I" gets his fee, and "B" gets his payment. Instead, if "I" denies his authorization, then "A" can redeem her deposit after block height  $d$ .

**Volatile deposits** So far, we have seen participants using persistent deposits, that are assimilated by the contract upon stipulation. Besides these, participants can also use volatile deposits, which are not assimilated upon stipulation. For instance:

```
(pre (deposit "A" 1 (ref (txA1)))
      (vol-deposit "A" x 1 (ref (txA2))))
```

gives "A" the possibility of contributing 1 $\text{₿}$  during the contract execution. However, "A" can choose instead to spend her volatile deposit outside the contract. The variable  $x$  is a handle to the volatile deposit, which can be used as follows:

```
(define (Pay?)
  (put (x) (withdraw "B")))
```

Since  $x$  is not paid upfront, there is no guarantee that  $x$  will be available when the contract demands it, as "A" can spend it for other purposes.

Volatile deposits can be exploited within more complex contracts, to handle situations where a participant wants to add some funds to the contract. For instance, assume a scenario where "A1" and "A2" want to give "B" 2 $\text{₿}$  as a present, paying 1 $\text{₿}$  each. However, "A2" is not sure a priori she will be able to pay, because she may need her 1 $\text{₿}$  for more urgent purposes: in this case, "A1" is willing to pay an extra bitcoin. We can model this scenario as follows: "A1" puts 2 $\text{₿}$  as a persistent deposit, while "A2" makes available a volatile deposit  $x$  of 1 $\text{₿}$ :

```
(contract
  (pre (deposit "A1" 2 (ref (txA1)))
        (vol-deposit "A2" x 1 (ref (txA2))))
  (choice
    (put (x) (split (2 -> (withdraw "B"))
                   (1 -> (withdraw "A1"))))
    (after 700000 (withdraw "B"))))
```

In the first branch, "A2" puts 1 $\text{₿}$  in the contract, and the balance is split between "B" (who takes 2 $\text{₿}$ , as expected), and "A1" (who takes her extra deposit back). The second branch is enabled after  $d$ , and it deals with the case where "A2" has not put her deposit by such deadline. In this case, "B" can redeem 2 $\text{₿}$ , while "A2" loses the extra deposit. Note that, in both cases, "B" will receive 2 $\text{₿}$ .



**Revealing secrets** A useful feature of Bitcoin smart contracts is the possibility for a participant to choose a secret, and unblock some action only when the secret is revealed. Further, different actions can be enabled according to the length of the secret. Secrets must be declared in the contract precondition, as follows:

```
(pre (secret "A" a "f9292914bfd27c426a23465fc122322abdb63b7"))
```

where "A" is the participant who owns the secret, *a* is its *name*, and "f9292914bfd27c426a23465fc122322abdb63b7" is its `hash160` hash. We never denote the value of the secret itself. A basic contract which exploits this feature is the following:

```
(define (PaySecret)
  (revealif (a) (pred (> a 1)) (withdraw "A")))
```

This contract asks "A" to commit to a secret of length greater than one, as stated in the predicate `(pred (> a 1))`. After revealing *a*, it allows "A" to redeem 1 ₿ upon revealing the secret. Until then, the deposit is frozen.

To reveal a secret without imposing a predicate use `(reveal)`. e.g.: `(reveal (a) (withdraw "A"))`

Note that we never refer to the value itself of the secret, rather we use its length. After compiling to Bitcoin, the actual length of the secret will be increased by  $\eta$ , where  $\eta$  is a security parameter, large enough to avoid brute-force preimage attack.

## 6.5.2 Verifying contracts with the BitML toolchain

Other than compiling contracts to transactions, the BitML toolchain allows to verify them before their execution.

A desirable property of smart contracts is *liquidity*, which requires that the contract balance is always eventually transferred to some participant. In a non-liquid contract, funds can be frozen forever, unavailable to anyone, hence effectively destroyed. There are many possible flavours of liquidity, depending e.g. on which participants are assumed to be honest, and on which are their strategies.

The toolchain can also verify arbitrary security properties, expressed as LTL queries.

**Liquidity** In the following contract, "A" and "B" contribute 1  $\text{฿}$  each for a donation of 2  $\text{฿}$  to either "C" or "D". We want to check if the contract is liquid or not, without supplying any strategy, i.e. without knowing which branch "A" and "B" will authorize.

This flavour of liquidity is called *strategy-less*. Intuitively, it corresponds to check if the contract is liquid for any possible strategy of any participants, whether they are honest or not.

To check the liquidity of the following contract, we add (check-liquid) at its end.

```
(participant "A" "0339bd7fade9167e09681...2bde1d57247fbe1")
(participant "B" "034a7192e922118173906...da4f85701a5f809")
(participant "C" "034f5ca30056b9dd89132...fea9c3467631e6b")
(participant "D" "037b60c121050e1fa6e7d...da1c63f0e2a157e")

(debug-mode)

(contract
  (pre
    (deposit "A" 1 "txid:2e647d8566f00a08d27...b2965e3501")
    (deposit "A" 1 "txid:625bc69c467b33e2abj...bc71016800"))

  (choice
    (auth "A" "B" (withdraw "C"))
    (auth "A" "B" (withdraw "D")))

  (check-liquid))
```

During the compilation of the contract, the tool-chain checks if it is liquid. The result is printed before the transactions in a comment-box.

```
/*=====
Model checking result for (check-liquid)

Result: false
counterexample({[0 | nil | 'xconf U empty | empty] < ( A, B) :
  withdraw C + (A, B) : withdraw D, 100000000 satoshi > 'xconf,
  'C-LockAuthControl} {[A lock withdraw C in 'xconf][0 | nil | '
  xconf U empty | empty] < Lock((A, B) : withdraw C) + (
  A, B) : withdraw D, 100000000 satoshi >
  'xconf, 'Rif1} {[A lock withdraw D in 'xconf][0 | nil | 'xconf
  U empty | empty] < Lock((A, B) : withdraw C) + Lock((A
  , B) : withdraw D), 100000000 satoshi > 'xconf, '
  Finalize}, {[0 | nil | 'xconf U empty | empty] < Lock((A,
  B) : withdraw C) + Lock((A, B) : withdraw D),
  100000000 satoshi > 'xconf, solution})
Model checking time: 143.0 ms
=====*/
```

The output indicates that the contract is not liquid. In fact, In order to unlock the funds, "A" and "B" must agree on the recipient of the donation, by giving their authorization on the same branch. This contract would be

liquid only by assuming the cooperation between "A" and "B": indeed, "A" alone cannot guarantee that the 2 ₿ will eventually be donated, as "B" can choose a different recipient, or even refuse to give any authorization.

A possible way to modify the contract to handle this situations by adding a timeout branch with (`after 700000 (split (1 -> (withdraw "A"))(1 -> (withdraw "B")))`). The new branch locks the contract until the block number 700000 is appended to the blockchain, modelling a delay. After the corresponding time passes, it unlocks and returns their deposits to "A" and "B".

```
(participant "A" "0339bd7fade9167e09681...2bde1d57247fbe1")
(participant "B" "034a7192e922118173906...da4f85701a5f809")
(participant "C" "034f5ca30056b9dd89132...fea9c3467631e6b")
(participant "D" "037b60c121050e1fa6e7d...da1c63f0e2a157e")

(debug-mode)

(contract
  (pre
    (deposit "A" 1 "txid:2e647d8566f00a08d27...b2965e35@1")
    (deposit "A" 1 "txid:625bc69c467b33e2abj...bc710168@0")))

  (choice
    (auth "A" "B" (withdraw "C"))
    (auth "A" "B" (withdraw "D")))
    (after 700000 (split (1 -> (withdraw "A")) (1 -> (withdraw "B"))))
  )

  (check-liquid))
```

Now the contract is liquid, and the toolchain confirms it.

```
/*=====
Model checking result for (check-liquid)

Result: true
Model checking time: 322.0 ms
=====*/
```

**Liquidity with strategies** In the following contract, "A" can reveal her secret and redeem its deposit. Otherwise, after a certain amount of time the block number 700000 will be appended to the blockchain, "B" can redeem "A"'s deposit, after providing his authorization to do so.

```
(participant "A" "0339bd7fade9167e09681...2bde1d57247fbe1")
(participant "B" "034a7192e922118173906...da4f85701a5f809")
(participant "C" "034f5ca30056b9dd89132...fea9c3467631e6b")
(participant "D" "037b60c121050e1fa6e7d...da1c63f0e2a157e")

(debug-mode)
```

```

(contract
  (pre
    (deposit "A" 1 "txid:2e647d8566f00a08d27...b2965e3501")
    (secret "A" a "f9292914bfd27c426a23465fc122322abdb63b7")))

  (choice
    (reveal (a) (withdraw "A"))
    (auth "B" (after 700000 (withdraw "B"))))

  (check-liquid))

```

Below, we check the strategy-less liquidity, detecting that the contract is not liquid. This is because if neither "A" reveals her secret nor "B" gives his authorization, the funds will be stuck forever.

```

/*=====
  Model checking result for (check-liquid)

Result: false
Secrets: a:1

counterexample({[0 | 700000 | 'xconf U empty | B, A] < B : after
  700000 : withdraw B + put empty reveal a if True . withdraw A,
  100000000 satoshi > 'xconf | {A : a # 1}, 'C-LockAuthRev} {{A
  lock-reveal a}[0 | 700000 | 'xconf U empty | B, A] Lock({A : a
  # 1}) | < B : after 700000 : withdraw B + put empty reveal a
  if True . withdraw A, 100000000 satoshi > 'xconf, 'Rifl} {{B
  lock after 700000 : withdraw B in 'xconf}[0 | 700000 | 'xconf
  U empty | B, A] Lock({A : a # 1}) | < Lock(B : after 700000 :
  withdraw B) + put empty reveal a if True . withdraw A, 100000000
  satoshi > 'xconf, 'Rifl} {{delta 700000}[700000 | nil | 'xconf
  U empty | B, A] Lock({A : a # 1}) | < Lock(B : after 700000 :
  withdraw B) + put empty reveal a if True . withdraw A,
  100000000 satoshi > 'xconf, 'Finalize}, {[700000 | nil | 'xconf
  U empty | B, A] Lock({A : a # 1}) | < Lock(B : after 700000 :
  withdraw B) + put empty reveal a if True . withdraw A,
  100000000 satoshi > 'xconf, solution}})
Model checking time: 104.0 ms
=====*/

```

The BitML toolchain allows us to specify the intended behaviour of a participant, called *strategy*. The security propriety is verified with respect to the specified strategies.

We check if the contract is liquid if the strategy of "A" consists in revealing her secret, expressed by (strategy "A"(do-reveal a)) as parameter of (check-liquid Strategy ...).

We also check the liquidity if "A" authorizes the second branch of the contract, with the strategy (strategy "B"(do-auth (auth "B"(after 700000 (withdraw "B"))))).

```
(participant "A" "0339bd7fade9167e09681...2bdeid57247fbe1")
(participant "B" "034a7192e922118173906...da4f85701a5f809")
(participant "C" "034f5ca30056b9dd89132...fea9c3467631e6b")
(participant "D" "037b60c121050e1fa6e7d...da1c63f0e2a157e")

(debug-mode)

(contract
  (pre
    (deposit "A" 1 "txid:2e647d8566f00a08d27...b2965e35@1")
    (secret "A" a "f9292914bfd27c426a23465fc122322abdb63b7")))

  (choice
    (reveal (a) (withdraw "A"))
    (auth "B" (after 700000 (withdraw "B")))))

  (check-liquid
    (strategy "A" (do-reveal a)))

  (check-liquid
    (strategy "B" (do-auth (auth "B" (after 700000 (withdraw "B"))))))
)
```

For both strategies, the contract is liquid.

```
/*=====
Model checking result for (check-liquid (strategy A (do-reveal a)))

Result: true

/*=====
Model checking result for (check-liquid (strategy B (do-auth (auth B (
  after 700000 (withdraw B))))))

Result: true
Model checking time: 270.0 ms
=====*/
```

**Quantitative liquidity** The previous flavours of liquidity require that no funds remain frozen within the contract. However, in some cases a participant could accept the fact that a portion of the funds remain frozen, especially when these funds would be ideally assigned to other participants.

In the following contract, "A" and "B" put 1 $\text{B}$  each. Each of them will get their own  $\text{B}$  back if they reveal their secret.

```
(participant "A" "0339bd7fade9167e...5bbb84211fd12bdeid57247fbe1")
(participant "B" "034a7192e9221181...5f7c7f403094da4f85701a5f809")

(debug-mode)

(contract
  (pre
```

```

(deposit "A" 1 "txid:2e647d8566f00a08...161c2078963d8deb2965e3501"
)
(deposit "B" 1 "txid:0f795bda36ac661f...69f5fb7ccc4c3d767db9f3401"
)
(secret "A" a "f9292914bfd27c426a23465fc122322abdb63b7")
(secret "B" b "9804ebb0fc4a8329981dd33aaff32b6cb579580a")

(split
  (1 -> (reveal (a) (withdraw "A")))
  (1 -> (reveal (b) (withdraw "B"))))

(check "A" has-more-than 1
  (strategy "A" (do-reveal a))))

```

In this setting, "A" is interested in checking if she will get back her bitcoin, assuming that she reveals her secret. We check it using `(check "A" has-more-than 1 (strategy "A" (do-reveal a)))`.

```

/*=====
Model checking result for (check A has-more-than 1 (strategy A (do-
  reveal a)))

Result: true
Model checking time: 134.0 ms
=====*/

```

**Custom LTL queries** The following contract is a *timed commitment*, where "A" wants to choose a secret `a`, and reveal it before the deadline `d`; if "A" does not reveal the secret within `d`, "B" can redeem the 1  $\text{฿}$  deposit as a compensation.

```

(participant "A" "029c5f6f5ef0095f547799...289fbc90c70209c1cccd")
(participant "B" "022c3afb0b654d3c2b0e2f...2f86f7647fa7b817af30")

(debug-mode)

(define (d) 700000)

(contract
  (pre
    (deposit "A" 1 "txA@0")
    (secret "A" a "f9292914bfd27c426a23465fc122322abdb63b7")))

  (choice
    (reveal (a) (withdraw "A"))
    (after (ref (d)) (withdraw "B"))))

  (check-query
    "[ ] <> (a revealed => A has-deposit >= 100000000 satoshi)")

  (check-query
    "[ ] <> (a revealed \\/\ B has-deposit >= 100000000 satoshi)"))

```

The toolchain allows us to check custom LTL properties, tailored specifically for the contract being verified, using (`check-query "query"`).

In the timed commitment contract, we want the following two properties to be satisfied.

- If "A" reveal her secret, she will get back her deposit. We check this property with (`check-query "[] <> (a revealed => A has-deposit >= 100000000 satoshi)"`).
- Either "B" gets to know the secret, or he will get the bitcoin as compensation. We check this property with (`check-query "[] <> (a revealed \\ / B has-deposit >= 100000000 satoshi)"`).

Due to the internal representation of numbers in the model checker, all  $\text{B}\$$  values have to be expressed in *satoshi* when checking custom LTL queries.

The result is true for both queries:

```

/*=====
Model checking result for (check-query [] (a revealed => <> A has-
  deposit >= 100000000 satoshi))

Result: true

/*=====
Model checking result for (check-query [] <> (a revealed \\ / B has-
  deposit >= 100000000 satoshi))

Result: true
Model checking time: 408.0 ms
=====*/

```

The first LTL property has the same semantic as checking the quantitative liquidity of 1  $\text{B}\$$  if the strategy of "A" is to reveal her secret, or (`check "A" has-more-than 1 (strategy "A"(do-reveal a))`). Instead, the second LTL property cannot be expressed as a combination of liquidity and strategies.

Other than `revealed` and `has-deposit >=`, you can express your LTL properties with `has-deposit`, and `has-deposit <=`.





## Chapter 7

# On-chain Fungible Tokens on Bitcoin

One of the main applications of blockchain technologies is the exchange of custom crypto-assets, called *tokens*. Token transfers currently involve  $\sim 50\%$  of the transactions on the Ethereum blockchain [107], and they are at the basis of many protocols built on top of that platform [6, 44]. Broadly, tokens are classified as *fungible* or *non-fungible*. Fungible tokens can be split into smaller units: different units of the same token can be used interchangeably. Further, users can join units of the same fungible token, and exchange them with other crypto-assets. Instead, non-fungible tokens cannot be split or joined.

Historically, the first implementations of tokens were developed before Ethereum, on top of Bitcoin. Some of them (e.g., [105]) used small bitcoin fractions to represent the token value; some others (e.g., [118, 100, 103]) embedded the token value in other transaction fields [13], to cope with the fluctuating bitcoin price. All these implementations have a common drawback: the correctness of the token actions is *not* guaranteed by the consensus protocol of the blockchain. In fact, the blockchain is used just to notarize the actions that manipulate tokens, but not to check that these actions are actually permitted. Typically, the owners of these tokens must resort to *off-chain* mechanisms (e.g., trusted authorities) to have some guarantees on the correct use of tokens, e.g. that they are not double-spent, or that distinct tokens are not joined.

By contrast, modern blockchain platforms support *on-chain* tokens, whose correctness is guaranteed by the consensus protocol of the blockchain. Some platforms (e.g., Algorand [79]) natively support to-

kens, while some others (e.g., Ethereum) encode them as smart contracts. Bitcoin, instead, does not support tokens natively, and its limited script language is not expressive enough to implement them as smart contracts. Since adding native tokens to Bitcoin appears to be out of reach, given the resilience of the Bitcoin community to radical changes [102], the only viable alternative is to devise a small, efficient extension of the script language which increases the expressiveness of Bitcoin enough to support tokens.

We implement a token contract as a single, succinct script which exploits neighbourhood covenants. We define a symbolic model of token actions, and a computational model, where performing these actions corresponds to appending transactions to the Bitcoin blockchain.

## 7.1 Overview of the approach

In this section we summarize our approach: in particular, we sketch our implementation of Bitcoin tokens, motivating the use of neighbourhood covenants to guarantee their security.

**Tokens** We propose a symbolic model of fungible tokens. Since non-fungible tokens are the special case of fungible ones where each token is generated exactly in one unit, hereafter we consider the general case of fungible tokens. The basic element of our model is the *deposit*, i.e. a term of the form:

$$\langle \mathbf{A}, v : \tau \rangle_x \quad (v \in \mathbb{N})$$

which represents the fact that a user  $\mathbf{A}$  owns  $v$  units of a token  $\tau$ , where  $\tau$  may denote either user-defined tokens or bitcoins ( $\mathfrak{B}$ ). The index  $x$  uniquely identifies the term within a *configuration*, i.e. a composition of deposits, e.g.:

$$\langle \mathbf{A}, 1 : \tau \rangle_x \mid \langle \mathbf{A}, 2 : \tau \rangle_y \mid \langle \mathbf{B}, 3 : \mathfrak{B} \rangle_z$$

We define a few actions to mint and manipulate tokens. First, any user  $\mathbf{A}$  can mint  $v$  units of a new token, spending a deposit of  $0 \mathfrak{B}$ . Performing this action (say, with  $v = 10$ ) is modelled as a state transition, whose labels records the performed action:

$$\langle \mathbf{A}, 0 : \mathfrak{B} \rangle_{x_0} \xrightarrow{gen} \langle \mathbf{A}, 10 : \tau \rangle_{x_1} \quad (7.1)$$

where the identifier  $x_1$  of the new deposit and the identifier  $\tau$  of the minted token are *fresh*. After performing the action,  $\mathbf{A}$  owns a deposit of ten units of the token  $\tau$ . As said before, one of the peculiar properties of

fungible tokens is that they can be *split*. When splitting her deposit in two smaller deposits, **A** can choose the owner of one of the new deposits, e.g.:

$$\langle \mathbf{A}, 10 : \tau \rangle_{x_1} \xrightarrow{\text{split}} \langle \mathbf{A}, 8 : \tau \rangle_{x_2} \mid \langle \mathbf{B}, 2 : \tau \rangle_{x_3} \quad (7.2)$$

A user can transfer the ownership of any of her deposits to another user. For instance, **A** can *give* her deposit  $x_2$  to **B**:

$$\langle \mathbf{A}, 8 : \tau \rangle_{x_2} \xrightarrow{\text{give}} \langle \mathbf{B}, 8 : \tau \rangle_{x_4} \quad (7.3)$$

After that, **B** owns a total of 10 units of  $\tau$  in two separate deposits, one with 8 units, and the other one with 2 units. This reflects the UTXO nature of Bitcoin: by contrast, in account-based blockchains like Ethereum, **B** would have a single account storing 10 units of  $\tau$ . Now, **B** can *join* his two deposits, obtaining a single deposit with 10 units of  $\tau$ . When performing the *join* action, **B** can also choose the owner of the new deposit, in this case transferring it back to **A**:

$$\langle \mathbf{B}, 8 : \tau \rangle_{x_4} \mid \langle \mathbf{B}, 2 : \tau \rangle_{x_3} \xrightarrow{\text{join}} \langle \mathbf{A}, 10 : \tau \rangle_{x_5} \quad (7.4)$$

A crucial property of the *join* operation is that only deposits of the same token can be joined together. Thus, two deposits of  $\tau$  and  $\tau'$  with  $\tau \neq \tau'$  cannot be joined:

$$\langle \mathbf{B}, 8 : \tau \rangle_{x_4} \mid \langle \mathbf{A}, 2 : \tau' \rangle_{x_6} \not\xrightarrow{\text{join}}$$

In this configuration, if both **A** and **B** agree, they can *exchange* the ownership of their tokens:

$$\langle \mathbf{B}, 8 : \tau \rangle_{x_4} \mid \langle \mathbf{A}, 2 : \tau' \rangle_{x_6} \xrightarrow{\text{xchg}} \langle \mathbf{A}, 8 : \tau \rangle_{x_7} \mid \langle \mathbf{B}, 2 : \tau' \rangle_{x_8}$$

The *xchg* operation also supports the exchange between bitcoins and other tokens, representing the trade of tokens. For instance, **A** can buy 2 units of  $\tau'$  from **B** for 1 $\mathfrak{B}$ :

$$\langle \mathbf{B}, 2 : \tau' \rangle_{x_8} \mid \langle \mathbf{A}, 1 : \mathfrak{B} \rangle_{x_9} \xrightarrow{\text{xchg}} \langle \mathbf{A}, 2 : \tau' \rangle_{x_{10}} \mid \langle \mathbf{B}, 1 : \mathfrak{B} \rangle_{x_{11}}$$

**Bitcoin** Although Bitcoin does not support user-defined tokens, it implements all the operations discussed above on its native crypto-currency. Intuitively, each deposit corresponds to a transaction output, and performing actions corresponds to appending a suitable transaction that redeems it.

For instance, minting bitcoins is obtained through coinbase transactions, which are used in Bitcoin to pay rewards to miners. We represent a coinbase transaction as follows:

$T_{7.1}$
$\text{in}(1): \perp$ $\text{wit}(1): \perp$ $\text{out}(1): \{\text{scr} : \text{versig}(pk_A, \text{rtx.wit}), \text{val} : 10\text{฿}\}$

In general, the  $\text{in}$  field points to a previous transaction on the blockchain, that the current one is trying to *spend*. Here, the “undefined” value  $\perp$  characterizes  $T_{7.1}$  as a coinbase, since it mints bitcoins without spending any transaction. The  $\text{out}$  field is a record, where  $\text{scr}$  is a *script*, and  $\text{val}$  is the amount of bitcoins that will be redeemed by a subsequent transaction which points to  $T_{7.1}$  and satisfies its script. Here, the script  $\text{versig}(pk_A, \text{rtx.wit})$  verifies a signature on the redeeming transaction ( $\text{rtx}$ , excluding its  $\text{wit}$  field) against  $A$ ’s public key  $pk_A$ . This signature is retrieved from the  $\text{wit}$  field of  $\text{rtx}$ . Since  $A$  is the only user who can redeem  $T_{7.1}$ , we can say that  $T_{7.1}$  is the *computational counterpart* of the deposit  $\langle A, 10 : \text{฿} \rangle_{x_1}$ .

To perform the *split* action (7.2) on  $\tau = \text{฿}$ , we can spend  $T_{7.1}$  with a transaction  $T_{7.2}$  with *two* outputs:

$T_{7.2}$
$\text{in}(1): (T_{7.1}, 1)$ $\text{wit}(1): \text{sig}_{sk_A}(T_1)$ $\text{out}(1): \{\text{scr} : \text{versig}(pk_A, \text{rtx.wit}), \text{val} : 8\text{฿}\}$ $\text{out}(2): \{\text{scr} : \text{versig}(pk_B, \text{rtx.wit}), \text{val} : 2\text{฿}\}$

The first output, that we denote by  $(T_{7.2}, 1)$ , corresponds to the deposit  $\langle A, 8 : \text{฿} \rangle_{x_2}$  in (7.2). Instead, the output  $(T_{7.2}, 2)$  corresponds to  $\langle B, 2 : \text{฿} \rangle_{x_3}$ . These outputs can be spent independently. For instance, performing the *give* action in (7.3) corresponds to appending a transaction which spends  $(T_{7.2}, 1)$ :

$T_{7.3}$
$\text{in}(1): (T_{7.2}, 1)$ $\text{wit}(1): \text{sig}_{sk_A}(T_{7.3})$ $\text{out}(1): \{\text{scr} : \text{versig}(pk_B, \text{rtx.wit}), \text{val} : 8\text{฿}\}$

At this point, we have two unspent outputs on the blockchain:  $(T_{7.2}, 2)$  and  $(T_{7.3}, 1)$ . We can perform the *join* action in (7.4) by spending both of

them *simultaneously* with the following transaction, which has *two* inputs:

$T_{7.4}$	
$\text{in}(1): (T_{7.2}, 2)$	$\text{in}(2): (T_{7.3}, 1)$
$\text{wit}(1): \text{sig}_{sk_B}(T_{7.4})$	$\text{wit}(2): \text{sig}_{sk_B}(T_{7.4})$
$\text{out}: \{\text{scr} : \text{versig}(pk_B, \text{rtx.wit}), \text{val} : 10\text{B}\}$	

**Implementing Bitcoin tokens with covenants** Although the Bitcoin script language is a bit more flexible than shown above, it does not allow to implement *on-chain* tokens. One of the first techniques to embed on-chain tokens in an *extended* version of Bitcoin was described in [63]. The technique relies on *covenants*, an extension of Bitcoin scripts which allows transactions to constrain the scripts of the redeeming ones. Roughly, a transaction output containing the script:

$$e \text{ and } \text{verrec}(\text{rtxo}(n))$$

where  $e$  is an arbitrary script, can only be redeemed by a transaction which makes  $e$  evaluate to true, and whose script in the  $n$ -th output is syntactically equal to  $e \text{ and } \text{verrec}(\text{rtxo}(n))$ . Using covenants, we can mint a token by appending the transaction  $T$  below, where the extra field **arg** is syntactic sugar for a sequence of values accessible by the script:

$T$	
...	
$\text{out}(1): \{\text{arg} : pk_A,$	
$\text{scr} : \text{versig}(\text{ctxo.arg}, \text{rtx.wit}) \text{ and}$	// verify signature
$\text{rtxo}(1).\text{val} = 1 \text{ and}$	// preserve value
$\text{verrec}(\text{rtxo}(1)),$	// preserve script
$\text{val} : 1\text{B}\}$	

The **arg** field identifies **A** as the owner of the token: to transfer the ownership to **B**, **A** must spend  $T$  with a transaction  $T'$ , setting its **arg** to **B**'s public key. For this to be possible,  $T'$  must satisfy the conditions specified in  $T$ 's script: (i) the **wit** field must contain the signature of the current owner; (ii) the output at index 1 must have  $1\text{B}$  value, to preserve the value of the token; (iii) the script at index 1 in  $T'$  must be equal to that in  $T$ . Once  $T'$  is on the blockchain, **B** can transfer the token to another user, by appending a transaction which redeems  $T'$ .

Note that the transaction  $T$  above actually mints a *non-fungible* token, which can be transferred from one user to another, but whose value cannot be *split* (further, the token has a subtle flaw related to *join* actions: we will

say more on this). The first step to turn the token into a fungible one is to support the *split* action. We can achieve this by adding a second element to the *arg* sequence, to represent the number of token units deposited in the transaction output. Using the notation  $w.i$  to access the  $i$ -th element of a sequence  $w$  (for  $1 \leq i \leq |w|$ ), we can implement a splittable token as follows:

$T_{split}$
<pre> ... out(1): {arg : pk<sub>A</sub> v,          scr : versig(ctxo.arg.1, rtx.wit) and               rtxo(1).arg.2 + rtxo(2).arg.2 = ctxo.arg.2 and               verrec(rtxo(1)) and verrec(rtxo(2)) and               outlen(rtx) = 2          val : ... }</pre>

The last two lines of the script ensure that any transaction which redeems  $T_{split}$  has exactly two outputs, each one with the same script of  $T_{split}$ . The second line ensures that the *split* preserves the number of token units (here, *val* is immaterial).

Now, let  $e_{split}$  be the script used in  $T_{split}$ . To extend the token with the *join* action, first we need to add a third element to the *arg* sequence, to encode the action performed by a transaction (say,  $G$  for *gen*,  $S$  for *split*, and  $J$  for *join*). The extended script could have the following form:

$$e \triangleq \text{if rtxo}(1).\text{arg}.3 = S \text{ then } e_{split} \text{ else } e_{join}$$

where  $e_{join}$  implements the join functionality, i.e.: (i) verify the signature on the redeeming transaction; (ii) check that the redeeming transaction has exactly *two* inputs and one output; (iii) ensure that the token units are preserved; (iv) ensure that the joined transactions represent units of the *same* token.

For instance, consider the transactions in Figure 7.1, where  $T_4$  and  $T_3$  represent, respectively, the deposits  $\langle B, 8 : \tau \rangle_{x_4}$  and  $\langle B, 2 : \tau \rangle_{x_3}$ . To perform the *join* action in (7.4), we must spend  $T_4$  and  $T_3$  with the transaction  $T_5$ : this requires to satisfy the script  $e$  in  $T_4$  and  $T_3$ . For condition (iii), the script must ensure that the 10 token units redeemed by  $T_5$  are the sum of the 8 units in  $T_4$  and the 2 units in  $T_3$ . For condition (iv), the script in  $T_4$  should check that it is the same as that in  $T_3$ , and *viceversa*. Hence, to implement conditions (iii)-(iv), the script in a transaction output must be able to access the fields in its *sibling*, i.e. the transaction output which is redeemed together (e.g.,  $(T_3, 1)$  is the sibling

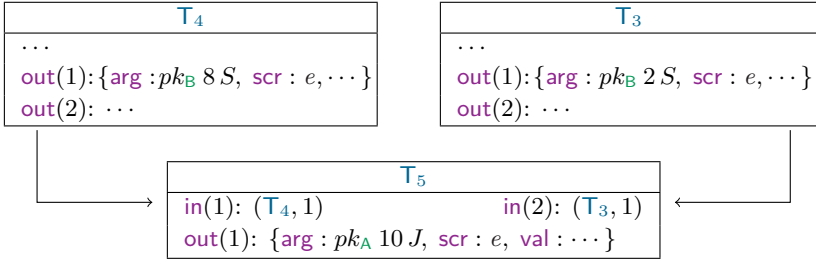


Figure 7.1: A transaction  $T_5$  attempting to join  $T_4$  and  $T_3$ .

of  $(T_4, 1)$  when appending  $T_5$ ). However, neither Bitcoin nor its extensions with covenants [18, 63, 66, 70] allow scripts to access the siblings.

**An insecure implementation of join** To implement the *join* action, we start by extending Bitcoin scripts with an operator to access the sibling transaction outputs:

$$stxo(n) \triangleq \text{output redeemed by the } n\text{-th input of rtx}$$

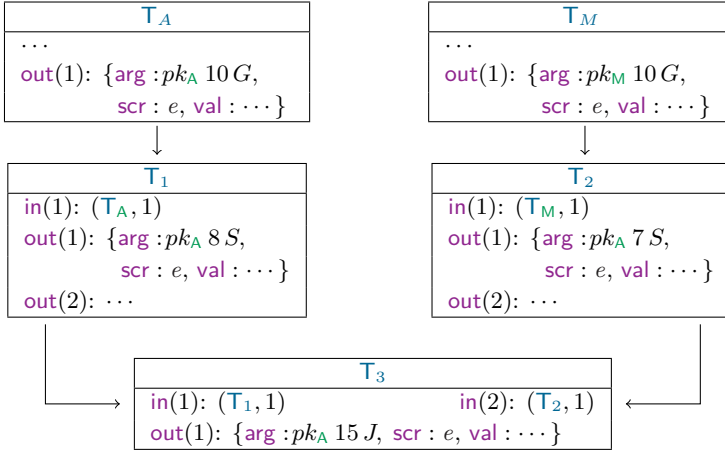
Using this new operator, we can encode the conditions (iii) and (iv) in  $e_{join}$  as follows:

$$rtxo(1).arg.2 = stxo(1).arg.2 + stxo(2).arg.2 \quad \text{(iii)}$$

$$verrec(stxo(1)) \text{ and } verrec(stxo(2)) \quad \text{(iv)}$$

Although this implementation of  $e_{join}$  correctly encodes the conditions, it introduces a security vulnerability: an adversary can join two deposits of *different* tokens. The attack is exemplified in Figure 7.2. The transactions  $T_A$  and  $T_M$  mint 10 units of *different* tokens, and transaction  $T_3$  joins them into a single deposit of the *same* token. Ideally, to counter this attack,  $e_{join}$  should check not only the sibling, but also its ancestors until the minting transaction, and verify that it corresponds to the minting ancestor of the current transaction output. Although this would be possible by adding script operators that can go up the transaction graph at an arbitrary depth, this would be highly inefficient from the point of view of miners, who should record the *whole* transaction graph, instead of just the set of unspent transactions (UTXO).

**A secure implementation with neighbourhood covenants** To address this issue, we use an operator which can go up the transaction graph



**Figure 7.2:** A join attack merging two different tokens.

only one level, introduced by neighborhood covenants in Chapter 4. By exploiting this new covenant, we can thwart the *join* attack of Figure 7.2, and eventually obtain a secure and efficient implementation of fungible tokens. We now sketch the script  $e_{\text{TOK}}$  which implements tokens. First, we add a fourth element to the **arg** sequence, to record in each transaction output the identifier of the token deposited in that output. As an identifier, we use the hash of the *parent* of the minting transaction, which we access through the script  $\text{txid}(\text{ptxo}(1))$ . When we evaluate the script contained in the minting transaction (e.g.,  $T_A$  and  $T_M$  in Figure 7.2), we require that its **arg** field actually contains the token identifier:

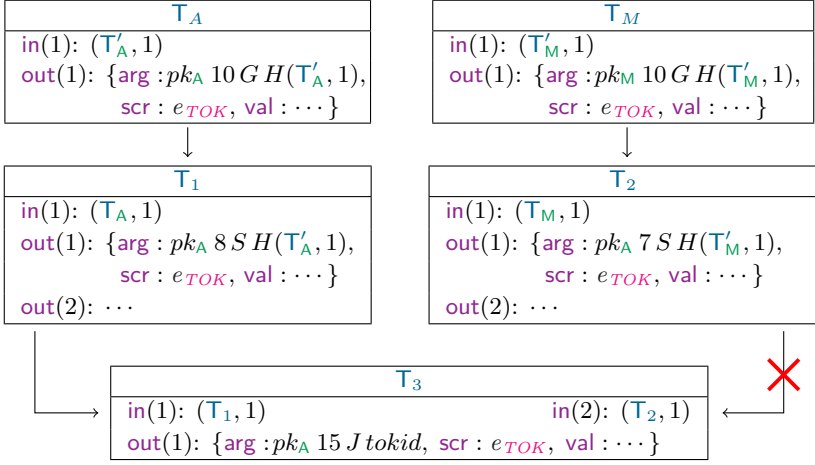
$$e_{\text{gen}} \triangleq \dots \text{ and } \text{ctxo.arg.4} = \text{txid}(\text{ptxo}(1))$$

Then, in the sub-scripts corresponding to all other token actions, we check that the redeeming transaction preserves the token identifier. E.g., in the *join* sub-script, besides checking conditions (iii) and (iv) as shown before, we add the condition:

$$e_{\text{join}} \triangleq \dots \text{ and } \text{ctxo.arg.4} = \text{rtxo}(1).\text{arg.4}$$

In Figure 7.3 we show how this resolves the attack of Figure 7.2. In order to append the malicious *join* transaction  $T_3$ , we must satisfy the script  $e_{\text{TOK}}$  in both  $T_1$  and  $T_2$ . These scripts check that the *tokid* in the redeeming transaction  $T_3$  is equal to the identifiers of the two branches,





**Figure 7.3:** Thwarting the join attack.

$H(T'_A, 1)$  and  $H(T'_M, 1)$ : by collision resistance of the hash function, this is not possible.

The discussion above shows how to counter an attack which attempts to *join* different tokens. However, the adversary could devise more ingenious attacks, e.g. forging units of an existing token. For instance, if  $T_M$  in Figure 7.3 were storing  $\mathfrak{B}$ , its owner could spend it to create a transaction  $T_2$  with *arbitrary* scripts and arguments. In particular,  $T_2$  could use  $e_{\text{TOK}}$  and any token identifier in  $\text{arg}.4$ , e.g.,  $H(T'_A, 1)$ , effectively *forging* new units of a pre-existing token. Our full  $e_{\text{TOK}}$  script exploits neighbourhood covenants to prevent these kinds of attacks as well.

## 7.2 A symbolic model of tokens

Let  $A, B, \dots$  range over *users*, and let  $\tau, \tau', \dots$  range over *tokens*, encompassing both user-defined ones and bitcoins ( $\mathfrak{B}$ ). A term  $\langle A, v : \tau \rangle_x$  represents a *deposit* of  $v \in \mathbb{N}$  units of the token  $\tau$  owned by  $A$  (the index  $x$  is a unique identifier of the deposit). A term  $A \triangleright_x \alpha$  represents  $A$ 's *authorization* to perform the *action*  $\alpha$  on the deposit  $x$ . The possible token actions are the following:

- $gen(x, v)$  represents the act of spending a bitcoin deposit  $x$  to mint  $v$  units of a new token. The owner of these units is the user who owned the deposit  $x$ .

- $burn(x, y)$  represents the act of destroying a sequence of deposits  $x$ , moving them to an unspendable deposit  $y$ .
- $split(x, v, \mathbf{B})$  represents the act of splitting a deposit  $x$  (say, containing  $v + v'$  units of a token  $\tau$ ) in two deposits of  $\tau$ . The first one of these deposits is owned by the same owner of  $x$ , and contains  $v$  token units. The second one is owned by  $\mathbf{B}$ , and contains the remaining  $v'$  units.
- $join(x, y, \mathbf{C})$  represents the joining of two deposits  $x$  and  $y$  of the same token, into a new deposit, owned by  $\mathbf{C}$ .
- $xchg(x, y)$  represents the act of atomically exchanging the owners of the two deposits  $x$  and  $y$  (not both of bitcoins). In particular, when one of the deposits stores  $\mathfrak{B}$ , this action represents buying/selling tokens for bitcoins.
- $give(x, \mathbf{B})$  represents a donation of the deposit  $x$  to  $\mathbf{B}$ .

Users follow a common pattern to perform token actions: (i) first, the involved users grant their authorization on the action; (ii) once all the needed authorizations have been granted, the action can actually be performed.

A *configuration*  $\Gamma$  is a composition of deposits and authorizations. We assume that configurations form a commutative monoid under the composition operator  $|$ , and we use  $\mathbf{0}$  to denote the empty configuration. We require that if  $\langle \mathbf{A}, v : \tau \rangle_x$  and  $\langle \mathbf{B}, v' : \tau' \rangle_{x'}$  both occur in  $\Gamma$ , then  $x \neq x'$ . We define a transition semantics between configurations in Figure 7.4. Transitions are decorated with labels, which describe the performed actions.

Rule  $[\text{GEN}]$  consumes a bitcoin deposit  $x$  owned by  $\mathbf{A}$  to generate  $v$  units of a new token  $\tau$ , which are stored in a fresh deposit  $y$ . Note that  $\mathbf{A}$ 's authorization is required to perform the action. For simplicity, we assume that minting tokens has no cost: it would be straightforward to adapt the rule to require a minting fee. Rule  $[\text{BURN}]$  removes from the configuration a single token deposit (when  $n = 1$  and  $\tau_1 \neq \mathfrak{B}$ ), or atomically removes a sequence of bitcoin deposits (when  $\tau_i = \mathfrak{B}$  for all  $i$ ). Rule  $[\text{SPLIT}]$  divides a deposit  $x$  in two fresh deposits  $y$  and  $z$ , preserving the number of token units. Rule  $[\text{JOIN}]$  allows  $\mathbf{A}$  and  $\mathbf{B}$  to merge two deposits of a token  $\tau$ , preserving the amount of token units, and transferring the new deposit to  $\mathbf{C}$ . Rule  $[\text{XCHG}]$  allows  $\mathbf{A}$  and  $\mathbf{B}$  to swap two deposits, containing either user-defined tokens or bitcoins. Finally, rule  $[\text{GIVE}]$  allows  $\mathbf{A}$  to donate one of her deposits to another user.

The transition relation  $\rightarrow$  is non-deterministic, because of the fresh

$$\begin{array}{c}
\frac{\Gamma = \mathbf{A} \triangleright_x \text{gen}(x, v) \mid \Gamma' \quad v > 0 \quad y, \tau \text{ fresh}}{\langle \mathbf{A}, 0 : \mathfrak{B} \rangle_x \mid \Gamma \xrightarrow{\text{gen}(x, v)} \langle \mathbf{A}, v : \tau \rangle_y \mid \Gamma'} \text{[GEN]} \\
\Gamma = \left( \parallel_{i \in 1..n} \mathbf{A}_i \triangleright_{x_i} \text{burn}(x_1 \cdots x_n, y) \right) \mid \Gamma' \quad n = 1 \vee (n \geq 1 \wedge \forall i : \tau_i = \mathfrak{B}) \\
\frac{\left( \parallel_{i \in 1..n} \langle \mathbf{A}_i, v_i : \tau_i \rangle_{x_i} \right) \mid \Gamma \xrightarrow{\text{burn}(x_1 \cdots x_n, y)} \Gamma'}{\Gamma = \mathbf{A} \triangleright_x \text{split}(x, v, \mathbf{B}) \mid \Gamma' \quad v, v' \geq 0 \quad y, y' \text{ fresh}} \text{[SPLIT]} \\
\frac{\langle \mathbf{A}, (v + v') : \tau \rangle_x \mid \Gamma \xrightarrow{\text{split}(x, v, \mathbf{B})} \langle \mathbf{A}, v : \tau \rangle_y \mid \langle \mathbf{B}, v' : \tau \rangle_{y'} \mid \Gamma'}{\Gamma = \mathbf{A} \triangleright_x \text{join}(x, y, \mathbf{C}) \mid \mathbf{B} \triangleright_y \text{join}(x, y, \mathbf{C}) \mid \Gamma' \quad z \text{ fresh}} \text{[JOIN]} \\
\frac{\langle \mathbf{A}, v : \tau \rangle_x \mid \langle \mathbf{B}, v' : \tau \rangle_y \mid \Gamma \xrightarrow{\text{join}(x, y, \mathbf{C})} \langle \mathbf{C}, (v + v') : \tau \rangle_z \mid \Gamma'}{\Gamma = \mathbf{A} \triangleright_x \text{xchg}(x, y) \mid \mathbf{B} \triangleright_y \text{xchg}(x, y) \mid \Gamma' \quad \tau \neq \mathfrak{B} \quad x', y' \text{ fresh}} \text{[XCHG]} \\
\frac{\langle \mathbf{A}, v : \tau \rangle_x \mid \langle \mathbf{B}, v' : \tau' \rangle_y \mid \Gamma \xrightarrow{\text{xchg}(x, y)} \langle \mathbf{A}, v' : \tau' \rangle_{x'} \mid \langle \mathbf{B}, v : \tau \rangle_{y'} \mid \Gamma'}{\Gamma = \mathbf{A} \triangleright_x \text{give}(x, \mathbf{B}) \mid \Gamma' \quad y \text{ fresh}} \text{[GIVE]} \\
\frac{\langle \mathbf{A}, v : \tau \rangle_x \mid \Gamma \xrightarrow{\text{give}(x, \mathbf{B})} \langle \mathbf{B}, v : \tau \rangle_y \mid \Gamma'}{}
\end{array}$$

Figure 7.4: Semantics of token actions.

names generated for deposits and tokens: however, given a transition  $\Gamma \xrightarrow{\alpha} \Gamma'$  the label  $\alpha$  is uniquely determined from  $\Gamma$  and  $\Gamma'$ . A *symbolic run*  $\mathcal{S}$  is a (possibly infinite) sequence  $\Gamma_0 \Gamma_1 \cdots$ , where  $\Gamma_0$  contains only  $\mathfrak{B}$  deposits, and for all  $i \geq 0$  there exists some (unique)  $\alpha_i$  such that  $\Gamma_i \xrightarrow{\alpha_i} \Gamma_{i+1}$ . For all  $i \geq 0$ , we denote with  $\mathcal{S}_i$  the  $i$ -th element of the run, when this element exists. If  $\mathcal{S}$  is finite, we denote its length as  $|\mathcal{S}|$ , and we write  $\Gamma_{\mathcal{S}}$  for its last configuration, i.e.  $\mathcal{S}_{|\mathcal{S}|-1}$ .

**Definition 7.1** (Token balance). *We define the balance of a token  $\tau \neq \mathfrak{B}$  in a configuration  $\Gamma$  inductively as follows:*

$$\begin{array}{ll}
\text{bal}_{\tau}(\mathbf{0}) = 0 & \text{bal}_{\tau}(\Gamma \mid \Gamma') = \text{bal}_{\tau}(\Gamma) + \text{bal}_{\tau}(\Gamma') \\
\text{bal}_{\tau}(\langle \mathbf{A}, v : \tau \rangle_x) = v & \text{bal}_{\tau}(\langle \mathbf{A}, v : \tau' \rangle_x) = 0 \quad (\tau' \neq \tau)
\end{array}$$

The following theorem establishes a basic preservation property: the balance of a token after a run is equal to the *minted* value minus the *burnt*

value, defined as:

$$\begin{aligned} \text{minted}_\tau(\mathcal{S}) &= v \text{ if } \exists i : \mathcal{S}_i \xrightarrow{\text{gen}(x,v)} \mathcal{S}_{i+1}, \text{ and} \\ &\quad \tau \text{ occurs in } \mathcal{S}_{i+1} \text{ but not in } \mathcal{S}_i \\ \text{burnt}_\tau(\mathcal{S}) &= \sum \left\{ v \mid \exists i : \mathcal{S}_i \xrightarrow{\text{burn}(x,y)} \mathcal{S}_{i+1}, \text{ and} \right. \\ &\quad \left. \mathcal{S}_i = \Gamma \mid \langle \mathbf{A}, v : \tau \rangle_x \right\} \end{aligned}$$

**Lemma 7.2.1.** Let  $\mathcal{S}$  be a finite symbolic run. For all  $\tau \neq \mathfrak{B}$ :

$$\text{bal}_\tau(\mathcal{S}) = \text{minted}_\tau(\mathcal{S}) - \text{burnt}_\tau(\mathcal{S})$$

*Proof.* By induction on  $\mathcal{S}$ , and by case analysis on each step. Inspecting each symbolic semantics rule, we can see that each step preserves the amount of token units, except for minting ( $[\text{GEN}]$ ) and burning ( $[\text{BURN}]$ ), which are explicitly taken into account by the equation.  $\square$

### 7.3 Implementing tokens in Bitcoin

In this section we show how to implement token actions in Bitcoin. To this purpose, we define a *computational model*, which describes the interactions of users who exchange messages and append transactions to the Bitcoin blockchain. A *computational run*  $\mathcal{C}$  is a sequence of bitstrings  $\gamma$ , each of which encodes one of the following actions: (i)  $\mathbf{A} \rightarrow * : m$ , denoting the broadcast of a bitstring  $m$ ; (ii)  $\mathbf{T}$ , denoting the appending of a transaction  $\mathbf{T}$  to the blockchain. A computational run always starts from a coinbase transaction  $\mathbf{T}_0$ . By extracting the transactions from a run  $\mathcal{C}$ , we obtain a blockchain  $\mathbf{B}_{\mathcal{C}}$ .

We simulate each symbolic token action in Figure 7.4 by appending a suitable transaction to the blockchain. We use the **arg** part of transaction outputs to record the token data:

1. **op** is the action implemented by the transaction:  $0 = \text{gen}$ ,  $1 = \text{burn}$ ,  $2 = \text{split}$ ,  $3 = \text{join}$ ,  $4 = \text{xchg}$ ,  $5 = \text{give}$ ;
2. **owner** is the (public key of the) user who owns of the token units controlled by the tx output;
3. **tkval** is the number of units controlled by the tx output;
4. **tkid** is the unique token identifier.

Symbolic authorization steps correspond to computational broadcasts of signatures.

We implement the token actions as a single script  $e_{\text{TOKEN}}$ , which uses a switch on the `op` value to jump to the first instruction of the requested action. Since the script is quite complex, we present independently the parts corresponding to each action, postponing the complete script to Appendix 7.4. All the transactions implementing token actions use *exactly* the same script, which we preserve throughout executions by using recursive covenants. To improve readability, we refer to the elements of the `arg` sequence by name rather than by index, e.g. we write `o.op` rather than `o.arg.1`.

**Gen** We implement the *gen* action by the following script:

```

1 if not verrec(ptxo(1)) // ctxo is a gen
2 then ctxo.tkid = txid(ptxo(1)) // token id
3   and ptxo(1).val = 0 // spent txo has 0 BTC
4   and outlen(ctxo) = 1 // gen has 1 out
5   and ctxo.tkval > 0 // positive token val
6 else ... // the other branches must preserve token id

```

Recall that *gen* produces a symbolic step of the form:

$$\langle \mathbf{A}, 0 : \mathfrak{B} \rangle_x \xrightarrow{\text{gen}(x,v)} \langle \mathbf{A}, v : \tau \rangle_y \quad (v > 0)$$

To translate this symbolic action into a computational one, we must spend a transaction output corresponding to  $\langle \mathbf{A}, 0 : \mathfrak{B} \rangle_x$ , and produce a fresh output corresponding to  $\langle \mathbf{A}, v : \tau \rangle_y$ . Assuming that the deposit  $x$  corresponds to an unspent transaction output  $(\mathbf{T}', 1)$  on the blockchain, this requires to append a transaction  $\mathbf{T}$  redeeming  $(\mathbf{T}', 1)$ , and ensuring that:

1. the parent transaction output  $(\mathbf{T}', 1)$  is not a token deposit, but just a plain  $\mathfrak{B}$  deposit (line 1);
2. `tkid` is the identifier of the parent tx output (line 2). This corresponds to identifying the fresh name  $\tau$  with the deposit name  $x$  of the redeemed  $\mathfrak{B}$  deposit;
3.  $0\mathfrak{B}$  are redeemed from the parent transaction (line 3);
4.  $\mathbf{T}$  has exactly one output (line 4);
5. `tkval` is positive (line 5), corresponding to the constraint  $v > 0$  in the symbolic semantics.

Notice that the first time the script is evaluated is when *redeeming*  $\mathbf{T}$  (not when appending it). At that time, `ctxo` will evaluate to  $(\mathbf{T}, 1)$ , and `ptxo(1)` to  $(\mathbf{T}', 1)$ . The script ensures that, when  $\mathbf{T}$  is redeemed, its `tkid` will contain a unique identifier of the token. Crucially, the scripts which

implement the other token actions will preserve this identifier in the `tkid` parameter. This identifier is essential to guarantee the correctness of the *join* and *xchg* actions.

**Burn** We implement the *burn* action by the script:

```

1 versig(ctxo.owner,rtx.wit) and // check owner
2 verscr(false,rtxo(1)) and // make rtx unspendable
3 inlen(rtxo(1)) = 1 and // rtx has 1 in
4 outlen(rtxo(1)) = 1 // rtx has 1 out

```

Recall that *burn* produces a symbolic step:

$$\left( \parallel_{i \in 1..n} \langle A_i, v_i : \tau_i \rangle_{x_i} \right) \xrightarrow{\text{burn}(x_1 \dots x_n, y)} \mathbf{0}$$

There are two cases, according to whether we are burning a single token deposit, or one or more  $\mathfrak{B}$  deposits. In the first case, assuming that the computational counterpart of  $x_1$  is the output  $(T', 1)$ , the corresponding computational step is to append a transaction  $T$  redeeming  $(T', 1)$ . The witness of  $T$  must carry a signature of the owner  $A$ , and its output script is *false*, making it unspendable. In the second case, it suffices to append a transaction  $T$  which redeems all the transaction outputs corresponding to  $x_1, \dots, x_n$ , and has a *false* script.

**Split** We implement the *split* action by the script:

```

1 versig(ctxo.owner,rtx.wit) // check owner
2 and verrec(rtxo(1)) // covenants on rtx
3 and verrec(rtxo(2))
4 and inlen(rtxo(1)) = 1 // rtx has 1 in
5 and outlen(rtxo(1)) = 2 // rtx has 2 outs
6 and rtxo(1).tkval >= 0 // positive token value
7 and rtxo(2).tkval >= 0
8 and rtxo(1).owner = ctxo.owner // preserve owner
9 and rtxo(1).tkid = ctxo.tkid // preserve tkid
10 and rtxo(2).tkid = ctxo.tkid
11 and rtxo(1).tkval + rtxo(2).tkval = ctxo.tkval

```

Recall that *split* produces a symbolic step of the form:

$$\langle A, (v + v') : \tau \rangle_x \xrightarrow{\text{split}(x, v, B)} \langle A, v : \tau \rangle_y \mid \langle B, v' : \tau \rangle_z$$

Assuming that  $x$  corresponds to an unspent transaction output  $(T', 1)$ , performing this step in Bitcoin requires to append a transaction  $T$  redeeming  $(T', 1)$ , and ensuring that:

1. the witness of  $T$  carries a signature of the owner (line 1);
2.  $T$  has only one input and two outputs, both containing the same script of  $(T', 1)$  (line 2-5);

3. the `tkval` of  $\mathsf{T}$ 's outputs `rtxo(1)` and `rtxo(2)` are  $\geq 0$  (line 6-7), corresponding to the precondition  $v, v' \geq 0$  in `[SPLIT]`;
4. `tkid` of  $\mathsf{T}$ 's outputs is the same of  $(\mathsf{T}', 1)$  (line 9-10);
5. the sum of token values `tkval` of the outputs of  $\mathsf{T}$  is equal to the token value of  $(\mathsf{T}', 1)$  (line-11).

**Join** We implement the *join* action by the script:

```

1 inlen(rtxo(1)) = 2 // rtx has 2 ins
2 and outlen(rtxo(1)) = 1 // rtx has 1 out
3 and verrec(rtxo(1)) // covenant on rtx
4 and verrec(stxo(2)) // covenants on both inputs
5 and verrec(stxo(1))
6 and ctxo.tkid = rtxo(1).tkid // preserve token id
7 and versig(ctxo.owner, rtx.wit) // check sig of owner
8 and rtxo(1).tkval = stxo(1).tkval + stxo(2).tkval

```

Recall that *join* produces a symbolic step of the form:

$$\langle \mathsf{A}, v : \tau \rangle_x \mid \langle \mathsf{B}, v' : \tau \rangle_y \xrightarrow{\text{join}(x,y,C)} \langle \mathsf{C}, (v + v') : \tau \rangle_z$$

Assume that  $x$  and  $y$  correspond to the unspent transaction outputs  $(\mathsf{T}', 1)$  and  $(\mathsf{T}'', 1)$ . To perform the corresponding computational step we append a transaction  $\mathsf{T}$  redeeming  $(\mathsf{T}', 1)$  and  $(\mathsf{T}'', 1)$ , and ensuring that, for both inputs:

1.  $\mathsf{T}$  has two inputs and one output, containing the same script of  $(\mathsf{T}', 1)$  and  $(\mathsf{T}'', 1)$  (line 1-5);
2. the token identifier `tkid` of the output of  $\mathsf{T}$  (`rtxo(1)`) is the same of  $(\mathsf{T}', 1)$  (line 6);
3. the witness of  $\mathsf{T}$  carries a signature of the owner (line 7);
4. the sum of token values `tkval` of both inputs is equal to the token value of  $(\mathsf{T}, 1)$  (line 8).

Note that the script in  $(\mathsf{T}', 1)$  ensures that the one in  $(\mathsf{T}'', 1)$  is the same, and vice-versa. In this way, we prevent joining tokens with bitcoins. To also prevent joining tokens of different type, the script checks that the `tkid` of the current transaction is the same as the one of the first output of the redeeming transaction. This is done by both inputs. In other words, `stxo(1).tkid = rtxo(1).tkid` and `stxo(2).tkid = rtxo(1).tkid`, implying that `stxo(1).tkid = stxo(2).tkid`.

**Exchange** We implement the *xchg* action by the script:

```

1 inlen(rtxo(1)) = 2 // rtx has 2 ins
2 and outlen(rtxo(1)) = 2 // rtx has 2 outs
3 and verrec(stxo(1)) // covenant on input 1
4 and verrec(rtxo(1)) // covenant on rtx(1)
5 and versig(ctxo.owner, rtx.wit) // check owner
6 and rtxo(1).owner = stxo(2).owner // exchange owner
7 and rtxo(2).owner = stxo(1).owner
8 and rtxo(1).tkval = stxo(1).tkval // preserve value
9 and rtxo(1).tkid = stxo(1).tkid // preserve tkid
10 if verrec(stxo(2)) then // exchange token/token
11     verrec(rtxo(2)) // covenant on rtx(2)
12     and rtxo(2).tkval = stxo(2).tkval // preserve tkval
13     and rtxo(2).tkid = stxo(2).tkid // preserve tkid
14 else // exchange token/BTC
15     verscr(versig(ctxo.owner, rtx.wit), rtxo(2))
16     and rtxo(2).val = stxo(2).val // preserve BTC

```

Recall that the symbolic  $xchg$  step has the form:

$$\langle A, v : \tau \rangle_x \mid \langle B, v' : \tau' \rangle_y \xrightarrow{xchg(x,y)} \langle A, v' : \tau' \rangle_{x'} \mid \langle B, v : \tau \rangle_{y'}$$

where  $\tau$  must be a token, while  $\tau'$  is either a  $\mathfrak{B}$  or a token.

Assume that  $x$  and  $y$  correspond to the unspent transaction outputs  $(T', 1)$  and  $(T'', 1)$ . To perform the corresponding computational step we append a transaction  $T$  redeeming  $(T', 1)$  and  $(T'', 1)$ , and ensuring that, for both inputs:

1.  $T$  has two inputs and two output (lines 1-2);
2. the first input and the first output of  $T$  must contain the same script of  $(T', 1)$  and  $(T'', 1)$  (lines 3-4);
3. the witness of  $T$  carries a signature of the owner (line 5);
4. the owner in the first output of  $T$  ( $rtxo(1)$ ) must be equal to the owner in the second input  $(T'', 1)$  (line 6);
5. dually, the owner in the second output of  $T$  ( $rtxo(2)$ ) must be equal to the owner in the first input  $(T', 1)$  (line 7);
6. the token value and identifier of the first output of  $T$  ( $rtxo(1)$ ) must be equal to those of  $(T', 1)$  (line 8-9).

Furthermore, if  $verrec(stxo(2))$  is true, i.e. we are exchanging a token with a token. In this case, we require that:

1. the second output of  $T$  must contain the same script of  $(T', 1)$  and  $(T'', 1)$  (line 11);
2. the token value and identifier of the second output of  $T$  ( $rtxo(2)$ ) must be equal to those of  $(T'', 1)$  (line 12-13);



When exchanging a token with a  $\mathfrak{B}$  deposit, we require:

1. the script of the second output of  $\mathbf{T}$  ( $\text{rtxo}(2)$ ) to be `versig(ctxo.owner, rtx.wit)` (line 15);
2. the value of the second output of  $\mathbf{T}$  to be equal to the value of  $(\mathbf{T}', 1)$  (line 16).

**Give** We implement the *give* action by the script:

```

1 inlen(rtxo(1)) = 1 // rtx has 1 in
2 and outlen(rtxo(1)) = 1 // rtx has 1 out
3 and versig(ctxo.owner, rtx.wit) // check owner
4 and verrec(rtxo(1)) // covenant on rtx(1)
5 and rtxo(1).tkid = ctxo.tkid // preserve tkid
6 and rtxo(1).tkval = ctxo.tkval // preserve value

```

Recall that *give* produces a symbolic step of the form:

$$\langle \mathbf{A}, v : \tau \rangle_x \xrightarrow{\text{give}(x, \mathbf{B})} \langle \mathbf{B}, v : \tau \rangle_y$$

Assuming that  $x$  corresponds to an unspent transaction output  $(\mathbf{T}', 1)$ , performing this step in Bitcoin requires to append a transaction  $\mathbf{T}$  redeeming  $(\mathbf{T}', 1)$ , and ensuring that:

1.  $\mathbf{T}$  has only one input and one output (lines 1-2);
2. the witness of  $\mathbf{T}$  carries a signature of the owner (line 3);
3. the output of  $\mathbf{T}$  ( $\text{rtxo}(1)$ ) contains the same script of  $(\mathbf{T}', 1)$  (line 4);
4. the token value and identifier of the first output of  $\mathbf{T}$  ( $\text{rtxo}(1)$ ) is the same of those of  $(\mathbf{T}', 1)$  (lines 5-6).

**Efficiency of the implementation** To estimate the efficiency of the implementation, we consider the number of cryptographic operations, as their execution cost is an order of magnitude greater than the other operations. In particular, performing `verrec` and `verscr` requires to compute the hash of a script (once this is done, the cost of comparing two hashes is negligible). This cost can be reduced by incentivising nodes to cache scripts. The most expensive token action is *xchg*, which, having two inputs, needs to verify 2 signatures and execute at most 10 covenants operations, which overall require to compute at most 6 script hashes. If nodes cache scripts, the cost of the action is not dissimilar to the one required to append a standard transaction with two inputs.

Note that, even though *e<sub>TOK</sub>* is a non-standard script, it could be used in a standard P2SH transaction, as in [81], if it did not exceed the

520-bytes limit. Taproot [121] would mitigate this issue: for scripts with multiple disjoint branches, Taproot allows the witness of the redeeming transaction to reveal just the needed branch. Therefore, the 520-bytes limit would apply to branches instead of the whole script.

# Conclusions

In this thesis we defined new formal models to improve the security of Bitcoin smart contracts, and developed the BitML toolchain.

In particular, we presented (i) a toolchain for developing BitML contracts (ii) a formal model of Bitcoin which is the foundation for (iii) a new process algebra for defining Bitcoin smart contracts (iv) a new extension to Bitcoin which extends its expressiveness as a smart contract platform, which has been exploited to implement (v) fungible tokens on Bitcoin, complete with a symbolic model.

## Main results and future works

### A formal model of Bitcoin transactions

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.

An alternative model of transactions in blockchain systems has been proposed in [99]. Roughly, blockchains are represented as directed acyclic graphs, where edges denote transfers of assets. This model is quite abstract, so that it can be instantiated to different blockchains (e.g., Bitcoin, Ethereum, and Hyperledger Fabric). Differently from ours, the model in [99] does not capture some peculiar features of Bitcoin, like e.g. transaction signatures and signature modifiers, output scripts, multi-signature verification, and Segregated Witnesses.

Our 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 enables their automated verification and analysis, which are of crucial importance in a context where design flaws may result in loss of money. For instance,

our model has been exploited in [9] to present a comprehensive survey of Bitcoin smart contracts.

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

In Definition 3.3, 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: (i)  $\mu = \{\text{wit} \mapsto \perp\}$  since the activation of the SegWit feature; (ii)  $\mu = \perp$ , beforehand. Consequently, the condition  $(\mathbb{T}, i, t) \overset{v}{\rightsquigarrow} (\mathbb{T}', j, t')$  item (a) of Definition 3.11 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 3.1 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 3.12 models blockchains as sequences of transactions, while 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 new blocks 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. Extending our model with coinbase transactions would falsify Theorem 3.1.5, since the overall value in the blockchain would no longer be decreasing. Definition 3.12 requires the timestamp of each transaction to increase monotonically. Instead, in Bitcoin a timestamp is valid if it is greater than the median timestamp of previous 11 blocks.

In Definitions 3.3 and 3.11, 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.

**Related works** Several works have proposed to use Bitcoin beyond the sole purpose of exchanging currency, by exploiting the flexibility of its scripting language. They propose to implement smart contracts, intended as sets of protocols of the participants involved in them.

Smart contracts requiring external state, namely oracles [87] and escrows [84], can be easily implemented using multi-signature transactions. Such implementations, however, rely on trusted third parties. The work [4] showed that Bitcoin can be used to implement timed commitments through deposit transactions. The commitments are then used to perform multiparty computations [1], such as calculating the winner of a lottery. The main drawback of this approach is indeed the deposit, which grows quadratically with the number of participants. More recently, [62] and [23] have proposed lottery smart contracts that require, respectively, zero and constant ( $\geq 0$ ) deposit. However, [62] requires the computation of an exponential number of signature w.r.t. the number of participants, while [23] only a quadratic one. The work [12] proposed a contingent payment protocol that can be implemented relying only on standard Bitcoin transaction. It allows to sell solutions for a class of NP problems (e.g. the factorization of a number), the use of zero-knowledge proofs ensure its correctness to the buyer.

## Extending Bitcoin with Neighborhood Covenants

We have proposed a formalisation of neighborhood covenants, an extension of traditional Bitcoin covenants in literature. we have exploited our formalisation to present a series of use cases which appear to be unfeasible in pure Bitcoin. We have introduced high-level contract primitives that exploit covenants to enable recursion, and allow contracts to receive new funds and parameters at runtime.

The first proposals of covenants in Bitcoin date back at least to 2013 [114]. Nevertheless, their inclusion into the official Bitcoin protocol is still uncertain, mainly because of the extremely cautious approach to implement changes to Bitcoin [102]. Still, the emerging of Bitcoin layer-2 protocols, like e.g. the Lightning Network [122], has revived the interest in covenants, as witnessed by a recent Bitcoin Improvement Proposal (BIP 119 [123]). The work in [63] propose a new opcode `CheckOutputVerify`, to explicitly access and constraint the outputs of the redeeming transaction. In contrast, [66] implements covenants exploiting the current implementation of `versig`, which checks a signature on data that is build by

implicitly accessing the redeeming transaction, to define a new operator `CheckSigFromStack`. Both implementations can require the script of the redeeming transaction to contain the same covenant of the spent one, enable recursive covenants. Another approach is to implement covenants without adding new opcodes [70], which would require a change in the consensus protocol, employing signatures to commit to transaction templates. The drawback of the approach is the need to delete the cryptographic keys used to sign the templates, which adds assumptions to the security model. As noted by [66], recursive covenants would allow to implement Bitcoin contracts that execute state machines, by appending transactions to trigger state transitions. Ergo [35] implement contracts as state machines on a UTXO blockchain using transaction trees. Differently from BitML [10], where transactions of a contract are pre-signed by its users, Ergo can use covenants to enforce the next transaction in the contract execution, enabling loops in the state transition. The work in [18] propose a formal model of covenants, which can be implemented in Bitcoin with modifications similar to the ones in [63, 66]. It exploits the model to specify some complex Bitcoin contracts, and discuss how to exploit covenants to design high-level language primitives for Bitcoin contracts.

This added flexibility can be exploited to design expressive high-level contract languages like Marlowe [68] and Plutus [30].

**Known limitations** Most of the scripts crafted in our use cases would produce non-standard transactions, that are rejected by Bitcoin nodes. To produce standard transactions from non-standard scripts, we can exploit P2SH [88]. This requires the transaction output to commit to the hash of the script, while the actual script is revealed in the witness of the redeeming transaction. Since, to check its hash, the script needs to be pushed to the stack, and the maximum size of a stack element is 520 bytes, longer scripts would be rejected. This clearly affects the expressiveness of contracts, as already observed in [10]. In particular, since the size of a script grows with the number of contract states (see e.g. Figure 4.10), contracts with many states would easily violate the 520 bytes limit. The introduction of Taproot [121] would mitigate this limit. For scripts with multiple disjoint branches, Taproot allows the witness of the redeeming transaction to reveal just the needed branch. Therefore, the 520 bytes limit would apply to branches, instead of the whole script. Another expressiveness limit derives from the fact that covenants can only constrain the scripts of the redeeming transaction. While this is enough to express non-fungible tokens (see Section 4.2.2), fungible ones seem to require more powerful mechanisms, because of the join operation. An alternative tech-

nique to enhancing covenants is to implement fungible tokens natively [33, 32], or to enforce their logic through a sidechain [117].

**Other extensions to Bitcoin.** Indeed, the Bitcoin scripting language features a very limited set of operations [89], and over the years many useful (and apparently harmless) opcodes have been disabled without a clear understanding of their alleged insecurity [101]. This is the case e.g., of bitwise logic operators, shift operators, integer multiplication, division and modulus. For this reason some developers proposed to re-enable some disabled opcodes [78], and some works in the literature proposed extensions to the Bitcoin scripting language so to enhance the expressiveness of smart contracts.

*Secure cash distribution with penalties* [57, 4, 26] is a cryptographic primitive which allows a set of participants to make a deposit, and then provide inputs to a function whose evaluation determines how the deposits are distributed among the participants. This primitive guarantees that dishonest participants (who, e.g., abort the protocol after learning the value of the function) will pay a penalty to the honest participants. This primitive does not seem to be directly implementable in Bitcoin, but it becomes so by extending the scripting language with the opcode CHECKSIGFROMSTACK discussed above. Secure cash distribution with penalties can be instantiated to a variety of smart contracts, e.g. lotteries [4] poker [57], and contingent payments. The latter smart contract can also be obtained through the opcode CHECKKEYPAIRVERIFY in [38], which checks if the two top elements of the stack are a valid key pair.

Another new opcode, called MULTIINPUT [62] consumes from the stack a signature  $\sigma$  and a sequence of `in` values  $(T_1, j_1) \cdots (T_n, j_n)$ , with the following two effects: (i) it verifies the signature  $\sigma$  against the redeeming transaction  $T$ , neglecting  $T.in$ ; (ii) it requires  $T.in$  to be equal to some of the  $T_i$ . Exploiting this opcode, [62] devise a fair  $N$ -party lottery which requires zero deposit, and  $O(N^2)$  off-chain signed transaction. The first one of these effects can be alternatively obtained by extending, instead of the scripting language, the signature modifiers. More specifically, [23] introduces a new signature modifier, which can set to  $\perp$  all the inputs of a transaction (i.e., no input is signed). In this way they obtain a fair multi-party lottery with similar properties to the one in [62].

Another way to improve the expressiveness of smart contracts is to replace the Bitcoin scripting language, e.g. with the one in [65]. This would also allow to establish bounds on the computational resources needed to run scripts.

Unfortunately, none of the proposed extensions has been yet included in the main branch of the Bitcoin Core client, and nothing suggests that they will be considered in the near future. Indeed, the development of Bitcoin is extremely conservative, as any change to its protocol requires an overwhelming consensus of the miners. So far, new opcodes can only be empirically assessed

through the Elements alpha project<sup>1</sup>, a testnet for experimenting new Bitcoin features. A significant research challenge would be that of formally proving that new opcodes do not introduce vulnerabilities, exploitable e.g. by Denial-of-Service attacks. For instance, unconstrained uses of the opcode `CAT` may cause an exponential space blow-up in the verification of transactions.

**Verification** Although designing contracts in the UTXO model seems to be less error-prone than in the shared memory model, e.g. because of the absence of reentrancy vulnerabilities (like the one exploited in the Ethereum DAO attack [124]), Bitcoin contracts may still contain security flaws. Therefore, it is important to devise verification techniques to detect security issues that may lead to the theft or freezing of funds. Recursive covenants make this task harder than in pure Bitcoin, since they can encode infinite-state transition systems, as in most of our use cases. Hence, model-checking techniques based on the exploration of the whole state space, like the one used in [3], cannot be applied.

**High-level Bitcoin contracts** The compiler of our extension of BitML is just sketched in Section 4.3, and we leave as future work its formal definition, as well as the extension of the computational soundness results of [22], ensuring the correspondence between the symbolic semantics of BitML and the underlying computational level of Bitcoin. Continuing along this line of research, it would be interesting to study new linguistic primitives that fully exploit the expressiveness of Bitcoin covenants, and to extend accordingly the verification technique of [24]. Note that our extension of the UTXO model is more restrictive than the one in [31], as the latter abstracts from the script language, just assuming that scripts denote any pure functions [126]. This added flexibility can be exploited to design expressive high-level contract languages like Marlowe [68] and Plutus [30].

## Bitcoin smart contracts as endpoint protocols

The formal model of smart contracts we have proposed in Chapter 5 is based on the current mechanisms of Bitcoin; indeed, this makes it possible to translate endpoint protocols into actual implementations interacting with the Bitcoin blockchain. However, constraining smart contracts to perfectly adhere to Bitcoin greatly reduces their expressiveness.

As witnessed in Section 5.2, designing secure smart contracts on Bitcoin is an error-prone task, similarly to designing secure cryptographic protocols. The reason lies in the fact that, to devise a secure contract, a designer has to anticipate any possible (mis-)behaviour of the other participants. The side effect is that endpoint protocols may be quite convoluted, as they must include compensations at all the points where something can go wrong. Therefore, tools

---

<sup>1</sup><https://elementsproject.org/elements/opcodes/>



to automate the analysis and verification of smart contracts may be of great help.

Recent works [2] propose to verify Bitcoin smart contracts by modelling the behaviour of participants as timed automata, and then using UPPAAL [25] to check properties against an attacker. This approach correctly captures the time constraints within the contracts. The downside is that encoding this UPPAAL model into an actual implementation with Bitcoin transactions is a complex task. Indeed, a designer without a deep knowledge of Bitcoin technicalities is likely to produce an UPPAAL model that can *not* be encoded in Bitcoin. A relevant research challenge is to study specification languages for Bitcoin contracts (like e.g. the one in Section 5.1), and techniques to *automatically* encode them in a model that can be verified by a model checker.

Remarkably, the verification of security properties of smart contracts requires to deal with non-trivial aspects, like temporal constraints and probabilities. This is the case, e.g., for the verification of fairness of lotteries (like e.g. the one discussed in Section 5.2.8); a further problem is that fairness must hold against any adversarial strategy. It is not clear whether in this case it is sufficient to consider a “most powerful” adversary, like e.g. in the symbolic Dolev-Yao model. In case a contract is not secure against arbitrary (PTIME) adversaries, one would like to verify that, at least, it is secure against *rational* ones [45], which is a relevant research issue. Additional issues arise when considering more concrete models of the Bitcoin blockchain, respect to the one in Section 1.1. This would require to model *forks*, i.e. the possibility that a recent transaction is removed from the blockchain. This could happen with rational (but dishonest) miners [59].

**DSLs for smart contracts.** As witnessed in Section 5.2, modelling Bitcoin smart contracts is complex and error-prone. A possible way to address this complexity is to devise high-level domain-specific languages (DSLs) for contracts, to be compiled in low-level protocols (e.g., the ones in Section 5.1). Indeed, the recent proliferation of non-Turing complete DSLs for smart contracts [97, 43, 28] suggests that this is an emerging research direction.

A first proposal of an high-level language implemented on top of Bitcoin is Typecoin [37]. This language allows to model the updates of a state machine as affine logic propositions. Users can “run” this machine by putting transactions on the Bitcoin blockchain. The security of the blockchain guarantees that only the legit updates of the machine can be triggered by users. A downside of this approach is that liveness is guaranteed only by assuming cooperation among the participants, i.e., a dishonest participant can make the others unable to complete an execution. Note instead that the smart contracts in Section 5.2 allow honest participants to terminate, regardless of the behaviours of the environment. In some cases, e.g. in the lottery in Section 5.2.8, abandoning the contract may even result in penalties (i.e., loss of the deposit paid upfront to stipulate the contract).

## A toolchain for BitML

While the business of smart contracts has flourished on platforms like Ethereum and Cardano, it never caught on Bitcoin. One of the main reasons is that, unlike the other platforms, Bitcoin has neither high-level contract languages, nor related development and verification tools.

A downside of using platforms with expressive, Turing-complete languages, is that they may expose contracts to a wider attack surface: indeed, a series of language-induced vulnerabilities of Ethereum contracts [8] has caused losses of hundreds of millions of USD [124, 120, 77].

Although our benchmarks witness a rich variety of contracts expressible in BitML, there is room for improvement. BitML is not *Bitcoin-complete*, i.e. some contracts executable in Bitcoin are not expressible in BitML. The main sources of this incompleteness are three: (i) all the transactions obtained by the compiler must be signed *before* stipulation by all the involved participants (only the signatures for authorizations can be provided at run-time); (ii) all transaction fields must be taken into account when computing signatures, while partial signatures (e.g. those obtained through `SIGHASH_ANYONECANPAY` and `SIGHASH_SINGLE`) are not used; (iii) off-chain interactions are limited to revealing secrets and providing authorizations. The first constraint is required to ensure that honest participants can always perform, at the Bitcoin level, the moves enabled in the corresponding BitML contract, regardless of the behaviour of the others. In this respect, BitML follows the standard assumption that participants are *non-cooperative*, i.e. at any moment after stipulation they can stop interacting (unlike TypeCoin [37], which assumes cooperation, allowing dishonest participants to make a contract deadlock). Yet, cooperation can be incentivized, by punishing misbehaviour with penalties, like e.g. in the timed commitment of Section 6.1. As a consequence of the design choices above, contracts with a dynamically-defined set of players (e.g., crowdfunding), or an unbounded number of iterations (e.g., micro-payment channels), are not expressible in BitML.

The limitations of BitML (and of Bitcoin) could be overcome in various ways. For instance, using Bitcoin “as-is”, it would be possible to relax constraint (iii) above, so to allow e.g. zero-knowledge off-chain protocols. This would enable to extend BitML with primitives to express *contingent payments* contracts, where participants trade solutions of a class of NP problems [12, 115]. Similarly, by relaxing constraint (i), we could extend BitML to enable dynamic stipulation of subcontracts, requiring that all the involved participants provide their signatures at run-time. This would allow to model e.g. micro-payment channels in BitML. Together with the use of `SIGHASH_ANYONECANPAY` (relaxing constraint (ii)), this would also allow for modelling crowdfunding contracts. As before, this extension could be implemented without modifying Bitcoin.

Other extensions of BitML would require extensions of Bitcoin. For instance, *covenants* [63, 66] would allow for implementing arbitrary finite-state machines. Controlled *input malleability* would allow to efficiently implement tournaments

in multi-player gambling games, like e.g. lotteries [23]. This can also be achieved through a new opcode that checks if the redeeming transaction belongs to a given set [62]. Contingent payments without zero-knowledge proofs can be achieved by exploiting a new opcode that checks the validity of key pairs [38]. A new opcode which checks signatures for arbitrary messages would allow for expressing general fair multiparty computations [57]. Further, fair and robust multiparty computations can be achieved using more complex transactions [54]. A more radical approach would be to replace the Bitcoin scripting language with a more expressive one, like e.g. Simplicity [65].

Compared with the tools for analysing Ethereum contracts [60, 73, 48, 50, 67, 49, 27, 69], whose precision is subject to the limitations derived by the Turing-completeness of the underlying languages, our toolchain features a sound and complete verification technique.

## On-chain Fungible Tokens on Bitcoin

We have proposed a secure and efficient implementation of fungible tokens on Bitcoin, exploiting neighbourhood covenants, a powerful yet simple extension of the Bitcoin script language. We formalise fungible tokens, including a symbolic model for fungible token, which can be applied to UTXO-based blockchains, and we prove some correctness properties of the model.

To keep the presentation simple, we have limited the functionality of tokens a bit, making split/join/exchange actions operate on just two deposits, and omitting time constraints. Removing these restrictions would only affect the size, but not the complexity, of our technical development. Further, it would allow to use tokens *as is* within high-level languages for Bitcoin contracts, e.g. BitML [22], simplifying the design of financial contracts which manage tokens and bitcoins. For instance, we would express as follows a basic zero-coupon bond [53] where an investor **A** pays upfront to a bank **B** 5 units of token  $\tau$ , and receives back  $1\text{฿}$  after a maturity date  $t$ :

$$\text{split } (5\tau \rightarrow \text{withdraw } \mathbf{B} \mid 2\text{฿} \rightarrow \text{after } t : \text{withdraw } \mathbf{A})$$

A research question arising from our work is how to exploit neighbourhood covenants in the design of high-level languages for Bitcoin contracts. Besides enhancing the expressiveness of these languages, neighbourhood covenants would enable a simpler compilation technique, compared e.g. that used by BitML, reducing the off-chain exchange of signatures.

**Related work** In Ethereum, similarly to our approach, tokens are implemented on top of the platform using custom code, following the ERC-20 and ERC-712 standards [106, 104]. Instead, the works in [33, 32] propose a generalisation of the EUTXO model [31] that natively supports custom tokens. Similarly, the work in [127] proposes an UTXO model that natively support tokens in the same way it supports the main cryptocurrency. In the model, each token or currency have the same status: fees for a token transaction are

paid directly with some units of the token itself, making the currencies mutually independent.

# List of Figures

1.1	Simplified Blockchain. . . . .	10
1.2	Two Bitcoin transactions. . . . .	11
2.1	Syntax of BitML contract preconditions. . . . .	14
2.2	Syntax of BitML contracts. . . . .	14
3.1	Signature modifiers. . . . .	28
3.2	Semantics of script expressions. . . . .	31
3.3	Example of three transactions. . . . .	33
3.4	Example of three transactions for Example 3.1.4 . . . . .	33
3.5	Three transactions for Examples 3.1.5 to 3.1.7. . . . .	35
3.6	Transactions of the chain contract. . . . .	39
4.1	Accessing transaction outputs through a script. . . . .	46
4.2	Semantics of neighbourhood covenants (extending Figure 3.2). . . . .	47
4.3	Transactions for the crowdfunding contract. . . . .	49
4.4	<b>A</b> creates a token with $T_0$ , and transfers it to <b>B</b> with $T_1$ . . . . .	50
4.5	<b>A</b> exploits the flaw to destroy a token, redeeming its value. . . . .	50
4.6	Transactions for the basic vault. . . . .	52
4.7	Transactions for the recursive vault. . . . .	53
4.8	Transactions for the pyramid scheme. . . . .	54
4.9	Transactions for King of the Ether Throne. . . . .	55
4.10	Script for King of the Ether Throne, obtained by compiling BitML. . . . .	57
5.1	Semantics of contract expressions. . . . .	62
5.2	Semantics of endpoint protocols. . . . .	64
5.3	Transactions of the naïve escrow contract. . . . .	65
5.4	Transactions of a contract relying on an oracle. . . . .	67
5.5	Transactions of the escrow contract. . . . .	69
5.6	Transactions of the Bitcoin mixing contract. . . . .	70
5.7	Transactions of the intermediated payment contract. . . . .	71
5.8	Transactions of the timed commitment. . . . .	73

---

5.9	Transactions of the timed commitment without signature exchange.	74
5.10	Transactions of the micropayment channel contract. . . . .	75
5.11	Transactions of the fair lottery with deposit. . . . .	77
5.12	Transactions of the contingent payment. . . . .	79
6.1	Toolchain architecture. . . . .	83
7.1	A transaction $T_5$ attempting to join $T_4$ and $T_3$ . . . . .	109
7.2	A <i>join</i> attack merging two different tokens. . . . .	110
7.3	Thwarting the <i>join</i> attack. . . . .	111
7.4	Semantics of token actions. . . . .	113

# List of Tables

3.1	Summary of notation. . . . .	25
6.1	Benchmarks for the BitML toolchain. . . . .	88





# Bibliography

## References

- [1] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. “Fair Two-Party Computations via Bitcoin Deposits”. In: *Financial Cryptography Workshops*. Vol. 8438. LNCS. Springer, 2014, pp. 105–121. DOI: [10.1007/978-3-662-44774-1\\_8](https://doi.org/10.1007/978-3-662-44774-1_8).
- [2] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. “Modeling bitcoin contracts by timed automata”. In: *International Conference on Formal Modeling and Analysis of Timed Systems*. Springer. 2014, pp. 7–22.
- [3] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. “Modeling Bitcoin contracts by timed automata”. In: *FORMATS*. Vol. 8711. LNCS. Springer, 2014, pp. 7–22. DOI: [10.1007/978-3-319-10512-3\\_2](https://doi.org/10.1007/978-3-319-10512-3_2).
- [4] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. “Secure Multiparty Computations on Bitcoin”. In: *IEEE S & P*. 2014, pp. 443–458. DOI: [10.1109/SP.2014.35](https://doi.org/10.1109/SP.2014.35).
- [5] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. “Secure multiparty computations on Bitcoin”. In: *Commun. ACM* 59.4 (2016), pp. 76–84. DOI: [10.1145/2896386](https://doi.org/10.1145/2896386).
- [6] Monika Di Angelo and Gernot Salzer. “Tokens, Types, and Standards: Identification and Utilization in Ethereum”. In: *DAPPS*. IEEE, 2020, pp. 1–10. DOI: [10.1109/DAPPS49028.2020.00001](https://doi.org/10.1109/DAPPS49028.2020.00001).

- [7] Nicola Atzei, Massimo Bartoletti, and Tiziana Cimoli. “A Survey of Attacks on Ethereum Smart Contracts SoK”. In: *Proceedings of the 6th International Conference on Principles of Security and Trust - Volume 10204*. New York, NY, USA: Springer-Verlag New York, Inc., 2017, pp. 164–186. ISBN: 978-3-662-54454-9. DOI: [10.1007/978-3-662-54455-6\\_8](https://doi.org/10.1007/978-3-662-54455-6_8). URL: [https://doi.org/10.1007/978-3-662-54455-6\\_8](https://doi.org/10.1007/978-3-662-54455-6_8).
- [8] Nicola Atzei, Massimo Bartoletti, and Tiziana Cimoli. “A Survey of Attacks on Ethereum Smart Contracts (SoK)”. In: *Principles of Security and Trust (POST)*. Vol. 10204. LNCS. Springer, 2017, pp. 164–186. DOI: [10.1007/978-3-662-54455-6\\_8](https://doi.org/10.1007/978-3-662-54455-6_8).
- [9] Nicola Atzei, Massimo Bartoletti, Tiziana Cimoli, Stefano Lande, and Roberto Zunino. “SoK: unraveling Bitcoin smart contracts”. In: *POST*. Vol. 10804. LNCS. Springer, 2018, pp. 217–242. DOI: [10.1007/978-3-319-89722-6](https://doi.org/10.1007/978-3-319-89722-6).
- [10] Nicola Atzei, Massimo Bartoletti, Stefano Lande, Nobuko Yoshida, and Roberto Zunino. “Developing secure Bitcoin contracts with BitML”. In: *ESEC/FSE*. 2019. DOI: <https://doi.org/10.1145/3338906.3341173>.
- [11] Nicola Atzei, Massimo Bartoletti, Stefano Lande, and Roberto Zunino. “A formal model of Bitcoin transactions”. In: *Financial Cryptography and Data Security*. Vol. 10957. LNCS. Springer, 2018. DOI: [10.1007/978-3-662-58387-6](https://doi.org/10.1007/978-3-662-58387-6).
- [12] Wacław Banasik, Stefan Dziembowski, and Daniel Malinowski. “Efficient Zero-Knowledge Contingent Payments in Cryptocurrencies Without Scripts”. In: *ESORICS*. Vol. 9879. LNCS. Springer, 2016, pp. 261–280. DOI: [10.1007/978-3-319-45741-3\\_14](https://doi.org/10.1007/978-3-319-45741-3_14).
- [13] Massimo Bartoletti, Bryn Bellomy, and Livio Pompianu. “A Journey into Bitcoin Metadata”. In: *J. Grid Comput.* 17.1 (2019), pp. 3–22. DOI: [10.1007/s10723-019-09473-3](https://doi.org/10.1007/s10723-019-09473-3).
- [14] Massimo Bartoletti, Salvatore Carta, Tiziana Cimoli, and Roberto Saia. “Dissecting Ponzi schemes on Ethereum: identification, analysis, and impact”. In: *Future Gener. Comput. Syst.* 102 (2020), pp. 259–277. DOI: [10.1016/j.future.2019.08.014](https://doi.org/10.1016/j.future.2019.08.014).
- [15] Massimo Bartoletti, Tiziana Cimoli, and Roberto Zunino. “Fun with Bitcoin Smart Contracts”. In: *ISoLA*. 2018, pp. 432–449. DOI: [10.1007/978-3-030-03427-6\\_32](https://doi.org/10.1007/978-3-030-03427-6_32).

- [16] Massimo Bartoletti, Stefano Lande, Maurizio Murgia, and Roberto Zunino. “Verification of recursive Bitcoin contracts”. In: *Submitted to LMCS* (2020). arXiv: [2011.14165](https://arxiv.org/abs/2011.14165).
- [17] Massimo Bartoletti, Stefano Lande, and Alessandro Sebastian Podda. “A Proof-of-Stake protocol for consensus on Bitcoin sub-chains”. In: *International Conference on Financial Cryptography and Data Security*. Springer, 2017, pp. 568–584.
- [18] Massimo Bartoletti, Stefano Lande, and Roberto Zunino. “Bitcoin Covenants Unchained”. In: *ISoLA*. To appear, 2020.
- [19] Massimo Bartoletti, Stefano Lande, and Roberto Zunino. “Computationally sound Bitcoin tokens”. In: *IEEE Computer Security Foundations Symposium*. To appear, 2021. arXiv: [2010.01347](https://arxiv.org/abs/2010.01347).
- [20] Massimo Bartoletti, Maurizio Murgia, and Roberto Zunino. “Renegotiation and recursion in Bitcoin contracts”. In: *COORDINATION*. Vol. 12134. LNCS. Springer, 2020, pp. 261–278. DOI: [10.1007/978-3-030-50029-0\\_17](https://doi.org/10.1007/978-3-030-50029-0_17).
- [21] Massimo Bartoletti and Livio Pompianu. “An Analysis of Bitcoin OP\_RETURN Metadata”. In: *Financial Cryptography and Data Security*. Ed. by Michael Brenner, Kurt Rohloff, Joseph Bonneau, Andrew Miller, Peter Y.A. Ryan, Vanessa Teague, Andrea Bracciali, Massimiliano Sala, Federico Pintore, and Markus Jakobsson. Cham: Springer International Publishing, 2017, pp. 218–230. ISBN: 978-3-319-70278-0.
- [22] Massimo Bartoletti and Roberto Zunino. “BitML: a calculus for Bitcoin smart contracts”. In: *ACM CCS*. 2018. DOI: [10.1145/3243734.3243795](https://doi.org/10.1145/3243734.3243795).
- [23] Massimo Bartoletti and Roberto Zunino. “Constant-deposit multi-party lotteries on Bitcoin”. In: *Financial Cryptography Workshops*. Vol. 10323. LNCS. Springer, 2017. DOI: [10.1007/978-3-319-70278-0](https://doi.org/10.1007/978-3-319-70278-0).
- [24] Massimo Bartoletti and Roberto Zunino. “Verifying liquidity of Bitcoin contracts”. In: *POST*. Vol. 11426. LNCS. Springer, 2019, pp. 222–247.
- [25] Gerd Behrmann, Alexandre David, and Kim G Larsen. “A tutorial on Uppaal”. In: *Formal methods for the design of real-time systems*. Vol. 3185. LNCS. Springer, 2004, pp. 200–236. DOI: [10.1007/978-3-540-30080-9\\_7](https://doi.org/10.1007/978-3-540-30080-9_7).

- [26] Iddo Bentov and Ranjit Kumaresan. “How to Use Bitcoin to Design Fair Protocols”. In: *CRYPTO*. Vol. 8617. LNCS. Springer, 2014, pp. 421–439. DOI: [10.1007/978-3-662-44381-1\\\_24](https://doi.org/10.1007/978-3-662-44381-1\_24).
- [27] Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cedric Fournet, Anitha Gollamudi, Georges Gonthier, Nadim Kobeissi, Aseem Rastogi, Thomas Sibut-Pinote, Nikhil Swamy, and Santiago Zanella-Beguelin. “Formal Verification of Smart Contracts”. In: *PLAS*. 2016.
- [28] Alex Biryukov, Dmitry Khovratovich, and Sergei Tikhomirov. “Findel: Secure Derivative Contracts for Ethereum”. In: *Financial Cryptography Workshops*. Vol. 10323. LNCS. Springer, 2017, pp. 453–467. DOI: [10.1007/978-3-319-70278-0\\\_28](https://doi.org/10.1007/978-3-319-70278-0\_28).
- [29] Dan Boneh and Moni Naor. “Timed Commitments”. In: *CRYPTO*. Vol. 1880. LNCS. Springer, 2000, pp. 236–254. DOI: [10.1007/3-540-44598-6](https://doi.org/10.1007/3-540-44598-6).
- [30] Lars Brünjes and Murdoch James Gabbay. “UTxO- vs account-based smart contract blockchain programming paradigms”. In: *CoRR* abs/2003.14271 (2020). arXiv: [2003.14271](https://arxiv.org/abs/2003.14271).
- [31] Manuel M.T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian, Michael Peyton Jones, and Philip Wadler. “The Extended UTXO Model”. In: *Financial Cryptography Workshops*. To appear. 2020.
- [32] Manuel M.T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian, Jann Müller, Michael Peyton Jones, Polina Vinogradova, and Philip Wadler. “Native Custom Tokens in the Extended UTXO Model”. In: *ISoLA*. To appear. 2020.
- [33] Manuel M.T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian, Jann Müller, Michael Peyton Jones, Polina Vinogradova, Philip Wadler, and Joachim Zahnentferner. “UTXO\$.ma\$: UTXO with Multi-Asset Support”. In: *ISoLA*. To appear. 2020.
- [34] Weili Chen, Zibin Zheng, Jiahui Cui, Edith Ngai, Peilin Zheng, and Yuren Zhou. “Detecting Ponzi Schemes on Ethereum: Towards Healthier Blockchain Technology”. In: *WWW*. ACM, 2018, pp. 1409–1418. DOI: [10.1145/3178876.3186046](https://doi.org/10.1145/3178876.3186046).
- [35] Alexander Chepurnoy and Amitabh Saxena. “Multi-stage Contracts in the UTXO Model”. In: *Cryptocurrencies and Blockchain Technology*. Vol. 11737. LNCS. Springer, 2019, pp. 244–254. DOI: [10.1007/978-3-030-31500-9\\\_16](https://doi.org/10.1007/978-3-030-31500-9\_16).

- [36] Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and José F. Quesada. “Maude: Specification and Programming in Rewriting Logic”. In: *TCS* (2001).
- [37] Karl Crary and Michael J. Sullivan. “Peer-to-peer affine commitment using Bitcoin”. In: *ACM Conf. on Programming Language Design and Implementation*. 2015, pp. 479–488. DOI: [10.1145/2737924.2737997](https://doi.org/10.1145/2737924.2737997).
- [38] Sergi Delgado-Segura, Cristina Pérez-Solà, Guillermo Navarro-Arribas, and Jordi Herrera-Joancomartí. “A fair protocol for data trading based on Bitcoin transactions”. In: *Future Generation Computer Systems* (2017). DOI: [10.1016/j.future.2017.08.021](https://doi.org/10.1016/j.future.2017.08.021).
- [39] Kevin Delmolino, Mitchell Arnett, Andrew Millerand Ahmed Kosba, and Elaine Shi. “Step by Step Towards Creating a Safe Smart Contract: Lessons and Insights from a Cryptocurrency Lab”. In: (2016).
- [40] Steven Eker, José Meseguer, and Ambarish Sridharanarayanan. “The Maude LTL Model Checker”. In: *Electr. Notes Theor. Comput. Sci.* 71 (2002), pp. 162–187. DOI: [10.1016/S1571-0661\(05\)82534-4](https://doi.org/10.1016/S1571-0661(05)82534-4).
- [41] Patrick C. Fischer, Albert R. Meyer, and Arnold L. Rosenberg. “Counter machines and counter languages”. In: *Mathematical systems theory* 2.3 (1968), pp. 265–283. DOI: [10.1007/BF01694011](https://doi.org/10.1007/BF01694011).
- [42] Matthew Flatt. “Creating languages in Racket”. In: *Commun. ACM* 55.1 (2012), pp. 48–56. DOI: [10.1145/2063176.2063195](https://doi.org/10.1145/2063176.2063195).
- [43] C. K. Frantz and M. Nowostawski. “From Institutions to Code: towards Automated Generation of Smart Contracts”. In: *eCAS Workshop*. 2016.
- [44] Michael Fröwis, Andreas Fuchs, and Rainer Böhme. “Detecting Token Systems on Ethereum”. In: *Financial Cryptography and Data Security*. Vol. 11598. LNCS. Springer, 2019, pp. 93–112. DOI: [10.1007/978-3-030-32101-7\\_7](https://doi.org/10.1007/978-3-030-32101-7_7).
- [45] Juan A. Garay, Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. “Rational Protocol Design: Cryptography against Incentive-Driven Adversaries”. In: *FOCS*. 2013, pp. 648–657. DOI: [10.1109/FOCS.2013.75](https://doi.org/10.1109/FOCS.2013.75).

- [46] Oded Goldreich, Silvio Micali, and Avi Wigderson. “How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority”. In: *ACM Symposium on Theory of Computing*. ACM, 1987, pp. 218–229. DOI: [10.1145/28395.28420](https://doi.org/10.1145/28395.28420).
- [47] David M. Goldschlag, Stuart G. Stubblebine, and Paul F. Syverson. “Temporarily hidden bit commitment and lottery applications”. In: *Int. J. Inf. Sec.* 9.1 (2010), pp. 33–50. DOI: [10.1007/s10207-009-0094-1](https://doi.org/10.1007/s10207-009-0094-1).
- [48] Ilya Grishchenko, Matteo Maffei, and Clara Schneidewind. “A Semantic Framework for the Security Analysis of Ethereum Smart Contracts”. In: *Principles of Security and Trust (POST)*. Vol. 10804. LNCS. Springer, 2018, pp. 243–269. DOI: [10.1007/978-3-319-89722-6\\_10](https://doi.org/10.1007/978-3-319-89722-6_10).
- [49] Ilya Grishchenko, Matteo Maffei, and Clara Schneidewind. “Foundations and Tools for the Static Analysis of Ethereum Smart Contracts”. In: *CAV*. Vol. 10981. LNCS. Springer, 2018, pp. 51–78. DOI: [10.1007/978-3-319-96145-3\\_4](https://doi.org/10.1007/978-3-319-96145-3_4).
- [50] Everett Hildenbrandt et al. “KEVM: A Complete Formal Semantics of the Ethereum Virtual Machine”. In: *IEEE Computer Security Foundations Symposium (CSF)*. IEEE Computer Society, 2018, pp. 204–217. DOI: [10.1109/CSF.2018.00022](https://doi.org/10.1109/CSF.2018.00022).
- [51] Everett Hildenbrandt et al. “KEVM: A Complete Semantics of the Ethereum Virtual Machine”. In: *2018 IEEE 31st Computer Security Foundations Symposium*. IEEE, 2018, pp. 204–217.
- [52] Yoichi Hirai. “Defining the Ethereum Virtual Machine for Interactive Theorem Provers”. In: *Financial Cryptography Workshops*. Vol. 10323. LNCS. Springer, 2017, pp. 520–535. DOI: [10.1007/978-3-319-70278-0\\_33](https://doi.org/10.1007/978-3-319-70278-0_33).
- [53] Simon L. Peyton Jones, Jean-Marc Eber, and Julian Seward. “Composing contracts: an adventure in financial engineering, functional pearl”. In: *International Conference on Functional Programming (ICFP)*. 2000, pp. 280–292. DOI: [10.1145/351240.351267](https://doi.org/10.1145/351240.351267).
- [54] Aggelos Kiayias, Hong-Sheng Zhou, and Vassilis Zikas. “Fair and Robust Multi-party Computation Using a Global Transaction Ledger”. In: *EUROCRYPT*. Vol. 9666. LNCS. Springer, 2016, pp. 705–734. DOI: [10.1007/978-3-662-49896-5\\_25](https://doi.org/10.1007/978-3-662-49896-5_25).
- [55] Ranjit Kumaresan and Iddo Bentov. “Amortizing Secure Computation with Penalties”. In: *ACM CCS*. 2016, pp. 418–429. DOI: [10.1145/2976749.2978424](https://doi.org/10.1145/2976749.2978424).

- [56] Ranjit Kumaresan and Iddo Bentov. “How to Use Bitcoin to Incentivize Correct Computations”. In: *ACM CCS*. 2014, pp. 30–41. DOI: [10.1145/2660267.2660380](https://doi.org/10.1145/2660267.2660380).
- [57] Ranjit Kumaresan, Tal Moran, and Iddo Bentov. “How to Use Bitcoin to Play Decentralized Poker”. In: *ACM CCS*. 2015, pp. 195–206. DOI: [10.1145/2810103.2813712](https://doi.org/10.1145/2810103.2813712).
- [58] Ranjit Kumaresan, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. “Improvements to Secure Computation with Penalties”. In: *ACM CCS*. 2016, pp. 406–417. DOI: [10.1145/2976749.2978421](https://doi.org/10.1145/2976749.2978421).
- [59] Kevin Liao and Jonathan Katz. “Incentivizing Blockchain Forks via Whale Transactions”. In: *Financial Cryptography Workshops*. Vol. 10323. LNCS. Springer, 2017, pp. 264–279. DOI: [10.1007/978-3-319-70278-0\\_17](https://doi.org/10.1007/978-3-319-70278-0_17).
- [60] Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. “Making Smart Contracts Smarter”. In: *ACM CCS*. 2016, pp. 254–269. DOI: [10.1145/2976749.2978309](https://doi.org/10.1145/2976749.2978309).
- [61] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. “A fistful of Bitcoins: characterizing payments among men with no names”. In: *Commun. ACM* 59.4 (2016), pp. 86–93. DOI: [10.1145/2896384](https://doi.org/10.1145/2896384).
- [62] Andrew Miller and Iddo Bentov. “Zero-Collateral Lotteries in Bitcoin and Ethereum”. In: *EuroS&P Workshops*. 2017, pp. 4–13. DOI: [10.1109/EuroSPW.2017.44](https://doi.org/10.1109/EuroSPW.2017.44).
- [63] Malte Möser, Ittay Eyal, and Emin Gün Sirer. “Bitcoin covenants”. In: *Financial Cryptography Workshops*. Vol. 9604. LNCS. Springer, 2016, pp. 126–141. DOI: [10.1007/978-3-662-53357-4\\_9](https://doi.org/10.1007/978-3-662-53357-4_9).
- [64] Xavier Nicollin and Joseph Sifakis. “An Overview and Synthesis on Timed Process Algebras”. In: *CAV*. Vol. 575. LNCS. 1991, pp. 376–398. DOI: [10.1007/3-540-55179-4\\_36](https://doi.org/10.1007/3-540-55179-4_36).
- [65] Russell O’Connor. “Simplicity: A New Language for Blockchains”. In: *PLAS*. 2017. URL: <http://arxiv.org/abs/1711.03028>.
- [66] Russell O’Connor and Marta Piekarska. “Enhancing Bitcoin transactions with covenants”. In: *Financial Cryptography Workshops*. Vol. 10323. LNCS. Springer, 2017. DOI: [10.1007/978-3-319-70278-0\\_12](https://doi.org/10.1007/978-3-319-70278-0_12).

- [67] Daejun Park, Yi Zhang, Manasvi Saxena, Philip Daian, and Grigore Rosu. “A formal verification tool for Ethereum VM bytecode”. In: *ACM ESEC/SIGSOFT FSE*. 2018, pp. 912–915. DOI: [10.1145/3236024.3264591](https://doi.org/10.1145/3236024.3264591).
- [68] Pablo Lamela Seijas and Simon J. Thompson. “Marlowe: Financial Contracts on Blockchain”. In: *ISoLA*. Vol. 11247. LNCS. Springer, 2018, pp. 356–375. DOI: [10.1007/978-3-030-03427-6\\_27](https://doi.org/10.1007/978-3-030-03427-6_27).
- [69] Ilya Sergey, Amrit Kumar, and Aquinas Hobor. “Scilla: a Smart Contract Intermediate-Level Language”. In: *CoRR* abs/1801.00687 (2018).
- [70] Jacob Swambo, Spencer Hommel, Bob McElrath, and Bryan Bishop. “Bitcoin Covenants: Three Ways to Control the Future”. In: *CoRR* abs/2006.16714 (2020).
- [71] Paul F. Syverson. “Weakly Secret Bit Commitment: Applications to Lotteries and Fair Exchange”. In: *IEEE CSFW*. 1998, pp. 2–13. DOI: [10.1109/CSFW.1998.683149](https://doi.org/10.1109/CSFW.1998.683149).
- [72] Nick Szabo. “Formalizing and Securing Relationships on Public Networks”. In: *First Monday* 2.9 (1997). <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/548>.
- [73] Petar Tsankov, Andrei Marian Dan, Dana Drachslor Cohen, Arthur Gervais, Florian Bueznli, and Martin T. Vechev. “Securify: Practical Security Analysis of Smart Contracts”. In: *CoRR* abs/1806.01143 (2018).
- [74] Marie Vasek and Tyler Moore. “There’s No Free Lunch, Even Using Bitcoin: Tracking the Popularity and Profits of Virtual Currency Scams”. In: *Financial Cryptography and Data Security*. 2015, pp. 44–61. DOI: [10.1007/978-3-662-47854-7\\_4](https://doi.org/10.1007/978-3-662-47854-7_4).
- [75] Andrew Chi-Chih Yao. “How to Generate and Exchange Secrets (Extended Abstract)”. In: *27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986*. IEEE Computer Society, 1986, pp. 162–167. DOI: [10.1109/SFCS.1986.25](https://doi.org/10.1109/SFCS.1986.25).

## Online references

- [76] *520-byte limitation on serialized script size*. URL: <https://github.com/bitcoin/bips/blob/master/bip-0016>.



- mediawiki#520-byte-limitation-on-serialized-script-size.
- [77] *A Postmortem on the Parity Multi-Sig Library Self-Destruct*. <https://goo.gl/Kw3gXi>. 11/2017.
- [78] *A proposal to reintroduce the disabled script opcodes*. URL: <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2017-May/014356.html>.
- [79] *Algorand Developer Docs - Assets*. <https://developer.algorand.org/docs/features/asa/>. 2020.
- [80] Adam Back and Iddo Bentov. *Note on fair coin toss via Bitcoin*. <http://www.cs.technion.ac.il/~idddo/cointossBitcoin.pdf>. 2013.
- [81] *BALZaC: Bitcoin Abstract Language, analyZer and Compiler*. <https://blockchain.unica.it/balzac/>. 2018.
- [82] *Bitcoin Average Transaction Fee historical chart*. URL: <https://bitinfocharts.com/comparison/bitcoin-transactionfees.html>.
- [83] *Bitcoin Core integration/staging tree*. URL: <https://github.com/bitcoin/bitcoin> (visited on 09/25/2018).
- [84] *Bitcoin developer guide - Escrow and arbitration*. <https://bitcoin.org/en/developer-guide#escrow-and-arbitration>. 2015.
- [85] *Bitcoin standard transactions*. URL: <https://developer.bitcoin.org/devguide/transactions.html#standard-transactions>.
- [86] *Bitcoin wiki - Contracts - Assurance contracts*. [https://en.bitcoin.it/wiki/Contract#Example\\_3:\\_Assurance\\_contracts](https://en.bitcoin.it/wiki/Contract#Example_3:_Assurance_contracts). 2012.
- [87] *Bitcoin wiki - Contracts - Using external state*. [https://en.bitcoin.it/wiki/Contract#Example\\_4:\\_Using\\_external\\_state](https://en.bitcoin.it/wiki/Contract#Example_4:_Using_external_state). 2012.
- [88] *Bitcoin wiki - Pay-to-Script Hash*. [https://en.bitcoinwiki.org/wiki/Pay-to-Script\\_Hash](https://en.bitcoinwiki.org/wiki/Pay-to-Script_Hash). 2017.
- [89] *Bitcoin wiki script*. URL: <https://en.bitcoin.it/wiki/Script> (visited on 10/01/2018).
- [90] BitFury group. *Smart Contracts on Bitcoin Blockchain*. <http://bitfury.com/content/5-white-papers-research/contracts-1.1.1.pdf>. 2015.

- [91] *BitML benchmarks contrats*. URL: <https://github.com/bitml-lang/bitml-compiler/tree/master/examples/benchmarks>.
- [92] *BitML compiler*. URL: <https://github.com/bitml-lang/bitml-compiler/>.
- [93] *BitML examples*. URL: <https://github.com/bitml-lang/bitml-compiler/tree/master/examples>.
- [94] *BitML Online Documentation*. URL: <https://blockchain.unica.it/bitml/docs/index.html>.
- [95] *BitML toolchain repository*. URL: <https://github.com/bitml-lang/>.
- [96] *BitML verifier*. URL: <https://github.com/bitml-lang/bitml-maude>.
- [97] Richard G. Brown, James Carlyle, Ian Grigg, and Mike Hearn. *Conda: An Introduction*. <http://r3cev.com/s/conda-introductory-whitepaper-final.pdf>. 2016.
- [98] Vitalik Buterin. *Ethereum: a next generation smart contract and decentralized application platform*. <https://github.com/ethereum/wiki/wiki/White-Paper>. 2013.
- [99] Christian Cachin, Angelo De Caro, Pedro Moreno-Sanchez, Björn Tackmann, and Marko Vukolić. *The Transaction Graph for Modeling Blockchain Semantics*. Cryptology ePrint Archive, Report 2017/1070. 2017. URL: <https://eprint.iacr.org/2017/1070>.
- [100] *Colu website*. <https://www.colu.com/>. 2020.
- [101] *CVE-2010-5141*. URL: [https://en.bitcoin.it/wiki/Common\\_Vulnerabilities\\_and\\_Exposures#CVE-2010-5141](https://en.bitcoin.it/wiki/Common_Vulnerabilities_and_Exposures#CVE-2010-5141).
- [102] Luke Dashjr. *BIP 0002*. [https://en.bitcoin.it/wiki/BIP\\_0002](https://en.bitcoin.it/wiki/BIP_0002). 2016.
- [103] Robby Dermody, Adam Krellenstein, Ouziel Slama, and Evan Wagner. *CounterParty: Protocol Specification*. 2014. URL: <http://counterparty.io/docs/protocol%5Fspecification/>.
- [104] William Entriken, Dieter Shirley, Jacob Evans, and Nastassia Sachs. *EIP 721: ERC-721 Non-Fungible Token Standard*. <https://eips.ethereum.org/EIPS/eip-721>.
- [105] *EPOBC protocol specification*. [https://github.com/chromaway/ngccbase/wiki/EPOBC\\_simple](https://github.com/chromaway/ngccbase/wiki/EPOBC_simple). 2020.
- [106] *ERC-20 Token Standard*. <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-20.md>. 2015.

- [107] *Ethereum token dynamics*. <https://stat.bloxy.info/superset/dashboard/tokens>.
- [108] Mike Hearn. *Rapidly-adjusted (micro)payments to a pre-determined party*. 2013. URL: [bitcointalk.org](http://bitcointalk.org).
- [109] *Hyperledger - Open Source Blockchain Technologies*. URL: <https://www.hyperledger.org/> (visited on 09/25/2018).
- [110] Rosco Kalis. *CashScript — Writing covenants*. 2019. URL: <https://cashscript.org/docs/guides/covenants/>.
- [111] *King of the Ether Throne: Post mortem investigation*. URL: <https://www.kingoftheether.com/postmortem.html>.
- [112] Johnson Lau. *Upgrading Bitcoin: Segregated Witness*. URL: <https://www.bitcoinhk.org/media/presentations/2016-03-16/2016-03-16-Segregated%5FWitness.pdf> (visited on 09/25/2018).
- [113] Eric Lombrozo, Johnson Lau, and Pieter Wuille. *Segregated Witness (Consensus layer)*. BIP 141, <https://github.com/bitcoin/bips/blob/master/bip-0141.mediawiki>. 2015.
- [114] Gregory Maxwell. *CoinCovenants using SCIP signatures, an amusingly bad idea*. <https://bitcointalk.org/index.php?topic=278122.0>. 2013.
- [115] Gregory Maxwell. *The first successful Zero-Knowledge Contingent Payment*. <https://bitcoincore.org/en/2016/02/26/zero-knowledge-contingent-payments-announcement/>. 2016.
- [116] Satoshi Nakamoto. *Bitcoin: a peer-to-peer electronic cash system*. <https://bitcoin.org/bitcoin.pdf>. 2008.
- [117] Jonas Nick, Andrew Poelstra, and Gregory Sanders. *Liquid: a Bitcoin Sidechain*. <https://blockstream.com/assets/downloads/pdf/liquid-whitepaper.pdf>. 2020.
- [118] *OpenAssets protocol*. <https://github.com/OpenAssets>. 2020.
- [119] *Oracalize website*. URL: <https://www.oracalize.it>.
- [120] *Parity Wallet Security Alert*. <https://paritytech.io/blog/security-alert.html>. 07/2017.
- [121] Anthony Towns Pieter Wuille Jonas Nick. *Taproot: SegWit version 1 spending rules*. BIP 341, <https://github.com/bitcoin/bips/blob/master/bip-0341.mediawiki>. 2020.

- 
- [122] Joseph Poon and Thaddeus Dryja. *The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments*. 2015. URL: <https://lightning.network/lightning-network-paper.pdf>.
- [123] Jeremy Rubin. *CHECKTEMPLATEVERIFY*. BIP 119, <https://github.com/bitcoin/bips/blob/master/bip-0119.mediawiki>. 2020.
- [124] *Understanding the DAO attack*. <http://www.coindesk.com/understanding-dao-hack-journalists/>. 06/2016.
- [125] Gavin Wood. *Ethereum: a secure decentralised generalised transaction ledger*. [gavwood.com/paper.pdf](http://gavwood.com/paper.pdf). 2014.
- [126] Joachim Zahmentferner. *An Abstract Model of UTXO-based Cryptocurrencies with Scripts*. 2018. URL: <https://eprint.iacr.org/2018/469>.
- [127] Joachim Zahmentferner. *Multi-Currency Ledgers*. Cryptology ePrint Archive, Report 2020/895. <https://eprint.iacr.org/2020/895>. 2020.

# Appendices



## 7.4 Supplementary material for Chapter 7

**Full token script** The full token script is the following:

```

1  if not verrec(ptxo(1)) then           // ctxo is a token generator
2      ctxo.op = 0                       // gen action
3  and ctxo.tkid = txid(ptxo(1))        // token id
4  and ptxo(1).val = 0                  // spent txo has 0 BTC
5  and outlen(ctxo) = 1                 // gen has 1 out
6  and ctxo.tkval > 0                   // positive token value
7  else
8  if rtxo.op = 1 then                  /****** BURN *****/
9      versig(ctxo.owner, rtx.wit)      // check owner
10     and verscr(false, rtxo(1))       // make rtx unspendable
11     and inlen(rtxo(1)) = 1           // rtx has 1 in
12     and outlen(rtxo(1)) = 1         // rtx has 1 out
13 else if rtxo.op = 2 then             /****** SPLIT *****/
14     versig(ctxo.owner, rtx.wit)      // check owner
15     and verrec(rtxo(1))              // covenants on rtx
16     and verrec(rtxo(2))
17     and inlen(rtxo(1)) = 1           // rtx has 1 in
18     and outlen(rtxo(1)) = 2         // rtx has 2 outs
19     and rtxo(1).tkval >= 0          // positive token value
20     and rtxo(2).tkval >= 0
21     and rtxo(1).owner = ctxo.owner   // preserve owner
22     and rtxo(1).tkid = ctxo.tkid     // preserve tkid
23     and rtxo(2).tkid = ctxo.tkid
24     and rtxo(1).tkval + rtxo(2).tkval = ctxo.tkval // preserve value
25 else if rtxo.op = 3 then             /****** JOIN *****/
26     inlen(rtxo(1)) = 2               // rtx has 2 in
27     and outlen(rtxo(1)) = 1         // rtx has 1 out
28     and verrec(rtxo(1))              // covenant on rtx
29     and verrec(stxo(2))              // covenants on both inputs
30     and verrec(stxo(1))
31     and ctxo.tkid = rtxo(1).tkid     // preserve token id
32     and versig(ctxo.owner, rtx.wit)  // check owner
33     and rtxo(1).tkval = stxo(1).tkval + stxo(2).tkval //preserve value
34 else if rtxo.op = 4 then             /****** XCHG *****/
35     inlen(rtxo(1)) = 2               // rtx has 2 ins
36     and outlen(rtxo(1)) = 2         // rtx has 2 outs
37     and versig(ctxo.owner, rtx.wit)  // check owner
38     and verrec(stxo(1))              // covenant on input 1
39     and verrec(rtxo(1))              // covenant on rtx(1)
40     and rtxo(1).owner = stxo(2).owner // swap owners
41     and rtxo(2).owner = stxo(1).owner
42     and rtxo(1).tkval = stxo(1).tkval // preserve value
43     and rtxo(1).tkid = stxo(1).tkid // preserve tkid
44     if verrec(stxo(2)) then          /****** EXCHANGE TOKEN/TOKEN *****/
45         verrec(rtxo(2))              // covenant on rtx(2)
46         and rtxo(2).tkval = stxo(2).tkval // preserve tkval
47         and rtxo(2).tkid = stxo(2).tkid // preserve tkid
48     else                               /****** EXCHANGE TOKEN/BTC *****/
49         verscr(versig(ctxo.owner, rtx.wit), rtxo(2))
50         and rtxo(2).val = stxo(2).val // preserve BTC
51 else if rtxo.op = 5 then             /****** GIVE *****/
52     inlen(rtxo(1)) = 1               // rtx has 1 in
53     and outlen(rtxo(1)) = 1         // rtx has 1 out
54     and versig(ctxo.owner, rtx.wit)  // check owner
55     and verrec(rtxo(1))              // covenant on rtx(1)
56     and rtxo(1).tkid = ctxo.tkid     // preserve tkid
57     and rtxo(1).tkval = ctxo.tkval   // preserve value
58 else false

```