

VERIFYING LIQUIDITY OF RECURSIVE BITCOIN CONTRACTS

MASSIMO BARTOLETTI^a, STEFANO LANDE^a, MAURIZIO MURGIA^b, AND ROBERTO ZUNINO^b

^a University of Cagliari, Italy
e-mail address: {bart,lande}@unica.it

^b University of Trento, Italy
e-mail address: {maurizio.murgia,roberto.zunino}@unitn.it

ABSTRACT. Smart contracts — computer protocols that regulate the exchange of crypto-assets in trustless environments — have become popular with the spread of blockchain technologies. A landmark security property of smart contracts is *liquidity*: in a non-liquid contract, it may happen that some assets remain frozen, i.e. not redeemable by anyone. The relevance of this issue is witnessed by recent liquidity attacks to Ethereum, which have frozen hundreds of USD millions. We address the problem of verifying liquidity on BitML, a DSL for smart contracts with a secure compiler to Bitcoin, featuring primitives for currency transfers, contract renegotiation and consensual recursion. Our main result is a verification technique for liquidity. We first transform the infinite-state semantics of BitML into a finite-state one, which focusses on the behaviour of a chosen set of contracts, abstracting from the moves of the context. With respect to the chosen contracts, this abstraction is sound, i.e. if the abstracted contract is liquid, then also the concrete one is such. We then verify liquidity by model-checking the finite-state abstraction. We implement a toolchain that automatically verifies liquidity of BitML contracts and compiles them to Bitcoin, and we assess it through a benchmark of representative contracts.

1. INTRODUCTION

Smart contracts — computer protocols that regulate the exchange of assets in trustless environments — have become popular with the growth of interest in blockchain technologies. Mainstream blockchain platforms like Ethereum, Tezos and Cardano, feature expressive high-level languages for programming smart contracts. This flexibility has a drawback in that it may open the door to attacks that steal or tamper with the assets controlled by vulnerable contracts [ABC17, LCO⁺16].

An alternative approach is to sacrifice the expressiveness of smart contracts to reduce the attack surface. This approach was pursued first by Bitcoin, where transactions can specify simple conditions on how to redeem them, using a limited set of logic, arithmetic, and cryptographic operators. Despite the limited expressiveness of these conditions, it is possible to encode a variety of smart contracts, e.g. gambling games, escrow services, crowdfunding systems, by suitably chaining transactions [ADMM14a, ADMM14c, ADMM16, ABC⁺18, BDM16, BZ17, BK14, KB14, KB16, KMB15, KVV16, MB17]. The common trait

Key words and phrases: Bitcoin, BitML, blockchain, smart contracts, liquidity, verification.

of these works is that they render contracts as cryptographic protocols, where participants can exchange/sign messages, read the blockchain, and append transactions. Verifying the correctness of these protocols is hard, since it requires to reason in a computational model, where participants can manipulate arbitrary bitstrings, only being constrained to use PPTIME algorithms.

Departing from this approach, BitML [BZ18a] allows to write Bitcoin contracts in a high-level, process-algebraic language. BitML features a compiler that translates contracts into sets of standard Bitcoin transactions. The compiler enjoys a computational soundness property, which guarantees that the execution of the compiled contract is coherent with the semantics of the source BitML specification, even in the presence of adversaries.

In this paper we address the problem of verifying BitML contracts, in an extension of BitML with renegotiation and recursion [BMZ20]. In particular, we focus on a landmark property of smart contracts, called *liquidity*, which ensures that funds cannot remain frozen within a contract. Before discussing our main contributions, we overview below BitML and our analysis technique.

1.1. BitML overview. In BitML, any participant can broadcast a *contract advertisement* $\{G\}C$, where C is the contract, specifying the rules to transfer bitcoins (₿), while G is a set of *preconditions* to its stipulation. Preconditions may require participants to deposit some ₿ in the contract, or to commit to some secret. Once $\{G\}C$ 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*. Stipulated contracts have a *balance*, initially set to the sum of the deposits required by its preconditions. This balance is updated when participants execute the contract, e.g. by depositing/withdrawing funds to/from the contract.

A contract C is a *choice* among zero or more branches. Each branch is a *guarded contract*, consisting of one action, and zero or more continuations. The guarded contract

withdraw A

transfers the whole balance to A , and then terminates. The guarded contract

split $\parallel_{i=1}^n w_i \rightarrow C_i$

decomposes the contract into n parallel components C_i , distributing the balance according to the weights w_i . The guarded contract

reveal \vec{a} **if** p

checks that all the secrets \vec{a} have been revealed and satisfy the predicate p (of course, a secret can be revealed only by the participant who has chosen it, as we will see in Section 2).

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 $A : D$, which requires the authorization of A , and the decoration **after** $t : D$, which requires to wait until time t .

Finally, the guarded contract

***** : **rngt** $X(\vec{\mathcal{E}})$

allows participants to renegotiate the contract. This requires first a definition for the variable X , which is given by an equation of the form $X(\vec{\beta}) = \{G\}C$. To finalize the renegotiation, all the participants involved in the current contract must accept the new contract C ,¹ by

¹We use the asterisk in ***** : **rngt** $X(\vec{\mathcal{E}})$ to stress that *all* the participants must accept the renegotiation.

satisfying its precondition G (similarly to stipulation). When this happens, the control passes to C , where the formal parameters $\vec{\beta}$ are instantiated to the actual parameters $\vec{\mathcal{E}}$. Note that C could refer to X , so enabling recursion.

We exemplify BitML by specifying the *timed commitment* contract, a basic protocol to construct more complex contracts, like e.g. lotteries and other games [ADMM14c]. Assume that a participant A wants to choose a secret, and promises to reveal it before some time t . The contract must ensure that if A does not reveal the secret in time, then she will pay a penalty of 1฿ to B (e.g., the opponent in a game). In BitML, this is modelled as follows:

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

The precondition requires A to pay upfront 1฿ , deposited in a transaction x , and to commit to a secret a . The contract is a guarded choice between two branches: any participant can fire an enabled guard, and make the contract evolve to its continuation. The guard of the left branch is `reveal a` , which is enabled only after A reveals the secret. Its continuation `withdraw A` allows anyone to transfer 1฿ to A . The guard of the right branch is `after t :withdraw B` , which is enabled only after time t and allows anyone to transfer 1฿ to B (here there is no continuation, since the `withdraw` terminates). So, before time t , A has the option to reveal a (avoiding the penalty), or to keep it secret (paying the penalty). If no branch is taken by time t , the first participant who fires its `withdraw` gets 1฿ .

1.2. Liquidity. The *liquidity* property requires that the contract balance can always be transferred eventually to some participant: in a non-liquid contract, funds can be frozen forever, unavailable to anyone, hence effectively destroyed². A simple form of liquidity could just require that participants can always cooperate to unfreeze funds. However, this notion would contrast with the setting of smart contracts, where participants are mutually untrusted, and may refuse to cooperate. For instance, consider a contract where A and B contribute 1฿ each for a donation of 2฿ to either C or D :

$$\{A:1@x \mid B:1@y\} (A:B:\text{withdraw } C + A:B:\text{withdraw } D)$$

As in the timed commitment example, this contract is a choice between two branches, both decorated with `A : B : ...`. This means that taking any branch requires the authorization of both users: if A and B disagree on the branch to take, the funds are frozen. When A and B agree on the recipient of the donation, the funds in the contract are unlocked, and they can be transferred to the chosen recipient.

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. Consequently, unless A trusts B , it makes sense to consider this contract as *non-liquid* from A 's point of view (and for similar reasons, also from that of B).

Consider again the timed commitment contract (we omit the preconditions for brevity):

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

This contract is liquid from A 's point of view, *even* if B is dishonest: indeed, A can reveal the secret and then redeem the funds from the contract. The timed commitment is

²To the best of our knowledge, the use of the term “liquidity” to refer to a contract property was first introduced in [TDD⁺18], in the setting of Ethereum contracts.

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, where **A** and **B** have to exchange their secrets or pay a 1฿ penalty, achieving liquidity is a bit more challenging. We first consider a wrong attempt:

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

This contract is liquid for **B**, but not for **A**. Indeed, if **A** performs the `reveal a` action, **B** could refuse to reveal b , making the contract stuck. Instead, **B** can wait time t and then fire `withdraw \mathbf{B}` ; if, in the meanwhile, **A** has fired `reveal a` , **B** can reveal his secret and fire `reveal b` , and then liquidate the contract.

To make the contract liquid for both participants, we amend it as follows:

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

Now, if **A** has fired `reveal a` but **B** refuses to reveal b , after time $t + 1$ she can liquidate the contract by performing the `withdraw \mathbf{A}` .

As a more involved example, consider a recursive variant of the timed commitment:

$$\begin{aligned} \mathbf{X}(n) &= \{\mathbf{A} : \text{secret } a \mid \mathbf{A} : 1\text{฿ } d\} \mathbf{C} \\ \mathbf{C} &= * : \text{rngt } \mathbf{X}(n + 1) + \text{reveal } a. \text{withdraw } \mathbf{A} + \text{after } (t + n) : \text{withdraw } \mathbf{B} \end{aligned}$$

This contract is a toy example of a recursive contract, where $\mathbf{X}(n)$ can be renegotiated, transferring its balance to $\mathbf{X}(n + 1)$. When this happens, **A** must commit to a new secret a , and provide an additional deposit d of 1฿ . Beyond renegotiation, $\mathbf{X}(n)$ allows **A** to reveal her secret and withdraw all the bitcoins deposited in the contract so far. If she does not reveal, **B** can fire the last branch after time $t + n$, transferring the whole balance to himself. This contract is liquid for both **A** and **B**. In every reachable state, a participant can stop renegotiating the current contract $\mathbf{X}(n)$. Then, anyone can liquidate the contract by either waiting until time $t + n$ and then performing the `withdraw \mathbf{B}` , or firing the `withdraw \mathbf{A}` when this action is enabled.

The examples above, albeit elementary, show that detecting if a contract is liquid is not straightforward, in general. Automatic techniques for the verification can be useful tools for the developers of smart contracts.

1.3. Verifying liquidity. One of the main contributions of this paper is a verification technique for the liquidity of BitML contracts. Our technique is based on a more general result, i.e. a correspondence between the concrete semantics of BitML and a new abstract semantics, which is finite-state and correctly approximates the concrete semantics. To obtain a finite-state abstraction, we need to cope with several sources of infiniteness of the concrete semantics: the unbounded passing of time, the stipulation and renegotiation of contracts, and the operations on bitcoin deposits. When studying the liquidity of a set of contracts X from the point of view of a participant **A**, we abstract away all this, by just recording the actions which can be performed on the descendants of X , distinguishing between the actions doable by **A** alone from those which require the cooperation of other participants. This abstraction produces a finite-state transition system, which we model-check for liquidity.

1.4. **Contributions.** We summarise our main contributions as follows:

- We introduce an extension of BitML featuring the renegotiation primitive $* : \text{rngt } X$. Compared to the version in [BZ18a], the current language is more expressive: besides allowing participants to provide new deposits and secrets at run-time, it also allows for *unbounded* recursion, still admitting compilation to Bitcoin.
- We formalize a notion of liquidity (Definition 3.3). With respect to a participant A , a contract is liquid when A alone can ensure that funds do not remain frozen within the contract, even in the presence of adversaries.
- We introduce an abstraction of the semantics of BitML which is finite-state, and sound with respect to the concrete (infinite-state) semantics. Building upon this abstraction, we devise a sound verification technique for liquidity in BitML (Theorem 5.7).
- We develop a toolchain for writing and verifying the liquidity of BitML contracts, and for deploying them on Bitcoin. The toolchain is based on a BitML embedding in Racket [Fla12], which allows for programming BitML contracts within the DrRacket IDE. The toolchain also implements a compiler from BitML contracts to standard Bitcoin transactions.
- We implement a collection of BitML contracts, which we use as a benchmark to evaluate our toolchain. This collection contains a variety of complex contracts, including financial services, auctions, lotteries, and other gambling games.
- We discuss alternative renegotiation primitives, which allow participants to choose some parameters (e.g. the amounts to be deposited) at renegotiation time, to change the set of participants involved in the renegotiated contract, and to renegotiate contracts without the consent of all participants.
- We discuss alternative notions of liquidity, e.g. taking into account participants' strategies.

1.5. **Comparison with previous work.** This paper borrows and extends the contributions of some past papers of ours. BitML was originally introduced in [BZ18a], in a version without renegotiation and recursion. A main limitation of this version was that the participants could not renegotiate the terms of a stipulated contract: this prevented from expressing common financial contracts, where funds have to be added by participants at run-time. Renegotiation and recursion were added in [BMZ20], where we showed that, despite the increased expressiveness, it was still possible to execute BitML on standard Bitcoin, preserving the security guarantees of BitML. These papers did not deal with verification of contracts, and with liquidity. This notion was introduced in [BZ19], in the original version of BitML, i.e. without renegotiation and recursion. The BitML toolchain was first presented in [ABL⁺19], supporting the compilation and verification of contracts in the original version of BitML.

The current paper is the first one which studies liquidity in the full BitML. Recursion adds significant complexity to verification, as it makes the calculus Turing-complete. Because of this, the abstraction in [BZ19] is no longer usable, so the current paper devises an alternative verification technique. The current paper also improves the BitML toolchain, extending the compiler and the liquidity verifier to contracts with renegotiation and recursion.

2. BITML WITH RENEGOTIATION AND RECURSION

We assume a set of *participants*, ranged over by A, B, \dots , a set of *deposit names* x, y, \dots , a set of *deposit variables* d, e, \dots , and a set of *secret names* a, b, \dots . We use χ, χ', \dots to range over deposits (both names and variables), and v, v', w to range over non-negative rational

values. We denote with \mathbf{Hon} the set of the *honest* participants. We denote with \mathcal{S}_A the set of secret names usable by A , requiring that $\mathcal{S}_A \cap \mathcal{S}_B = \emptyset$ if $A \neq B$.

Definition 2.1 (Contract precondition). Contract preconditions have the following syntax (the deposits χ in a contract precondition G must be distinct):

$$\begin{array}{ll}
 G ::= A : v @ \chi & \text{deposit of } v\mathfrak{B} \text{ put by } A \\
 \quad | A : \mathbf{secret} a & \text{secret committed by } A \ (a \in \mathcal{S}_A) \\
 \quad | G \mid G & \text{composition} \quad \diamond
 \end{array}$$

The precondition $A : v @ \chi$ requires A to own $v\mathfrak{B}$ in a deposit χ , and to spend it for stipulating the contract. The precondition $A : \mathbf{secret} a$ requires A to generate a secret a , and commit to it before the contract starts. After stipulation, A can choose whether to disclose the secret a , or not.

To define contracts, we assume a finite set of *recursion variables*, ranged over by X, Y, \dots , and a language of *static expressions* $\mathcal{E}, \mathcal{E}', \dots$, formed by integer constants k , integer variables β, β', \dots , and the usual arithmetic operators. We omit to define the syntax and semantics of static expressions, since they are standard. We assume that a closed static expression evaluates to an integer value. We use the $\vec{\cdot}$ notation for finite sequences.

Definition 2.2 (Contract). Contracts are terms with the syntax in Figure 1, where: (i) the summation $\sum_{i \in I} D_i$ is over a finite set of indices I ; (ii) each recursion variable X has a unique defining equation $X(\vec{\beta}) = \{G\}C$; (iii) renegotiations $* : \mathbf{rngt} X(\vec{\mathcal{E}})$ have the correct number of arguments; (iv) the names \vec{a} in $\mathbf{reveal} \vec{a} \text{ if } p$ are distinct, and they include those in p . We denote with 0 the empty sum. The order of decorations is immaterial, e.g., $\mathbf{after} \mathcal{E} : A : B : D$ is equivalent to $B : A : \mathbf{after} \mathcal{E} : D$. \diamond

A contract C is a choice among guarded contracts. The guarded contract $\mathbf{withdraw} A$ transfers the whole balance to A . A guarded contract $\mathbf{reveal} \vec{a} \text{ if } p. C'$ continues as C' once all the secrets \vec{a} have been revealed and satisfy the predicate p . The guarded contract $\mathbf{split} (w_1 \rightarrow C_1 \mid \dots \mid w_n \rightarrow C_n)$ divides the contract into n contracts C_i , splitting the balance according to the weights w_i . To restrict *who* can execute a branch and *when*, one can use the decoration $A : D$, requiring to wait for A 's authorization, and the decoration $\mathbf{after} \mathcal{E} : D$, requiring to wait until the time specified by the static expression \mathcal{E} . The guarded contract $* : \mathbf{rngt} X(\vec{\mathcal{E}})$ allows the participants involved in the contract to renegotiate it. Intuitively, if $X(\vec{\beta}) = \{G\}C$, then the contract continues as $C\{\vec{\mathcal{E}}/\vec{\beta}\}$ if all the participants mentioned in G give their authorization, and satisfy the precondition G .

Definition 2.3 (Contract advertisement). A contract advertisement is a term $\{G\}^\circ C$, such that: (i) \circ is either empty or a deposit name; (ii) each secret name in C occurs in G ; (iii) G requires a deposit from each A in $\{G\}^\circ C$; (iv) each $* : \mathbf{rngt} X(\vec{\mathcal{E}})$ in C refers to a defining equation $X(\vec{\beta}) = \{G'\}C'$ where the participants in G' are the same as those in G . \diamond

Intuitively, $\{G\}C$ is the advertisement of a contract C with preconditions G , while $\{G\}^x C$ is the advertisement of a renegotiation of an existing contract x . Condition (iii) is used to guarantee that the contract is stipulated only if *all* the involved participants give their consent: namely, A 's consent is rendered as A 's authorization to spend one of her deposits. The last condition is only used to simplify the technical development: we outline in Section 8 how to relax it, by allowing renegotiations to exclude some participants, or to include new ones, which were not among those who originally stipulated the contract.

$C ::= \sum_{i \in I} D_i$	contract	$p ::= true$	truth
$D ::=$	guarded contract	$ p \wedge p$	conjunction
$withdraw\ A$	transfer the balance to A	$ \neg p$	negation
$ reveal\ \vec{a}\ if\ p.\ C$	reveal secrets (if p is true)	$ E = E$	equality
$ split\ _i (w_i \rightarrow C_i)$	split the balance	$ E < E$	less than
$ A : D$	wait for A 's authorization	$E ::= \mathcal{E}$	static expression
$ after\ \mathcal{E} : D$	wait until time \mathcal{E}	$ a$	secret
$ * : rngt\ X \langle \vec{\mathcal{E}} \rangle$	renegotiate the contract	$ E + E$	addition
		$ E - E$	subtraction

Figure 1: Syntax of BitML contracts.

$\Gamma ::=$	configuration	$\xi ::=$	authorization to ...
0	empty	$\# \triangleright \{G\}^\circ C$	commit to $\{G\}^\circ C$
$ \{G\}^\circ C$	contract advertisement	$ x \triangleright \{G\}^\circ C$	spend x for $\{G\}^\circ C$
$ \langle C, v \rangle_x$	active contract storing $v\mathfrak{B}$	$ x \triangleright D$	take branch D in contract x
$ \langle A, v \rangle_x$	A 's deposit of $v\mathfrak{B}$	$ x, y \triangleright \langle A, v \rangle$	join deposits x and y
$ A[\xi]$	A 's authorization for ξ	$ x \triangleright \langle A, v \rangle, \langle A, v' \rangle$	split deposit x in two
$ \{A : a\#N\}$	committed secret of A	$ x \triangleright B$	donate deposit x to B
$ A : a\#N$	revealed secret of A	$ \vec{x}, i \triangleright y$	destroy i -th deposit in \vec{x}
$ t$	global time		
$ \Gamma \mid \Gamma'$	composition		

Figure 2: Configurations and authorizations.

We now define the semantics of BitML, starting from its configurations.

Definition 2.4 (Configuration). Configurations are terms with the syntax in Figure 2, where: (i) in a committed secret, $N \in \mathbb{N} \cup \{\perp\}$ (where \perp denotes an ill-formed commitment); (ii) in a revealed secret, $N \in \mathbb{N}$; (iii) in a configuration there are no duplicate authorizations; (iv) in a configuration containing $\langle \dots \rangle_x$ and $\langle \dots \rangle_y$, it must be $x \neq y$; (v) there exists at most one term t . We assume that $(\mid, 0)$ is a commutative monoid, and we denote indexed parallel compositions with $||_i$. We say that Γ is *initial* when it contains only deposits (i.e., terms $\langle A, v \rangle_x$), and that it is a *timed configuration* when it contains a term t . We denote with $\text{cn}(\Gamma)$ the set of contract names x such that Γ contains $\langle C, v \rangle_x$, for some C and v . \diamond

The intuition behind the various terms in configurations is the following:

- $\langle C, v \rangle_x$ is a stipulated contract storing $v\mathfrak{B}$, uniquely identified by the name x ;
- $\langle A, v \rangle_x$ is a deposit of $v\mathfrak{B}$ owned by A , and uniquely identified by the name x ;
- $A[\xi]$ is A 's *authorizations* to perform some action ξ ;
- $\{A : a\#N\}$ represents A 's commitment to a secret N , identified by a ;
- $A : a\#N$ represents a secret N , identified by a , and revealed by A .

Definition 2.5 (BitML semantics). The semantics of BitML is a Labelled Transition System (LTS) between timed configurations. In the rest of this section we describe the reduction

rules of the LTS, which implicitly define the labels ℓ . A *concrete run* \mathcal{R} is a sequence $\Gamma_0 \xrightarrow{\ell_0}_b \Gamma_1 \xrightarrow{\ell_1}_b \dots$, where Γ_0 is timed. If \mathcal{R} is finite, we write $\Gamma_{\mathcal{R}}$ for the untimed part of its last configuration.

Below we gently introduce the BitML semantics, first illustrating each construct through examples, and then giving the general rule. Labels represent the actions performed by participants. A decoration $\mathbf{A} : \dots$ in the label means that the action can be performed only by \mathbf{A} . The absence of such a decoration means that the action can be performed by anyone. Note that labels are not instrumental to define the BitML semantics: yet, they are essential to the definition of liquidity, since there we need to associate actions to the participants who can perform them. In the examples, we will omit the labels.

Deposits. A deposit $\langle \mathbf{A}, v \rangle_x$ can be reduced in several ways: it can be split into smaller deposits, joined with another deposit, transferred to another participant, or destroyed. In all cases, its owner \mathbf{A} must first authorise the action. The reduction rules for deposits are detailed in Figure 3. Rule [DEP-AUTHJOIN] allows \mathbf{A} to authorize the merge of two deposits x, y into a single one, creating the needed authorization. The label of the form $\mathbf{A} : \dots$ records that only \mathbf{A} can perform this move. Rule [DEP-JOIN] uses this authorization to create a single deposit z of \mathbf{A} . The rules [DEP-AUTHDIVIDE] and [DEP-DIVIDE] act similarly, allowing a deposit of \mathbf{A} to be divided in two parts. The rules [DEP-AUTHDONATE] and [DEP-DONATE] allow \mathbf{A} to transfer one of her deposits to another participant. The rules [DEP-AUTHDESTROY] and [DEP-DESTROY] allow a set of participants to destroy a set of deposits $x_1 \dots x_n$. To do that, first each participant \mathbf{A}_i must provide the needed authorization $\mathbf{A}_i[\vec{x}, i \triangleright y]$ for their own deposit x_i . When all the authorizations have been collected, rule [DEP-DESTROY] eliminates the deposits. The last two rules in Figure 3 are needed to properly represent the fact that computational participants can create (and put on the ledger) transactions without a counterpart in the symbolic model. To achieve a meaningful correspondence between the symbolic and the computational models, putting on the ledger such transactions is rendered with the rule [DEP-DESTROY].

Stipulation: advertisement. Any participant can broadcast a contract advertisement $\{\mathbf{G}\}\mathbf{C}$, provided that all the deposits mentioned in \mathbf{G} exist in the current configuration, and that the names of the secrets in \mathbf{G} are fresh. This is formalised by the following rule:

$$\frac{\forall \mathbf{A}: v @ x \text{ in } \mathbf{G} : \langle \mathbf{A}, v \rangle_x \text{ in } \Gamma \quad \text{all secrets in } \mathbf{G} \text{ fresh}}{\Gamma \xrightarrow{\text{adv}(\{\mathbf{G}\}\mathbf{C})}_b \{\mathbf{G}\}\mathbf{C} \mid \Gamma} \text{[C-ADV]}$$

We exemplify this and the following rules through a running example. Let:

$$\mathbf{G} = \mathbf{A}: 1 @ x \mid \mathbf{B}: 1 @ d \mid \mathbf{A}: \text{secret } a$$

This precondition requires both \mathbf{A} and \mathbf{B} to deposit $1\mathfrak{B}$, but \mathbf{A} 's deposit name x is known, while \mathbf{B} 'one is not known yet, so we refer to it through a deposit variable d . Let \mathbf{C} be an arbitrary contract involving only \mathbf{A} and \mathbf{B} , and let $\Gamma_0 = \Gamma \mid \langle \mathbf{A}, 1 \rangle_x \mid \langle \mathbf{B}, 1 \rangle_y$ for some Γ . By rule [C-ADV], the configuration Γ_0 can take the transition:

$$\Gamma_0 \rightarrow_b \Gamma_0 \mid \{\mathbf{G}\}\mathbf{C} = \Gamma_1$$

$$\begin{array}{c}
\frac{\frac{\langle \mathbf{A}, v \rangle_x \mid \langle \mathbf{A}, v' \rangle_y \mid \Gamma \xrightarrow{\mathbf{A}:x,y} \langle \mathbf{A}, v \rangle_x \mid \langle \mathbf{A}, v' \rangle_y \mid \mathbf{A}[x, y \triangleright \langle \mathbf{A}, v + v' \rangle] \mid \Gamma}{\Gamma = \mathbf{A}[x, y \triangleright \langle \mathbf{A}, v + v' \rangle] \mid \mathbf{A}[y, x \triangleright \langle \mathbf{A}, v + v' \rangle] \mid \Gamma' \quad z \text{ fresh}}}{\langle \mathbf{A}, v \rangle_x \mid \langle \mathbf{A}, v' \rangle_y \mid \Gamma \xrightarrow{\text{join}(x,y)} \langle \mathbf{A}, v + v' \rangle_z \mid \Gamma'} \text{[DEP-AUTHJOIN]} \\
\\
\frac{\frac{\langle \mathbf{A}, v + v' \rangle_x \mid \Gamma \xrightarrow{\mathbf{A}:x,v,v'} \langle \mathbf{A}, v + v' \rangle_x \mid \mathbf{A}[x \triangleright \langle \mathbf{A}, v \rangle, \langle \mathbf{A}, v' \rangle] \mid \Gamma}{\Gamma = \mathbf{A}[x \triangleright \langle \mathbf{A}, v \rangle, \langle \mathbf{A}, v' \rangle] \mid \Gamma' \quad y, y' \text{ fresh}}}{\langle \mathbf{A}, v + v' \rangle_x \mid \Gamma \xrightarrow{\text{divide}(x,v,v')} \langle \mathbf{A}, v \rangle_y \mid \langle \mathbf{A}, v' \rangle_{y'} \mid \Gamma'} \text{[DEP-AUTHDIVIDE]} \\
\\
\frac{\frac{\langle \mathbf{A}, v \rangle_x \mid \Gamma \xrightarrow{\mathbf{A}:x,B} \langle \mathbf{A}, v \rangle_x \mid \mathbf{A}[x \triangleright B] \mid \Gamma}{\Gamma = \mathbf{A}[x \triangleright B] \mid \Gamma' \quad y \text{ fresh}}}{\langle \mathbf{A}, v \rangle_x \mid \Gamma \xrightarrow{\text{donate}(x,B)} \langle B, v \rangle_y \mid \Gamma'} \text{[DEP-AUTHDONATE]} \\
\\
\frac{\frac{\vec{x} = x_1 \cdots x_n \quad j \in 1..n \quad y \text{ fresh (except in destroy authorizations for } \vec{x})}{(\|_{i=1}^n \langle \mathbf{A}_i, v_i \rangle_{x_i}) \mid \Gamma \xrightarrow{\mathbf{A}_j:\vec{x},j} (\|_{i=1}^n \langle \mathbf{A}_i, v_i \rangle_{x_i}) \mid \mathbf{A}_j[\vec{x}, j \triangleright y] \mid \Gamma}{\vec{x} = x_1 \cdots x_n \quad \Gamma = (\|_{i=1}^n \mathbf{A}_i[\vec{x}, i \triangleright y]) \mid \Gamma'} \text{[DEP-AUTHDESTROY]} \\
\\
\frac{\Gamma = (\|_{i=1}^n \mathbf{A}_i[\vec{x}, i \triangleright y]) \mid \Gamma'}{(\|_{i=1}^n \langle \mathbf{A}_i, v_i \rangle_{x_i}) \mid \Gamma \xrightarrow{\text{destroy}(\vec{x})} \Gamma'} \text{[DEP-DESTROY]}
\end{array}$$

Figure 3: Semantics of Bitcoin deposits.

Stipulation: commitment. To stipulate an advertised contract, all the participants mentioned in it must fulfill the preconditions, by making available the required deposits, and committing to the required secrets. In our example, \mathbf{A} has one secret to commit, so she can perform the following step:

$$\Gamma_1 \rightarrow_b \Gamma_1 \mid \{\mathbf{A} : a \# N\} \mid \mathbf{A}[\# \triangleright \{\mathbf{G}\}\mathbf{C}] = \Gamma_2 \quad (2.1)$$

where the term $\{\mathbf{A} : a \# N\}$ represents \mathbf{A} 's commitment to the secret N , while $\mathbf{A}[\# \triangleright \{\mathbf{G}\}\mathbf{C}]$ represents finalising the commitment phase for \mathbf{A} . Participant \mathbf{B} has no secrets to commit, but he must choose one of his deposits (e.g., $\langle \mathbf{B}, 1 \rangle_y$) to fulfill the precondition $\mathbf{B}:1 @ d$:

$$\Gamma_2 \rightarrow_b \Gamma_2 \mid \mathbf{B} : d \leftarrow y \mid \mathbf{B}[\# \triangleright \{\mathbf{G}\}\mathbf{C}] = \Gamma_3 \quad (2.2)$$

In general, these steps are formalised by the following rule:

$$\frac{\begin{array}{l}
\Delta_s = \|\|_{i=1}^k \{\mathbf{A} : a_i \# N_i\} \\
a_1 \cdots a_k \text{ are all the secrets of } \mathbf{A} \text{ in } \mathbf{G} \\
\forall i \in 1..k : \#N : \{\mathbf{A} : a_i \# N\} \text{ in } \Gamma \\
\forall i \in 1..k : \#N : \mathbf{A} : a_i \# N \text{ in } \Gamma \\
\forall i \in 1..k : N_i \in \begin{cases} \mathbb{N} & \text{if } \mathbf{A} \in \text{Hon} \\ \mathbb{N} \cup \{\perp\} & \text{otherwise} \end{cases}
\end{array} \quad \begin{array}{l}
\Delta_d = \|\|_{i=1}^h \mathbf{A} : d_i \leftarrow x_i \\
d_1 \cdots d_h \text{ are all the deposits of } \mathbf{A} \text{ in } \mathbf{G} \\
\forall i \in 1..h : \mathbf{A} : v_i @ d_i \text{ in } \mathbf{G} \\
\forall i \in 1..h : \#x : (\mathbf{A} : d_i \leftarrow x) \text{ in } \Gamma \\
\forall i \in 1..h : \langle \mathbf{A}, v_i \rangle_{x_i} \text{ in } \Gamma \\
\forall i \neq j \in 1..h : x_i \neq x_j \\
\forall x, v, i \in 1..h : \mathbf{A} : v @ x \text{ in } \mathbf{G} \implies x_i \neq x
\end{array} \text{[C-AUTHCOMMIT]} \\
\hline
\{\mathbf{G}\}^\circ \mathbf{C} \mid \Gamma \xrightarrow{\mathbf{A}:\{\mathbf{G}\}^\circ \mathbf{C}, \Delta_s | \Delta_d} \{\mathbf{G}\}^\circ \mathbf{C} \mid \Gamma \mid \Delta_s \mid \Delta_d \mid \mathbf{A}[\# \triangleright \{\mathbf{G}\}^\circ \mathbf{C}]
\end{array}$$

The rule precondition ensures that the final configuration fulfills the conditions required by G . We use the notation $\{G\}^\circ C$ to refer to a contract advertisement of the form $\{G\}C$ or $\{G\}^x C$. The case $\{G\}C$ corresponds to the original contract stipulation, while the other case corresponds to renegotiation. In this way, we can use the same rule in both situations. Note that condition (iii) in Definition 2.4 ensures that rule [C-AUTHCOMMIT] cannot be used more than once to generate the same authorization. The same is true for all the other rules that generate authorizations.

Stipulation: authorization. Back to our example, in the configuration Γ_3 of (2.2), A and B must perform an additional sequence of steps to authorize the transfer of their deposits x , y to the contract:

$$\Gamma_3 \rightarrow_b \Gamma_3 \mid A[x \triangleright \{G\}C] \rightarrow_b \Gamma_3 \mid A[x \triangleright \{G\}C] \mid B[y \triangleright \{G\}C] = \Gamma_4 \quad (2.3)$$

where the terms $A[x \triangleright \{G\}C]$ and $B[y \triangleright \{G\}C]$ represent the authorizations to spend x and y for stipulation. In general, these steps are obtained through the following rule:

$$\frac{\forall B \text{ in } G : B[\# \triangleright \{G\}^\circ C] \text{ in } \Gamma \quad A : v \otimes \chi \text{ in } G \quad \Gamma \vdash \chi = x}{\{G\}^\circ C \mid \Gamma \xrightarrow{A:\{G\}^\circ C, x}_b \{G\}^\circ C \mid \Gamma \mid A[x \triangleright \{G\}^\circ C]} \text{[C-AUTHINITDEP]}$$

The first premise requires that all participants have finalised the commitment phase. While rule [C-AUTHCOMMIT] allows a participant to add all her commitments to the configuration in a single step, each application of rule [C-AUTHINITDEP] allows one to authorize the spending of a single deposit. If a deposit variable d was used in G , the relation $\Gamma \vdash d = x$ ensures that the configuration contains a binding $A : d \leftarrow x$. The relation \vdash is defined as follows:

$$\Gamma \vdash x = x \quad \Gamma \mid A : d \leftarrow x \vdash d = x$$

Stipulation: initialization. In the configuration Γ_4 of (2.3) all the needed authorizations have been granted, so the advertisement can be turned into an active contract. This step consumes the deposits and the authorizations, and it initializes the new contract, with a fresh name z , and with a balance corresponding to the sum of all the consumed deposits:

$$\Gamma_4 \rightarrow_b \Gamma \mid \{A : a \# N\} \mid \langle C, 2 \rangle_z$$

In general, this step is obtained through the following rule:

$$\frac{\begin{array}{l} G = (\parallel_{i \in I} A_i : v_i \otimes x_i) \mid (\parallel_{i \in J} B_i : v'_i \otimes d_i) \mid (\parallel_{i \in K} C_i : \text{secret } a_i) \quad z \text{ fresh} \\ \Delta = (\parallel_{i \in I} \langle A_i, v_i \rangle_{x_i}) \mid (\parallel_{i \in J} \langle B_i, v'_i \rangle_{x'_i}) \mid (\parallel_{i \in J} B_i : d_i \leftarrow x'_i) \mid \\ (\parallel_{i \in I} A_i [x_i \triangleright \{G\}C]) \mid (\parallel_{i \in J} B_i [x'_i \triangleright \{G\}C]) \mid (\parallel_{A \in G} A[\# \triangleright \{G\}C]) \end{array}}{\{G\}C \mid \Delta \mid \Gamma \xrightarrow{\text{init}(G, C)}_b \langle C, \sum_{i \in I} v_i + \sum_{i \in J} v'_i \rangle_z \mid \Gamma} \text{[C-INIT]}$$

Note that the part $\{G\}C \mid \Delta$ of the configuration contains all the terms that are consumed by the step. The following rules define the behaviour of a contract after stipulation.

Withdraw. Executing `withdraw A` terminates the contract, and transfers its balance to `A`:

$$\langle \text{withdraw } A, v \rangle_x \rightarrow_b \langle A, v \rangle_y$$

After the contract x is terminated, a fresh deposit of $v\mathfrak{B}$ owned by `A` is created. The general rule is the following:

$$\frac{y \text{ fresh}}{\langle \text{withdraw } A, v \rangle_x \mid \Gamma \xrightarrow{\text{withdraw}(A, v, x)}_b \langle A, v \rangle_y \mid \Gamma} \text{[C-WITHDRAW]}$$

The case where the action `withdraw A` has an alternative branch is dealt with by the rule [C-BRANCH], discussed below.

Split. The `split` primitive divides the contract balance in parts, each one controlled by its own contract. For instance:

$$\langle (\text{split } 2 \rightarrow C_1 \mid 3 \rightarrow C_2), 5 \rangle_x \rightarrow_b \langle C_1, 2 \rangle_y \mid \langle C_2, 3 \rangle_z$$

After this step, the new spawned contracts C_1 and C_2 are executed concurrently. The general rule is the following:

$$\frac{w = \sum_{i=1}^k w_i \quad v_i = (v \cdot w_i) / w \quad y_1 \cdots y_k \text{ fresh}}{\langle \text{split } \parallel_{i=1}^k (w_i \rightarrow C_i), v \rangle_x \mid \Gamma \xrightarrow{\text{split}(x)}_b (\parallel_{i=1}^k \langle C_i, v_i \rangle_{y_i}) \mid \Gamma} \text{[C-SPLIT]}$$

Note that the weights w_i in the `split` do not represent actual \mathfrak{B} values, but the proportion w.r.t. the contract balance. For instance:

$$\langle (\text{split } 2 \rightarrow C_1 \mid 3 \rightarrow C_2), 10 \rangle_x \rightarrow_b \langle C_1, 4 \rangle_y \mid \langle C_2, 6 \rangle_z$$

Revealing secrets. Any participant can reveal one of her secrets, using the rule:

$$\frac{N \neq \perp}{\{A : a \# N\} \mid \Gamma \xrightarrow{A:a}_b A : a \# N \mid \Gamma} \text{[C-AUTHREV]}$$

The premise $N \neq \perp$ is needed to avoid the case where a participant does not know the secret she has committed to. Indeed, at the level of Bitcoin, commitments are represented as cryptographic hashes of bitstrings, and revealing a secret amounts to broadcasting a preimage, i.e. a value whose hash is equal to the committed value. If a participant commits to a random value, then with overwhelming probability she will not be able to provide a preimage. The label $A : a$ represents the fact that only `A`, the participant who performed the commitment, can fire the transition.

Reveal. The prefix `reveal \vec{a} if p` can be fired if all the committed secrets \vec{a} have been revealed, and satisfy the guard p . For instance, if $\Gamma = A : a \# N \mid B : b \# M$:

$$\langle \text{reveal } ab \text{ if } a = b. C, v \rangle_x \mid \Gamma \rightarrow_b \langle C, v \rangle_y \mid \Gamma \quad \text{if } M = N$$

The general rule is the following:

$$\frac{\vec{a} = a_1 \cdots a_n \quad \Delta = \parallel_{i=1}^n B_i : a_i \# N_i \quad \llbracket p \rrbracket_\Delta = \text{true} \quad y \text{ fresh}}{\langle \text{reveal } \vec{a} \text{ if } p. C, v \rangle_x \mid \Gamma \mid \Delta \xrightarrow{\text{rev}(\vec{a}, y)}_b \langle C, v \rangle_y \mid \Gamma \mid \Delta} \text{[C-REV]}$$

where the semantics of predicates $\llbracket p \rrbracket_\Delta$ is defined by the following equations:

$$\begin{aligned} \llbracket true \rrbracket_\Delta &= true & \llbracket p_1 \wedge p_2 \rrbracket_\Delta &= \llbracket p_1 \rrbracket_\Delta \text{ and } \llbracket p_2 \rrbracket_\Delta & \llbracket \neg p \rrbracket_\Delta &= \text{not } \llbracket p \rrbracket_\Delta \\ \llbracket a \rrbracket_\Delta &= N \text{ if } \Delta \text{ contains } \mathbf{A} : a \# N & \llbracket E_1 \bullet E_2 \rrbracket_\Delta &= \llbracket E_1 \rrbracket_\Delta \bullet \llbracket E_2 \rrbracket_\Delta & (\bullet \in \{+, -\}) \\ \llbracket N \rrbracket_\Delta &= N & \llbracket E_1 \circ E_2 \rrbracket_\Delta &= \llbracket E_1 \rrbracket_\Delta \circ \llbracket E_2 \rrbracket_\Delta & (\circ \in \{=, <\}) \end{aligned}$$

Authorizing branches. A branch $\mathbf{A} : D$ can be taken only provided that \mathbf{A} has granted her authorization. This can be done through the following rule:

$$\frac{}{\langle \mathbf{A} : D + C, v \rangle_x \mid \Gamma \xrightarrow{\mathbf{A}:(x, \mathbf{A} : D)} \langle \mathbf{A} : D + C, v \rangle_x \mid \mathbf{A}[x \triangleright \mathbf{A} : D] \mid \Gamma} \text{[C-AUTHBRANCH]}$$

Reducing branches. Once all the authorizations for a branch occur in the configuration, anyone can trigger the transition, provided that the time constraints (if any) are respected. For instance, if $D = \mathbf{A} : \text{after } 1000 : \text{withdraw } \mathbf{B}$, we have the transition:

$$\langle D + C, v \rangle_x \mid \mathbf{A}[x \triangleright D] \mid 1051 \rightarrow_b \langle \mathbf{B}, v \rangle_y \mid 1051$$

The general rule is the following:

$$\frac{\langle D', v \rangle_x \mid \Gamma \xrightarrow{\ell} \Gamma' \quad D = \mathbf{A}_1 : \dots : \mathbf{A}_k : \text{after } t_1 : \dots : \text{after } t_m : D' \quad D' \neq \mathbf{A} : \dots \quad x \notin \text{cn}(\Gamma') \quad t \geq t_1, \dots, t_m \quad D' \neq \text{after } t' : \dots}{\langle D + C, v \rangle_x \mid \parallel_{i=1}^k \mathbf{A}_i[x \triangleright D] \mid \Gamma \mid t \xrightarrow{\ell} \Gamma' \mid t} \text{[C-BRANCH]}$$

Delaying. In any configuration, we always allow time to advance:

$$\frac{\delta > 0}{\Gamma \mid t \xrightarrow{\delta} \Gamma \mid t + \delta} \text{[C-DELAY]}$$

Renegotiation: advertisement. Contract renegotiation is similar to stipulation, including advertisement, commitment, authorization, and contract initialization. We illustrate these phases through a running example. Consider a configuration:

$$\Gamma_0 = \langle * : \text{rngt } \mathbf{X}(k) + C_{alt}, v \rangle_x \mid \Gamma \quad \text{where } \mathbf{X}(\beta) = \{G\}C$$

where C_{alt} contains the branches alternative to rngt , $G = \mathbf{A} : 1 @ d \mid \mathbf{B} : 1 @ e \mid \mathbf{A} : \text{secret } a$, and C is an arbitrary contract involving only \mathbf{A} and \mathbf{B} , possibly containing the integer variable β in static expressions. The renegotiation is advertised as follows:

$$\Gamma_0 \rightarrow_b \Gamma_0 \mid \{G'\}^x C' = \Gamma_1 \quad (2.4)$$

where $\{G'\}^x C'$ is obtained by transforming $\{G\}C$ as follows: (i) variables d, e are renamed into fresh ones d', e' , and similarly the secret name a into a' , (ii) the static expressions in C are evaluated, assuming $\beta = k$, and replaced with their results. The superscript x in

$\{G'\}^x C'$ is used to record that, when the renegotiation is concluded, the contract x must be closed. In general, this step is defined through the following rule:

$$\frac{\begin{array}{l} \{G\}C \equiv \mathbf{X}\langle\vec{\mathcal{E}}\rangle \\ \text{all secrets in } G \text{ fresh} \\ \forall A: v @ d \text{ in } G : d \text{ fresh} \end{array} \quad \begin{array}{l} \forall \{G'\}^x C' \text{ in } \Gamma : \{G'\}C' \not\equiv \mathbf{X}\langle\vec{\mathcal{E}}\rangle \\ \forall A_i: v_i @ x_i \text{ in } G : \langle A_i, v_i \rangle_{x_i} \text{ in } \Gamma \end{array}}{\begin{array}{l} \langle * : \text{rngt } \mathbf{X}\langle\vec{\mathcal{E}}\rangle + C_{\text{alt}}, v \rangle_x \mid \Gamma \xrightarrow{\text{advRngt}(\{G\}^x C)} \rangle_b \\ \langle * : \text{rngt } \mathbf{X}\langle\vec{\mathcal{E}}\rangle + C_{\text{alt}}, v \rangle_x \mid \Gamma \mid \{G\}^x C \end{array}} \quad [\text{C-ADV RNGT}]$$

The relation $\{G\}C \equiv \mathbf{X}\langle\vec{\mathcal{E}}\rangle$ used in the rule premise holds when, for some G', C' , there exists a defining equation $\mathbf{X}\langle\vec{\beta}\rangle = \{G'\}C'$ such that $\{G\}C$ is the transformation of $\{G'\}C'$ obtained by instantiating the formal parameters with the actual ones, and by α -converting the secret names and deposit variables, as done e.g. in Equation (2.4). Before formalising this transformation below in Equation (2.5), we need a few auxiliary notions. We denote with $\llbracket \mathcal{E} \rrbracket$ the evaluation of a closed static expression \mathcal{E} (the actual definition is standard, so we omit it). We overload $\llbracket - \rrbracket$ to contracts and contract advertisements: namely, $\llbracket C \rrbracket$ is the contract obtained by substituting all the occurring static expressions with their valuation, and $\llbracket \{G\}C \rrbracket = \{G\}\llbracket C \rrbracket$. Then, we denote with $\{G\}C \equiv_\alpha \{G'\}C'$ the α -equivalence between two contract advertisements w.r.t. secret names and deposit variables. Finally, we define:

$$\{G\}C \equiv \mathbf{X}\langle\vec{\mathcal{E}}\rangle \iff \exists G', C' : \mathbf{X}\langle\vec{\beta}\rangle = \{G'\}C' \text{ and } \llbracket \{G'\}C' \llbracket \llbracket \vec{\mathcal{E}} \rrbracket / \vec{\beta} \rrbracket \rrbracket \equiv_\alpha \{G\}C \quad (2.5)$$

Renegotiation: commitment. In the subsequent steps participants choose the actual deposit names, and **A** commits to her secret. If **A** owns in Γ a deposit $\langle A, 1 \rangle_y$, she can choose $d' = y$ to satisfy the precondition G . Similarly, **B** can choose $e' = z$ if he owns such a deposit in Γ . These choices are performed as follows:

$$\begin{array}{l} \Gamma_1 \rightarrow_b \Gamma_1 \mid A : d' \leftarrow y \mid \{A : a' \# N\} \mid A[\# \triangleright \{G'\}^x C'] \quad = \Gamma_2 \\ \Gamma_2 \rightarrow_b \Gamma_2 \mid B : e' \leftarrow z \mid B[\# \triangleright \{G'\}^x C'] \quad = \Gamma_3 \end{array}$$

The general case is defined by the rule [C-AUTHCOMMIT] described before.

Renegotiation: authorization. At this point, participants must authorise to spend their deposits and the balance of the contract at x . This is done through a series of steps:

$$\begin{array}{l} \Gamma_3 \rightarrow_b \Gamma_3 \mid A[y \triangleright \{G'\}^x C'] \\ \rightarrow_b \Gamma_3 \mid A[y \triangleright \{G'\}^x C'] \mid A[x \triangleright \{G'\}^x C'] \quad = \Gamma_4 \\ \rightarrow_b \Gamma_4 \mid B[z \triangleright \{G'\}^x C'] \\ \rightarrow_b \Gamma_4 \mid B[z \triangleright \{G'\}^x C'] \mid B[x \triangleright \{G'\}^x C'] \quad = \Gamma_5 \end{array}$$

The general case is defined by the rule [C-AUTHINITDEP] described above.

Renegotiation: initialization. Finally, the renegotiated contract is stipulated. This amounts to closing the old contract x , consuming the deposits y and z , and transferring the balance v of the old contract to the new one, which is given a fresh name x' :

$$\Gamma_5 \rightarrow_b \langle C', v + 2 \rangle_{x'} \mid \Gamma'$$

where Γ' is Γ without the deposits y and z . Note that the branches in C_{alt} are discarded only in the last step above, where we complete the renegotiation. Before this step, it would have been possible to take one of the branches in C_{alt} , aborting the renegotiation.

The initialization step is defined by the following rule, which is analogous to [C-INIT]:

$$\frac{\begin{array}{l} \{G\}C \equiv \mathbf{X}(\vec{\mathcal{E}}) \quad y \text{ fresh} \\ G = (\|_{i \in I} A_i : v_i @ x_i) \mid (\|_{i \in J} B_i : v'_i @ d_i) \mid (\|_{i \in K} C_i : \text{secret } a_i) \\ \Delta = (\|_{i \in I} \langle A_i, v_i \rangle_{x_i}) \mid (\|_{i \in J} \langle B_i, v'_i \rangle_{x'_i}) \mid (\|_{i \in J} B_i : d_i \leftarrow x'_i) \mid \\ (\|_{i \in I} A_i [x_i \triangleright \{G\}^x C]) \mid (\|_{i \in J} B_i [x'_i \triangleright \{G\}^x C]) \mid \\ (\|_{A \in G} A [\# \triangleright \{G\}^x C] \mid A [x \triangleright \{G\}^x C]) \end{array}}{\langle * : \text{rngt } \mathbf{X}(\vec{\mathcal{E}}) + C_{alt}, v \rangle_x \mid \{G\}^x C \mid \Gamma \mid \Delta \xrightarrow{\text{rngt}(x, G, C)}_b \langle C, v + \sum_{i \in I} v_i + \sum_{i \in J} v'_i \rangle_y \mid \Gamma} \text{ [C-RNGT]}$$

The main difference between [C-INIT] and [C-RNGT] is that the latter transfers the balance v of the old contract to the new one.

Executing BitML on Bitcoin. Stipulating or renegotiating a BitML contract C in Bitcoin requires each participant to invoke the BitML compiler, which was first introduced in [BZ18a] and then extended in [BMZ20] with renegotiation and recursion. In particular, each participant has to (i) generate a key pair for each subcontract of C , (ii) exchange the generated public keys with the other participants, (iii) sign each subcontract of C , and (iv) exchange all the signatures. The whole protocol is detailed in [BZ18b] (Definition 21). When dealing with a contract C among N participants and having M subcontracts, such protocol requires $O(NM)$ broadcasts. After stipulation/renegotiation, executing a step of the BitML semantics corresponds to appending a transaction to the Bitcoin blockchain.

3. LIQUIDITY

In this section we formalise a notion of contract liquidity. Aiming at generality, we parameterise this notion over:

- an LTS \rightarrow , which models the contract behaviour;
- a subset \mathcal{L}_A of the labels of the LTS \rightarrow , which represents the moves that can be performed by an *honest* participant A , without requiring the cooperation of the other participants.

Once these parameters are fixed, we define when, in a configuration Γ of the LTS, a set X of contracts in Γ is liquid. Roughly, this happens when the honest participant A can always make the funds stored by the contracts X be transferred to some participant. In the meanwhile, the other participants may play against her, e.g. by not revealing some secrets, or by not granting their authorizations for some branch. Note that by suitably instantiating the parameters \rightarrow and \mathcal{L}_A , we will be able to use the same notion of liquidity both with the concrete and with the abstract BitML semantics. For instance, for the concrete semantics we choose $\rightarrow = \rightarrow_b$, and as \mathcal{L}_A all the labels except those of the form $B : \ell$ with $B \neq A$.

$$\begin{aligned}
orig_{\Gamma}(\Gamma, x) &= x \quad \text{if } x \in \text{cn}(\Gamma) \\
orig_{\Gamma}(\mathcal{R} \xrightarrow{\ell} \Gamma', x) &= \begin{cases} orig_{\Gamma}(\mathcal{R}, x) & \text{if } x \in \text{cn}(\Gamma_{\mathcal{R}}) \cap \text{cn}(\Gamma') \\ orig_{\Gamma}(\mathcal{R}, y) & \text{if } x \in \text{cn}(\Gamma') \setminus \text{cn}(\Gamma_{\mathcal{R}}) \text{ and } \{y\} = \text{cn}(\Gamma_{\mathcal{R}}) \setminus \text{cn}(\Gamma') \end{cases}
\end{aligned}$$

Figure 4: Origin of a contract name within a run.

We start by introducing an auxiliary partial function $orig_{\Gamma}(\mathcal{R}, x)$ that, given a contract name x and a run \mathcal{R} starting from Γ , determines the (unique) ancestor y of x in Γ , if any. Intuitively, $orig_{\Gamma}(\mathcal{R}, x) = y$ means that y has evolved along the run \mathcal{R} , eventually leading to x , and possibly to other contracts. In BitML, this happens in one of the following cases. First, a **split** can spawn new contracts, e.g.:

$$\langle \text{split } (v_1 \rightarrow C_1 \mid v_2 \rightarrow C_2), v_1 + v_2 \rangle_x \xrightarrow{\text{split}(x)}_b \langle C_1, v_1 \rangle_{y_1} \mid \langle C_2, v_2 \rangle_{y_2}$$

Here, both y_1 and y_2 have x as ancestor. Second, **reveal** reduces as follows:

$$\langle \text{reveal } a.C, v \rangle_x \mid \dots \xrightarrow{\text{reveal}(a,x)}_b \langle C, v \rangle_y \mid \dots$$

In this case, the ancestor of y is x . Third, a **rngt** evolves as follows:

$$\langle * : \text{rngt } X(\vec{\mathcal{E}}), v \rangle_x \mid \dots \xrightarrow{\text{init}(x,G,C)}_b \langle C, v + \dots \rangle_y \mid \dots$$

Also in this case, the ancestor of y is x .

Definition 3.1 (Origin). Let \mathcal{R} be a run starting from Γ , and let x be a contract name. We define the partial function $orig_{\Gamma}(\mathcal{R}, x)$ by induction on the length of \mathcal{R} in Figure 4.

Note that, in Figure 4, the condition $x \in \text{cn}(\Gamma') \setminus \text{cn}(\Gamma_{\mathcal{R}})$ checks that the name x has been introduced in the last transition of the run. Dually, $\{y\} = \text{cn}(\Gamma_{\mathcal{R}}) \setminus \text{cn}(\Gamma')$ checks that y has been consumed. The use of a singleton $\{y\}$ is justified by the fact that in BitML each transition can consume at most one contract name.

Example 3.2. Let $\Gamma = \langle C_1, v \rangle_y \mid \langle A, v \rangle_z$, and let \mathcal{R} be the following run starting from Γ , where the contracts C_1 and C_2 are immaterial, but for the fact that they enable the displayed moves:

$$\begin{aligned}
&\langle C_1, v \rangle_y \mid \langle A, v \rangle_z \rightarrow_b \langle C_1, v \rangle_y \mid \langle A, v \rangle_z \mid \{G\}C_2 \rightarrow_b^* \langle C_1, v \rangle_y \mid \langle C_2, v \rangle_x \\
&\quad \xrightarrow{\text{split}(x)}_b \langle C_1, v \rangle_y \mid \langle C_2', v \rangle_{x'} \\
&\quad \xrightarrow{\text{split}(y)}_b \langle C_1', v' \rangle_{y'} \mid \langle C_1'', v - v' \rangle_{y''} \mid \langle C_2', v \rangle_{x'}
\end{aligned}$$

We have that $orig_{\Gamma}(\mathcal{R}, y') = orig_{\Gamma}(\mathcal{R}, y'') = y$, since the corresponding contracts have been obtained through a split of the ancestor y , which was in Γ . Instead, $orig_{\Gamma}(\mathcal{R}, x')$ is undefined, because its ancestor x is not in Γ . Further, $orig_{\Gamma}(\mathcal{R}, y) = y$, while $orig_{\Gamma}(\mathcal{R}, x)$ is undefined.

Before formalising liquidity, we give some further intuition. Assume that A is an honest participant, who cares about the liquidity of a set of contracts X in Γ . After an arbitrary sequence of transitions $\Gamma \rightarrow \dots \rightarrow \Gamma'$, where any participant may perform actions, A wants to liquidate all the contracts in Γ' originating from X , transferring their funds to participants' deposits. We want A to be able to liquidate contracts *without* the help of the other participants.

Definition 3.3 (Liquidity). Let \rightarrow be an LTS, let \mathcal{L}_A be a subset of its labels, let Γ be a configuration of the LTS, and let x be a contract name. We say that:

- x is *liquidable* in Γ if there exists a run $\mathcal{R} = \Gamma \xrightarrow{\ell_1} \dots \xrightarrow{\ell_n} \Gamma'$ such that:
 - (1) for all $i \in 1..n$, $\ell_i \in \mathcal{L}_A$;
 - (2) there exists no $z \in \text{cn}(\Gamma')$ such that $\text{orig}_\Gamma(\mathcal{R}, z) = x$.
- x is *liquid* in Γ if, for all runs $\mathcal{R} = \Gamma \rightarrow \dots \rightarrow \Gamma'$, all the contract names y with $\text{orig}_\Gamma(\mathcal{R}, y) = x$ are liquidable in Γ' .

We extend this to sets X of contracts names: X is liquid in Γ iff all $x \in X$ are liquid in Γ .

Intuitively, a contract x is liquidable by A when A can perform a sequence of transitions which eventually lead to a configuration containing no contract names originated from x . Consequently, all the funds in the contract x have been transferred to participants' deposits. When liquidating x , A 's moves can not reveal secrets of other participants, or generate authorizations for them: A must be able to unfreeze the funds on her own, performing actions $\ell_i \in \mathcal{L}_A$. Note that if $x \notin \text{cn}(\Gamma)$, then x is trivially liquidable in Γ .

The notion of liquidity is based upon that of liquidability. A contract x is liquid in a configuration Γ when, after an arbitrary sequence of moves performed by *any* participant, the contract names originated by x are liquidable by A .

We remark that, although Definition 3.3 will be instantiated with the semantics of BitML, the basic concepts it relies upon (runs, origin of contracts, moves of a participant) are quite general. Hence, our notion of liquidity, as well as the variants proposed in Section 7, can be applied to other languages for smart contracts, using their transition semantics.

Example 3.4. Recall the timed commitment contract TC from Section 1:

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

Let $\Gamma = \langle TC, 1 \rangle_x \mid \{A: a \# N\}$ be a configuration where the contract has been stipulated. We show that x is liquid in Γ with respect to any participant. In the configurations reachable from Γ , the contract TC has not progressed, or it has reduced to $\text{withdraw } A$. In the first case, x is liquidable by anyone, by firing $\text{withdraw } B$ after time t (alternatively, A may reveal the secret and then fire $\text{withdraw } A$). Instead, if the contract has reduced to $\text{withdraw } A$, anyone can liquidate it. Since all the descendants of x are liquidable, x is liquid.

Example 3.5. Let $G = A:1@y \mid B:1@z \mid A:\text{secret } a$. Consider the following contracts, where p is an arbitrary predicate on a :

$$C_1 = \text{reveal } a \text{ if } p.\text{withdraw } A + \text{reveal } a \text{ if } \neg p.\text{withdraw } B$$

$$C_2 = \text{split } (1 \rightarrow \text{reveal } a \text{ if } p.\text{withdraw } A \mid 1 \rightarrow \text{reveal } a \text{ if } \neg p.\text{withdraw } B)$$

$$C_3 = \text{withdraw } A + B:\text{split } (1 \rightarrow B:\text{withdraw } A \mid 1 \rightarrow \text{withdraw } B)$$

For $i \in \{1, 2, 3\}$, let $\Gamma_i = \langle C_i, 2 \rangle_x \mid \{A: a \# N\}$, with $N \neq \perp$. We have that:

- x is liquid in Γ_1 for A , but not for any other participant. C_1 has three reducts: C_1 itself, $\text{withdraw } A$, and $\text{withdraw } B$. The last two contracts are trivially liquidable by anyone. Instead C_1 is liquidable only by A , by revealing a and firing a reveal branch (since their guards are p and $\neg p$, one of them will be enabled), and finally firing the corresponding withdraw . Instead, from B 's viewpoint x is not liquid, because A could refuse to reveal.
- x is not liquid in Γ_2 for anyone. Indeed, none of the reducts of C_2 is liquidable, because one of the two reveal branches is stuck.

- x is liquid in Γ_3 for \mathbf{B} , but not for \mathbf{A} . Indeed, if \mathbf{B} authorizes and performs the `split`, then the reduct is not liquidable by \mathbf{A} , since \mathbf{B} could deny his second authorization.

Note that liquidability in Definition 3.3 requires that the moves $\ell_1 \dots \ell_n$ are performed atomically, effectively forbidding the adversary to interfere. Atomicity might be realistic in some blockchains, but not in others. For instance, Algorand features atomic sequences of transactions natively [BBL⁺21], while in Ethereum it is possible to perform atomically a sequence $\ell_1 \dots \ell_n$ of calls by deploying a new contract with a suitable function which calls $\ell_1 \dots \ell_n$ in sequence. Bitcoin, instead, does not support atomic sequences of transactions: an honest participant could start to perform the sequence, but at some point in the middle the adversary can interfere. Repeated interference could lead to an infinite run, where each attempt by the honest participant is hindered by the adversary. To illustrate the issue, consider a (not-BitML) LTS with states 0, 1, 2 and transitions $0 \xrightarrow{s} 1$, $1 \xrightarrow{s} 2$, and $1 \xrightarrow{p} 0$, where 0 is the initial state and 2 is the final state, where the funds have been liquidated. If atomic sequences of moves were allowed, an honest participant could always reach the final state by firing the atomic sequence $s s$. Otherwise, an adversary could always prevent the participant from reaching the final state, by firing p after each s . Hence, this LTS would be considered liquid only by assuming atomic sequences of moves.

However, infinite adversary interference as those shown above are not reproducible in BitML, for the following reason. Our notion of liquidity requires that any descendent of the contract must be liquidable by \mathbf{A} *alone*. This means that \mathbf{A} can do that by performing a sequence of moves which do not include any renegotiation, since renegotiations can be finalized only with the cooperation of all the other participants. Without renegotiation, the contract eventually terminates, so an adversary can interfere at most a finite number of times. After the last interference, \mathbf{A} has still a way of terminating the contract, and she can do this alone, satisfying item (2) of Definition 3.3.

4. CASE STUDIES

In this section we illustrate BitML and liquidity through a few example contracts. The automatic verification of these contracts will be discussed later on in Section 6.

A fair lottery. Consider a lottery between two players. The preconditions require \mathbf{A} and \mathbf{B} to commit to one secret each (a and b , respectively), and to put a deposit of $3\mathfrak{B}$ each ($1\mathfrak{B}$ as a bet, and $2\mathfrak{B}$ as a penalty for dishonest behaviour):

$$\begin{aligned} \text{Lottery}(\text{Win}) &= \text{split}(\ \\ &\quad 2 \rightarrow (\text{reveal } b \text{ if } 0 \leq b \leq 1. \text{ withdraw } \mathbf{B}) + (\text{after } t : \text{ withdraw } \mathbf{A}) \\ &\quad | 2 \rightarrow (\text{reveal } a. \text{ withdraw } \mathbf{A}) + (\text{after } t : \text{ withdraw } \mathbf{B}) \\ &\quad | 2 \rightarrow \text{Win}) \\ \text{Win} &= \text{reveal } a \text{ b if } a = b. \text{ withdraw } \mathbf{A} + \text{reveal } a \text{ b if } a \neq b. \text{ withdraw } \mathbf{B} \end{aligned}$$

The contract splits the balance in three parts, of $2\mathfrak{B}$ each. The first part allows \mathbf{B} to reveal b and then redeem $2\mathfrak{B}$; otherwise, after the deadline \mathbf{A} can redeem \mathbf{B} 's penalty (as in the timed commitment). Similarly, the second part allows \mathbf{A} to redeem $2\mathfrak{B}$ by revealing a . To determine the winner we compare the secrets, in the subcontract *Win*: \mathbf{A} wins if the secrets are the same, otherwise \mathbf{B} wins. This lottery is *fair*, since: (i) if both players are honest, then they will reveal their secrets within the deadline (redeeming $2\mathfrak{B}$ each), and then they

will have a $1/2$ probability of winning³; (ii) if a player is dishonest, not revealing the secret, then the other player has a positive payoff, since she can redeem 4€ .

Although fair, the lottery is *not* liquid, neither from **A**'s nor from **B**'s point of view, because if one of the two players does not reveal her secret, then the *Win* subcontract is not liquidable, and so the 2€ stored therein are frozen. We can recover liquidity by replacing *Win* with the following contract, where $t' > t$:

$$\text{Win}_2 = \text{Win} + (\text{after } t' : \text{reveal } a. \text{ withdraw } \mathbf{A}) + (\text{after } t' : \text{reveal } b. \text{ withdraw } \mathbf{B})$$

Now, even if one of the two players does not reveal, the honest player fire her *reveal* at time t' , liquidating the 2€ stored in *Win*₂.

Zero-coupon bonds. A zero-coupon bond [JES00] is a financial contract where an investor **A** pays 1€ upfront to a bank **B**, and receives back 2€ after a maturity date (say, year 2030). We can express this contract in BitML as follows. The contract precondition requires **A** to provide a deposit x of 1€ , and **B** to provide a deposit of 2€ . The contract is as follows:

$$\text{ZCB} = \text{split} (1 \rightarrow \text{withdraw } \mathbf{B} \mid 2 \rightarrow \text{after } 2030 : \text{withdraw } \mathbf{A})$$

Upon stipulation, all the deposits required in the preconditions pass under the control of *ZCB*, and can no longer be spent by **A** and **B**. The contract splits these funds in two parts: 1€ , that can be withdrawn by **B** at any moment, and 2€ , that can be withdrawn by **A** after the maturity date.

Although *ZCB* correctly implements the functionality of zero-coupon bounds, it is quite impractical: for the whole period from the stipulation to the maturity date, 2€ are frozen within the contract, and cannot be used by the bank in any way. Although this is a desirable feature for the investor, since it guarantees that he will receive 2€ even if the bank bankrupts, it is quite undesirable for the bank. In the real world, the bank would be free to use its own funds, together with those of investors, to make further financial transactions through which to repay the investments. The risk that the bank bankrupts is mitigated by external mechanisms, like insurances or government intervention.

To overcome this issue, we can exploit renegotiation. We first revise the precondition, which now requires only **A**'s deposit. The revised contract is:

$$\begin{aligned} \text{ZCB}_2 &= \text{split} (1 \rightarrow \text{withdraw } \mathbf{B} \mid 0 \rightarrow * : \text{rngt } \mathbf{X}) \\ \mathbf{X} &= \{\mathbf{B} : 2 @ d\} \text{ after } 2030 : \text{withdraw } \mathbf{A} \end{aligned}$$

As before, the bank can withdraw 1€ at any moment after stipulation. In the second part of the *split*, the participants renegotiate the contract: if they both agree, 0€ pass under the control of the contract **X**. The precondition of **X** requires the bank to provide 2€ in a fresh deposit; upon renegotiation, **A** can withdraw 2€ after the maturity date. The crucial difference with *ZCB* is that the deposit variable d is instantiated at *renegotiation* time, unlike x , which must be fixed at *stipulation* time.

The revised contract *ZCB*₂ solves the problem of *ZCB*, in that it no longer freezes 2€ for the whole duration of the bond: the bank could choose to renegotiate the contract, paying 2€ , just before the maturity date. This flexibility comes at a cost, since **A** loses the

³Note that **B** could increase his probability to win the lottery by choosing a secret $b > 1$, since doing so would increase the chances that $a \neq b$. We require that $0 \leq b \leq 1$, so that if **B** chooses a secret outside that range he will lose his 2€ deposit in the first part of *split*, and so **B**'s *average* payoff would be negative. Instead, **A** can only decrease decrease her probability to win by choosing a secret $a \notin \{0, 1\}$. For this reason, it is not necessary to require that $0 \leq a \leq 1$.

guarantee to eventually receive $2\mathfrak{B}$. To address this issue we need to add, as in the real world, an external mechanism. More specifically, we assume an insurance company I that, for an annual premium of $p\mathfrak{B}$ paid by the bank, covers a face amount of $f\mathfrak{B}$ (with $2 > f > 10p$):

$$A:1 @ x_1 \mid B:p @ x_2 \mid I:f @ x_3$$

We revise the bond contract as follows:

$$\begin{aligned} ZCB3 &= \text{split } (1 \rightarrow \text{withdraw } B \\ &\quad | p \rightarrow \text{withdraw } I \\ &\quad | f \rightarrow *: \text{rngt } X\langle 1 \rangle + \text{after } 2021 : \text{withdraw } A) \\ X\langle n \in 1..9 \rangle &= \{B:p @ d\} \\ &\quad \text{split } (p \rightarrow \text{withdraw } I \\ &\quad | f \rightarrow *: \text{rngt } X\langle n + 1 \rangle + \text{after } (2021 + n) : \text{withdraw } A) \\ X\langle 10 \rangle &= \{B:2 @ d\} \\ &\quad \text{split } (f \rightarrow \text{withdraw } I \\ &\quad | 2 \rightarrow \text{after } 2030 : \text{withdraw } A) \end{aligned}$$

The contract starts by transferring $1\mathfrak{B}$ to the bank, and the first year of the premium to the insurer. The remaining $f\mathfrak{B}$ are transferred to the renegotiated contract $X\langle 1 \rangle$, or, if the renegotiation is not completed by 2021, to the investor.

The contracts $X\langle n \rangle$, for $n \in 1..9$, allow the insurer to receive the annual premium until 2030: if the bank does not renegotiate the contract for the following year (paying the corresponding premium), then the investor can redeem the face amount of $f\mathfrak{B}$. Finally, the contract $X\langle 10 \rangle$ can be triggered if the bank deposits the $2\mathfrak{B}$: when this happens, the face amount is given back to the insurer, and the investor can redeem $2\mathfrak{B}$ after the maturity date.

Compared to $ZCB2$, the contract $ZCB3$ offers more protection to the investor. To see why, we must evaluate A 's payoff for all the possible behaviours of the other participants. If B and I are both honest, then A will redeem $2\mathfrak{B}$, as in the ideal contract ZCB . Instead, if either B or I do not accept to renegotiate some $X\langle n \rangle$, then A can redeem $f\mathfrak{B}$ as a partial compensation (unlike in $ZCB2$, where A just loses $1\mathfrak{B}$). In the real world, A could use this compensation to cover the legal fee to sue the bank in court; also, I could e.g. increase the premium for future interactions with B . By further refining the contract, we could model these real-world mechanisms as oracles, which sanction dishonest participants according to the evidence collected in the blockchain and in messages broadcast by participants. For instance, if B and I accept the renegotiation $X\langle n \rangle$ but A does not, then the oracle would be able to detect A 's dishonesty by inspecting the authorizations broadcast in year $2021 + n$. The sanction could consist e.g. in preventing A from buying other bonds from B .

Despite the apparent complexity, the contract is liquid for any participant: indeed, any potentially blocking renegotiation has an alternative timeout branch, which allows to liquidate the contract.

A fair recursive coin flipping game. Consider a simple game where two players repeatedly flip coins, and the one who wins two consecutive flips takes the pot. The precondition requires each player to deposit $3\mathfrak{B}$ and choose a secret:

$$A:3 @ x \mid A:\text{secret } a \mid B:3 @ y \mid B:\text{secret } b$$

```

CFG = reveal b if  $0 \leq b \leq 1$ .(
  reveal ab if  $a = b$ . (* : rngt  $X_A\langle 1 \rangle$  + after 3: SplitA)
  + reveal ab if  $a \neq b$ . (* : rngt  $X_B\langle 1 \rangle$  + after 3: SplitB)
  + after 2: withdraw B)
  + after 1: withdraw A
 $X_A\langle n \rangle$  = {A:secret a | B:secret b}
  reveal b if  $0 \leq b \leq 1$ .(
  reveal ab if  $a = b$ . withdraw A
  + reveal ab if  $a \neq b$ . (* : rngt  $X_B\langle n+1 \rangle$  + after  $(3n+3)$ : SplitB)
  + after  $(3n+2)$ : withdraw B)
  + after  $(3n+1)$ : withdraw A
 $X_B\langle n \rangle$  = {A:secret a | B:secret b}
  reveal b if  $0 \leq b \leq 1$ .(
  reveal ab if  $a = b$ . (* : rngt  $X_A\langle n+1 \rangle$  + after  $(3n+3)$ : SplitA)
  + reveal ab if  $a \neq b$ . withdraw B
  + after  $(3n+2)$ : withdraw B)
  + after  $(3n+1)$ : withdraw A
SplitA = split (4 → withdraw A | 2 → withdraw B)
SplitB = split (4 → withdraw B | 2 → withdraw A)

```

Figure 5: A recursive coin flipping game.

The contract *CFG* (Figure 5) asks *B* to reveal his secret first: if *B* waits too much, *A* can withdraw the contract funds after time 1. Then, it is *A*'s turn to reveal (before time 2, otherwise *B* can withdraw the funds). The current flip winner is *A* if the secrets of *A* and *B* are equal, otherwise it is *B*. At this point, the contract can be renegotiated as $X_A\langle 1 \rangle$ or $X_B\langle 1 \rangle$, depending on the flip winner (the parameter 1 represents the round). If players do not agree on the renegotiation, then the funds are split fairly, according to the current expected win. The contract $X_A\langle n \rangle$ requires *A* and *B* to generate fresh secrets for the n -th round. If *A* wins again, she can withdraw the pot, otherwise the contract can be renegotiated as $X_B\langle n+1 \rangle$. If the players do not agree on the renegotiation, the pot is split fairly between them. The contract X_B is similar.

This game is fair, i.e. the expected payoff of a rational player is always non-negative, notwithstanding the behaviour of the other player. Rational players must choose random secrets in $\{0, 1\}$, since non uniformly distributed secrets can make the adversary bias the coin flip in her favour. Further, choosing a secret different from 0 or 1 would be irrational: if done by *B*, this would prevent himself from revealing (by the predicate in the *reveal b*), and so *A* could win after the timeout; if done by *A*, this would make *B* win the round (since *B* wins when the secrets are different). Rationality also requires to reveal secrets in time (before the alternative *after* branch is enabled), and to take the *Split* branch if restipulation does not occur in time. This ensures that, when renegotiation happens, there is still time to reveal the round secrets. Indeed, a late renegotiation could enable the other player to win by timeout. To show fairness, first consider the case where renegotiation always happens. A rational player wins each coin flip with probability $1/2$, at least: so, the probability of winning the whole game is also $1/2$, at least. In the general case, the renegotiation at the end of each round may fail. When this happens, the rational player takes the *Split* branch, distributing the pot according to the expected payoff in the *current* game state,

thus ensuring the fairness of the game. The player who won the last coin flip is expected to win $p\mathfrak{B}$, with $p = 1/2 \cdot 6 + 1/2 \cdot (1/2 \cdot p + 1/2 \cdot 0)$, giving $p = 4$. Accordingly, the *Split* contracts transfer $4\mathfrak{B}$ to the winner of the last flip and $(6 - 4)\mathfrak{B} = 2\mathfrak{B}$ to the other player.

The contract *CFG* is liquid, because the **after** branches always offer liquidable alternatives to potentially blocking branches.

5. A SAFE ABSTRACTION OF THE BITML SEMANTICS

The concrete BitML semantics is infinite-state, because participants can always create new deposits, stipulate new contracts and renegotiate them, and can advance the current time. In this section we introduce an abstract semantics of BitML, which reduces the state space to a finite one, safely approximating liquidity. We construct our abstraction in three steps:

- First, we abstract concrete configurations Γ as *abstract configurations* $\alpha_{A,X}(\Gamma)$, where $A \in \mathbf{Hon}$, and X is the (finite) set of contract names under observation. Roughly, $\alpha_{A,X}(\Gamma)$ discards all the terms in the configuration, except for the contracts X , which are abstracted as follows. We remove **reveal** actions, authorizations, and time constraints, only recording whether reducing a contract D requires cooperation from some participant different from A . Further, the abstraction discards the actual parameters \mathcal{E} of renegotiations.
- We define a semantics \rightarrow_{\sharp} of abstract configurations. This semantics partitions the moves in two sets: the moves x , which represent a reduction of the contract x that can be performed by A *alone*, and the moves $* : x$, which instead represent actions that may require cooperation from other participants. Intuitively, the moves of the first kind are those required by liquidity. Although this abstract semantics substantially simplifies the concrete one, it is still infinite-state. We establish that the abstract semantics is an over-approximation of the concrete one (Theorem 5.4). Further, we show that abstract runs containing only moves of A alone under-approximate concrete runs (Theorem 5.5).
- We define a finite-state refinement $\rightarrow_{\sharp fin}$ of \rightarrow_{\sharp} , which allows for model-checking liquidity. Liquidity w.r.t. $\rightarrow_{\sharp fin}$ is proved to be equivalent to liquidity w.r.t. \rightarrow_{\sharp} (Theorem 5.10).

We start by defining the abstraction of configurations. Hereafter, we assume that all the abstractions are done w.r.t. the same honest participant A .

Definition 5.1 (Abstraction of configurations). For all sets of contract names X , we define the contract abstraction function α_A and the configuration abstraction function $\alpha_{A,X}$ in Figure 6. We abstract each defining equation $\mathbf{X}(\vec{\beta}) = \{G\}C$ as $\mathbf{X} =_{\sharp} \alpha_A(C)$. The syntax of abstract configurations is implicitly given by the equations in Figure 6. We further allow recursion variables \mathbf{X} within abstract configurations. We identify $* : * : D$ with $* : D$.

The equations in Figure 6 follow the intuition described above. The prefix τ used in **reveal** is used to abstract the fact that A alone can unconditionally reveal some of her secrets. Instead, we abstract as $* : \tau$ the case where some of the secrets to be revealed do not belong to A , or the truth of the predicate p is unknown. Since the secrets have been removed from configurations, this is a conservative (but safe) abstraction.

We now describe the abstract semantics of BitML. In the relation \rightarrow_{\sharp} between abstract configurations, the rules to advertise contracts, for deposits, and for delays are removed. There is a rule for making a contract $\langle \mathbf{withdraw} A, v \rangle_x$ reduce to a deposit $\langle A, v \rangle_y$ is replaced so that $\langle \mathbf{withdraw} A, v \rangle_x$ reduces to 0 (the empty configuration).

$$\begin{aligned}
\alpha_A(\sum_{i \in I} D_i) &= \sum_{i \in I} \alpha_A(D_i) & \alpha_A(\text{withdraw } B) &= \text{withdraw } B \\
\alpha_A(\text{reveal } \vec{a} \text{ if } p. C) &= \begin{cases} \tau. \alpha_A(C) & \text{if } p = \text{true} \text{ and } \vec{a} \subseteq \mathcal{S}_A \\ * : \tau. \alpha_A(C) & \text{otherwise} \end{cases} \\
\alpha_A(\text{split } (\parallel_{i \in I} v_i \rightarrow C_i)) &= \text{split } (\parallel_{i \in I} \alpha_A(C_i)) \\
\alpha_A(B : D) &= \begin{cases} \alpha_A(D) & \text{if } B = A \\ * : \alpha_A(D) & \text{otherwise} \end{cases} & \alpha_A(\text{after } t : D) &= \alpha_A(D) \\
\alpha_A(* : \text{rngt } X(\vec{\mathcal{E}})) &= * : \text{rngt } X \\
\alpha_{A,X}(\Delta \mid \Delta') &= \alpha_{A,X}(\Delta) \mid \alpha_{A,X}(\Delta') \\
\alpha_{A,X}(\langle C, v \rangle_x) &= \langle \alpha_A(C) \rangle_x \quad \text{if } x \in X & \alpha_{A,X}(\Delta) &= 0 \quad \text{otherwise}
\end{aligned}$$

Figure 6: Abstraction of contracts and configurations.

$$\begin{array}{c}
\frac{}{\langle \text{withdraw } B \rangle_x \mid \Gamma^\# \xrightarrow{x} \Gamma^\#} \text{[A-WITHDRAW]} \qquad \frac{y \text{ fresh}}{\langle \tau.C \rangle_x \mid \Gamma^\# \xrightarrow{x} \langle C \rangle_y \mid \Gamma^\#} \text{[A-REV]} \\
\frac{y_i \text{ fresh}}{\langle \text{split } \parallel_i C_i \rangle_x \mid \Gamma^\# \xrightarrow{x} \parallel_i \langle C_i \rangle_{y_i} \mid \Gamma^\#} \text{[A-SPLIT]} \\
\frac{\langle D \rangle_x \mid \Gamma^\# \xrightarrow{\ell} \Gamma'^\# \quad x \in \text{cn}(\ell)}{\langle D + C \rangle_x \mid \Gamma^\# \xrightarrow{\ell} \Gamma'^\#} \text{[A-BRANCH]} \qquad \frac{\langle D \rangle_x \mid \Gamma^\# \xrightarrow{\ell} \Gamma'^\# \quad x \in \text{cn}(\ell)}{\langle * : D \rangle_x \mid \Gamma^\# \xrightarrow{*:\ell} \Gamma'^\#} \text{[A-EXT]} \\
\frac{X =_\# C \quad y \text{ fresh}}{\langle * : \text{rngt } X \rangle_x \mid \Gamma^\# \xrightarrow{*:x} \langle C \rangle_y \mid X \mid \Gamma^\#} \text{[A-RNGT]}
\end{array}$$

Figure 7: Abstract BitML semantics.

Definition 5.2 (Abstract semantics). We define the relation $\rightarrow_\#$ between abstract configurations in Figure 7. We use ℓ to range over labels, which have the form x or the form $* : x$, where x is a contract name. We identify $* : * : \ell$ with $* : \ell$. We define $\mathcal{L}^\#$ as the subset of labels of the form x . An *abstract run* $\mathcal{R}^\#$ is a sequence $\Gamma_0 \rightarrow_\# \Gamma_1 \rightarrow_\# \dots$.

We briefly comment the rules in Figure 7, which define the abstract semantics $\rightarrow_\#$ of BitML. Most rules are straightforward. In rule [A-EXT], we record in the label $* : \ell$ the fact that reducing a contract $* : D$ might require the cooperation from some other participant. In rule [A-RNGT], we record through the term X in the configuration the fact that we have unfolded the recursion variable X . This is not strictly needed to make the abstract semantics safely approximate the concrete one: rather, it is a technical expedient to build the finite-state semantics $\rightarrow_{\#fin}$ on-top of $\rightarrow_\#$.

Correspondence between the semantics. We now establish a correspondence between the abstract and the concrete semantics of BitML. In all the statements below, we assume that the concrete configuration Γ is reachable from an initial configuration.

We introduce below the notion of *descendants* of a set of contracts, which is dual to that of originator in Definition 3.1. This notion is exploited whenever we reduce an abstract configuration $\alpha_{A,X}(\Gamma)$ mimicking the steps of a concrete run \mathcal{R} from Γ to Γ' . Intuitively, the resulting configuration will be an abstraction of Γ' against the descendants of X .

Definition 5.3 (Descendants). For all concrete configurations Γ , runs $\mathcal{R} = \Gamma \rightarrow_b^* \Gamma'$, and set of contract names X , we define the set of contract names $desc_\Gamma(\mathcal{R}, X)$ as follows:

$$desc_\Gamma(\mathcal{R}, X) = \{x \in \text{cn}(\Gamma') \mid \text{orig}_\Gamma(\mathcal{R}, x) \in X\}$$

Theorem 5.4 establishes that each concrete run $\mathcal{R} = \Gamma \rightarrow_b^* \Gamma'$ has a corresponding abstract run, whose first element is the abstraction of Γ , and the last element is the abstraction of Γ' , plus recursion variables.

Theorem 5.4 (Over-approximation). *Let $X \subseteq \text{cn}(\Gamma)$, and let $\mathcal{R} = \Gamma \rightarrow_b^* \Gamma'$ be a concrete run. Then, there exist $\mathbf{X}_1, \mathbf{X}_2, \dots$ such that:*

$$\alpha_{A,X}(\Gamma) \rightarrow_{\sharp}^* \alpha_{A,desc_\Gamma(\mathcal{R},X)}(\Gamma') \mid \parallel_i \mathbf{X}_i$$

Proof. (sketch) By induction on the number of concrete moves. We check that each move $\Delta \rightarrow_b \Delta'$ is matched by zero or one abstract moves. First, note that the abstract configuration contains only contracts in X . We have the following cases, according to the concrete move:

- If $\Delta \rightarrow_b \Delta'$ does neither consume nor create a contract, then the abstractions of Δ and Δ' are the same, so we match the concrete move with zero abstract moves.
- If $\Delta \rightarrow_b \Delta'$ creates contracts without consuming any contract, then it is an instance of rule [C-INIT]. In this case, the created contract is not a descendant of any contract in X , so it does not occur in the abstraction of Δ' . As in the previous case, we match the concrete move with zero abstract moves.
- If $\Delta \rightarrow_b \Delta'$ consumes a contract $x \notin X$, then the abstraction discards the contract x and its descendants, and so we match the concrete move with zero abstract moves.
- If $\Delta \rightarrow_b \Delta'$ consumes a contract $x \in X$, then one of the rules [C-SPLIT], [C-REV], [C-WITHDRAW] or [C-RNGT] have been used. The moves [C-SPLIT], [C-WITHDRAW] and [C-RNGT] are matched by the corresponding [A-*] abstract moves. The move [C-REV] is matched either by [A-REV] or [A-EXT], depending on how the **reveal** was abstracted. Note that [C-RNGT] may instantiate the formal parameters in the contract with actual values. However, these values are abstracted away by α_A , so this instantiation is immaterial. The term \mathbf{X} introduced by [A-RNGT] is accounted for by the statement (it will belong to $\parallel_i \mathbf{X}_i$). \square

Theorem 5.5 associates some abstract runs with concrete runs. More specifically, it considers an abstract run starting from the abstraction of a concrete configuration Γ (plus recursion variables), and whose labels represent actions performable by \mathbf{A} (e.g., the actions in \mathcal{L}^\sharp). The theorem constructs a concrete run starting from Γ , whose last configuration, once abstracted, matches the last configuration in the abstract run. Further, the labels in the concrete run belong to the set of concrete actions performable by \mathbf{A} (which we denote by \mathcal{L}_A^b , i.e. all the concrete labels except those of the form $\mathbf{B} : \ell$ with $\mathbf{B} \neq \mathbf{A}$).

Theorem 5.5 (Under-approximation). *Let $X \subseteq \text{cn}(\Gamma)$, and assume that:*

$$\alpha_{A,X}(\Gamma) \mid \parallel_i \mathbf{X}_i \rightarrow_{\sharp}^* \Gamma' \quad \text{with labels in } \mathcal{L}^\sharp$$

Then, there exist \mathcal{R} and Γ' such that:

$$\mathcal{R} = \Gamma \rightarrow_b^* \Gamma' \quad \text{with labels in } \mathcal{L}_A^b, \text{ and } \Gamma'_{\#} = \alpha_{A, \text{desc}_{\Gamma}(\mathcal{R}, X)}(\Gamma') \mid \parallel_i \mathbf{X}_i$$

Proof. (sketch) By induction on the number of moves. Since each abstract move must be in $\mathcal{L}^{\#}$, it must be derived by one of the rules [A-WITHDRAW], [A-REV], or [A-SPLIT], possibly as a premise of [A-BRANCH]. To fall into these cases, the concrete contract must have one of the following forms:

- **withdraw B** or **split** \dots . Here, the abstract move is matched by the concrete moves [C-WITHDRAW] or [C-SPLIT], the labels of which belong to \mathcal{L}_A^b .
- **reveal \vec{a} if true** where \vec{a} are secrets of **A**, the participant w.r.t. whom we abstract. In this case, the abstract move [A-REV] is matched by [C-REV], the label of which belongs to \mathcal{L}_A^b . If the secrets have not been revealed, yet, we perform one or more [C-AUTHREV] moves (the labels of which also belong to \mathcal{L}_A^b) before the [C-REV] move. Note that the premise $N \neq \perp$ in rule [C-AUTHREV] is satisfied, because we are assuming that the participant **A** against whom we are abstracting is honest ($A \in \text{Hon}$), and rule [C-AUTHCOMMIT] ensures that honest participants have committed to secrets $N \neq \perp$.
- any of the above, constrained by **after** $t : \dots$ or **B** : \dots with $B = A$. Here, if needed, we perform a [C-DELAY] to advance the current time until time t (a move in \mathcal{L}_A^b). Then, if needed, **A** can perform a [C-AUTHBRANCH] to authorize the branch (this move is also in \mathcal{L}_A^b). After that, the constraints are satisfied, so we proceed as in the previous items. \square

Safe approximation of liquidity. Theorem 5.7 below establishes the soundness of our abstraction w.r.t. liquidity. It exploits the following corollary for liquidability.

Corollary 5.6 (Abstraction soundness against liquidability). *Let $x \in X \subseteq \text{cn}(\Gamma)$. If x is liquidable in $\alpha_{A, X}(\Gamma) \mid \vec{Y}$ w.r.t. $\rightarrow_{\#} / \mathcal{L}^{\#}$, then x is liquidable in Γ w.r.t. $\rightarrow_b / \mathcal{L}_A^b$.*

Proof. Since x is liquidable in $\alpha_{A, X}(\Gamma) \mid \vec{Y}$, there exists an abstract run with labels in $\mathcal{L}^{\#}$ and leading to a configuration $\Gamma'_{\#}$ without contracts originated from x . By Theorem 5.5, there is a corresponding concrete run with labels in \mathcal{L}_A^b starting from Γ and leading to a configuration Γ' the abstraction of which is $\Gamma'_{\#}$. Since the abstraction preserves all the contracts originated from x , in Γ' there are none. Therefore, x is liquidable in Γ . \square

Theorem 5.7 (Abstraction soundness against liquidity). *Let $x \in X \subseteq \text{cn}(\Gamma)$. If x is liquid in $\alpha_{A, X}(\Gamma)$ w.r.t. $\rightarrow_{\#} / \mathcal{L}^{\#}$, then x is liquid in Γ w.r.t. $\rightarrow_b / \mathcal{L}_A^b$.*

Proof. Let $\mathcal{R} = \Gamma \rightarrow \dots \rightarrow \Gamma'$ be a concrete run, and let y be such that $\text{orig}_{\Gamma}(\mathcal{R}, y) = x$. By Theorem 5.4, there exists a corresponding abstract run:

$$\Gamma_{\#} = \alpha_{A, X}(\Gamma) \rightarrow_{\#}^* \Gamma'_{\#} = \alpha_{A, \text{desc}_{\Gamma}(\mathcal{R}, X)}(\Gamma') \mid \parallel_i \mathbf{X}_i$$

Since x is liquid in $\Gamma_{\#}$ and y is a descendant of x in the abstract run, y is liquidable in $\Gamma'_{\#}$. By Corollary 5.6, y is liquidable in Γ' . Therefore, x is liquid in Γ . \square

A finite-state abstraction of BitML. The abstract semantics \rightarrow_{\sharp} is infinite-state, even up-to renaming of contract names. Indeed, each application of rule [A-RNGT] unfolds a contract \mathbf{X} , which can spawn other parallel contracts before recursing. Therefore, the number of parallel contracts in reachable abstract configurations can grow unboundedly. This hinders verification based on model-checking the whole state space.

We now introduce another abstract semantics, called $\rightarrow_{\sharp fin}$, which has a *finite* number of reachable configurations (up-to renaming of contract names), and which preserves the liquidity w.r.t. \rightarrow_{\sharp} : more specifically, a contract is liquid w.r.t. \rightarrow_{\sharp} if and only if it is liquid w.r.t. $\rightarrow_{\sharp fin}$. We define the transition relation $\rightarrow_{\sharp fin}$ by the following rule:

$$\frac{\Gamma \xrightarrow{\ell}_{\sharp} \Gamma' \quad \nexists \Delta, \mathbf{X} : \Gamma' = \Delta \mid \mathbf{X} \mid \mathbf{X}}{\Gamma \xrightarrow{\ell}_{\sharp fin} \Gamma'}$$

Intuitively, $\rightarrow_{\sharp fin}$ can mimic any move of \rightarrow_{\sharp} , except for the moves [A-RNGT] which renegotiate a contract which has already been renegotiated. Technically, this is ensured by constraining the configuration to contain at most one occurrence of each recursion variable \mathbf{X} .

Theorem 5.8. *Starting from any abstract configuration, the set of states reachable through $\rightarrow_{\sharp fin}$ is finite, up-to renaming of contract names.*

Proof. Direct consequence of the definition of the abstract semantics. \square

The following lemma establishes the equivalence of \rightarrow_{\sharp} and $\rightarrow_{\sharp fin}$ against liquidity, and it is instrumental to prove that they also agree on liquidity (Theorem 5.10).

Lemma 5.9. *Let Γ^{\sharp} be an abstract configuration without multiple occurrences of any \mathbf{X} . Then, for all contract names x :*

$$x \text{ liquidable in } \Gamma^{\sharp} \text{ w.r.t. } \rightarrow_{\sharp} / \mathcal{L}^{\sharp} \iff x \text{ liquidable in } \Gamma^{\sharp} \text{ w.r.t. } \rightarrow_{\sharp fin} / \mathcal{L}^{\sharp}$$

Proof. Recall that liquidity only considers moves in \mathcal{L}^{\sharp} . The only difference between \rightarrow_{\sharp} and $\rightarrow_{\sharp fin}$ is that the latter forbids, in some configurations, the application of rule [A-RNGT]. Since the label of this rule is not in \mathcal{L}^{\sharp} , this difference is immaterial for liquidity. \square

Theorem 5.10. *Let Γ^{\sharp} be an abstract configuration without occurrences of any \mathbf{X} . Then, for all contract names x :*

$$x \text{ liquid in } \Gamma^{\sharp} \text{ w.r.t. } \rightarrow_{\sharp} / \mathcal{L}^{\sharp} \iff x \text{ liquid in } \Gamma^{\sharp} \text{ w.r.t. } \rightarrow_{\sharp fin} / \mathcal{L}^{\sharp}$$

Proof. (sketch) For \Rightarrow , consider a run \mathcal{R}^{\sharp} of $\rightarrow_{\sharp fin}$ starting from Γ^{\sharp} and leading to some Γ'^{\sharp} , and let y be a descendant of x . Since \rightarrow_{\sharp} includes $\rightarrow_{\sharp fin}$, \mathcal{R}^{\sharp} is also a run of \rightarrow_{\sharp} . Since x is liquid in Γ^{\sharp} w.r.t. \rightarrow_{\sharp} , then y must be liquidable in Γ'^{\sharp} w.r.t. \rightarrow_{\sharp} . By Lemma 5.9, it follows that y is liquidable also in $\rightarrow_{\sharp fin}$. Therefore, x is liquid in Γ^{\sharp} w.r.t. $\rightarrow_{\sharp fin}$.

For \Leftarrow , take any contract $\langle \mathbf{C} \rangle_y$ which originates from the contract x in a run of \rightarrow_{\sharp} starting from Γ^{\sharp} . By contradiction, assume that y is not liquidable for \rightarrow_{\sharp} . In the \rightarrow_{\sharp} run, $\langle \mathbf{C} \rangle_y$ might be reached after several renegotiations involving some defined contracts $\mathbf{X}_1, \mathbf{X}_2, \dots$ more than once. In such case, it is also possible to reach \mathbf{C} without renegotiating the same \mathbf{X}_i more than once, because each use of rule [A-RNGT] on \mathbf{X}_i spawns the same contract, and contracts in parallel do not interact. Hence, there exists some \rightarrow_{\sharp} run leading to a non liquidable $\langle \mathbf{C} \rangle_z$ and whose configurations never include any \mathbf{X}_i more than once. This run is therefore also a $\rightarrow_{\sharp fin}$ run, leading to a non liquidable z according to Lemma 5.9 — contradiction with the liquidity of x w.r.t. $\rightarrow_{\sharp fin}$. \square

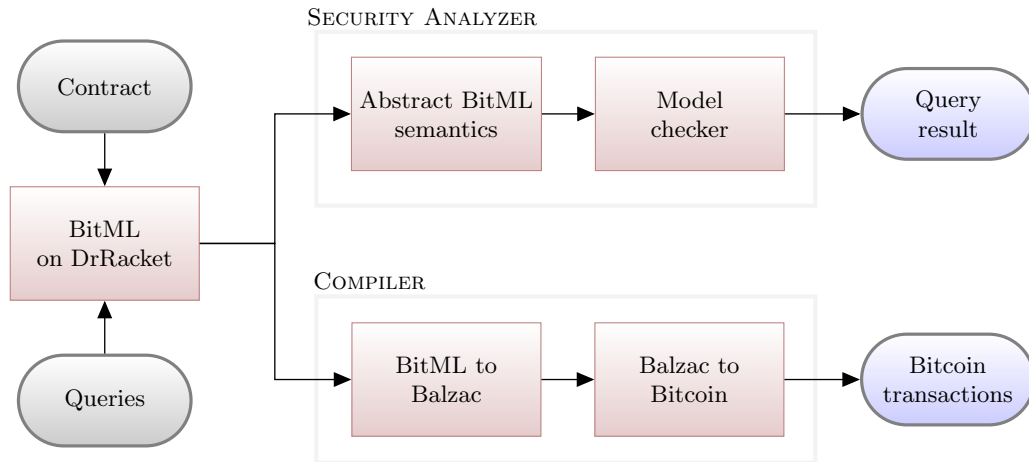


Figure 8: Toolchain architecture.

6. THE BITML TOOLCHAIN

We now describe the BitML toolchain, whose architecture is displayed in Figure 8. The development workflow is the following: (a) write the BitML contract in the DrRacket IDE; (b) verify that the contract is liquid w.r.t. the given participant; (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 verifier implements the abstract BitML semantics in Maude, a model-checking framework based on rewriting logic [CDE⁺02]. The toolchain is open-source⁴; a tutorial is available online⁵, including references to our experiments on the Bitcoin testnet.

Benchmarks. To evaluate our toolchain, we use a benchmark of representative use cases⁶ including financial contracts [ST18, BKT17], auctions, lotteries [ADMM16, MB17] and gambling games. For each contract in the benchmark, we display in Table 1 the number of involved participants, the number of transactions obtained by the compiler, and the number of Maude rewrites for checking liquidity. Notably, the tool automatically verifies that all the contracts in the benchmark are liquid. The verification time for all the benchmarks is in the order of milliseconds on a consumer-grade laptop, except for *ZCB3*, which requires ~ 1 s.

We compare the performance of our tool against [ADMM14b], which models Bitcoin contracts in Uppaal, a model-checking framework based on Timed Automata. The most complex contract modelled in [ADMM14b] is the mutual timed commitment with 2 participants: this requires ~ 30 s to be verified in Uppaal, while our tool verifies the same property in ~ 1 ms. This speedup is due to the higher abstraction level of BitML over [ADMM14b], which operates at the (lower) level of Bitcoin transactions. Note that increasing the number of participants in the mutual timed commitment significantly affects verification time. This is because the size of the contract increases and there are more committed secrets in play. Indeed, a larger number of secrets increases the branching in the (abstract) BitML LTS,

⁴<https://github.com/bitml-lang>

⁵<https://blockchain.unica.it/bitml>

⁶<https://github.com/bitml-lang/bitml-compiler/tree/master/examples/benchmarks>

Contract	Participants	Transactions	Rewrites
Mutual timed commitment	2	15	72
Mutual timed commitment	3	34	207
Mutual timed commitment	4	75	644
Mutual timed commitment	5	164	2097
Escrow (early fees)	3	12	104
Escrow (late fees)	3	11	110
Zero Coupon Bond	3	8	189
Coupon Bond	3	18	9101
Future(C)	3	$5 + T_C$	$136 + R_C$
Option(C, D)	3	$14 + T_C + T_D$	$162 + R_C + R_D$
Lottery (quadratic collateral)	2	15	1466
Lottery (0 collateral)	2	8	31
Lottery (0 collateral)	4	587	167
Rock-Paper-Scissors	2	23	2322
Morra game	2	40	89
Shell game	2	23	48
Auction (2 turns)	2	42	218
Coin flipping game	2	32	563
Zero coupon bond (v3)	3	44	1196813

Table 1: Benchmarks for the BitML toolchain.

since in each state any secret not revealed so far can be revealed. This increases the size of the state space, and consequently the complexity of model checking.

We observe that the number of Maude rewritings for *ZCB3* is particularly large. This is due to the contract generating many parallel components, which cause the explosion of the state space. More specifically, *ZCB3* performs 10 steps of recursion, each one involving a split, which creates several parallel sub-contracts (half of them with a choice).

The exponential blow-up due to parallel components is a common performance issue of general model checking, and so it also occurs in our setting. Indeed, the benchmarks in Table 1 only include the contracts for which we have managed to verify liquidity. We have not explicitly looked for the simplest examples for which the verification is practically unfeasible. However, by the discussion above, we expect that increasing the number of recursion steps by ~ 20 units in the *ZCB3* contract would be enough to make the model checker exhaust the available resources. Besides increasing the number of recursion steps, there are many other ways to design a contract for which verifying liquidity is unfeasible. For instance, if the contract contains a split of N withdraw actions, the size of the state space is at least $O(2^N)$. A similar size is obtained for a contract with N reveal actions in sequence, since at each point of the computation one has the option of revealing any of their secrets. The very same problem is witnessed by a contract requiring N authorizations, e.g. as in $A_1 : \dots : A_N : D$.

Limitations. One of the main difficulties that we have encountered in developing contracts is that some complex BitML specifications can not be compiled to pure Bitcoin, because of the 520-byte limit on the size of each value pushed to the evaluation stack [And19]. In some cases, we managed to massage the BitML contract so to make its compilation respect the constraint. For instance, a pattern that easily violates the 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 p_0 and p_1 , 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:

```
(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 choices. The scripts of these transactions encode the disjunction of *two* logical conditions, corresponding to the two branches of the choices. Using this workaround we have managed to compile a 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 predicates⁷) allowed us to compile all the contracts in Table 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.

7. VARIANTS OF LIQUIDITY

We now discuss some variants of the notion of liquidity of Section 3.

Liquidity under a strategy. The runs $\mathcal{R} = \Gamma \rightarrow \dots \rightarrow \Gamma'$ in Definition 3.3 allow any participant to perform *any* enabled move. For instance, consider the contract:

$$\{A:1@x \mid A:\text{secret } a \mid A:\text{secret } a'\}(\text{reveal } a.\text{withdraw } A + \text{reveal } a'.B:\text{withdraw } A)$$

This contract is *not* liquid for A : indeed, if A performs $\text{reveal } a'$, the reduct $B:\text{withdraw } A$ is not liquidable by A alone, since B may refuse to give his authorization. To overcome this issue, A can follow the strategy of always performing the $\text{reveal } a$. In this way, she is sure to be able liquidate the contract from any reachable state. A possible extension of the notion of liquidity in Definition 3.3 is to make it parametric on A 's strategy, and consider only the runs $\mathcal{R} = \Gamma \rightarrow \dots \rightarrow \Gamma'$ which are coherent with it.

Multiparty liquidity. Definition 3.3 requires that A *alone* can liquidate each descendent of the contract. We can relax it by considering a set of collaborative participants. For instance, consider an escrow contract between A and B , involving also a mediator M :

$$\begin{aligned} \text{Escrow} &= A:\text{withdraw } B + B:\text{withdraw } A + A:\text{Resolve} + B:\text{Resolve} \\ \text{Resolve} &= \text{split}(0.1 \rightarrow \text{withdraw } M \\ &\quad \mid 0.9 \rightarrow M:\text{withdraw } A + M:\text{withdraw } B) \end{aligned}$$

⁷<https://github.com/bitml-lang/bitml-compiler/blob/master/bitml/exp.rkt>

After the contract has been stipulated, A can choose to pay B , by authorizing the first branch. Similarly, B can allow A to take her money back, by authorizing the second branch. If they do not agree, any of them can invoke a mediator M to resolve the dispute, invoking a *Resolve* branch. There, the initial deposit (say, of $v\text{฿}$) is split in two parts: $0.1v$ goes to the mediator, while $0.9v$ is assigned either to A and B , depending on M 's choice. This contract is *not* liquid for A , because B can invoke the mediator, who can refuse to act, freezing the funds within the contract (similarly for B). Instead, assuming that M is collaborative, the contract is liquid for both A and B . Indeed, a collaborative M will always authorize either the *withdraw A* or the *withdraw B* to unlock $0.9v$. Multiparty liquidity where all participants are collaborative was used e.g. in [TDD⁺18] in the context of Ethereum contracts.

Quantitative liquidity. Definition 3.3 requires that no funds remain frozen within the contract. However, in some cases A could accept the fact that a portion of the funds remain frozen, especially when these funds would be assigned to other participants. We could define a contract *v-liquid* for A if at least v bitcoins are guaranteed to be redeemable by anyone. For instance, *Lottery(Win)* of Section 4 is non-liquid for A , but it is 4฿ -liquid. Instead, *Lottery(Win₂)* is 6฿ -liquid, and then also liquid, under this strategy. A refinement of this notion could require that at least $v\text{฿}$ are transferred to A , rather than to any participant. Under this notion, both *Lottery(Win)* and *Lottery(Win₂)* would be 2฿ -liquid for A .

Other variants of liquidity may take into account the time when funds become liquid, the payoff of strategies (e.g., ruling out irrational adversaries), or fairness issues.

8. VARIANTS OF CONTRACT PRIMITIVES

The renegotiation primitive we have proposed for BitML is motivated by its simplicity, and by the possibility of compiling into standard Bitcoin transactions. By adding some degree of complexity, we can devise more general primitives, which could be useful in certain scenarios. We discuss below some alternatives.

Renegotiation-time parameters. The primitive $* : \text{rngt } X\langle \vec{\mathcal{E}} \rangle$ allows participants to choose at run-time only the deposit variables used in the renegotiated contracts, and to commit to new secrets. A possible extension is to allow participants to choose at run-time *arbitrary* values for the renegotiation parameters $\vec{\mathcal{E}}$.

For instance, consider a mortgage payment, where a buyer A must pay 10฿ to a bank B in 10 installments. After A has paid the first five installments (of 1฿ each), the bank might propose to renegotiate the contract, varying the amount of the installment. Using the BitML renegotiation primitive presented in Section 2, we could not model this contract, since the new amount and the number of installments are unknown at the time of the original stipulation. Technically, the issue is that the primitive $* : \text{rngt } X\langle \vec{\mathcal{E}} \rangle$ only involves static expressions \mathcal{E} , the value of which is determined at stipulation time.

To cope with non-statically known values, we could extend guarded contracts with terms of the form $* : \text{rngt } X\langle B : v \rangle$, declaring that the value v is to be chosen by B at renegotiation time. For instance, this would allow to model our installments payment plan as $\text{IPP}\langle 1 \rangle$,

with the following defining equations:

$$\begin{aligned} \text{IPP}\langle\beta < 5\rangle &= \{\mathbf{A}:1@d\}(\text{split } 1 \rightarrow \text{withdraw } \mathbf{B} \mid 0 \rightarrow * : \text{rngt } \text{IPP}\langle\beta + 1\rangle) \\ \text{IPP}\langle 5\rangle &= \{\mathbf{A}:1@d\}(\text{split } 1 \rightarrow \text{withdraw } \mathbf{B} \mid 0 \rightarrow * : \text{rngt } \mathbf{Y}\langle\mathbf{B} : k, \mathbf{B} : v\rangle) \\ \mathbf{Y}\langle\beta \neq 1, \beta'\rangle &= \{\mathbf{A}:\beta'@d\}(\text{split } \beta' \rightarrow \text{withdraw } \mathbf{B} \mid 0 \rightarrow * : \text{rngt } \mathbf{Y}\langle\beta - 1, \beta'\rangle) \\ \mathbf{Y}\langle 1, \beta'\rangle &= \{\mathbf{A}:\beta'@d\} \text{withdraw } \mathbf{B} \end{aligned}$$

where in $\text{IPP}\langle 5\rangle$, the bank chooses the number of installments k , as well as the amount v of each installment. Note that if \mathbf{A} does not agree with these values, the renegotiation fails. A more refined version of the contract should take this possibility into account, by adding suitable compensation branches. Although adding the new primitive would moderately increase the complexity of the semantics and of the compiler, this extension can still be implemented on top of standard Bitcoin.

Renegotiation with a given set of participants. As we have remarked in Section 2, a renegotiation can be performed only if *all* the participants of the contract agree. We could relax this, by just requiring the agreement of a *given* set of participants (possibly, not among those who originally stipulated the contract).

For instance, consider an escrow service between a buyer \mathbf{A} and a seller \mathbf{B} for the purchase of an item worth $1\mathfrak{B}$. The normal case is when the \mathbf{A} authorizes the transfer of $1\mathfrak{B}$ after receiving the item, but it may happen that a dishonest \mathbf{B} never ships the item, or that a dishonest buyer never authorizes the payment. To cope with these cases, the participants could renegotiate the contract, including an escrow service \mathbf{M} which mediates the dispute:

$$\begin{aligned} \mathbf{A} : \text{withdraw } \mathbf{B} + \mathbf{B} : \text{withdraw } \mathbf{A} + \mathbf{A} : \mathbf{M} : \text{rngt } \text{Refd}_{\mathbf{A}} + \mathbf{B} : \mathbf{M} : \text{rngt } \text{Refd}_{\mathbf{B}} \\ \text{Refd}_{\mathbf{P}} = \{\mathbf{P}:0.1@d\} \text{split } (0.1 \rightarrow \text{withdraw } \mathbf{M} \mid 1 \rightarrow \text{withdraw } \mathbf{P}) \end{aligned}$$

where $\mathbf{A} : \mathbf{M} : \text{rngt } \text{Refd}_{\mathbf{A}}$ means that only \mathbf{A} and \mathbf{M} need to agree in order for the contract $\text{Refd}_{\mathbf{A}}$ to be executed, resolving the dispute. In this case it is crucial that the renegotiation is possible even without the agreement between \mathbf{A} and \mathbf{B} . Indeed, if \mathbf{M} decides to refund \mathbf{A} (by authorizing $\text{Refd}_{\mathbf{A}}$), it is not to be expected that also \mathbf{B} agrees. Similarly to the one discussed before, also this extension can be implemented on-top of Bitcoin.

Non-consensual renegotiation. In the variants of $* : \text{rngt}$ discussed before, renegotiation requires one or more participants to agree. Hence, each use of $* : \text{rngt}$ must include suitable alternative branches, to be fired in case the renegotiation fails. In certain scenarios, we may want to renegotiate the contract without the participants having to agree. To this purpose, we can introduce a new primitive $\text{call } \mathbf{X}$, which continues as \mathbf{X} without requiring anyone to agree. We assume that the defining equations of this primitive have the form $\mathbf{X}(\vec{\beta}) = \{v\}C$, where v represents the amount of \mathfrak{B} added to the contract, by anyone.

We exemplify the new primitive in a two-players game which starts with a bet of $1\mathfrak{B}$ from \mathbf{A} , and a bet of $2\mathfrak{B}$ from \mathbf{B} . Then, starting from \mathbf{A} , players take turns adding $2\mathfrak{B}$ each to the pot. The first one who is not able to provide the additional $2\mathfrak{B}$ within a given time loses the game, allowing the other player to take the whole pot. The contract is as follows:

$$\begin{aligned} C &= \{\mathbf{A}:1@x \mid \mathbf{B}:2@y\}(\text{call } \mathbf{X}_{\mathbf{A}}\langle 2\rangle + \text{after } 1 : \text{withdraw } \mathbf{B}) \\ \mathbf{X}_{\mathbf{A}}\langle n\rangle &= \{2\}(\text{call } \mathbf{X}_{\mathbf{B}}\langle n+1\rangle + \text{after } n : \text{withdraw } \mathbf{A}) \\ \mathbf{X}_{\mathbf{B}}\langle n\rangle &= \{2\}(\text{call } \mathbf{X}_{\mathbf{A}}\langle n+1\rangle + \text{after } n : \text{withdraw } \mathbf{B}) \end{aligned}$$

Unlike `* : rngt`, the action `call` can be fired without the authorizations of all the players: it just requires that the authorization to gather $2\mathfrak{B}$ is provided, by anyone. Even though the sender of these $2\mathfrak{B}$ is not specified in the contract, it is implicit in the game mechanism: for instance, when $X_A\langle n \rangle$ calls $X_B\langle n + 1 \rangle$, only participant `B` is incentivized to add $2\mathfrak{B}$, since not doing so will make `A` win.

Implementing the `call` primitive on top of Bitcoin seems unfeasible: even if it were possible to use complex off-chain multiparty computation protocols [GMR⁺19], doing so might be impractical. Rather, we would like to extend Bitcoin as much as needed for the new primitive. In our implementation of BitML, we compile contracts to sets of transactions and make participants sign them. In standard BitML this is doable since, at stipulation time, we can finitely over-approximate the reducts of the original contract. Recursion can make this set infinite, e.g. $X_A\langle 2 \rangle, X_A\langle 3 \rangle, \dots$, hence impossible to compile and sign statically. A way to cope with this is to extend Bitcoin with *malleable* signatures which only cover the part of the transaction not affected by the parameter n in $X_B\langle n \rangle$. Further, signatures must not cover the in fields of transactions, since they change as recursion unfolds. In this way, the same signature can be reused for each call.

Adding malleability provides flexibility, but poses some risks. For instance, instead of redeeming the transaction corresponding to $X_A\langle n \rangle$ with the transaction of $X_B\langle n + 1 \rangle$ one could instead use the transaction of $X_B\langle n + 100 \rangle$, since the two transactions have the same signature. To overcome this problem, we could add a new opcode to allow the output script of $X_B\langle n \rangle$ to access the parameter in the redeeming transaction, so to verify that it is indeed $n + 1$ as intended. Similarly, to check that we have $2\mathfrak{B}$ more in the new transaction, an opcode could provide the value of the new output. The same goal could be achieved by exploiting *covenants* [MES16, OP17, BLZ20].

9. CONCLUSIONS

We have investigated linguistic primitives to renegotiate BitML contracts, and their implementation on standard Bitcoin. More expressive primitives could be devised by relaxing this constraint, e.g. assuming the extended UTXO model [CCM⁺20].

Our verification technique is based on a sound abstraction of the state space of contracts. Since this abstraction is finite-state, it can be model-checked to verify the required properties. If we assume that integers are unbounded, and that participants always accept renegotiations, the extension of BitML presented in Section 2 can simulate a counter machine, so making BitML Turing-complete. Hence, any verification technique for BitML cannot be sound and complete. Alternative techniques to model checking (e.g., type-based approaches [DBHP19]) could be used to analyse relevant contract properties.

Acknowledgements. Massimo Bartoletti is partially supported by Aut. Reg. Sardinia projects *Sardcoin*, *Smart collaborative engineering*, and Conv. Fondazione di Sardegna & Atenei Sardi project F74I19000900007 *ADAM*. Maurizio Murgia and Roberto Zunino are partially supported by MIUR PON *Distributed Ledgers for Secure Open Communities*.

REFERENCES

- [ABC17] Nicola Atzei, Massimo Bartoletti, and Tiziana Cimoli. A survey of attacks on Ethereum smart contracts (SoK). In *Principles of Security and Trust (POST)*, volume 10204 of *LNCS*, pages 164–186. Springer, 2017. doi:10.1007/978-3-662-54455-6_8.

- [ABC⁺18] Nicola Atzei, Massimo Bartoletti, Tiziana Cimoli, Stefano Lande, and Roberto Zunino. SoK: unraveling Bitcoin smart contracts. In *POST*, volume 10804 of *LNCS*, pages 217–242. Springer, 2018. doi:10.1007/978-3-319-89722-6.
- [ABL⁺19] Nicola Atzei, Massimo Bartoletti, Stefano Lande, Nobuko Yoshida, and Roberto Zunino. Developing secure Bitcoin contracts with BitML. In *ESEC/FSE*, pages 1124–1128. ACM, 2019. doi:<https://doi.org/10.1145/3338906.3341173>.
- [ADMM14a] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Fair two-party computations via Bitcoin deposits. In *Financial Cryptography Workshops*, volume 8438 of *LNCS*, pages 105–121. Springer, 2014. doi:10.1007/978-3-662-44774-1_8.
- [ADMM14b] 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 (FORMATS)*, volume 8711 of *LNCS*, pages 7–22. Springer, 2014. doi:10.1007/978-3-319-10512-3_2.
- [ADMM14c] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on Bitcoin. In *IEEE S & P*, pages 443–458, 2014. First appeared on Cryptology ePrint Archive, <http://eprint.iacr.org/2013/784>. doi:10.1109/SP.2014.35.
- [ADMM16] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on Bitcoin. *Commun. ACM*, 59(4):76–84, 2016. doi:10.1145/2896386.
- [And19] Gavin Andresen. Bitcoin script size limit, 2019. BIP 16, <https://github.com/bitcoin/bips/blob/master/bip-0016.mediawiki#520-byte-limitation-on-serialized-script-size>.
- [BBL⁺21] Massimo Bartoletti, Andrea Bracciali, Cristian Lepore, Alceste Scalas, and Roberto Zunino. A formal model of Algorand smart contracts. In *Financial Cryptography and Data Security*, volume 12674 of *LNCS*, pages 93–114. Springer, 2021. doi:10.1007/978-3-662-64322-8_5.
- [BDM16] Wacław Banasik, Stefan Dziembowski, and Daniel Malinowski. Efficient zero-knowledge contingent payments in cryptocurrencies without scripts. In *ESORICS*, volume 9879 of *LNCS*, pages 261–280. Springer, 2016. doi:10.1007/978-3-319-45741-3_14.
- [BK14] Iddo Bentov and Ranjit Kumaresan. How to use Bitcoin to design fair protocols. In *CRYPTO*, volume 8617 of *LNCS*, pages 421–439. Springer, 2014. doi:10.1007/978-3-662-44381-1_24.
- [BKT17] Alex Biryukov, Dmitry Khovratovich, and Sergei Tikhomirov. Findel: Secure derivative contracts for Ethereum. In *Financial Cryptography Workshops*, volume 10323 of *LNCS*, pages 453–467. Springer, 2017. doi:10.1007/978-3-319-70278-0_28.
- [BLZ20] Massimo Bartoletti, Stefano Lande, and Roberto Zunino. Bitcoin covenants unchained. In *ISoLA*, volume 12478 of *ISOLA*, pages 25–42. Springer, 2020. doi:10.1007/978-3-030-61467-6_3.
- [BMZ20] Massimo Bartoletti, Maurizio Murgia, and Roberto Zunino. Renegotiation and recursion in Bitcoin contracts. In *Proc. COORDINATION*, volume 12134 of *LNCS*, pages 261–278. Springer, 2020. doi:10.1007/978-3-030-50029-0_17.
- [BZ17] Massimo Bartoletti and Roberto Zunino. Constant-deposit multiparty lotteries on Bitcoin. In *Financial Cryptography Workshops*, volume 10323 of *LNCS*, pages 231–247. Springer, 2017. doi:10.1007/978-3-319-70278-0.
- [BZ18a] Massimo Bartoletti and Roberto Zunino. BitML: a calculus for Bitcoin smart contracts. In *ACM CCS*, 2018. doi:10.1145/3243734.3243795.
- [BZ18b] Massimo Bartoletti and Roberto Zunino. BitML: a calculus for Bitcoin smart contracts. *IACR Cryptol. ePrint Arch.*, page 122, 2018. URL: <http://eprint.iacr.org/2018/122>.
- [BZ19] Massimo Bartoletti and Roberto Zunino. Verifying liquidity of Bitcoin contracts. In *POST*, volume 11426 of *LNCS*. Springer, 2019.
- [CCM⁺20] Manuel M.T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian, Michael Peyton Jones, and Philip Wadler. The extended UTXO model. In *Workshop on Trusted Smart Contracts*, 2020.
- [CDE⁺02] Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Jose F. Quesada. Maude: specification and programming in rewriting logic. *Theor. Comput. Sci.*, 285(2):187–243, 2002. doi:10.1016/S0304-3975(01)00359-0.
- [DBHP19] Ankush Das, Stephanie Balzer, Jan Hoffmann, and Frank Pfenning. Resource-aware session types for digital contracts. *CoRR*, abs/1902.06056, 2019.
- [Fla12] Matthew Flatt. Creating languages in Racket. *Commun. ACM*, 55(1):48–56, 2012. doi:10.1145/2063176.2063195.

- [GMR⁺19] Lewis Gudgeon, Pedro Moreno-Sanchez, Stefanie Roos, Patrick McCorry, and Arthur Gervais. Sok: Off the chain transactions. *IACR Cryptology ePrint Archive*, 2019:360, 2019.
- [JES00] 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)*, pages 280–292, 2000. doi:10.1145/351240.351267.
- [KB14] Ranjit Kumaresan and Iddo Bentov. How to use Bitcoin to incentivize correct computations. In *ACM CCS*, pages 30–41, 2014. doi:10.1145/2660267.2660380.
- [KB16] Ranjit Kumaresan and Iddo Bentov. Amortizing secure computation with penalties. In *ACM CCS*, pages 418–429, 2016. doi:10.1145/2976749.2978424.
- [KMB15] Ranjit Kumaresan, Tal Moran, and Iddo Bentov. How to use Bitcoin to play decentralized poker. In *ACM CCS*, pages 195–206, 2015. doi:10.1145/2810103.2813712.
- [KVV16] Ranjit Kumaresan, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. Improvements to secure computation with penalties. In *ACM CCS*, pages 406–417, 2016. doi:10.1145/2976749.2978421.
- [LCO⁺16] Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. Making smart contracts smarter. In *ACM CCS*, pages 254–269, 2016. doi:10.1145/2976749.2978309.
- [MB17] Andrew Miller and Iddo Bentov. Zero-collateral lotteries in Bitcoin and Ethereum. In *EuroS&P Workshops*, pages 4–13, 2017. doi:10.1109/EuroSPW.2017.44.
- [MES16] Malte Möser, Ittay Eyal, and Emin Gün Sirer. Bitcoin covenants. In *Financial Cryptography Workshops*, volume 9604 of *LNCS*, pages 126–141. Springer, 2016. doi:10.1007/978-3-662-53357-4_9.
- [OP17] Russell O’Connor and Marta Piekarska. Enhancing Bitcoin transactions with covenants. In *Financial Cryptography Workshops*, volume 10323 of *LNCS*. Springer, 2017. doi:10.1007/978-3-319-70278-0_12.
- [ST18] Pablo Lamela Seijas and Simon J. Thompson. Marlowe: Financial contracts on blockchain. In *ISoLA*, volume 11247 of *LNCS*, pages 356–375. Springer, 2018. doi:10.1007/978-3-030-03427-6_27.
- [TDD⁺18] Petar Tsankov, Andrei Marian Dan, Dana Drachler-Cohen, Arthur Gervais, Florian Bünzli, and Martin T. Vechev. Securify: Practical Security Analysis of Smart Contracts. In *ACM CCS*, pages 67–82, 2018. doi:10.1145/3243734.3243780.