

BitML: a calculus for Bitcoin smart contracts

Massimo Bartoletti
University of Cagliari
Cagliari, Italy
bart@unica.it

Roberto Zunino
University of Trento
Trento, Italy
roberto.zunino@unitn.it

Abstract—We propose a domain-specific language for smart contracts, which allows participants to transfer cryptocurrency according to agreed contract terms. We define a symbolic and a computational model for reasoning about their security. In the symbolic model, participants act according to the semantics of the domain-specific language. Instead, in the computational model they exchange bitstrings, and publish transactions on the Bitcoin blockchain. A compiler is provided to translate smart contracts into standard Bitcoin transactions. We prove the correctness of our compiler, showing that computational attacks to compiled smart contracts are also observable in the symbolic model.

Index Terms—Bitcoin, smart contracts, process calculi

I. INTRODUCTION

Cryptocurrencies like Bitcoin and Ethereum have revived the idea of *smart contracts*, agreements between untrusted parties that can be automatically enforced without a trusted intermediary [1]. These agreements define protocols which regulate cryptocurrency exchanges among participants. For instance, a smart contract can be a lottery which collects bets from a set of players, determines the winner in a fair manner, and then transfers the whole pot to the winner.

Disintermediation is made possible by the *blockchain*, a public, append-only record of transactions, and by the *consensus protocol* of the cryptocurrency [2]. The execution of smart contracts relies on the blockchain to log all the participants’ moves; further, the underlying logic of transactions is exploited to enable all and only the moves permitted by the contract. The consensus protocol is used to consistently update the blockchain: economic incentives are provided to ensure that the nodes of the cryptocurrency network have the same view of the blockchain. In this way, the state of each contract (and consequently, the asset of each user) is uniquely determined by the sequence of its transactions on the blockchain.

Smart contracts have different incarnations, depending on the cryptocurrency on which they are based. In Ethereum, contracts are expressed as programs in a Turing-equivalent language. Any user can publish a contract on the blockchain. This makes the contract available to other users, who can then run it by calling its functions (concretely, by publishing suitable transactions on the blockchain). Such openness comes at the price of a wide attack surface: attackers may exploit vulnerabilities in the implementation of contracts, or may publish themselves Trojan-horse contracts with hidden vulnerabilities, to steal or tamper with the assets controlled by contracts. Indeed, a series of vulnerabilities in Ethereum contracts [3] have been exploited, causing money losses in the order of hundreds of millions of dollars [4]–[6].

Unlike Ethereum, Bitcoin does not provide a language for smart contracts: rather, they are described as cryptographic protocols where participants send/receive messages, verify signatures, and put/search transactions on the blockchain. Lotteries [7]–[10], Poker [11], contingent payments [12]–[14], and more general fair multi-party computations [15], [16] witness the variety of smart contracts supported by Bitcoin.

Describing smart contracts at this level of abstraction is complex and error-prone, for two reasons. First, they often rely on advanced features of Bitcoin, (e.g., transaction scripts, signature modifiers, temporal constraints), whose actual behaviour relies on low-level implementation details. Second, establishing the correctness of each smart contract requires to prove the computational security of its protocol. This task requires the skills of expert cryptographers, and even in this case it is a significant effort. By contrast, working in an high-level symbolic model would relieve smart contract programmers from (most of) this burden, since the much higher level of abstraction would allow security proofs to be carried out with automatic tools.

Contributions: We propose BitML, a domain-specific language for Bitcoin smart contracts. BitML is a process calculus, with primitives to stipulate contracts, and to exchange currency according to the contract terms, possibly involving temporal constraints. BitML can express most of the Bitcoin smart contracts proposed so far in the literature [17], e.g. escrow services, timed commitments, multi-player lotteries, etc.

We provide our calculus with a *symbolic semantics*, and a security model where participants use strategies to choose which moves to perform (among those allowed by the semantics). Each honest participant has her own strategy, while the dishonest ones are controlled by a single adversarial strategy. The adversary is based on the (symbolic) Dolev-Yao model: he schedules the participant moves, eavesdrops messages, and impersonates other participants. However, we prevent the adversary from dropping messages forever, coherently with the peer-to-peer nature of the Bitcoin communication network, where (proper) requests are always eventually processed. Since BitML features temporal constraints, our security model handles time. More specifically, we allow the adversary to let time pass, but only provided that the other participants are also willing to. In this way we rule out attacks where the adversary advances the time making a participant miss a deadline.

Our main contribution is a compiler which translates BitML contracts into standard Bitcoin transactions. Participants can

perform the actions of the contract by publishing the corresponding transactions on the blockchain. The main technical challenge is to guarantee the correctness of the compiler, i.e. that the symbolic execution of the contract matches the “computational” one performed on the Bitcoin network. This correspondence must hold also in the presence of computational adversaries: otherwise, attacks at the Bitcoin level could be unobservable at the level of the symbolic semantics.

To prove the correctness of the compiler, we start by defining a *computational model* of the honest participants and of the adversary. Both are equipped with computational strategies, which allow them to inspect the past run, and choose which actions to perform. In the computational model, actions can broadcast bitstrings, publish transactions on the blockchain, and make time pass. Compared to the symbolic level, the computational level extends the power of the adversary: he can arbitrarily manipulate bitstrings, being only subject to complexity bounds, and he can exploit any kind of valid Bitcoin transaction, beyond those obtained by the compiler.

Then, we introduce a relation between symbolic and computational runs, called *coherence*. Intuitively, coherence holds when both runs perform “morally” equivalent actions: e.g., the same amount of bitcoins is spent in both runs at the same time, for the same purpose. Defining coherence is technically demanding, because the symbolic model is far more abstract than the computational one; further, the computational adversary can perform actions with no symbolic counterpart.

To illustrate one of the subtleties in formalising coherence, suppose that a dishonest participant A owns $v\mathfrak{B}$, represented both in the symbolic run (as a term $\langle A, v \rangle$) and in the computational run (as a transaction T). Computationally, A can redeem T with a transaction T' which can not be represented symbolically (e.g., because it can not be produced by the compiler). In such case, a coherent symbolic step is to make $\langle A, v \rangle$ disappear, as if $v\mathfrak{B}$ were destroyed. In subsequent steps, coherence will not keep track of the descendants of T' in the symbolic run. This loss of information at the symbolic level could in principle allow for computational attacks without a symbolic counterpart. However, this is not the case, since computational attacks can always be adapted so to have a symbolic counterpart. Indeed, to attack honest participants, A has to stipulate contracts with them: this requires A to put a deposit, computationally represented as a transaction T'' . Instead of obtaining T'' from T' , which makes T'' untraceable at the symbolic level, A can perform symbolically-traceable actions to create from T a deposit T''' with the same value of T'' , to be used in the computational attack. This adaptation is feasible, because untraceable computational actions do not allow the adversary to artificially increase his wealth. Hence, the value of T'' can not exceed the value of T , so making it possible to produce T''' with symbolic actions.

Leveraging on the definition of coherence, we establish the correctness of the compiler, as a *computational soundness* result [18]. Specifically, we prove that if honest participants execute contracts using symbolic strategies and compiler-generated transactions, then the computational runs resulting

from their interaction with the adversary will have a symbolic counterpart (with overwhelming probability). Consequently, proofs of trace-based security properties carried out in the symbolic model can be lifted *for free* to the computational model. This result is crucial, since it enables the development of analysis and verification techniques *at the symbolic level*, which would be much more burdensome to obtain at the (far more concrete) computational level.

II. A CALCULUS FOR BITCOIN CONTRACTS

In BitML, contracts allow participants to interact according to the following workflow:

- 1) A participant broadcasts a *contract advertisement* $\{G\}C$. The component C is the contract, which specifies how funds can be transferred among participants; G is a set of *preconditions* to its execution. Roughly, G requires participants to deposit some \mathfrak{B} , either upfront or during the contract execution, and to commit to some secrets.
- 2) Participants can then choose whether to accept the advertisement, or not. When all the involved participants have accepted $\{G\}C$, satisfying its preconditions, the contract C becomes stipulated. Then, participants can transfer the deposited funds by acting as prescribed by C .

We assume a set Part of *participants*, ranged over by A, B, \dots , and we denote with $\text{Hon} \subseteq \text{Part}$ a non-empty set of *honest* participants. We also assume a set of names, of two kinds: x, y, \dots denote *deposits* of \mathfrak{B} , while a, b, \dots denote *secrets*. We denote with \bar{x} a finite sequence of deposit names, and we adopt a similar notation for sequences of other kinds.

The preconditions in G have three possible forms:

- $A : !v @ x$ (“persistent deposit”) requires A to own $v\mathfrak{B}$ in a deposit x , and to spend it for stipulating the contract.
- $A : ?v @ x$ (“volatile deposit”) only requires A to authorize the spending of x during the execution of C . Since x is not paid upfront, there is no guarantee that $v\mathfrak{B}$ will be available when C demands x , as A can spend it for other purposes. Intuitively, volatile deposits $A : ?v @ x$ give A the option of contributing $v\mathfrak{B}$ during the contract execution, but there is no obligation for A to do so.
- $A : \text{secret } a$ requires A to generate a random nonce a , and commit to it by publishing its hash $H(a)$ before C starts. During the execution of C , A can choose whether to reveal a to the other participants.

Definition 1 (Contract preconditions).

$G ::=$	contract precondition
$A : ?v @ x$	volatile deposit of $v\mathfrak{B}$, expected from A
$A : !v @ x$	persistent deposit of $v\mathfrak{B}$, expected from A
$A : \text{secret } a$	committed secret by A
$ G G$	composition ◇

Once C is stipulated, it starts its execution with a *balance*, initially set to the sum of the persistent deposits in its advertisement. The execution of C will affect this balance, when participants deposit/withdraw funds to/from the contract.

We formalise contracts using a process calculus. A contract C is a *choice* among branches. Intuitively, a branch D performs an action, and possibly proceeds with a continuation C' . The action `put x & reveal a if p` atomically performs the following: (i) spend all the volatile deposits x , adding their values to the current balance; (ii) check that all the secrets a have been revealed, and that they satisfy the predicate p .

The guarded contract `split $v_1 \rightarrow C_1 \mid \dots \mid v_n \rightarrow C_n$` divides the contract into n contracts C_i , each one with balance v_i . The sum of the v_i must be equal to the current balance.

The prefix `withdraw A` transfers the whole balance to A (to transfer only a part of it, one can perform a `split`).

Note that, when enabled, the above actions can be fired by anyone at anytime. To restrict *who* can execute a branch 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 .

Definition 2 (Contracts). We define the syntax of contracts in Figure 1. We denote with 0 the empty sum. We abbreviate `put x & reveal a if p` as: (i) `put x` when a is empty and p is *true*, (ii) `reveal a if p` when x is empty, (iii) τ when x and a are empty and p is *true*, and (iv) we omit “if p ” when the predicate p is true. In guarded contracts, we assume that the order of decorations is immaterial, e.g., we consider `after $t : A : B : D$` equivalent to `B : A : after $t : D$` . \diamond

Definition 3 (Contract advertisement). A contract advertisement is a term $\{G\}C$ such that the following conditions hold: (i) all the names in G are distinct; (ii) all the names in C occur in G ; (iii) all the names in `put x & reveal a if p` are distinct; further, all the names in p occur in a ; (iv) each participant in $\{G\}C$ has a persistent deposit in G . \diamond

The last condition is not restrictive in practice: we can craft a contract that allows a participant A to deposit just a small fraction of bitcoin, and then immediately transfer it back to A through a `split`. In the following examples, we do not show these dummy deposits in the contract preconditions.

Example 1 (Escrow). A buyer A wants to buy an item from a seller B , but they do not trust each other. So, they would like to use a contract to ensure that B will get paid if and only if A gets her item. We would like to guarantee that, even if either A or B are dishonest, exactly one of them will be able to redeem the money: if they disagree, an arbiter C will decide who gets the money.

We model this contract in BitML as follows. The precondition $G = A : !v @ x$ requires A to provide a persistent deposit of $v\text{฿}$. The contract C is a choice among four branches:

$$C = A : \text{withdraw } B + B : \text{withdraw } A \\ + C : \text{withdraw } A + C : \text{withdraw } B$$

In the first branch, `A : withdraw B` , participant B can redeem $v\text{฿}$ if A provides her authorization. The second branch is dual, and can be used e.g., if the seller wants to refund the buyer for a damaged item. The last two branches are used if neither

A nor B give their authorizations: in this case, the arbiter C chooses whether to authorize A or B to redeem the deposit.

We also propose a variant of the contract where C can issue a *partial* refund of $\zeta v\text{฿}$ to A , and of $(1-\zeta)v\text{฿}$ to B (similarly to [19], [20]). The possible values of ζ are given by a finite set Z in the range $[0, 1]$. We model the new contract as:

$$C' = A : \text{withdraw } B + B : \text{withdraw } A + \sum_{\zeta \in Z} C : D_\zeta \\ D_\zeta = \text{split } (\zeta v \rightarrow \text{withdraw } A \mid (1-\zeta)v \rightarrow \text{withdraw } B)$$

The case of full refunds is obtained with $Z = \{0, 1\}$. If $Z = \{0, 1/2, 1\}$, C can also choose to refund $v/2\text{฿}$ to both. \diamond

Example 2 (Intermediated payment). Assume that A wants to send an indirect payment of $v_C\text{฿}$ to C , routing it through an intermediary B who can choose whether to authorize the payment, in this case retaining a fee $v_B\text{฿}$, or not. Since A does not trust B , she wants to use a contract to guarantee that: (i) if B is honest, then $v_C\text{฿}$ are transferred to C ; (ii) if B is *not* honest, then A does not lose money.

In BitML, we use $G = A : ! (v_B + v_C) @ x$ as precondition, and the following contract:

$$C = B : \text{split } (v_B \rightarrow \text{withdraw } B \mid v_C \rightarrow \text{withdraw } C) \\ + \text{after } t : \text{withdraw } A$$

The first branch can only be taken if B authorizes the payment: in this case, B gets his fee, and C gets his payment. If B denies his authorization, after time t , A can redeem her deposit. \diamond

Example 3 (Timed commitment). Assume that A wants to choose a secret a , and reveal it after some time t — guaranteeing that the revealed value is the chosen secret. This can be obtained through a *timed commitment* protocol [15], [21]–[23]. As in [15], we force A to pay to B a penalty of $v\text{฿}$ if A does not reveal the secret within t . In BitML, we use the precondition $G = A : !v @ x \mid A : \text{secret } a$, and:

$$C = (\text{reveal } a. \text{withdraw } A) + (\text{after } t : \text{withdraw } B)$$

Only A can choose the first branch, by revealing a . After that, anyone can further reduce the contract, and transfer $v\text{฿}$ to A . Only after time t , B can choose the second branch, and collect A 's penalty. Before the deadline, A has the option to reveal a (avoiding the penalty), or to keep it secret (paying the penalty). If A reveals a after time t , a race condition occurs: in such case, the first one who fires the `withdraw` gets the money.

We can also model *mutual* timed commitment as follows:

$$G = A : !v @ x \mid A : \text{secret } a \mid B : !v @ y \mid B : \text{secret } b \\ C = \text{reveal } a. C' + (\text{after } t : \text{withdraw } B) \\ C' = \text{reveal } b. C'' + (\text{after } t : \text{withdraw } A) \\ C'' = \text{split } (v \rightarrow \text{withdraw } A \mid v \rightarrow \text{withdraw } B)$$

The contract C can reduce to C' if A reveals a , otherwise (after t) B can redeem $2v\text{฿}$. In C' , if B reveals b , then both participants can redeem their deposits, running C'' . Otherwise, A can redeem $2v\text{฿}$. \diamond

$C ::= \sum_{i \in I} D_i$	contract	$p ::=$	predicate	$E ::=$	arithmetic expression
$D ::=$	guarded contract	$true$	truth	N	32-bit constant
put x & reveal a if $p.C$	collect deposits x and secrets a	$ p \wedge p$	conjunction	$ a $	length of a secret
withdraw A	transfer the balance to A	$ \neg p$	negation	$ E + E$	addition
split $v \rightarrow C$	split the balance ($ v = C $)	$ E = E$	equality	$ E - E$	subtraction
$A : D$	wait for A 's authorization	$ E < E$	less than		
after $t : D$	wait until time t				

Fig. 1: Syntax of BitML contracts.

Example 4 (Two-players lottery). A multiparty lottery is a protocol where n players put their bets in a pot, and a winner — fairly chosen among the players — redeems the whole pot. Various contracts for multiparty lotteries on Bitcoin have been proposed in [8]–[10], [15], [24], [25].

We model the lottery in [15], for two players A and B who bet 1฿ each. The contract preconditions are the following:

$$G = A : ! 2 @ y_A \mid A : ? 1 @ x_A \mid A : \text{secret } a \\ \mid B : ! 2 @ y_B \mid B : ? 1 @ x_B \mid B : \text{secret } b$$

where a and b are committed secrets, y_A, y_B , are collaterals (used as compensations in case of dishonest behaviour), and x_A, x_B are the bets. The contract C is the following:

```
put  $x_A x_B$ . (split
  2  $\rightarrow$  reveal  $b$  if  $0 \leq |b| \leq 1$ . withdraw  $B$ 
  + after  $t'$ : withdraw  $A$ 
| 2  $\rightarrow$  reveal  $a$  if  $0 \leq |a| \leq 1$ . withdraw  $A$ 
  + after  $t'$ : withdraw  $B$ 
| 2  $\rightarrow$  reveal  $ab$  if  $|a| = |b|$ . withdraw  $A$ 
  + reveal  $ab$  if  $|a| \neq |b|$ . withdraw  $B$ )
+ after  $t$ : split (2  $\rightarrow$  withdraw  $A$  | 2  $\rightarrow$  withdraw  $B$ )
```

The players can put their bets by performing put; if they do not place their bets within t , each player can use the after branch to get their collateral back. After the bets are placed, the balance is split in three parts. Player B must reveal b before the deadline t' ; otherwise, A can redeem B 's collateral (note that this is a timed commitment, as in Example 3). Similarly, A must reveal a . To fairly determine the winner, we further require that the length of secrets is either 0 or 1 (after compiling to Bitcoin, the actual lengths are $\eta + 0$ and $\eta + 1$ so to avoid trivial preimage attacks). In this way, the third part of the split can compute the winner: A if the secrets have the same length, B otherwise. \diamond

We now define a reduction semantics of contracts. It is organised in two layers: a bottom layer, taking the form of an LTS between (untimed) configurations, and a top layer, in the form of a timed LTS between *timed configurations*.

An (untimed) configuration Γ is a parallel composition of:

- *contract advertisements* $\{G\}C$, representing a contract which has been proposed, but not stipulated yet.
- *active contracts* $\langle C, v \rangle_x$, representing a stipulated contract, holding a current balance of $v\text{฿}$. The name x uniquely identifies the active contract.

- *personal deposits* $\langle A, v \rangle_x$, representing a fund of $v\text{฿}$ owned by A , and with unique name x .
- *authorizations* $A[\chi]$, representing the consent of A to perform some operation χ .
- *committed secrets* $\{A : a \# N\}$, i.e. A has committed a random nonce a , by broadcasting its hash $H(a)$. The length of a , which is secret as well, is determined by N .
- *revealed secrets* $A[a \# N]$, representing the fact that A has revealed her secret a (hence, its length N).

Definition 4 (Configurations). We define configurations Γ, Δ, \dots through the syntax in Figure 2, where we stipulate that in a configuration there are no duplicate authorizations. We assume that $(|, 0)$ is a commutative monoid. Indexed parallel compositions are denoted with \parallel . We say that Γ is *initial* when it contains only terms of the form $\langle A, v \rangle_x$. \diamond

Definition 5 (LTS of untimed configurations). The LTS of configurations is defined in Figure 3 (advertisements and stipulation), Figure 4 (contracts), and Figure 8 (deposits). The latter is deferred to Section A, as for the semantics of predicates (Definition 24). The function cv from labels α to sets of names is defined as: $cv(\text{put}(x, a, y)) = cv(\text{withdraw}(A, v, y)) = cv(\text{split}(y)) = \{y\}$, otherwise $cv(\alpha)$ is empty. \diamond

When a participant A owns a deposit $\langle A, v \rangle_x$, she can employ that amount for several operations. For instance, she can divide the deposit into two smaller deposits, or join it with another deposit of hers to form a larger one. The deposit can also be transferred to another participant, or used to stipulate a contract. Before such operations can be performed, A must authorize it (at the Bitcoin level, this happens through suitable signatures of A). For this reason, we model both the authorization step and the actual operation using two distinct LTS moves, formalized by two separate semantics rules. For instance, rule [DEP-AUTHJOIN] authorizes the join of two deposits:

$$\langle A, v \rangle_x \mid \langle A, v' \rangle_y \mid \Gamma \rightarrow \langle A, v \rangle_x \mid \langle A, v' \rangle_y \mid A[\chi_x] \mid \Gamma$$

where $\chi_x = x, y \triangleright \langle A, v + v' \rangle$ authorizes to spend x . After we also obtain the dual authorization χ_y , rule [DEP-JOIN] performs the actual join. Given $\Gamma = A[\chi_x] \mid A[\chi_y] \mid \Gamma'$, we infer:

$$\langle A, v \rangle_x \mid \langle A, v' \rangle_y \mid \Gamma \rightarrow \langle A, v + v' \rangle_z \mid \Gamma$$

Rules for contracts are better explained through an example.

$\Gamma, \Delta ::=$	untimed configuration	$\chi ::=$	authorization to ...
0	empty	$x \triangleright \{G\}C$	accept an advertised contracts $\{G\}C$
$ \{G\}C$	contract advertisement	$x \triangleright D$	take branch D
$ \langle C, v \rangle_x$	an active contract containing $v\mathfrak{B}$	$x, y \triangleright \langle A, v \rangle$	join deposit x with y into a deposit for A
$ \langle A, v \rangle_x$	a deposit of $v\mathfrak{B}$ redeemable by A	$x \triangleright \langle A, v \rangle, \langle A, v' \rangle$	divide a deposit x in two deposits for A
$ \mathbf{A}[\chi]$	authorization of A to perform χ	$x \triangleright B$	donate deposit x to B
$ \{A : a \# N\}$	committed secret of A ($N \in \mathbb{N} \cup \{\perp\}$)	$\mathbf{x}, i \triangleright y$	destroy i -th deposits in \mathbf{x} through y
$ \mathbf{A}[a \# N]$	revealed secret of A ($N \in \mathbb{N}$)		
$ \Gamma \mid \Delta$	parallel composition		

Fig. 2: Configurations.

$$\begin{array}{c}
\frac{\{G\}C \text{ contains at least one participant in Hon} \quad \Gamma \text{ contains } \langle A_i, v_i \rangle_{x_i} \text{ for all } A_i : ! v_i \otimes x_i \text{ in } \{G\}C}{a \text{ fresh, for each } A : \text{secret } a \text{ in } G \quad \Gamma \text{ contains } \langle A_i, v_i \rangle_{x_i} \text{ for all } A_i : ? v_i \otimes x_i \text{ in } \{G\}C} \text{[C-ADVERTISE]} \\
\Gamma \xrightarrow{\text{advertise}(\{G\}C)} \{G\}C \mid \Gamma \\
\\
\frac{a_1 \cdots a_k \text{ secrets of } A \text{ in } G \quad \forall i \in 1..k : \mathfrak{B}N : \{A : a_i \# N\} \in \Gamma \quad \forall i \in 1..k : N_i \in \begin{cases} \mathbb{N} & \text{if } A \in \text{Hon} \\ \mathbb{N} \cup \{\perp\} & \text{otherwise} \end{cases}}{\Delta = \{A : a_1 \# N_1\} \mid \cdots \mid \{A : a_k \# N_k\}} \text{[C-AUTHCOMMIT]} \\
\{G\}C \mid \Gamma \xrightarrow{A : \{G\}C, \Delta} \{G\}C \mid \Gamma \mid \Delta \\
\\
\frac{\forall B : \forall a \text{ secret of } B \text{ in } G : \exists N : (\{B : a \# N\} \in \Gamma \text{ or } B[a \# N] \in \Gamma) \quad G \text{ contains a deposit } x \text{ of } A}{\{G\}C \mid \Gamma \xrightarrow{A : \{G\}C, x} \{G\}C \mid \Gamma \mid \mathbf{A}[x \triangleright \{G\}C]} \text{[C-AUTHINT]} \\
\\
\frac{G = (\|_{i \in I} A_i : ? v_i \otimes x_i) \mid (\|_{i \in J} B_i : ! v'_i \otimes y_i) \mid (\|_{i \in K} C_i : \text{secret } a_i) \quad x \text{ fresh}}{\Gamma = (\|_{i \in I} A_i[x_i \triangleright \{G\}C]) \mid (\|_{i \in J} B_i[y_i \triangleright \{G\}C]) \mid \Gamma'} \text{[C-INIT]} \\
\{G\}C \mid \Gamma \mid (\|_{i \in J} \langle B_i, v'_i \rangle_{y_i}) \xrightarrow{\text{init}(G, C)} \langle C, \sum_{i \in J} v'_i \rangle_x \mid \Gamma
\end{array}$$

Fig. 3: Semantics of untimed configurations: rules for advertisement and stipulation.

$$\begin{array}{c}
\frac{v = v_1 \cdots v_k \quad C = C_1 \cdots C_k \quad \sum_{i=1}^k v_k = v' \quad x_1 \cdots x_k \text{ fresh}}{\langle \text{split } v \rightarrow C, v' \rangle_y \mid \Gamma \xrightarrow{\text{split}(y)} (\|_{i=1}^k \langle C_i, v_i \rangle_{x_i}) \mid \Gamma} \text{[C-SPLIT]} \\
\\
\frac{N \neq \perp}{\{A : a \# N\} \xrightarrow{A : a} A[a \# N]} \text{[C-AUTHREV]} \quad \frac{x = x_1 \cdots x_m \quad \Gamma = \|\|_{i=1}^m \langle A_i, v_i \rangle_{x_i} \quad z \text{ fresh} \quad a = a_1 \cdots a_n \quad \Delta = \|\|_{i=1}^n B_i[a_i \# N_i] \quad \llbracket p \rrbracket_\Delta = \text{true}}{\langle \text{put } x \ \& \ \text{reveal } a \ \text{if } p. C, v \rangle_y \mid \Gamma \mid \Delta \mid \Gamma' \xrightarrow{\text{put}(x, a, y)} \langle C, v + \sum_{i=1}^m v_i \rangle_z \mid \Delta \mid \Gamma'} \text{[C-PUTREV]} \\
\\
\frac{x \text{ fresh}}{\langle \text{withdraw } A, v \rangle_y \mid \Gamma \xrightarrow{\text{withdraw}(A, v, y)} \langle A, v \rangle_x \mid \Gamma} \text{[C-WITHDRAW]} \quad \frac{D \equiv A : D'}{\langle D + C, v \rangle_x \mid \Gamma \xrightarrow{A : x, D} \langle D + C, v \rangle_x \mid \mathbf{A}[x \triangleright D]} \text{[C-AUTHCONTROL]} \\
\\
\frac{D \equiv A_1 : \cdots : A_k : \text{after } t_1 : \cdots : \text{after } t_m : D' \quad D' \not\equiv A : \cdots}{\langle D', v \rangle_x \mid (\|_{i=1}^k A_i[x \triangleright D]) \mid \Gamma \xrightarrow{\alpha} \Gamma' \quad x \in \text{cv}(\alpha) \quad D' \not\equiv \text{after } t : \cdots}{\langle D + C, v \rangle_x \mid (\|_{i=1}^k A_i[x \triangleright D]) \mid \Gamma \xrightarrow{\alpha} \Gamma'} \text{[C-CONTROL]}
\end{array}$$

Fig. 4: Semantics of untimed configurations: rules for contracts.

$$\frac{\Gamma \xrightarrow{\alpha} \Gamma' \quad \text{cv}(\alpha) = \emptyset}{(\Gamma, t) \xrightarrow{\alpha} (\Gamma', t)} \text{[ACTION]} \quad \frac{\delta > 0}{(\Gamma, t) \xrightarrow{\delta} (\Gamma, t + \delta)} \text{[DELAY]} \quad \frac{D \equiv \text{after } t_1 : \cdots : \text{after } t_m : D' \quad D' \not\equiv \text{after } t' : \cdots}{\langle D, v \rangle_x \mid \Gamma \xrightarrow{\alpha} \Gamma' \quad x \in \text{cv}(\alpha) \quad t \geq t_1, \dots, t_m} \text{[TIMEOUT]}$$

Fig. 5: Semantics of timed configurations.

Example 5. Recall the timed commitment in Example 3. Assume that $A \in \text{Hon}$, and let $\Gamma = \langle A, v \rangle_x$. A possible (untimed) computation, where A reveals her secret and then redeems the deposit, is the following:

$$\begin{aligned}
\Gamma &\rightarrow \Gamma \mid \{G\}C & (1) \\
&\rightarrow \Gamma \mid \{G\}C \mid \{A : a\#N\} & (2) \\
&\rightarrow \Gamma \mid \{G\}C \mid \{A : a\#N\} \mid A[x \triangleright \{G\}C] & (3) \\
&\rightarrow \langle C, v \rangle_{x_1} \mid \{A : a\#N\} \mid A[x \triangleright \{G\}C] = \Gamma_1 & (4) \\
&\rightarrow \langle C, v \rangle_{x_1} \mid A[a\#N] \mid A[x \triangleright \{G\}C] & (5) \\
&\rightarrow \langle \text{withdraw } A, v \rangle_{x_2} \mid A[a\#N] \mid A[x \triangleright \{G\}C] & (6) \\
&\rightarrow \langle A, v \rangle_{x_3} \mid A[a\#N] \mid A[x \triangleright \{G\}C] & (7)
\end{aligned}$$

Step (1) advertises $\{G\}C$, by rule [C-ADVISE]. This move is possible because A is honest, and her deposit is in Γ . At step (2), A commits to a secret, by rule [C-AUTHCOMMIT]. This is possible because there are no previous commitments of a in Γ ; the length N is a natural, since A is honest. At step (3), A gives her authorization to stipulate C , by rule [C-AUTHINIT]. This is possible since the configuration contains A 's deposit and committed secret. At step (4) the contract C becomes stipulated, by rule [C-INIT]. Stipulation is possible because A authorized it, and made available her persistent deposit x . After the move, the $v\mathbb{B}$ in x are transferred to the contract. At step (5), A reveals her secret by rule [C-AUTHREV], which verifies that $N \neq \perp$. After that, the action `reveal` a is performed at step (6), by rule [C-PUTREV], followed by [C-CONTROL] with $k = m = 0$. Finally, A withdraws $v\mathbb{B}$ at step (7) by rule [C-WITHDRAW], obtaining a fresh deposit x_3 . \diamond

We discuss a few subtleties in the rules. First, [C-ADVISE] requires as a side condition that at least one of the participants involved in each stipulation is honest (one of the weakest assumptions in cryptographic protocols). The same effect of running contracts among dishonest participants can still be obtained by redistributing funds through the rules for deposits. Hence, this side condition does not affect the power of the adversary. A further motivation for the side condition is that the correctness of our compiler will rely on this assumption.

Rule [C-AUTHCOMMIT] allows dishonest participants to choose an “invalid” length \perp for their committed secrets. This reflects the fact that, in the computational model, A commits to a secret by broadcasting a bitstring, meant to be the hash of the secret. If A is dishonest, she could instead broadcast an arbitrary bitstring w , preventing herself later on to reveal a preimage of w . Similarly, the length \perp prevents the reveal action in the symbolic model.

Definition 6 (LTS of timed configurations). Timed configurations are pairs (Γ, t) , where $t \in \mathbb{N}$ is the global time. The LTS between timed configurations is defined in Figure 5, where labels α are either untimed labels, or delays $\delta \in \mathbb{N}$. \diamond

Example 6. Recall the configuration Γ_1 from Example 5. Starting from $(\Gamma_1, 0)$ we can have an alternative timed computation, where B waits until $t' > t$ to redeem A 's deposit:

$$\begin{aligned}
(\Gamma_1, 0) &\rightarrow (\Gamma_1, t') & (8) \\
&\rightarrow (\langle B, v \rangle_y \mid \{A : a\#N\} \mid A[x \triangleright \{G\}C], t') & (9)
\end{aligned}$$

Step (8) lets the time pass, by rule [DELAY]. At step (9), B fires the prefix `withdraw` B within the `after`, and in this way he collects $v\mathbb{B}$ from A 's deposit. This is obtained by using rule [C-WITHDRAW] in the premise of [C-TIMEOUT]. \diamond

III. SYMBOLIC STRATEGIES AND ADVERSARIES

A symbolic strategy Σ^s is a PPTIME algorithm which allows participants to select which action(s) to perform, among those permitted by the semantics. We distinguish between two kinds of strategies: those for honest participants, and that for the adversary Adv (who also controls the dishonest participants). Strategies can inspect the past execution.

Definition 7 (Symbolic runs). A symbolic run R^s is a (possibly infinite) sequence $(\Gamma_0, t_0) \xrightarrow{\alpha_0} (\Gamma_1, t_1) \xrightarrow{\alpha_1} \dots$ where Γ_0 is initial, and $t_0 = 0$. If R^s is finite, we write Γ_{R^s} for its last untimed configuration, and δ_{R^s} for the last global time. We write $R^s \xrightarrow{\alpha} \dot{R}^s$ when \dot{R}^s extends R^s with the transition $(\Gamma_{R^s}, \delta_{R^s}) \xrightarrow{\alpha} (\Gamma_{\dot{R}^s}, \delta_{\dot{R}^s})$. \diamond

The strategy of a participant can inspect the past run, except for the (lengths of the) committed secrets which have not been revealed yet. The *stripping* of a symbolic run censors this information. Instead, revealed secrets are not censored.

Definition 8 (Stripping of symbolic runs). The stripping of a symbolic run, denoted $\text{strip}(R^s)$, is the sequence obtained from R^s by replacing each committed secret $\{A : a\#N\}$ with $\{A : a\#\perp\}$, and each label $A : \{G\}C, \Delta$ with $A : \{G\}C, 0$. We accordingly define label stripping. \diamond

A strategy receives as input a stripped run R_*^s . Further, the strategy of A has access to an infinite sequence $r_A \in \{0, 1\}^\omega$ of independent and uniformly random bits. Instead of modelling the access to r_A through an oracle, we simply pass r_A as input to the strategy.

Definition 9 (Randomness sources). A randomness source is a function r from participants $A \in \text{Hon} \cup \{\text{Adv}\}$ to infinite bitstrings $r(A) \in \{0, 1\}^\omega$. We usually write r_A for $r(A)$. \diamond

Definition 10 (Symbolic participant strategies). The symbolic strategy of a participant $A \in \text{Hon}$ is a PPTIME algorithm $\Sigma_A^s(R_*^s, r_A)$, taking as input a stripped symbolic run R_*^s and a random sequence r_A . The output is a finite sequence of α -moves such that the following constraints hold:

- 1) if $\alpha \in \Sigma_A^s(\text{strip}(R^s), r_A)$, then $R^s \xrightarrow{\alpha}$;
- 2) if $B : \dots \in \Sigma_A^s(\text{strip}(R^s), r_A)$, then $B = A$;
- 3) if $A : \{G\}C, \Delta$ and $A : \{G\}C, \Delta'$ in $\Sigma_A^s(\text{strip}(R^s), r_A)$, then $\Delta = \Delta'$;
- 4) if $\alpha \in \Sigma_A^s(\text{strip}(R^s), r_A)$ and $R^s \xrightarrow{\alpha_1} \dot{R}^s \xrightarrow{\alpha}$, then $\alpha \in \Sigma_A^s(\text{strip}(\dot{R}^s), r_A)$. \diamond

The constraints in Definition 10 are needed to rule out ill-formed strategies. (1) requires that Σ_A^s only chooses moves enabled by the semantics; (2) states that Σ_A^s cannot choose moves of B ; (3) guarantees that the lengths of secrets are chosen coherently (i.e., A cannot choose different lengths for the same secret); (4) requires the strategy to be *persistent*: if

on a run Σ_A^s chooses α , and α is not taken as the next step in the run (e.g., because some other participant acts earlier), then Σ_A^s must still choose α after that step, if still enabled.

The adversary Adv acts on behalf of all the dishonest participants. Her symbolic strategy has access to the current (stripped) run, and to a random sequence r_{Adv} . Following the Dolev-Yao model, Adv controls the scheduling among all participants. At each moment, Adv can see all the sequences of moves Λ_i^s chosen by each honest A_i , and can perform *exactly* one action, denoted as λ^s . This can be one of the actions in some Λ_i^s (represented as a pair (A_i, j) , where $j \leq |\Lambda_i^s|$), or one action α not requiring the authorization of an honest participant, or a delay δ if all the honest participants agree.

Definition 11 (Symbolic adversary strategies). A symbolic adversary strategy is a PPTIME algorithm $\Sigma_{\text{Adv}}^s(R_*^s, r_{\text{Adv}}, \Lambda^s)$, taking as input a stripped symbolic run R_*^s , a random sequence r_{Adv} , and a list $\Lambda^s = \Lambda_1^s \cdots \Lambda_k^s$ of sequences of stripped moves for each $A_i \in \text{Hon} = \{A_1, \dots, A_k\}$. The output is a single adversary action λ^s such that, for all symbolic runs R^s , if $\Sigma_{\text{Adv}}^s(\text{strip}(R^s), r_{\text{Adv}}, \Lambda^s) = \lambda^s$, one of the following holds:

- 1) $\lambda^s = (A_i, j)$, $\Lambda_i^s = \alpha_1 \cdots \alpha_m$, and $\alpha_j = A_i : \dots$;
- 2) $\lambda^s = \alpha$, $\alpha \neq A : \dots$ for any A , $\alpha \notin \mathbb{N}$, and $R^s \xrightarrow{\alpha}$;
- 3) $\lambda^s = B : \dots$, $\lambda^s \neq B : a$ for any a , $B \notin \text{Hon}$ and $R^s \xrightarrow{\lambda^s}$;
- 4) $\lambda^s = \delta$, where $\forall i \in 1..k : (\Lambda_i^s = \emptyset \text{ or } \exists \delta_i \in \Lambda_i^s : \delta_i \geq \delta)$.
- 5) $\lambda^s = B : a$, where (i) $B \notin \text{Hon}$; (ii) $\Gamma_{\text{strip}(R^s)}$ contains $\{B : a \# \perp\}$; (iii) for some prefix R_*^s of $\text{strip}(R^s)$, $\Sigma_{\text{Adv}}^s(R_*^s, r_{\text{Adv}}, \epsilon) = B : \{G\}C, \Delta$ where $\{B : a \# N\}$ in Δ , for some $N \neq \perp$ and $\{G\}C$;

If $\Sigma_{\text{Adv}}^s(\text{strip}(R^s), r_{\text{Adv}}, \Lambda^s) = B : \{G\}C, \Delta$ for $B \notin \text{Hon}$, we ask $\Sigma_{\text{Adv}}^s(\text{strip}(R^s), r_{\text{Adv}}, \epsilon) = B : \{G\}C, \Delta$. \diamond

Conditions (1)–(4) are straightforward; in (5) the adversary makes a dishonest participant B reveal a secret a . In such case, we require that the (stripped) run contains the corresponding committed secret $\{B : a \# \perp\}$, and that, at some previous point R_*^s in the (stripped) run, Σ_{Adv}^s chose a non- \perp length N for a . This requirement is achieved by considering $\Sigma_{\text{Adv}}^s(R_*^s, r_{\text{Adv}}, \epsilon)$, where the use of ϵ is due to the ignorance of the parameter Λ^s generated at the point R_*^s . We consequently require that the N chosen by Σ_{Adv}^s does not depend on Λ^s . This restriction does *not* limit the power of Adv , who can first perform a sequence of actions λ^s depending on Λ^s , appending them to R^s , and then choose N depending on such actions λ^s .

We now characterise the runs obtained under an adversary with strategy Σ_{Adv}^s taking as input the outputs of the participant strategies. This notion, called *conformance*, also involves a randomness source, fed as input to the strategies.

Definition 12 (Symbolic conformance). Let Σ^s be a set of symbolic strategies, comprising those of honest participants A_1, \dots, A_k and of the adversary, and let r be a randomness source. We say that a symbolic run R^s *conforms to* (Σ^s, r) when one of the following conditions hold:

- 1) $R^s = (\Gamma, 0)$, with Γ initial;
- 2) $\hat{R}^s \xrightarrow{\alpha} R^s$, where \hat{R}^s conforms to (Σ^s, r) , and, given $\Lambda_i^s = \Sigma_{A_i}^s(\text{strip}(\hat{R}^s), r_{A_i})$ for all $i \in 1..k$ and $\Lambda_*^s =$

$\text{strip}(\Lambda_1^s \cdots \Lambda_k^s)$, if $\lambda^s = \Sigma_{\text{Adv}}^s(\text{strip}(\hat{R}^s), r_{\text{Adv}}, \Lambda_*^s)$ then $\lambda^s = \alpha$, or $\lambda^s = (A_i, j)$ for $\alpha = \alpha_j$ and $\Lambda_i^s = \alpha_1 \cdots \alpha_m$. If Σ^s does not comprise the adversary strategy, we say that R^s conforms to (Σ^s, r) when there exists some strategy Σ_{Adv}^s such that R^s conforms to $(\Sigma^s \Sigma_{\text{Adv}}^s, r)$. \diamond

IV. COMPUTATIONAL MODEL

In this section we introduce our computational model, which will be the target of the BitML compiler. We rely on [26] for a formal model of Bitcoin transactions and its blockchain. Here we just briefly recap these notions, providing some intuition.

In Bitcoin, *transactions* describe transfers of currency (\mathbb{B}). The log of all transactions is maintained on a public, immutable and decentralised data structure called *blockchain*. We represent transactions as records, with fields *in*, *wit*, *out* and *absLock*. Consider e.g. the transactions T_1 and T_2 below:

T_1	T_2
in: \dots	in: $0 \mapsto (T_1, 0)$
wit: \dots	wit: $0 \mapsto \text{sig}_{K(A)}$
out: $0 \mapsto (\lambda c.\text{versig}_{K(A)}(c), v_0 \mathbb{B})$	out: $0 \mapsto (\lambda x.H(x) = k, v_0 \mathbb{B})$
$1 \mapsto (\lambda c.\text{versig}_{K(B)}(c), v_1 \mathbb{B})$	absLock: t

The transaction T_1 has two *outputs*: the $v_0 \mathbb{B}$ in $\text{out}(0)$ can be redeemed by any transaction T whose *in* field refers to $(T_1, 0)$, and whose *wit* field satisfies the predicate in in $\text{out}(0)$ (similarly for the other output). This is the case e.g. of the transaction T_2 above. Its *witness* $\text{sig}_{K(A)}$ is the signature of A on T_2 (excluding the *wit* field itself), as required by $T_1.\text{out}(0)$.

If T_1 is on the blockchain and its $\text{out}(0)$ is unspent, A can update the blockchain by appending T_2 . This moves $v_0 \mathbb{B}^1$ from T_1 to T_2 . The transaction T_2 has only one output, which can be redeemed by any transaction providing a witness having hash k . The time t in $T_2.\text{absLock}$ represents the earliest moment when T_2 can be put on the blockchain. A subsequent transaction can redeem $(v_1 + v_0) \mathbb{B}$ in a single shot. This requires two inputs, $(T_1, 1)$ and $(T_2, 0)$, and two witnesses. The witness associated to the first input is a signature of B ; the other is a preimage of k .

Definition 13 (Blockchain). A blockchain \mathbf{B} is a sequence $(T_0, t_0) \cdots (T_n, t_n)$, where T_0 is a *coinbase* transaction (i.e., $T_0.\text{in} = \perp$), and $t_i \leq t_j$ for all $0 \leq i < j \leq n$. \diamond

Hereafter, we assume that blockchains are *consistent*, i.e. obtained by appending transactions respecting the Bitcoin protocol (as formalised in [26]). Borrowing from [26], we write $\mathbf{B} \triangleright (T, t)$ when (T, t) consistently updates \mathbf{B} .

We now introduce the computational counterparts of symbolic runs and strategies. Our model uses PPTIME algorithms w.r.t. a security parameter $\eta \in \mathbb{N}$. We follow the random oracle model [27] for cryptographic hashes. Namely, a special entity $O \notin \text{Part}$ provides to every A (including Adv) the access to a hash functionality. To compute a hash, A sends a message $A \rightarrow O : m$, and waits for the reply $O \rightarrow A : H(m)$, where $H(m)$ comprises η random bits. We extend the randomness

¹In the actual Bitcoin, the value the outputs of T_2 must be strictly smaller than v_0 , and the difference is paid to the Bitcoin network. For simplicity, in this paper (as in [26]) we neglect these *transaction fees*.

source also to O , assuming $r_{\mathsf{O}} \in \{0, 1\}^\omega$ to be defined. For simplicity, we assume that all communications (except those involving O) are (reliable) broadcasts.

Definition 14 (Computational labels). A computational label λ^c is a bitstring, encoding one of the following actions, where $A \in \text{Part} \cup \{\text{Adv}\}$, and m is a bitstring:

$A \rightarrow * : m$	A broadcasts message m	
$A \rightarrow \mathsf{O} : m$	A queries the oracle with message m	
$\mathsf{O} \rightarrow A : m$	O answers to A with message m	
T	put on the ledger a Bitcoin transaction	
δ	perform a delay	◇

We associate each A with two key pairs: K_A for signing messages, and \hat{K}_A for redeeming deposits. These key pairs are generated through r_A , if $A \in \text{Hon}$. We write $K_A(r_A)$, $\hat{K}_A(r_A)$ for the key pairs generated using the first 2η bits of r_A . For the participants *not* in Hon , denoted with B_1, \dots, B_k , we write $K_{B_i}(r_{\text{Adv}})$, $\hat{K}_{B_i}(r_{\text{Adv}})$ for the key pairs generated from r_{Adv} using the i -th 2η bits. Given any key pair K , we denote with K^p its public part, and with K^s its private part.

Definition 15 (Computational runs). Let r be a randomness source. A computational run R^c is a finite sequence of computational labels, beginning with a prefix R_0^c such as:

$$\begin{aligned} & \mathsf{T}_0 \cdots A_i \rightarrow * : (K_{A_i}^p(r_{A_i}), \hat{K}_{A_i}^p(r_{A_i})) \cdots \quad (\forall A_i \in \text{Hon}) \\ & \cdots B_j \rightarrow * : (K_{B_j}^p(r_{\text{Adv}}), \hat{K}_{B_j}^p(r_{\text{Adv}})) \cdots \quad (\forall B_j \notin \text{Hon}) \end{aligned}$$

where T_0 is a coinbase transaction, and for each participant P (honest or dishonest), there exists an output of T_0 redeemable with the private key \hat{K}_P^s . We say that the run R_0^c is *initial*. We denote with δ_{R^c} the sum of all the delays in R^c . ◇

The labels $P_i \rightarrow * : (K_{P_i}^p, \hat{K}_{P_i}^p)$ represent a broadcast of P_i 's public keys (of both kinds) to all participants. In this way, each participant starts with some funds (possibly 0), and knows the public keys of the others.

Analogously to Definition 8, we define the stripping of a computational run. We denote with $\text{strip}_A(R^c)$ the run obtained by removing from R^c all the messages *not* visible by A , i.e. the messages between O and some other $B \neq A$.

Definition 16 (Stripping of computational runs). For each computational run R^c and participant $A \in \text{Part} \cup \{\text{Adv}\}$, we define $\text{strip}_A(R^c)$ as the computational run obtained from R^c by replacing each label λ^c with $\text{strip}_A(\lambda^c)$, defined as follows:

$$\text{strip}_A(\lambda^c) = \begin{cases} \epsilon & \text{if } \lambda^c = B \rightarrow \mathsf{O} : m, \text{ with } B \neq A \\ \epsilon & \text{if } \lambda^c = \mathsf{O} \rightarrow B : m, \text{ with } B \neq A \\ \lambda^c & \text{otherwise} \end{cases} \quad \diamond$$

By extracting the transactions from a computational run, we obtain a blockchain, whose consistency will be ensured by Definition 20. This will be used to define strategies.

Definition 17 (Blockchain of a computational run). For each R^c , we define the blockchain \mathbf{B}_{R^c} inductively as follows:

$$\mathbf{B}_{\mathsf{T}_0} = (\mathsf{T}_0, 0) \quad \mathbf{B}_{R^c \lambda^c} = \begin{cases} \mathbf{B}_{R^c}(\mathsf{T}, \delta_{R^c}) & \text{if } \lambda^c = \mathsf{T} \\ \mathbf{B}_{R^c} & \text{otherwise} \end{cases} \quad \diamond$$

Definition 18 (Computational participant strategies). The computational strategy of a participant $A \in \text{Hon}$ is a PPTIME algorithm $\Sigma_A^c(R_*^c, r_A)$, taking as input a (A -stripped) computational run R_*^c and a random sequence r_A . The output is a *finite* set Λ^c of computational labels, such that if $\lambda^c \in \Lambda^c$, then one of the following items holds:

- 1) $\lambda^c = A \rightarrow * : m$ or $\lambda^c = A \rightarrow \mathsf{O} : m$, for some m ;
- 2) $\lambda^c = \mathsf{T}$, where $\mathbf{B}_{R_*^c} \triangleright (\mathsf{T}, \delta_{R_*^c})$, where in R_*^c we can find (in this order): (i) all the inputs of T ; (ii) a message $B \rightarrow * : \mathsf{T}$, for some B ; (iii) for each witness w in T , a message $B \rightarrow * : w$, for some B .
- 3) $\lambda^c = \delta$.

We further require that participant strategies are persistent: i.e., if $\Lambda^c = \Sigma_A^c(R_*^c, r_A)$, then for all λ^c such that $\mathbf{B}_{R^c \lambda^c}$ is consistent, $\Sigma_A^c(R_*^c \lambda^c, r_A)$ includes the computational labels:

$$\begin{aligned} & \{\mathsf{T} \in \Lambda^c \mid \mathbf{B}_{R_*^c \lambda^c} \triangleright (\mathsf{T}, \delta_{R_*^c \lambda^c})\} \cup \\ & \{A \rightarrow * : m \in \Lambda^c \mid (A \rightarrow * : m) \neq \lambda^c\} \cup \\ & \{A \rightarrow \mathsf{O} : m \in \Lambda^c \mid (A \rightarrow \mathsf{O} : m) \neq \lambda^c\} \quad \diamond \end{aligned}$$

Persistency ensures that, if A at a certain point wants to perform some λ^c (sending m or putting T on the blockchain), she cannot change her mind in the future, until λ^c is performed. Condition (2) requires that, before being able to append T , one has to make both T and its witnesses public. This models the fact that, before T actually appears on the blockchain, it has to be broadcast on the Bitcoin network, potentially enabling an adversary to know T beforehand.

Definition 19 (Computational adversary strategies). A computational adversary strategy is a PPTIME algorithm $\Sigma_{\text{Adv}}^c(R_*^c, r_{\text{Adv}}, \Lambda^c)$, taking as input a (Adv -stripped) computational run R_*^c , a random sequence r_{Adv} , and a list Λ^c of sequences of moves of each $A_i \in \text{Hon} = \{A_1, \dots, A_k\}$. The output is a single computational label λ^c such that if $\Sigma_{\text{Adv}}^c(R_*^c, r_{\text{Adv}}, \Lambda^c) = \lambda^c$, one of the following holds:

- 1) $\lambda^c = A \rightarrow * : m$ or $\lambda^c = \text{Adv} \rightarrow \mathsf{O} : m$, for some m, A ;
- 2) $\lambda^c = \mathsf{T}$, where $\mathbf{B}_{R_*^c} \triangleright (\mathsf{T}, \delta_{R_*^c})$, where in R_*^c we can find (in this order): (i) all the inputs of T ; (ii) a message $B \rightarrow * : \mathsf{T}$, for some B ; (iii) for each witness w in T , a message $B \rightarrow * : w$, for some B .
- 3) $\lambda^c = \delta$, where $\forall i \leq k : (\Lambda_i^c = \emptyset \text{ or } \exists \delta_i \in \Lambda_i^c : \delta_i \geq \delta)$. ◇

Note that (1) allows Adv to impersonate any A , to use either A 's messages in Λ^c , or self-produced ones. The other cases allow Adv to (consistently) extend the blockchain, and to delay (if all the honest participants agree).

Given a set of strategies Σ^c for all the honest participants and for the adversary, and a randomness source r , we now define which runs R^c can result from making everyone interact. Such runs are said to *conform* to (Σ^c, r) .

Definition 20 (Computational conformance). Let Σ^c be a set of computational strategies, including those of the honest participants A_1, \dots, A_k and of the adversary, and let r be a randomness source. We say that a computational run R^c *pre-conforms to* (Σ^c, r) if one of the following conditions holds:

- 1) R^c is initial;
- 2) $R^c = \dot{R}^c \mathcal{X}^c$, where \dot{R}^c pre-conforms to (Σ^c, r) , and $\mathcal{X}^c = \Sigma_{\text{Adv}}^c(\text{strip}_{\text{Adv}}(\dot{R}^c), r_{\text{Adv}}, \Lambda^c)$, where $\Lambda^c = \Lambda_1^c \dots \Lambda_k^c$, $\Lambda_i^c = \Sigma_{A_i}^c(\text{strip}_{A_i}(\dot{R}^c), r_{A_i})$, \mathcal{O} not occurring in Λ^c , \mathcal{X}^c .
- 3) $R^c = \dot{R}^c(\text{Adv} \rightarrow \mathcal{O} : m)(\mathcal{O} \rightarrow \text{Adv} : h_m)$, where \dot{R}^c pre-conforms to (Σ^c, r) , and $\Sigma_{\text{Adv}}^c(\text{strip}_{\text{Adv}}(\dot{R}^c), r_{\text{Adv}}, \Lambda^c) = \text{Adv} \rightarrow \mathcal{O} : m$, where $\Lambda^c = \Lambda_1^c \dots \Lambda_k^c$, $\Lambda_i^c = \Sigma_{A_i}^c(\text{strip}_{A_i}(\dot{R}^c), r_{A_i})$, and \mathcal{O} does not occur in Λ^c .
- 4) $R^c = \dot{R}^c(A_j \rightarrow \mathcal{O} : m)(\mathcal{O} \rightarrow A_j : h_m)$, where \dot{R}^c pre-conforms to (Σ^c, r) , $\Lambda_i^c = \Sigma_{A_i}^c(\text{strip}_{A_i}(\dot{R}^c), r_{A_i})$, and $(A_j \rightarrow \mathcal{O} : m)$ is the first occurrence of a query to the oracle in $\Lambda_1^c \dots \Lambda_k^c$.

Further, in both Items 3 and 4, given n the number of distinct queries to \mathcal{O} in R^c , we require that if m was already requested, then h_m is its reply in \dot{R}^c ; otherwise, h_m is the portion of $r_{\mathcal{O}}$ of length η starting from $n\eta$.

We say that R^c *conforms to* (Σ^c, r) if R^c is a prefix of a run pre-conforming to (Σ^c, r) . \diamond

Above, in Items 3 and 4 we handle the queries to \mathcal{O} , modelling the hash functionality as in the random oracle model. We let queries have higher priority than other actions.

V. COMPILING CONTRACTS TO BITCOIN TRANSACTIONS

We compile a contract advertisement $\{G\}C$ into a sequence of standard Bitcoin transactions². Our compiler relies on the following parameters, which depend on G and C :

- **PartG** is the set of all participants occurring in G ;
- **part** maps deposit names in G to the corresponding participants (e.g., $\text{part}(x) = A$ if $A : ?v @ x$ in G);
- **txout** maps deposit names in G to the corresponding Bitcoin transaction output (T, o) ;
- **val** maps deposit names in G to the value contained in the deposit (e.g., $\text{val}(x) = v$ if $A : ?v @ x$ in G);
- **sechash** maps secret names in G to the corresponding committed hashes.

Further, we assume that participants generate the following key pairs, and exchange their public parts:

- $\mathbf{K}(A)$, for each $A \in \text{PartG}$;
- $\mathbf{K}(D, A)$, for each subterm D of C , and each $A \in \text{PartG}$.

For a set of participants $\mathcal{P} = \{A_1, \dots, A_n\}$, we denote with $\mathbf{K}(D, \mathcal{P})$ the set of key pairs $\{\mathbf{K}(D, A_1), \dots, \mathbf{K}(D, A_n)\}$.

Definition 21 (BitML compiler). The function $\mathbb{B}_{\text{adv}}(\cdot)$ from contract advertisements to sequences of Bitcoin transactions is defined by the rules in Figure 6. \diamond

²Standard Bitcoin transactions are obtained from those obtained by our compiler using the tool <https://github.com/bitcoin-transaction-model>. This is crucial, since the Bitcoin network currently discards non-standard transactions.

The function \mathbb{B}_{adv} produces all the transactions for $\{G\}C$. In particular, T_{init} will be the first to be put on the blockchain, representing the stipulation of C . Function \mathbb{B}_C assigns to a contract $C = \sum_i D_i$ a transaction T_C , which can be redeemed only by using the keys of the subterms D_i of C . The functions \mathbb{B}_D and \mathbb{B}_{out} handle the possible actions of each D . The action `split` is handled by \mathbb{B}_{par} . We provide more intuition through an example.

Example 7. We compile the timed commitment $\{G\}C$ of Example 3, where $G = A : !v @ x \mid A : \text{secret } a \mid B : !0 @ y$, and $C = D_1 + D_2$, with $D_1 = \text{reveal } a. \text{withdraw } A$ and $D_2 = \text{after } t : \text{withdraw } B$. Assume that: $A \in \text{Hon}$, $\text{txout}(x) = (T_A, 0)$ for some T_A whose output 0 has value v and is redeemable by A , $\text{sechash}(a) = h_a$, $\text{val}(x) = v$. Similarly, for B : $\text{txout}(y) = (T_B, 0)$, and $\text{val}(y) = 0$.

The compiler produces $\mathbb{B}_{\text{adv}}(\{G\}C) = T_{\text{init}} T_1 T_2$, where:

$$\begin{aligned} T_1 &= \mathbb{B}_D(D_1, D_1, T_{\text{init}}, 0, v, \{A, B\}, 0) \\ &= \mathbb{B}_C(\text{withdraw } A, D_1, T_{\text{init}}, 0, v, \emptyset, \{A, B\}, 0) \\ &= T' \mathbb{B}_D(\text{withdraw } A, \text{withdraw } A, T', 0, v, \{A, B\}, 0) = T' T'_A \\ T_2 &= \mathbb{B}_D(D_2, D_2, T_{\text{init}}, 0, v, \{A, B\}, 0) \\ &= \mathbb{B}_D(\text{withdraw } B, D_2, T_{\text{init}}, 0, v, \{A, B\}, t) = T'_B \end{aligned}$$

The obtained transactions are in Figure 7, where:

$$\begin{aligned} e_1 &= \mathbb{B}_{\text{out}}(D_1) \\ &= \text{versig}_{\mathbf{K}(D_1, \{A, B\})}(\text{SA}_{\text{SB}}) \wedge H(b) = h_a \wedge |b| \geq \eta \\ e_2 &= \mathbb{B}_{\text{out}}(D_2) = \text{versig}_{\mathbf{K}(D_2, \{A, B\})}(\text{SA}_{\text{SB}}) \\ e' &= \mathbb{B}_{\text{out}}(\text{withdraw } A) = \text{versig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}(\text{SA}_{\text{SB}}) \end{aligned}$$

For the sake of readability we do not use distinct variables for different signatures of the same participant. The participants start by generating the transactions off chain, and exchanging the signatures shown in Figure 7. Doing this, A commits to her secret, whose hash h_a occurs in the output script T_{init} . After that, both A and B sign T_{init} , and put it on the ledger, stipulating the contract. The transaction T_{init} can be redeemed either by T' or by T'_B . In the first case, T' has to add to her witness a value a such that $H(a) = h_a$ and $|a| \geq \eta$. After that, A can redeem her deposit (now in T') by putting T'_A on the blockchain. In the second case, T_{init} can be redeemed by T'_B : however, this transaction can be put on the blockchain only after t , because of the timelock in T'_B . \diamond

VI. SECURITY OF THE COMPILER

In this section we prove the security of our compiler. We start with two auxiliary definitions: coherence between symbolic and computational runs (Section VI-A), and a mapping from symbolic to computational strategies (Section VI-B). Our computational soundness result (Theorem 1) is in Section VI-C.

A. Relating symbolic and computational runs

Two runs R^s and R^c are coherent when each symbolic step in R^s is matched by its computational implementation in R^c . We provide some intuition by discussing a few cases. Advertising a contract (rule [C-ADVERTISE]) corresponds to a broadcast of a suitable bitstring, where symbolic deposit names

$$\begin{array}{l}
G = (\|_{i \in I} A_i : ? v_i \otimes x_i) \mid (\|_{i \in J} B_i : ! v'_i \otimes y_i) \mid (\|_{i \in K} C_i : \text{secret } a_i) \\
C = \sum_{i=1}^m D_i \quad v = \sum_{i \in J} v'_i \\
e_i = \mathbb{B}_{\text{out}}(D_i) \quad (\forall i \in 1..m) \quad \mathbf{x} = \biguplus_{i=1}^m \text{fv}(e_i) \\
\mathcal{T}_i = \mathbb{B}_{\text{D}}(D_i, D_i, \mathbb{T}_{\text{init}}, 0, v, \text{PartG}, 0) \quad (\forall i \in 1..m) \\
\hline
\mathbb{B}_{\text{adv}}(\{G\}C) = \mathbb{T}_{\text{init}} \mathcal{T}_1 \cdots \mathcal{T}_m
\end{array}$$

\mathbb{T}_{init}
in: $i \mapsto \text{txout}(y_i) \quad (\forall i \in J)$
wit: \perp
out: $(\lambda \mathbf{x}. \bigvee_{i=1}^m e_i, v)$

$$\begin{array}{l}
C = \sum_{i=1}^m D_i \quad I = \{z_1, \dots, z_k\} \\
e_i = \mathbb{B}_{\text{out}}(D_i) \quad (\forall i \in 1..m) \\
\mathbf{x} = \biguplus_{i=1}^m \text{fv}(e_i) \\
\mathcal{T}_i = \mathbb{B}_{\text{D}}(D_i, D_i, \mathbb{T}_C, 0, v, \text{PartG}, 0) \quad (\forall i \in 1..m) \\
\hline
\mathbb{B}_{\text{C}}(C, D_p, \mathbb{T}, o, v, I, \mathcal{P}, t) = \mathbb{T}_C \mathcal{T}_1 \cdots \mathcal{T}_m
\end{array}$$

\mathbb{T}_C
in: $0 \mapsto (\mathbb{T}, o), i \mapsto \text{txout}(z_i) \quad (\forall i \in 1..k)$
wit: $0 \mapsto \text{sig}_{\mathbf{K}(D_p, \mathcal{P})}, i \mapsto \text{sig}_{\mathbf{K}(\text{part}(z_i))} \quad (\forall i \in 1..k)$
out: $(\lambda \mathbf{x}. \bigvee_{i=1}^m e_i, v)$
absLock: t

$$\frac{D \neq A_1 : \dots : A_n : \text{after } t_1 : \dots : \text{after } t_k : \text{put } z \ \& \ \text{reveal } a \ \text{if } p. \ C \quad \varsigma \ \text{fresh}}{\mathbb{B}_{\text{out}}(D) = \text{versig}_{\mathbf{K}(D, \text{PartG})}(\varsigma)}$$

$$\frac{D \equiv A_1 : \dots : A_n : \text{after } t_1 : \dots : \text{after } t_k : \text{put } z \ \& \ \text{reveal } a \ \text{if } p. \ C \quad \mathbf{a} = a_1 \cdots a_m \quad \varsigma, b_1 \cdots b_m \ \text{fresh}}{\mathbb{B}_{\text{out}}(D) = \text{versig}_{\mathbf{K}(D, \text{PartG})}(\varsigma) \wedge \mathbb{B}(p) \wedge \bigwedge_{i=1}^m \text{H}(b_i) = \text{sechash}(a_i) \wedge |b_i| \geq \eta}$$

$$\frac{D = \text{withdraw } A}{\mathbb{B}_{\text{D}}(D, D_p, \mathbb{T}, o, v, \mathcal{P}, t) = \{\text{in} : (\mathbb{T}, o), \text{wit} : \text{sig}_{\mathbf{K}(D_p, \mathcal{P})}, \text{out} : (\lambda \varsigma. \text{versig}_{\mathbf{K}(A)}(\varsigma), v), \text{absLock} : t\}}$$

$$\frac{D = \text{put } z \ \& \ \text{reveal } a \ \text{if } p. \ C \quad v' = v + \sum_{x \in z} \text{val}(x)}{\mathbb{B}_{\text{D}}(D, D_p, \mathbb{T}, o, v, \mathcal{P}, t) = \mathbb{B}_{\text{C}}(C, D_p, \mathbb{T}, o, v', z, \mathcal{P}, t)} \quad \frac{D = \text{split } v \rightarrow C \quad v = v_1 \cdots v_k \quad \sum_{i=1}^k v_i \leq v}{\mathbb{B}_{\text{D}}(D, D_p, \mathbb{T}, o, v, \mathcal{P}, t) = \mathbb{B}_{\text{par}}(C, D_p, \mathbb{T}, o, v, \mathcal{P}, t)}$$

$$\frac{D = A : D'}{\mathbb{B}_{\text{D}}(D, D_p, \mathbb{T}, o, v, \mathcal{P}, t) = \mathbb{B}_{\text{D}}(D', D_p, \mathbb{T}, o, v, \mathcal{P} \setminus \{A\}, t)} \quad \frac{D = \text{after } t' : D'}{\mathbb{B}_{\text{D}}(D, D_p, \mathbb{T}, o, v, \mathcal{P}, t) = \mathbb{B}_{\text{D}}(D', D_p, \mathbb{T}, o, v, \mathcal{P}, \max\{t, t'\})}$$

$$\begin{array}{l}
C = C_1 \cdots C_n \quad C_i = \sum_{j=1}^{k_i} D_{i,j} \quad (\forall i \in 1..n) \\
v = v_1 \cdots v_n \quad e_{i,j} = \mathbb{B}_{\text{out}}(D_{i,j}) \quad (\forall i \in 1..n, j \in 1..k_i) \\
\mathbf{x}_i = \biguplus_{j=1}^{k_i} \text{fv}(e_{i,j}) \quad (\forall i \in 1..n) \\
\mathcal{T}_{i,j} = \mathbb{B}_{\text{D}}(D_{i,j}, D_{i,j}, \mathbb{T}_C, i-1, v_i, \text{PartG}, 0) \quad (\forall i \in 1..n, j \in 1..k_i) \\
\hline
\mathbb{B}_{\text{par}}(C, D_p, \mathbb{T}, o, v, \mathcal{P}, t) = \mathbb{T}_C(\mathcal{T}_{i,j})_{i \in 1..n, j \in 1..k_i}
\end{array}$$

\mathbb{T}_C
in: (\mathbb{T}, o)
wit: $\text{sig}_{\mathbf{K}(D_p, \mathcal{P})}$
out: $i-1 \mapsto (\lambda \mathbf{x}_i. \bigvee_{j=1}^{k_i} e_{i,j}, v_i) \quad (\forall i \in 1..n)$
absLock: t

$$\begin{array}{l}
\mathbb{B}(\text{true}) = \text{true} \quad \mathbb{B}(p_1 \wedge p_2) = \mathbb{B}(p_1) \wedge \mathbb{B}(p_2) \quad \mathbb{B}(\neg p) = \neg \mathbb{B}(p) \quad \mathbb{B}(e_1 \circ e_2) = \mathbb{B}(e_1) \circ \mathbb{B}(e_2) \\
\mathbb{B}(N) = N \quad \mathbb{B}(|a|) = |a| - \eta \quad \mathbb{B}(e_1 \bullet e_2) = \mathbb{B}(e_1) \bullet \mathbb{B}(e_2)
\end{array}$$

Fig. 6: Compiling contracts to Bitcoin transactions.

\mathbb{T}_{init}	\mathbb{T}'	\mathbb{T}'_A	\mathbb{T}'_B
in: $0 \mapsto (\mathbb{T}_A, 0), 1 \mapsto (\mathbb{T}_B, 0)$	in: $(\mathbb{T}_{\text{init}}, 0)$	in: $(\mathbb{T}', 0)$	in: $(\mathbb{T}_{\text{init}}, 0)$
wit: \perp	wit: $\text{sig}_{\mathbf{K}(D_1, \{A, B\})}$	wit: $\text{sig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}$	wit: $\text{sig}_{\mathbf{K}(D_2, \{A, B\})}$
out: $(\lambda \varsigma_{A \& B}. e_1 \vee e_2, v)$	out: $(\lambda \varsigma_{A \& B}. e', v)$	out: $(\lambda \varsigma. \text{versig}_{\mathbf{K}(A)}(\varsigma), v)$	out: $(\lambda \varsigma. \text{versig}_{\mathbf{K}(B)}(\varsigma), v)$
			absLock: t

Fig. 7: Transactions obtained by compiling the timed commitment contract.

have been encoded as Bitcoin transaction outputs. Committing a secret (rule $[C\text{-AUTHCOMMIT}]$) corresponds to a sequence of computational actions: generating a nonce, obtaining its hash (querying \mathcal{O}), signing and broadcasting it. Contract stipulation (rule $[C\text{-INIT}]$) corresponds to putting on the blockchain the transaction T_{init} obtained by the compiler. Moves related to active contracts (e.g., $[C\text{-PUTREV}]$) correspond to putting on the blockchain the related transactions, as generated by the compiler. Authorizing a move (rule $[C\text{-AUTHCONTROL}]$) corresponds to broadcasting the needed signature. Revealing a secret (rule $[C\text{-AUTHREVE}]$) corresponds to broadcasting the associated nonce.

As anticipated in Section I, computational moves without a symbolic representation are ignored by the coherence relation. For instance, the computational Adv can broadcast arbitrary bitstrings, with no symbolic meaning. Adv can also put on the ledger any transaction T allowed by Bitcoin, including those not generated by our compiler, which so has no symbolic counterpart. In such case, coherence might still hold, depending on the inputs of T . If no input of T is symbolically represented, coherence holds, but we associate no corresponding move in R^s . Another case when coherence holds is that in which T redeems transactions which implement deposits $\langle A, v \rangle$: here, we represent T in R^s using $[DEP\text{-DESTROY}]$. Instead, coherence does not hold if T redeems some transaction which implements an active contract $\langle C, v \rangle$. Consequently, when coherence holds, the computational behavior of transactions implementing contracts closely matches their symbolic semantics.

Defining coherence is rather long, since it has to deal with all the possible cases of the symbolic semantics, and technically intricate, since computational runs have no intrinsic structure. For these reasons, we defer the actual definition to Appendix D (Definition 25). This definition exploits three maps, $txout$, $sechash$ and κ , the role of which is similar to that of the compiler parameters. Hereafter we write $R^s \sim_r R^c$ when the two runs are coherent w.r.t. the randomness source r .

B. Mapping symbolic to computational strategies

Given a symbolic strategy Σ^s of an honest participant A , we translate it to a computational strategy Σ^c which emulates its behaviour. We first describe how to translate the stipulation of a contract advertisement $\{G\}C$, i.e. by providing a computational counterpart to the rules in Figure 3. In the computational model, this is performed by following a *stipulation protocol* (Definition 22), which will be exploited to construct Σ^c .

Note that, while in the symbolic setting we model contract advertisements as terms $\{G\}C$, containing names x for the deposits, in the computational setting we encode such advertisements as bitstrings, using transaction outputs (T, o) instead of deposit names x . In the stipulation protocol, we assume as given this representation \mathcal{C} , which we call *computational contract advertisement*, and the parameters of Σ_A^c , i.e. the computational run R^c and the random sequence r_A . After decoding \mathcal{C} as a symbolic advertisement $\{G\}C$, A interacts with the other participants to obtain all the parameters required to compile $\{G\}C$. Then, A computes $\mathbb{B}_{adv}(\{G\}C)$, and finally she puts the generated T_{init} transaction on the ledger.

Definition 22 (Stipulation protocol). Let $A \in \text{Hon}$, let R_*^c be a (A -stripped) computational run, and let r_A be a random sequence. The stipulation protocol for a computational contract advertisement \mathcal{C} is the following.

- 1) A decodes \mathcal{C} , constructing a symbolic contract advertisement $\{G\}C$; in doing this, A chooses distinct symbolic names for all the transaction outputs in \mathcal{C} . The mapping $txout$ is defined according to the used correspondence between names and transaction outputs.
- 2) A infers from G the parameters $part$, $PartG$, and val .
- 3) A uses r_A to obtain the key pairs $K_A(r_A)$ and $\hat{K}_A(r_A)$. The key $K_A(r_A)$ is used by A to sign all the protocol messages. Further, A reads, from the initial prefix of the run R_*^c , the public keys $K_B^p(r_B)$ and $\hat{K}_B^p(r_B)$ of the all $B \in PartG \setminus \{A\}$. The keys $K_B^p(r_B)$ are then used by A to filter out the incoming messages with incorrect signatures.
- 4) A defines compiler keys: $\mathbf{K}(A)$ is $\hat{K}_A(r_A)$, while the public part of $\mathbf{K}(B)$ is $\hat{K}_B^p(r_B)$ for $B \neq A$. The keys $\mathbf{K}(D, A)$ are generated by A consuming $\sim \eta$ fresh bits from r_A , and their public parts are shared with the others. Dually, A defines the public parts of keys $\mathbf{K}(D, B)$ using the first broadcasts by the other participants. Let k be the sequence of public keys known by A after this step.
- 5) A generates from r_A a secret nonce of the desired length (to be defined by Σ_A^c) for every $A : secret a$ in G . Then, A computes the hashes $h = h_1 \cdots h_k$ of secret nonces (by querying \mathcal{O}), and broadcasts all these hashes as a single message $m(\mathcal{C}, h, k)$. Dually, A receives the hashes h' from the other participants. When doing this, A defines $sechash$ using the first (correctly signed) $m(\mathcal{C}, h', k)$ in R_*^c which has no duplicate hashes, and has no hashes already occurring (signed) in R_*^c .
- 6) A computes $\mathbb{B}_{adv}(\{G\}C)$, generating a list of transactions. The signatures by A occurring in the witnesses are computed and shared with others, while the signatures of other participants are received and verified. Note that the T_{init} transaction, (whose wit field is left to \perp by the compiler), is not signed, yet.
- 7) Only at this point, after having verified all the other signatures, A signs the first transaction, adding the signatures for her persistent deposits using $\hat{K}_A(r_A)$. Dually, she receives the signatures from the other participants for their own persistent deposits.
- 8) Finally, A broadcasts the signed T_{init} , and its witnesses, and puts it on the ledger.

In order to define the mapping from symbolic to computational strategies, we first sketch how to parse a (stripped) computational run R_*^c so to obtain a (stripped) symbolic run R_*^s . This requires to find, when possible, a symbolic counterpart to computational actions, mapping, e.g., computational advertisements to symbolic ones, signatures to authorizations, transactions to symbolic actions (involving deposits or contracts), and secrets to reveal actions. This correspondence closely follows Definition 25. Indeed, to perform the parsing,

we can exploit maps similar to $txout$, $sechash$ and κ to keep track of the correspondence at each step of R_*^s .

Using such maps, we can detect when a transaction T in R_*^c has some input with a symbolic counterpart (i.e. in $\text{ran } txout$). When that happens, we can map T into its corresponding action in R_*^s . According to Definitions 18 and 19, T has to be preceded by the broadcasts of its witnesses w , which, in turn, must be preceded by the broadcast of T . This allows to parse the signatures in w , so to generate authorizations in R_*^s .

A few cases must be handled with more care. For instance, Adv could broadcast a signature *before* the signed transaction. When this happens, we simply ignore the message; duplicate signatures are ignored as well. Further, a computational contract advertisement could involve only dishonest participants: we ignore that as well. Adv can consume her own deposits (among those tracked by $txout$), to create an arbitrary transaction without a symbolic counterpart. In this case, in the semantics we use the $[DEP-DESTROY]$ move in R_*^s , preceded by its authorizations, making those deposits disappear from the symbolic world. When the hash of a secret is committed, it can not be parsed precisely as a $[C-AUTHCOMMIT]$ move, since the latter involves the length of the secrets, which can not be inferred from the hash. However, this is not needed, since the *stripped* run R_*^s does not involve such lengths.

Definition 23 (From symbolic to computational strategies).

Let Σ_A^s be a symbolic strategy, with $A \in \text{Hon}$. We define $\aleph(\Sigma_A^s) = \Sigma_A^c$ below. Given the parameters R_*^c, r_A of Σ_A^c , we:

- 1) parse the (stripped) run R_*^c , so to obtain a corresponding symbolic (stripped) run R_*^s , as sketched above;
- 2) halve the random sequence r_A as $(\pi_1(r_A), \pi_2(r_A))$;
- 3) evaluate $\Lambda^s = \Sigma_A^s(R_*^s, \pi_1(r_A))$;
- 4) convert the symbolic actions Λ^s into computational actions Λ^c , and define $\Sigma_A^c(R_*^c, r_A) = \Lambda^c$. When Λ^s contains $A : \{G\}C, \Delta$ or $A : \{G\}C, x$, their conversion follows the stipulation protocol (Definition 22), using $\pi_2(r_A)$. There, at item 5, we choose the length of each secret by adding η to the corresponding value N in Δ . \diamond

C. Computational soundness

Our main result follows. We assume that honest participants have an associated symbolic strategy, and that their computational strategy is consequently obtained through the mapping \aleph . We also assume that Bitcoin uses secure cryptographic primitives, namely ideal hash functions (according to the random oracle model), and a digital signature scheme which is robust against *existential forgery* attacks. We also assume that the Bitcoin blockchain is append-only, i.e. that there are no forks, and *consistent*, i.e. miners only append valid transactions [26]. This is a rough abstraction of the security properties established in [28]–[30]

Our computational soundness result establishes that any computational run conforming to the (computational) strategies, with overwhelming probability has a corresponding symbolic run conforming to the (symbolic) strategies.

Theorem 1 (Computational soundness). *Let Σ^s be a set of symbolic strategies for all $A \in \text{Hon}$. Let Σ^c be a set of computational strategies such that $\Sigma_A^c = \aleph(\Sigma_A^s)$ for all $A \in \text{Hon}$, and with an arbitrary adversary strategy Σ_{Adv}^c . Fix $k \in \mathbb{N}$. Then, the following set has overwhelming probability:*

$$\left\{ r \mid \begin{array}{l} \forall R^c \text{ conforming to } (\Sigma^c, r) \text{ with } |R^c| \leq \eta^k : \\ \exists R^s \text{ conforming to } (\Sigma^s, \pi_1(r)) \text{ with } R^s \sim_r R^c \end{array} \right\}$$

The proof of Theorem 1 is in Appendix D.

VII. CONCLUSIONS

Our computational soundness result bridges the gap between the cryptography community, where smart contracts have been investigated first, and the programming languages community. It guarantees that, if some safety properties are violated at the computational level, then they are also violated at the symbolic level. So, reachability-based symbolic analyses can be soundly used to prove safety properties of smart contracts.

BitML allows for writing contracts that never deadlock, e.g. the timed commitment protocol and the lottery in Section II. Intuitively, this symbolic property is preserved by the compiler, which requires all the needed signatures in the stipulation phase. We conjecture that this liveness property can be proved, using the same notion of coherence of Theorem 1. Further, we could strengthen Theorem 1 to take probabilities into account. More specifically, when in the symbolic model an event has probability p , then in the computational model the same event should have a probability $p \pm \varepsilon$ where ε is negligible.

A first proposal of an high-level language relying on Bitcoin is [31]. This language allows to model the updates of a state machine as affine logic propositions. Users can “run” this machine by putting transactions on the blockchain, with the guarantee that only the legit updates can be performed. Using a computational model similar to ours (including coherence) it would be possible to establish computational soundness also for [31]. A downside of [31] is that liveness is guaranteed only by assuming cooperative participants, i.e., a dishonest participant can make the others unable to complete an execution. Note instead that in BitML, honest participants can always make a contract progress, regardless of the behaviour of the environment. Cooperation is incentivized by punishing misbehaviour with penalties, like e.g. in the lottery in Example 4.

The expressiveness of BitML could be improved by replacing the Bitcoin scripting language, using e.g. the approach of [32].

REFERENCES

- [1] N. Szabo, “Formalizing and securing relationships on public networks,” *First Monday*, vol. 2, no. 9, 1997, <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/548>.
- [2] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten, “SoK: Research perspectives and challenges for Bitcoin and cryptocurrencies,” in *IEEE S & P*, 2015, pp. 104–121.
- [3] N. Atzei, M. Bartoletti, and T. Cimoli, “A survey of attacks on Ethereum smart contracts (SoK),” in *Principles of Security and Trust (POST)*, ser. LNCS, vol. 10204. Springer, 2017, pp. 164–186. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-54455-6_8
- [4] “Understanding the DAO attack,” June 2016, <http://www.coindesk.com/understanding-dao-hack-journalists/>.
- [5] “Parity Wallet security alert,” July 2017, <https://paritytech.io/blog/security-alert.html>.

- [6] “A Postmortem on the Parity Multi-Sig library self-destruct,” November 2017, <https://goo.gl/Kw3gXi>.
- [7] M. Andrychowicz, S. Dziembowski, D. Malinowski, and L. Mazurek, “Fair two-party computations via Bitcoin deposits,” in *Financial Cryptography Workshops*, ser. LNCS, vol. 8438. Springer, 2014, pp. 105–121.
- [8] I. Bentov and R. Kumaresan, “How to use Bitcoin to design fair protocols,” in *CRYPTO*, ser. LNCS, vol. 8617. Springer, 2014, pp. 421–439.
- [9] A. Miller and I. Bentov, “Zero-collateral lotteries in Bitcoin and Ethereum,” in *EuroS&P Workshops*, 2017, pp. 4–13.
- [10] M. Bartoletti and R. Zunino, “Constant-deposit multiparty lotteries on Bitcoin,” in *Financial Cryptography Workshops*, 2017, also in IACR Cryptology ePrint Archive 955/2016.
- [11] R. Kumaresan, T. Moran, and I. Bentov, “How to use Bitcoin to play decentralized poker,” in *ACM CCS*, 2015, pp. 195–206.
- [12] W. Banasik, S. Dziembowski, and D. Malinowski, “Efficient zero-knowledge contingent payments in cryptocurrencies without scripts,” in *ESORICS*, ser. LNCS, vol. 9879. Springer, 2016, pp. 261–280.
- [13] S. Delgado-Segura, C. Pérez-Solà, G. Navarro-Arribas, and J. Herrera-Joancomartí, “A fair protocol for data trading based on bitcoin transactions,” *Future Generation Computer Systems*, 2017.
- [14] G. Maxwell, “The first successful zero-knowledge contingent payment,” <https://bitcoincore.org/en/2016/02/26/zero-knowledge-contingent-payments-announcement/>, 2016.
- [15] M. Andrychowicz, S. Dziembowski, D. Malinowski, and L. Mazurek, “Secure multiparty computations on Bitcoin,” in *IEEE S & P*, 2014, pp. 443–458.
- [16] R. Kumaresan and I. Bentov, “How to use Bitcoin to incentivize correct computations,” in *ACM CCS*, 2014, pp. 30–41.
- [17] N. Atzei, M. Bartoletti, T. Cimoli, S. Lande, and R. Zunino, “SoK: unraveling Bitcoin smart contracts,” 2018, IACR.
- [18] M. Abadi and P. Rogaway, “Reconciling two views of cryptography (the computational soundness of formal encryption),” *J. Cryptology*, vol. 20, no. 3, p. 395, 2007.
- [19] “Bitcoin developer guide - escrow and arbitration,” <https://bitcoin.org/en/developer-guide#escrow-and-arbitration>.
- [20] BitFury group, “Smart contracts on Bitcoin blockchain,” 2015, <http://bitfury.com/content/5-white-papers-research/contracts-1.1.1.pdf>.
- [21] D. Boneh and M. Naor, “Timed commitments,” in *CRYPTO*, ser. LNCS, vol. 1880. Springer, 2000, pp. 236–254.
- [22] P. F. Syverson, “Weakly secret bit commitment: Applications to lotteries and fair exchange,” in *IEEE CSFW*, 1998, pp. 2–13.
- [23] D. M. Goldschlag, S. G. Stubblebine, and P. F. Syverson, “Temporarily hidden bit commitment and lottery applications,” *Int. J. Inf. Sec.*, vol. 9, no. 1, pp. 33–50, 2010.
- [24] M. Andrychowicz, S. Dziembowski, D. Malinowski, and L. Mazurek, “Secure multiparty computations on Bitcoin,” *Commun. ACM*, vol. 59, no. 4, pp. 76–84, 2016.
- [25] A. Back and I. Bentov, “Note on fair coin toss via Bitcoin,” <http://www.cs.technion.ac.il/~iddo/cointossBitcoin.pdf>, 2013.
- [26] N. Atzei, M. Bartoletti, S. Lande, and R. Zunino, “A formal model of Bitcoin transactions,” in *Financial Cryptography and Data Security*, 2018.
- [27] M. Bellare and P. Rogaway, “Random oracles are practical: A paradigm for designing efficient protocols,” in *ACM Conference on Computer and Communications Security*. ACM, 1993, pp. 62–73.
- [28] C. Badertscher, U. Maurer, D. Tschudi, and V. Zikas, “Bitcoin as a transaction ledger: A composable treatment,” in *CRYPTO*, 2017, pp. 324–356.
- [29] J. A. Garay, A. Kiayias, and N. Leonardos, “The Bitcoin backbone protocol: Analysis and applications,” in *EUROCRYPT*, ser. LNCS, vol. 9057. Springer, 2015, pp. 281–310.
- [30] A. E. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou, “Hawk: The blockchain model of cryptography and privacy-preserving smart contracts,” in *IEEE Symp. on Security and Privacy*, 2016, pp. 839–858.
- [31] K. Crary and M. J. Sullivan, “Peer-to-peer affine commitment using Bitcoin,” in *ACM Conf. on Programming Language Design and Implementation*, 2015, pp. 479–488.
- [32] R. O’Connor, “Simplicity: A new language for blockchains,” in *PLAS*, 2017. [Online]. Available: <http://arxiv.org/abs/1711.03028>

A. Supplementary material for Section II

Definition 24 (Semantics of predicates).

$$\begin{aligned}
\llbracket true \rrbracket_{\Gamma} &= true \\
\llbracket p_1 \wedge p_2 \rrbracket_{\Gamma} &= \llbracket p_1 \rrbracket_{\Gamma} \text{ and } \llbracket p_2 \rrbracket_{\Gamma} \\
\llbracket \neg p \rrbracket_{\Gamma} &= \text{not } \llbracket p \rrbracket_{\Gamma} \\
\llbracket E_1 \circ E_2 \rrbracket_{\Gamma} &= \llbracket E_1 \rrbracket_{\Gamma} \circ \llbracket E_2 \rrbracket_{\Gamma} \quad \circ \in \{=, <\} \\
\llbracket N \rrbracket_{\Gamma} &= N \\
\llbracket |a| \rrbracket_{\Gamma} &= N \quad \text{if } \Gamma \text{ contains } A[a\#N] \\
\llbracket E_1 \bullet E_2 \rrbracket_{\Gamma} &= \llbracket E_1 \rrbracket_{\Gamma} \bullet \llbracket E_2 \rrbracket_{\Gamma} \quad \bullet \in \{+, -\}
\end{aligned}$$

The above semantics is well-defined, provided that there exists a unique $A[a\#N]$ in Γ for each $|a|$ in the predicate. Our semantics of contracts will ensure that is indeed the case.

We now describe the rules of deposits, starting from the simplest ones: those which deal with deposits, without directly affecting contracts.

The rule $[\text{DEP-AUTHJOIN}]$ allows A to authorize the merge of two deposits x, y into a single one, creating the needed authorization. The label of the form $A : \dots$ records that only participant A can perform this move. The rule $[\text{DEP-JOIN}]$ uses this authorization to create a single deposit z of A . The rules $[\text{DEP-AUTHDIVIDE}]$ and $[\text{DEP-DIVIDE}]$ act similarly, allowing a deposit of A to be divided in two parts. The rules $[\text{DEP-AUTHDONATE}]$ and $[\text{DEP-DONATE}]$ allow A to transfer one of her deposits to another participant. The pair of rules $[\text{DEP-AUTHDESTROY}]$ and $[\text{DEP-DESTROY}]$ allow a set of participants to destroy a set of deposits $x_1 \dots x_n$. To do that, first each participant A_i must provide the needed authorization $A_i[x, i \triangleright y]$ for their own deposit x_i . When all the authorizations have been collected, rule $[\text{DEP-DESTROY}]$ eliminates the deposits. The last pair of rules is needed to properly represent the fact that computational participants can create (and put on the ledger) transactions which do not have a counterpart in our 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 $[\text{DEP-DESTROY}]$.

B. Supplementary material for Section III

The following lemma states that stripping preserves the symbolic transitions, except for those which reveal secrets.

Lemma 2. *Let $\alpha \neq A : a$, for any A and a . (i) if $R^s \xrightarrow{\alpha} \hat{R}^s$, then $\text{strip}(R^s) \xrightarrow{\alpha} \text{strip}(\hat{R}^s)$; (ii) if $\text{strip}(R^s) \xrightarrow{\alpha} R_*^s$, then $R^s \xrightarrow{\alpha} \hat{R}^s$, for some \hat{R}^s such that $\text{strip}(\hat{R}^s) = \text{strip}(R_*^s)$.*

Lemma 3. *Let $R_*^s = \text{strip}(R^s)$ for some symbolic run R^s . Also, let $\text{Hon} = \{A_1, \dots, A_k\}$, and let $\Lambda^s = \Lambda_1^s \dots \Lambda_k^s$, where $\Lambda_i^s = \Sigma_{A_i}^s(R_*^s, r_{A_i})$ for $i \in 1..k$. If $\Sigma_{\text{Adv}}^s(R_*^s, r, \Lambda^s) = \lambda^s$, then:*

- 1) if $\lambda^s = (A_i, j)$, then $R^s \xrightarrow{\alpha_j}$, where $\Lambda_i^s = \alpha_1 \dots \alpha_m$;
- 2) otherwise, $R^s \xrightarrow{\lambda^s}$.

$A, B, \dots \in \text{Part}$	Participants	Γ, Δ, \dots	Configuration (untimed)	r	Random generator
x, y, \dots	Deposit names	(Γ, t)	Configuration (timed)	r_A	Random nonce (of A)
a, a', b, \dots	Secrets names	α, α', \dots	Labels (including delays)	m, m', \dots	Bitstrings
$t, t', \delta, \dots \in \mathbb{N}$	Delays	R^s, \hat{R}^s, \dots	Symbolic runs	T, T', \dots	Transactions
$v, v' \dots \in 10^{-8}\mathbb{N}$	Currency values	$R_x^s, \hat{R}_x^s, \dots$	Stripped symb. runs	R^c, \hat{R}^c, \dots	Computational runs
G, G', \dots	Preconditions	λ^s	Symbolic label	R_*, \hat{R}_*, \dots	Stripped comp. runs
C, C', \dots	Contracts	Λ^s	Set of symbolic labels	λ^c	Computational label
D, D', \dots	Guarded contracts	Σ_A^s	Symbolic strategy	Λ^c	Set of computational labels
$\{G\}C$	Advertisement	$\mathfrak{B}_{adv}(\cdot)$	Compiler	Σ_A^c	Computational strategy
				$\mathfrak{N}(\cdot)$	Strategy mapping

TABLE I: Summary of notation.

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

Fig. 8: Semantics of untimed configurations: rules for deposits.

C. Supplementary material for Section IV

In Lemma 4 we will show that there exists exactly one such maximal run.

Lemma 4. *Let Σ^c be a set of computational strategies for the honest participants and for the adversary, and let r be a randomness source. There exist exactly one maximal run R^c which conforms to (Σ^c, r) .*

D. Supplementary material for Section VI

Hereafter, we assume that at least one of the participants in $\{G\}C$ has a secret a in G . Since rule [C-ADVERTISE] requires that secret names are fresh, this ensures that the same contract can not be advertised twice. Expressiveness is not affected by this assumption, since a needs not be used in C . In all our examples, we implicitly assume this condition to be respected.

Definition 25 (Coherence). We inductively define the relation $coher(R^s, R^c, r, txout, sechash, \kappa)$, where (i) R^s is a symbolic run, (ii) R^c is a computational run, (iii) r is a randomness source, (iv) $txout$ is an injective function from

names x (occurring in R^s) to transaction outputs (T, o) (where T occurs in R^c), respecting values; (v) $sechash$ is a mapping from secret names a (occurring in R^s) to bitstrings; (vi) κ maps triples $(\{G\}C, D, A)$, where D is a subterm of C , to public keys.

Base case: $coher(R^s, R^c, r, txout, sechash, \kappa)$ holds if all the following conditions hold: (i) $R^s = (\Gamma_0, 0)$, with Γ_0 initial; (ii) $R^c = T_0 \cdots$ initial; (iii) all the public keys in R^c are generated from r , according to Definition 15; (iv) $txout$ maps exactly the x of $\langle A, v \rangle_x$ in Γ_0 to an output in T_0 of value $v\mathfrak{B}$, and spendable with $\hat{K}_A(r_A)$; (v) $\text{dom } sechash = \emptyset$; (vi) $\text{dom } \kappa = \emptyset$.

Inductive case: $coher(\hat{R}^s \xrightarrow{\alpha} (\Gamma, t), \hat{R}^c, r, txout, sechash, \kappa)$ holds if $coher(\hat{R}^s, \hat{R}^c, r, txout', sechash', \kappa')$ and one of the following cases applies.

- 1) $\alpha = advertise(\{G\}C)$, $\lambda^c = A \rightarrow * : \mathcal{C}$, where \mathcal{C} is obtained by encoding $\{G\}C$ as a bitstring, representing each x in it as the transaction output $txout'(x)$. Further, $txout' = txout$, $sechash' = sechash$, and $\kappa' = \kappa$.

- 2) $\alpha = A : \{G\}C, \Delta$, where: (i) for some B , \dot{R}^c contains $B \rightarrow * : C$, where C is obtained from $\{G\}C$ and $txout'$ as in Item 1. Note that \dot{R}^c might contain several such messages; below, we let C represent the first occurrence. (ii) for some B , $\lambda^c = B \rightarrow * : (C, h, k)$ (signed by A), where h is a sequence comprising a bitstring h_i with $|h_i| = \eta$ for each secret a_i in Δ , and k is a sequence of keys, as the one produced by the stipulation protocol. We require that λ^c is the first occurrence, in the run \dot{R}^c , of such a message after C . (iii) Let N_i be the length of a_i fixed in Δ . If $N_i \neq \perp$, we require that \dot{R}^c contains, for some B , a query to the oracle $B \rightarrow O : m_i$, and a subsequent reply $O \rightarrow B : h_i$ such that $|m_i| = \eta + N_i$. Otherwise, if $N_i = \perp$, we require that h_i does not occur as a reply from O to any query of length $\geq \eta$. (iv) No hash is reused: the h_i are pairwise distinct, and also distinct from $sechash'(b)$ for any $b \in \text{dom}(sechash')$. (v) $txout = txout'$. (vi) $sechash$ extends $sechash'$ so that for each secret a_i we have $sechash(a_i) = h_i$. (vii) If $A \in \text{Hon}$, we define κ by extending κ' according to k , so to record the public keys of all participants occurring in G for each subterm D of C . If κ' already defines such keys, or $A \notin \text{Hon}$, we let $\kappa = \kappa'$.
- 3) $\alpha = A : \{G\}C, x$, where: (i) $\lambda^c = B \rightarrow * : m$ for some B , where m is the signature of the transaction T_{init} of $\dot{B}_{adv}(\{G\}C)$ relatively to the input x with $\dot{K}_A(r_A)$. The parameters of the compiler are set as follows: **part**, **PartG** and **val** are inferred from G , we let **txout** = $txout'$, **sechash** = $sechash'$, and $\mathbf{K}(B) = \dot{K}_B^p(r_B)$, $\mathbf{K}(D, B) = \kappa'(\{G\}C, D, B)$ for each B , and D subterm of C . (ii) For some B , we have $B \rightarrow * : T_{init}$ occurring in \dot{R}^c . (iii) λ^c is the first occurrence in \dot{R}^c of a broadcast of m after T_{init} . (iv) $txout = txout'$, $sechash = sechash'$, and $\kappa = \kappa'$.
- 4) $\alpha = \text{init}(G, C)$, where: (i) α consumes from \dot{R}^s the advertisement $\{G\}C$ and its persistent deposits to produce $\langle C, v \rangle_z$. (ii) $\lambda^c = T_{init}$ where T_{init} is the first transaction in $\dot{B}_{adv}(\{G\}C)$. The needed compiler parameters are obtained as in Item 3. (iii) $sechash = sechash'$, $\kappa = \kappa'$, and $txout$ extends $txout'$, mapping z to T_{init} .
- 5) $\alpha = A : x, D$, where: (i) \dot{R}^s contains $\langle C', v \rangle_x$ with $C' = D + \sum_i D_i$, for some $D = A : D'$. (ii) In \dot{R}^s , we find that $\langle C', v \rangle_x$ has $\{G\}C$ as its ancestor advertisement. (iii) $\lambda^c = B \rightarrow * : m$, where m is a signature with key $\kappa'(\{G\}C, D, A)$ of the first transaction T in $\dot{B}_D(D, D, T', o, v, \text{PartG}, 0)$, where $(T', o) = txout'(x)$. The compiler parameters are obtained as in Item 3. (iv) $txout = txout'$, $sechash = sechash'$, and $\kappa = \kappa'$. (v) \dot{R}^c contains $B \rightarrow * : T$ for some B , and m is the first signature of T in \dot{R}^c after the first broadcast of T .
- 6) $\alpha = \text{put}(x, a, y)$, where: (i) $x = x_1 \cdots x_k$. (ii) In $\Gamma_{\dot{R}^s}$, the action α consumes $\langle D + C, v \rangle_y$ and the deposits $\langle A_i, v_i \rangle_{x_i}$ to produce $\langle C', v' \rangle_{y'}$, where $D = \cdots : \text{put} \cdots \text{reveal} \cdots C'$. Let t be maximum deadline in an after in front of D . (iii) In \dot{R}^s , we find that $\langle D + C, v \rangle_y$ has $\{G\}C''$ as its ancestor advertise-
- ment, for some G and C'' . (iv) $\lambda^c = T$ where T is the first transaction of $\dot{B}_C(C', D, T', o, v', x, \text{PartG}, t)$, where $(T', o) = txout'(y)$. The compiler parameters are obtained as in Item 3. (v) $txout$ extends $txout'$ so that y' is mapped to $(T, 0)$, $sechash = sechash'$, and $\kappa = \kappa'$.
- 7) $\alpha = A : a$, where: (i) $\lambda^c = B \rightarrow * : m$ from some B with $|m| \geq \eta$. (ii) $\dot{R}^c = \cdots (B \rightarrow O : m)(O \rightarrow B : sechash'(a)) \cdots$, for some B . (iii) $txout = txout'$, $sechash = sechash'$ and $\kappa = \kappa'$. (iv) In \dot{R}^s we find an $A : \{G\}C, \Delta$ action, with a in G , with a corresponding broadcast in \dot{R}^c of $m' = (C, h, k)$. (v) λ^c is the first broadcast of m in \dot{R}^c after the first broadcast of m' .
- 8) $\alpha = \text{split}(y)$, where: (i) In \dot{R}^s , the action α consumes $\langle D + C, v \rangle_y$ to obtain $\langle C_0, v_0 \rangle_{x_0} \mid \cdots \mid \langle C_k, v_k \rangle_{x_k}$ where $D = \cdots : \text{split } v \rightarrow C$ and $C = C_0 \cdots C_k$. Let t be the maximum deadline in an after in front of D . (ii) In \dot{R}^s , we find that $\langle D + C, v \rangle_y$ has $\{G\}C'$ as its ancestor advertisement. (iii) $\lambda^c = T$ where T is the first transaction of $\dot{B}_{par}(C, D, T', o, \text{PartG}, t)$ where $(T', o) = txout'(y)$. The compiler parameters are obtained as for Item 3. (iv) $txout$ extends $txout'$ mapping each x_i to (T, i) , $sechash = sechash'$, and $\kappa = \kappa'$.
- 9) $\alpha = \text{withdraw}(A, v, y)$, where: (i) In \dot{R}^s , the action α consumes $\langle D + C, v \rangle_y$ to obtain $\langle A, v \rangle_x$, where $D = \cdots : \text{withdraw } A$. (ii) In \dot{R}^s , we find that $\langle D + C, v \rangle_y$ has $\{G\}C'$ as its ancestor advertisement. (iii) $\lambda^c = T$ where T is the first transaction of $\dot{B}_D(D, D, T', o, v, \text{PartG}, 0)$ where $(T', o) = txout'(y)$. The compiler parameters are obtained as for Item 3. (iv) $txout$ extends $txout'$ mapping x to $(T, 0)$, $sechash = sechash'$, and $\kappa = \kappa'$.
- 10) $\alpha = A : x, x'$, where: (i) In \dot{R}^s we find $\langle A, v \rangle_x$ and $\langle A, v' \rangle_{x'}$. (ii) In \dot{R}^c we find $B \rightarrow * : T$ for some B, T , where T has as its two inputs $txout'(x)$ and $txout'(x')$, and a single output of value $v + v'$ redeemable with $\dot{K}_A(r_A)$. (iii) $\lambda^c = B \rightarrow * : m'$ for some B, m' , where m' is the signature of T with $\dot{K}_A(r_A)$. (iv) λ^c is the first broadcast of m' in \dot{R}^c after the first broadcast of T . (v) $txout = txout'$, $sechash = sechash'$, and $\kappa = \kappa'$.
- 11) $\alpha = \text{join}(x, y)$, where: (i) In \dot{R}^s the action α spends $\langle A, v \rangle_x$ and $\langle A, v' \rangle_{x'}$ to obtain $\langle A, v + v' \rangle_y$. (ii) $\lambda^c = T$ is a transaction having as inputs $txout'(x)$ and $txout'(x')$, and having one output of value $v + v'$ redeemable with $\dot{K}_A(r_A)$. (iii) $txout$ extends $txout'$ mapping y to $(T, 0)$, $sechash = sechash'$, and $\kappa = \kappa'$.
- 12) $\alpha = A : x, v, v'$. Similar to Item 10.
- 13) $\alpha = \text{divide}(x, v, v')$. Similar to Item 11.
- 14) $\alpha = A : x, B$. Similar to Item 10.
- 15) $\alpha = \text{donate}(x, B)$. Similar to Item 11.
- 16) $\alpha = A : y, j$, where: (i) $y = y_1 \cdots y_k$. (ii) In \dot{R}^s we find $\langle B_i, v_i \rangle_{y_i}$ for $i \in 1..k$, with $B_j = A$. (iii) In \dot{R}^c we find $B \rightarrow * : T$ for some B, T , where T has as its inputs $txout'(y_i)$ for $i \in 1..k$, and possibly others not in $\text{ran } txout'$. (iv) $\lambda^c = B \rightarrow * : m$ from some B, m where m is a signature of T with $\dot{K}_A(r_A)$, corresponding to

the j -th input. (v) λ^c is the first broadcast of m in \dot{R}^c after the first broadcast of \mathbf{T} . (vi) λ^c does not correspond to any of the other cases, i.e. there is no other symbolic action α for which $\dot{R}^s \alpha$ would be coherent with $\dot{R}^c \lambda^c$. (vii) $txout = txout'$, $sechash = sechash'$, and $\kappa = \kappa'$.

- 17) $\alpha = \text{destroy}(\mathbf{x})$, where: (i) $\mathbf{x} = x_1 \cdots x_k$. (ii) In \dot{R}^s , α consumes $\langle \mathbf{A}_i, v_i \rangle_{x_i}$ to obtain 0. (iii) $\lambda^c = \mathbf{T}$ from some \mathbf{T} having as inputs $txout'(x_1), \dots, txout'(x_k)$, and possibly others not in $\text{ran } txout'$. (iv) λ^c does not correspond to any of the other cases, i.e. there is no other symbolic action α for which $\dot{R}^s \alpha$ would be coherent with $\dot{R}^c \lambda^c$. (v) $txout = txout'$, $sechash = sechash'$, and $\kappa = \kappa'$.
- 18) $\alpha = \delta = \lambda^c$, and $txout = txout'$, $sechash = sechash'$, and $\kappa = \kappa'$.

Inductive case 2: $\text{coher}(R^s, R^c \lambda^c, r, txout, sechash, \kappa)$ holds if $\text{coher}(R^s, R^c, r, txout, sechash, \kappa)$ and one of the following cases applies:

- 1) $\lambda^c = \mathbf{T}$ where no input of \mathbf{T} belongs to $\text{ran } txout$.
- 2) $\lambda^c = \mathbf{A} \rightarrow \mathbf{O} : m$ or $\lambda^c = \mathbf{O} \rightarrow \mathbf{A} : m$, for some \mathbf{A}, m .
- 3) $\lambda^c = \mathbf{A} \rightarrow * : m$, where λ^c does not correspond to any symbolic move, according to the first inductive case.

We write $R^s \sim_r R^c$ iff $\text{coher}(R^s, R^c, r, txout, sechash, \kappa)$ for some $txout, sechash$, and κ . \diamond

Proof of Theorem 1 (sketch): Assume that R^c satisfies the hypotheses, but has no corresponding R^s which is coherent (to R^c) and conforming (to the symbolic strategies). Consider the longest prefix \dot{R}^c of R^c having a corresponding \dot{R}^s which is coherent (to \dot{R}^c) and conforming (to the computational strategies). We have that $R^c = \dot{R}^c \lambda^c \cdots$. We now show that either $\dot{R}^c \lambda^c$ has a corresponding symbolic run $\dot{R}^s \dot{R}^s$ which is coherent and conforming to the symbolic strategies (contradicting the maximality of \dot{R}^c), or the adversary succeeded in a signature forgery, or in a preimage attack (which can happen only with negligible probability). Note that, by obtain coherence, \dot{R}^s must be either empty, or contain a single symbolic action.

We proceed by cases on λ^c :

- 1) $\lambda^c = \mathbf{B} \rightarrow * : m$. Then, coherence must hold for some \dot{R}^s . Indeed, the definition of coherence maps m to an authorization (if it is the first broadcast of a signature), a revealed secret (if it is the first broadcast of a preimage), or in *all* other cases it simply ignores m . So, we can choose \dot{R}^s as the corresponding move, or to be empty, and obtain coherence. In these cases, we also obtain conformance. Indeed, in the last case (\dot{R}^s empty) the run $\dot{R}^s \dot{R}^s = \dot{R}^s$ is trivially conforming. For the authorization or reveal cases, we note that if the computational Adv was able to generate m , it is either forged (with negligible probability), or it originated from some honest \mathbf{A} . Since $\Sigma_{\mathbf{A}}^c = \aleph(\Sigma_{\mathbf{A}}^s)$, it follows that, at some time in the past, $\Sigma_{\mathbf{A}}^s$ enabled the authorization or reveal. By persistency, it is also enabled at the end of \dot{R}^s , hence Σ_{Adv}^s can choose such action, and achieve conformance.

2) If $\lambda^c = \mathbf{T}$, we consider the following subcases according to the inputs of \mathbf{T} :

- a) If no input of \mathbf{T} belongs to $\text{ran } txout$, then coherence and conformance are achieved taking \dot{R}^s to be empty.
- b) Otherwise, if at least one of the inputs of \mathbf{T} belongs to $\text{ran } txout$, then we look in \dot{R}^s for all the deposits and active contracts corresponding to such inputs. By definition of computational strategy, we must find in \dot{R}^c a (first) broadcast $\mathbf{B} \rightarrow * : \mathbf{T}$ followed by a (first) broadcast $\mathbf{B} \rightarrow * : m$ for all witnesses m of \mathbf{T} . By the coherence of \dot{R}^c , in \dot{R}^s the messages $\mathbf{B} \rightarrow * : m$ correspond to suitable authorization/reveal moves for each of the (counterparts of the) inputs of \mathbf{T} . We consider the following subcases:

- i) If all the inputs are deposits, then we let \dot{R}^s perform the symbolic move corresponding to \mathbf{T} (e.g., init or join). Note that if \mathbf{T} can not be represented symbolically, we can choose \dot{R}^s to perform a destroy. Such moves are feasible symbolically since we already have their authorizations. Such \dot{R}^s leads to a coherent run, which is also conforming, since even if no honest strategy wants to perform the move, Adv can perform it on its own, having all the authorizations.

- ii) Otherwise, some input \mathbf{T}' of \mathbf{T} must correspond to an active contract. This must be originated from an advertisement, which has to involve at least one honest participant \mathbf{A} , by definition of the symbolic semantics (rule [C-ADVERTISE]). However, by construction, our compiler makes \mathbf{T}' require the signatures of all the participants, hence including \mathbf{A} . Since such signature must occur as a witness in \mathbf{T} , the adversary Adv must have forged it (with negligible probability), or must have obtained it from \mathbf{A} . In the latter case, \dot{R}^s contains an authorization for the symbolic move corresponding to λ^c . By choosing \dot{R}^s accordingly, we obtain a coherent and conforming run.

- 3) Finally, if $\lambda^c = \delta$, we simply choose \dot{R}^s to perform δ . Coherence trivially holds. For conformance, we note that by definition of computational strategy, all the honest participants must output a Λ^c which either contains some $\delta' \geq \delta$, or is empty. By definition of \aleph , this must also be the case in Λ^s , resulting in conformance. \square