

BitML: A Calculus for Bitcoin Smart Contracts

Massimo Bartoletti
University of Cagliari
bart@unica.it

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

ABSTRACT

We introduce BitML, a domain-specific language for specifying contracts that regulate transfers of bitcoins among participants, without relying on trusted intermediaries. We define a symbolic and a computational model for reasoning about BitML security. In the symbolic model, participants act according to the semantics of BitML, while in the computational model they exchange bitstrings, and read/append transactions on the Bitcoin blockchain. A compiler is provided to translate contracts into standard Bitcoin transactions. Participants can execute a contract by appending these transactions on the Bitcoin blockchain, according to their strategies. We prove the correctness of our compiler, showing that computational attacks on compiled contracts are also observable in the symbolic model.

CCS CONCEPTS

• **Security and privacy** → **Distributed systems security; Formal security models; Security protocols;**

KEYWORDS

Bitcoin; smart contracts; process calculi

1 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 [58]. These agreements regulate cryptocurrency exchanges among participants: for instance, a lottery collects bets from players, determines the winner in a fair manner, and then transfers the pot to the winner.

Disintermediation is made possible by the *blockchain*, a public, append-only record of transactions, and by the *consensus protocol* followed by the nodes to update the blockchain [27]. 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: suitable economic incentives ensure that the nodes of the 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 platform on which they are based. In Ethereum, they are expressed as programs in a Turing-equivalent bytecode 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-horses with hidden vulnerabilities, to steal or tamper with the assets controlled by contracts. Indeed,

a series of vulnerabilities in Ethereum contracts [14] have caused losses in the order of hundreds of millions of USD [4, 6, 7].

Unlike Ethereum, Bitcoin does not provide a language for smart contracts: rather, in literature they are expressed as cryptographic protocols where participants send/receive/sign messages, verify signatures, and put/search transactions on the blockchain [15]. Lotteries [11, 20, 22, 47], gambling games [42], micro-payment channels [31, 48, 54], contingent payments [18, 32, 46], and more general fair multi-party computations [12, 41] witness the variety of smart contracts supported by Bitcoin.

Describing smart contracts at this level of abstraction is complex and error-prone. Indeed, establishing the correctness of a smart contract requires to prove the *computational security* of a cryptographic protocol, where — besides the usual primitives — participants can craft Bitcoin transactions and interact with the Bitcoin network. Further, these protocols often rely on advanced features of Bitcoin (e.g., transaction scripts, signature modifiers, segregated witnesses), whose actual behaviour relies on low-level implementation details. The task of proving the security of such kind of protocols 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 introduce BitML (after “Bitcoin Modelling Language”), a domain-specific language for Bitcoin smart contracts. BitML is a process calculus, with primitives to stipulate contracts and to exchange currency according to the contract terms. In this respect, BitML departs from the current practice of representing Bitcoin contracts as cryptographic protocols: rather, BitML pioneers the “contracts-as-programs” paradigm for Bitcoin, by completely abstracting from Bitcoin transactions and cryptographic details. Despite the high level of abstraction, BitML can express most of the Bitcoin smart contracts proposed so far [15], e.g. escrow services, timed commitments, lotteries, gambling games, etc. The operational semantics of BitML allows for reasoning about the behaviour of these contracts in a *symbolic* setting, where the underlying cryptography and Bitcoin machinery are abstracted away.

One of our main contributions is a compiler to translate BitML contracts into standard Bitcoin transactions. Participants can perform the contract actions by publishing the corresponding transactions on the blockchain. The crucial 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 Bitcoin. 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.

We establish the correctness of the BitML compiler through a *computational soundness* theorem [9]. More specifically, we prove that if honest participants use compiler-generated transactions,

then the actual Bitcoin executions resulting from their interaction with computational adversaries will have a symbolic counterpart as a BitML execution (with overwhelming probability). Basically, this implies that computational attacks to compiled contracts in Bitcoin are also observable in the symbolic semantics of BitML. A practical consequence of this result is that 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.

2 OVERVIEW OF THE APPROACH

In this section we overview our approach, introducing its main components, and discussing the relations among them. In particular, we provide a gentle introduction to BitML, illustrating its expressiveness through a series of examples. We then discuss the main results of the paper, illustrating their practical consequences.

BitML in a nutshell. In BitML, contracts allow participants to interact according to the following workflow. First, a participant broadcasts a *contract advertisement* $\{G\}C$. The component C is the contract, which specifies the rules to transfer bitcoins ($\text{\$}$) among participants. The component G is a set of *preconditions* to its execution: roughly, it requires participants to deposit some $\text{\$}$, either upfront or during the contract execution, and to commit to some secrets. 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. Only at this point, participants can transfer the deposited funds by acting as prescribed by C . Once C is stipulated, it starts its execution with a *balance*, initially set to the sum of the deposits in its advertisement. The execution of C will affect this balance, when participants deposit/withdraw funds to/from the contract.

As a first example, assume a buyer A who wants to buy an item from a seller B . The participants want to use a contract to ensure that B will get paid if and only if A gets her item. Assuming the cost of the item is $1\text{\$}$, the contract precondition $G = A : ! 1\text{\$} @ x$ requires A to provide a $1\text{\$}$ deposit (where x is the deposit name), which will be transferred to B only after A 's consent. The contract has two mutually exclusive clauses (separated by $+$):

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

The first clause allows B to withdraw $1\text{\$}$ from the contract, if A provides her authorization (denoted by $A : \dots$). Instead, the second clause allows A to get back her deposit upon B 's authorization (e.g., in case B acknowledges a problem with the shipment).

Note that the above contract gives little guarantees when the participants dishonestly deny their authorization: in particular, A can receive the item and then prevent B from withdrawing the payment, while B can freeze A 's deposit without shipping the item. A less naïve contract should guarantee that, even if A or B are dishonest, exactly one of them will be able to redeem the deposit. To ensure this property, we resort to a mediator M who resolves disputes between A and B . Assuming that the mediator takes a fee

of $0.1\text{\$}$ (cut down from A 's deposit), we can craft the following:

$$\begin{aligned} \text{Escrow} &= \text{PayOrRefund} + A : \text{Resolve}_{0.1,0.9} + B : \text{Resolve}_{0.1,0.9} \\ \text{Resolve}_{v,v'} &= \text{split}(v\text{\$} \rightarrow \text{withdraw } M \\ &\quad | v'\text{\$} \rightarrow M : \text{withdraw } A + M : \text{withdraw } B) \end{aligned}$$

Besides the two clauses in *PayOrRefund*, the contract *Escrow* features two additional clauses, which allow A and B to trigger the dispute resolution, specified by *Resolve*. The two parallel clauses therein split the $1\text{\$}$ deposit in two parts: $0.1\text{\$}$ go to the mediator, while $0.9\text{\$}$ are assigned either to A and B , depending on M 's choice.

BitML contracts feature other primitives besides those seen so far: for instance, they can express time constraints, and allow participants to choose/reveal secrets. We show these features to model a *timed commitment* protocol [12, 26, 35, 57]. There, a participant A wants to choose a secret, and reveal it after some time t — guaranteeing that the revealed value is the chosen secret. We force A to pay to another participant B a penalty of $1\text{\$}$ if A does not reveal the secret within t . In the contract precondition, A declares a deposit of $1\text{\$}$, and a secret with name a . The contract is the following:

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

Only A can choose the first clause, by revealing a . When doing so, A can take her $1\text{\$}$ deposit back. After the deadline t , B can choose the second clause, and collect A 's penalty. Before t , A has the option to reveal a (avoiding the penalty), or to keep it secret (paying the penalty). As a borderline case, if A reveals a after t , a race condition occurs: the first one who makes a step gets the money.

Timed commitment contracts like the one above are the basis upon which constructing complex contracts which distribute bitcoins according to values chosen by participants (e.g., gambling games, lotteries, etc.). For instance, consider a simple “odds and evens” game between two players. The contract preconditions require A and B to commit to one secret each (a and b , respectively), and to put a deposit of $3\text{\$}$ each ($1\text{\$}$ as a bet, and $2\text{\$}$ as a penalty in case of dishonest behaviour). The contract is the following:

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

The balance is split in three parts. Player B must reveal b by the deadline t ; otherwise, A can redeem B 's penalty (as in the timed commitment). Similarly, A must reveal a . To determine the winner we compare the *lengths* of the secrets, in the third clause of *split*. The winner is A if the secrets have the same length, otherwise it is B . Checking that b 's length is either 0 or 1 (in the first clause) is needed to achieve fairness: indeed, B can increase his probability to win $2\text{\$}$ in the third clause by choosing a secret with length $N > 1$. However, doing so would make B lose his $2\text{\$}$ deposit, so overall B 's *average* payoff would be negative. So, a rational B would then choose a secret of length 0 or 1 — as well as a rational A , who otherwise decreases her probability to win. When both lengths are 0 or 1, A and B can redeem their $2\text{\$}$ penalty, and have a $1/2$ probability to win, if at least one of them chooses the length uniformly.

The last primitive we present is `put x`, which allows a participant to provide a “volatile” deposit x after stipulation. We illustrate it in a variant of *Escrow*, where both **A** and **B** contribute to the fee for the mediator, only if they do not agree within *PayOrRefund*:

$$\begin{aligned} \text{EscrowPut} &= \text{PayOrRefund} + \text{after } t : \text{withdraw } \mathbf{B} \\ &+ \text{put } x . (\text{put } y . \text{Resolve}_{0.2,1} + \text{after } t' : \text{withdraw } \mathbf{A}) \end{aligned}$$

As a precondition, we assume that **A** and **B** declare additional volatile $0.1\mathbb{B}$ deposits (respectively, x and y). If they disagree, **A** has to pay her deposit before time t , otherwise **B** can withdraw **A**’s $1\mathbb{B}$ deposit. The same for **B**, which must pay his $0.1\mathbb{B}$ deposit before $t' > t$. After that, they invoke *Resolve*, paying $0.2\mathbb{B}$ to the mediator, and assigning $1\mathbb{B}$ to the winner of the dispute.

Symbolic model. To reason about the security of BitML, we introduce two security models: a symbolic model and a computational one. The symbolic model is based on the semantics of BitML, which abstracts from Bitcoin and its blockchain. In this model, participants are represented as strategies (formally, PPTIME algorithms) which allow them to choose which actions to perform, in any given state. Each honest participant has its own strategy, while the dishonest ones are collectively represented as a single adversarial strategy. Strategies can read the current BitML *configuration* — a public, shared state — and output actions which determine the next state. Further, the adversary strategy can schedule the participants’ moves, eavesdrop messages, and impersonate other participants. Within a configuration, we can find e.g. *deposits* of the form $\langle \mathbf{A}, v\mathbb{B} \rangle$, modelling the ownership of $v\mathbb{B}$ by participant **A**. Participant can freely split or gather their own deposits, or transfer them to other participants. Further, deposits can be spent to stipulate contracts: technically, stipulation creates an *active contract*, i.e. a term of the form $\langle \mathbf{C}, v\mathbb{B} \rangle$, with an initial balance $v\mathbb{B}$ amounting to the sum of all the spent deposits. The contract **C** determines how the balance of $v\mathbb{B}$ can be distributed among participants, depending on their actual interaction. Contracts have an operational semantics, describing all their possible transitions: participants strategies (either honest or adversarial) can only choose among the transitions allowed by the semantics. For instance, a possible symbolic run of *Escrow* is the following (omitting the stipulation and the intermediate steps):

$$\langle \text{Escrow}, 1\mathbb{B} \rangle \rightarrow^* \langle \mathbf{A}, 0.9\mathbb{B} \rangle \mid \langle \mathbf{M}, 0.1\mathbb{B} \rangle$$

which models the case where the mediator has resolved a dispute in favour of **A** (there are also runs where the dispute is won by **B**, or the mediator is not invoked). The timed commitment *TC* has two kinds of runs, according to **A**’s strategy: those where **A** reveals the secret, and those where **B** redeems the deposit. For instance:

$$\langle \text{TC}, 1\mathbb{B} \rangle \mid t_0 \rightarrow^* \langle \mathbf{A}, 1\mathbb{B} \rangle \mid \mathbf{A} : a\#N \mid t_0$$

is an run where **A** has revealed the secret (represented by the term $\mathbf{A} : a\#N$, where N is a ’s length) at time $t_0 < t$. Instead, in the run:

$$\langle \text{TC}, 1\mathbb{B} \rangle \mid t_0 \rightarrow^* \langle \mathbf{B}, 1\mathbb{B} \rangle \mid t_0 + \delta$$

with $t_0 + \delta > t$, **A** has not revealed the secret, and **B** has collected **A**’s $1\mathbb{B}$ penalty. Finally, *OddsEvens* features symbolic runs where either **A** or **B** redeem the winnings, and they reveal or not their secrets. For instance, in the following run (where $t_0 < t$):

$$\langle \text{OddsEvens}, 6\mathbb{B} \rangle \mid t_0 \rightarrow^* \langle \mathbf{A}, 2\mathbb{B} \rangle \mid \langle \mathbf{B}, 4\mathbb{B} \rangle \mid \mathbf{A} : a\#0 \mid \mathbf{B} : b\#1 \mid t_0$$

players have revealed their secrets (and redeemed the $2\mathbb{B}$ deposits), and **B** has won the pot, since the lengths of the secrets are different.

Compiling BitML to Bitcoin. We design a compiler that, given a BitML contract, outputs a set of Bitcoin transactions through which the contract can be actually executed on Bitcoin. Our compiler only relies on standard features of Bitcoin transactions: for instance, the `withdraw` primitive relies on signature verification, `split` relies on transactions with multiple outputs, and `reveal` exploits the hash opcode of Bitcoin scripts. Sequencing and choice are obtained, respectively, by transaction cascading and by the absence of double-spending. We show several examples of compilation in Section 7.

A crucial question is: *how to preserve security properties of a BitML contract once it is compiled into Bitcoin?* Indeed, lifting symbolic security to computational security is a necessary result in order to justify the adoption of high-level languages for smart contracts: in the absence of such result, proving that a contract is secure in the symbolic model would not guarantee its security under a computational attacker. Proving such preservation result is challenging, as the symbolic model is a substantially higher-level than what can actually happen in Bitcoin. For instance, while the executions of contracts strictly follow the BitML semantics (so, the possible configurations are pre-determined), in Bitcoin adversaries can append arbitrary transactions to the blockchain (not necessarily those obtained by the compiler), crafting them with data in their knowledge, possibly sniffed over the network.

Computational model. In order to reason about the behaviour of participants in Bitcoin, we introduce a computational model, where attackers are only subject to the usual restrictions of standard computational models (i.e. they can only manipulate bitstrings using PPTIME algorithms). As in the symbolic model, also computational participants are rendered as strategies: these strategies can listen to network traffic and scan the blockchain, in order to decide their next actions. As an action, a participant can append a transaction to the blockchain: for this to happen, the bitstring generated by a strategy must be the encoding of a transaction which redeems some unspent transaction outputs on the blockchain. Another action is broadcasting a bitstring: this is used for off-chain communications, e.g., when participants need to exchange signatures, or to reveal their secrets. The adversary strategy defines the behaviour of the dishonest participants, and acts as a scheduler for all the participants’ moves, as in the symbolic model. The main difference, here, is that computational actions are bitstrings, instead of symbolic terms. Computational adversaries can obtain these bitstrings from their knowledge (e.g., sniffed messages), with the only limitation of crafting them through PPTIME algorithms. For simplicity, our model slightly abstracts from the actual low-level behaviour of Bitcoin (e.g., we assume the blockchain to be immutable and without forks, and we neglect transaction fees).

Relating symbolic and computational security. The main result of this paper is a *computational soundness* theorem [9]: it states that, with overwhelming probability, runs in the computational model have a corresponding run in the symbolic model. Therefore, computational attacks can be observed in the symbolic model. More precisely, the theorem states that we can securely execute BitML contracts on Bitcoin according to the following workflow. First, each honest participant chooses her symbolic strategy, which drives the

stipulation and execution of contracts according to the symbolic semantics. Then, they translate their strategies to computational ones; in this step, we leverage our compiler to map the involved contracts into Bitcoin transactions. Finally, participants execute their computational strategies, being subject to computational adversaries. Our computational soundness result ensures that computational strategies admit no attacks, unless the original symbolic strategies also admit the same attack. To state computational soundness, we define a correspondence relation (called *coherence*) between runs of contracts in the two models. Intuitively, (symbolic) deposits $\langle A, v\mathbb{B} \rangle$ correspond to (computational) unspent transaction outputs for $v\mathbb{B}$ which can be redeemed using only A 's signature. Active contracts $\langle C, v \rangle$ also correspond to unspent transaction outputs, but they involve more complex redeem conditions (e.g., they may require to provide a signature from *all* the contract participants, or to reveal a secret having a certain length). Performing a transition in the symbolic semantics of a contract C corresponds to appending to the blockchain one of the transactions obtained by compiling C .

Practical consequences of computational soundness. To illustrate the practical applications of our results, recall the *OddsEvens* contract introduced before. Assume that A 's symbolic strategy is to (i) stipulate the contract, depositing $3\mathbb{B}$ and committing to a randomly chosen secret a of length 0 or 1; (ii) reveal her secret at a certain time $t' < t$; (iii) withdraw her deposit at time t' ; (iv) collect the pot, if she is the winner. According to the symbolic semantics, in all the runs of *OddsEvens* which are *conformant* to A ' strategy (and to any adversarial strategy), after time t' , A will receive at least $2\mathbb{B}$ (i.e. her deposit) from the contract. Computational soundness guarantees that the same holds in the computational model, i.e. when executing the compiled contract on Bitcoin. Note that proving this result directly in the computational model is inconvenient, as the adversary — not being constrained to follow the structure of the contract — can play any sequence of actions that respects the computational model, and the consistency of the Bitcoin blockchain (where he can also append transactions crafted by himself, not obtained by the BitML compiler). In this setting, proving security against all adversaries would require to cope with a universal quantification over all possible PPTIME algorithms. By contrast, proving security in the symbolic model is significantly simpler: a verification algorithm should check if the desired property holds in all the reachable configurations in the contract runs conformant with A 's strategy. This simplification is similar to the one obtained when reasoning about the security of cryptographic protocols in symbolic models, instead of computational ones [9].

3 RELATED WORK

The first proposal to implement smart contracts on Bitcoin dates back (at least) to 2012 [1], and the first scientific paper to 2013 [12]. Since then, the research on smart contracts has evolved along different directions: (i) studying contracts that can be run directly on Bitcoin; (ii) increasing the expressiveness of contracts through Bitcoin extensions; (iii) developing high-level languages for Bitcoin contracts; (iv) studying new blockchain infrastructures for smart contracts. Below we briefly survey the literature along these lines.

Bitcoin smart contracts. Basic smart contracts which transfer bitcoins according to the external state [2, 3] can be implemented

using multi-signature transactions. Timed commitments for Bitcoin were originally introduced by [12], and then used to implement multiparty computations, like lotteries. The lottery in [12] requires each player to deposit a collateral which grows quadratically with the number of players. Subsequent works proposed Bitcoin extensions which allow for lotteries without collaterals. More general forms of fair multiparty computations on “pure” Bitcoin, exploiting SegWit [43], were proposed in [11, 22, 41]. Contingent payments for Bitcoin (i.e., contracts which allow to sell solutions for a class of NP problems) were introduced in [18, 46].

Extensions of the Bitcoin scripting language. Other works proposed extensions of Bitcoin to enhance the expressiveness of smart contracts. In [49] the Bitcoin scripting language is extended with *covenants*, a construct that can constrain the structure of the redeeming transaction. Another implementation of covenants is proposed in [53], exploiting a (currently disabled) opcode to concatenate arbitrary data, and introducing a new opcode to verify signatures against it. Covenants enable contracts that implement *vaults*, i.e. protocols which scatter a money transfer along a sequence of transactions, giving the ability to the owner of the vault to abort the transfer if he detects misbehaviour. More generally, recursive covenants allow to implement a state machine through a sequence of transaction that store its state. The collaterals of multiparty lotteries can be eliminated through Bitcoin extensions: e.g., [20] requires input malleability (i.e. the possibility of not signing any input), while [47] requires a new opcode that checks if the redeeming transaction belongs to a predetermined set. The work [32] proposed a contingent payment protocol that does not rely on zero-knowledge proofs, but instead requires a new opcode to check if the two top elements of the stack are a valid key pair. The work [42] captures general multiparty, interactive, stateful computations by exploiting a new opcode (similar to the one in [53]), to check signatures for arbitrary messages.

Languages for Bitcoin smart contracts. Only a few languages for Bitcoin contracts have been proposed so far. TypeCoin [30] allows to model the updates of a state machine as affine logic propositions. Users “run” this machine by putting transactions on the blockchain, with the guarantee that only legit updates can be performed. A downside of [30] is that liveness is guaranteed only when participants cooperate, i.e., an adversary can prevent the others from completing the contract. Note instead that in BitML, honest participants can always make a contract progress. The other languages we are aware of, IVY [5], BALZaC [8] and Simplicity [52], are high-level alternatives to the Bitcoin scripting language, that can be compiled into Bitcoin scripts. In order to implement a smart contract using these languages, one still needs to design it as a protocol involving message exchanges and transactions (although with more readable scripts). Note that these languages do not allow to describe the whole contract, but only the *individual* transactions used in the associated protocol. Compared to these approaches, BitML has two main advantages: first, it can express the *whole* contract within a single term; second, it relieves the designer from the burden of Bitcoin transactions, which, instead, can be automatically generated by our compiler. The loss of expressiveness caused by the abstraction from low-level Bitcoin details (discussed in Section 10),

is repaid by a gain of elegance in the specifications of contracts, and by the simplification of upcoming verification techniques.

Smart contracts beyond Bitcoin. After Bitcoin, other platforms and languages for smart contracts have been created [19]. Currently, the most popular ones are Ethereum [28] (in the permissionless setting) and Hyperledger Fabric [10] (in the permissioned one). The recent attacks on Ethereum contracts have given rise to extensive research on how to make it more secure. A few papers study EVM, the bytecode language which is the target of compilation of Ethereum contracts. Among them, [44] formalises the semantics of EVM, and develops a symbolic execution of EVM contracts, to detect some vulnerability patterns. Another approach based on the analysis of vulnerability patterns on dependency graphs is pursued by [59], which develop a tool called Securify. A more detailed formalisation of EVM (validated against the official Ethereum test suite) is in [37], which also proposes a set of general security properties relevant for avoiding classic vulnerabilities in Ethereum contracts. The work [36] develops EtherTrust, a framework for the static verification of Ethereum smart contracts at the EVM level, which can establish the absence of re-entrancy vulnerabilities. Also [23] detects vulnerabilities of Ethereum contracts, by translating Solidity and EVM code into F* [56]. Further, given a Solidity program and an alleged compilation of it into EVM bytecode, [23] verifies that the two pieces of code have equivalent behaviours. The work [38] uses the Isabelle/HOL proof assistant [51] to verify the EVM obtained by compiling the Solidity code of “Deed”, a contract which is part of the Ethereum Name Service. In particular, it proves that, upon an invocation of the contract, only its owner can decrease the balance. The work [55] proposes Scilla, a strongly typed intermediate language where contracts are represented as Communicating Automata. Compared to EVM, Scilla is more structured: this simplifies formal reasoning, and makes contracts more amenable to verification. Ongoing work aims at a Coq formalization of Scilla. Other papers propose domain-specific languages for Ethereum contracts. Among them, [45] represents smart contracts as finite state automata, where state transitions can be constrained according to contract variables and inputs; a tool is provided to translate these automata into Solidity code. The work [24] compiles to Solidity a fragment of the language for financial contracts introduced by [39]. While the previous works address qualitative properties of contracts, the work [29] develops a framework for their *quantitative* analysis, by transforming contracts (specified in an abstract language) into state-based games. This allows one to compute the worst-case guaranteed utility resulting from interacting with a contract, which can be helpful to detect vulnerabilities.

4 THE BITML CALCULUS

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 \mathbb{B} , while a, b, \dots denote *secrets*. We denote with \vec{x} a finite sequence of deposit names, and we adopt a similar notation for sequences of other kinds.

Definition 1 (Contract preconditions).

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

The precondition $A : ! v @ x$ requires A to own $v \mathbb{B}$ in a deposit x , and to spend it for stipulating a contract C . Instead, $A : ? v @ x$ only requires A to pre-authorize the spending of x . Since x is not spent upfront, there is no guarantee that $v \mathbb{B}$ will be available when C demands x , as A can spend it for other purposes. Finally, $A : \text{secret } a$ requires A to generate a random nonce a , and commit to it before C starts. During the execution of C , A can choose whether to reveal a to the other participants, or not.

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

A contract C is a *choice* among branches. Intuitively, a branch D performs an action, and possibly proceeds with a continuation C' . The action $\text{put } \vec{x} \ \& \ \text{reveal } \vec{a} \ \text{if } p$ atomically performs the following: (i) spend all the volatile deposits \vec{x} , adding their values to the current balance; (ii) check that all the secrets \vec{a} have been revealed, and that they satisfy the predicate p . The guarded contract $\text{split } v_1 \rightarrow C_1 \mid \dots \mid v_n \rightarrow C_n$ 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 $\text{withdraw } A$ 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 $\text{after } t : D$, which requires to wait until time t .

Definition 3 (Contract advertisement). A contract advertisement is a term $\{G\}C$ satisfying the following conditions: (i) the names in G are distinct; (ii) each name in C occurs in G ; (iii) the names in $\text{put } \vec{x} \ \& \ \text{reveal } \vec{a} \ \text{if } p$ are distinct; and, each name in p occurs in \vec{a} ; (iv) each A in $\{G\}C$ has a persistent deposit in G . ◇

The last condition will be used to guarantee that the contract is stipulated only if *all* the involved participants give their authorizations. Indeed, in order to transform a contract advertisement $\{G\}C$ into an active contract, our semantics requires only the authorizations of participants with persistent deposits in G . So, condition (iv), makes the participants involved in $\{G\}C$ equal to those with persistent deposits in G , causing the stipulation to happen only when everyone agrees. Requiring exactly the authorizations to spend persistent deposits in the symbolic semantics is key to implement contract stipulation in Bitcoin: indeed, to record that a contract has been stipulated, in Bitcoin one has to append a suitable transaction, say T_{init} . When doing this, the Bitcoin network cannot check conditions corresponding to committing to a secret, or to pre-authorize a volatile deposit, since these actions are performed *off-chain*. The

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

Figure 1: Syntax of BitML contracts.

only condition that can be checked is that persistent deposits are spent contextually with appending T_{init} : this can be obtained by setting the inputs of T_{init} to these deposits. Note also that condition (iv) 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`.

Semantics We introduce a reduction semantics of BitML. Because of space limitations, here we provide the underlying intuitions, relegating the full formalisation to Appendix A.2. The rules of the semantics are grouped into four sets: (i) rules for managing deposits; (ii) rules for advertising contracts and stipulating them; (iii) rules for executing active contracts; (iv) rules for handling time. The untimed rules follow a common pattern, in order to perform an operation: first, the involved participants give their authorization; then, the operation is actually performed. Both cases are rendered as transitions in the semantics. Transitions are decorated with labels, which describe the performed actions. For simplicity, we ignore these labels in our informal description below.

The configurations of the semantics contain the following terms: (i) *contract advertisements* $\{G\}C$ represent a contract which has been proposed, but not stipulated yet; (ii) *active contracts* $\langle C, v \rangle_x$ represent a stipulated contract, holding a current balance of $v\mathbb{B}$. The name x uniquely identifies the active contract; (iii) *personal deposits* $\langle A, v \rangle_x$ represent a fund of $v\mathbb{B}$ owned by A , and with unique name x ; (iv) *authorizations* $A[\chi]$ represent the consent of A to perform some operation χ ; (v) *committed secrets* $\{A : a\#N\}$, represent A who has committed a random nonce a of (secret) length N , by broadcasting its hash $H(a)$; (vi) *revealed secrets* $A : a\#N$ represent the fact that A has revealed her secret a (hence, its length N).

Definition 4 (Configurations). The syntax of configurations is:

$\Gamma ::= 0$	empty
$ \{G\}C$	contract advertisement
$ \langle C, v \rangle_x$	an active contract containing $v\mathbb{B}$
$ \langle A, v \rangle_x$	a deposit of $v\mathbb{B}$ redeemable by A
$ A[\chi]$	authorization of A to perform χ
$ \{A : a\#N\}$	committed secret of A ($N \in \mathbb{N} \cup \{\perp\}$)
$ A : a\#N$	revealed secret of A ($N \in \mathbb{N}$)
$ \Gamma \Gamma'$	parallel composition

Further, $\Gamma | t$ is a *timed* configuration, where $t \in \mathbb{N}$ is a global time.

We now illustrate the BitML semantics through a series of examples, which cover all the primitives. When time is immaterial, we will only show the steps of the untimed semantics. A full execution of the timed commitment contract is shown in Appendix A.2.

Deposits. When a participant A owns a deposit $\langle A, v \rangle_x$, she can employ that amount for several operations: she can divide the deposit into two smaller deposits, or join it with another deposit of hers to form a larger one; the deposit can also be transferred to another participant, or destroyed. For instance, to authorize the join of two deposits, A can perform the following step:

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

where $\chi_x = x, y \triangleright \langle A, v + v' \rangle$ means that A authorizes to spend x . After A also provides the dual authorization χ_y , any participant can perform the actual join as follows:

$$\langle A, v \rangle_x | \langle A, v' \rangle_y | A[\chi_x] | A[\chi_y] \rightarrow \langle A, v + v' \rangle_z$$

Advertisement. Any participant, at any time, can advertise a new contract C (with preconditions G) by performing the following step:

$$\Gamma \rightarrow \Gamma | \{G\}C$$

The rule requires that all the deposits mentioned in G exist in Γ , that secrets names are fresh, and that at least one of the participants in G is honest. The last condition, useful to obtain computational soundness, does not limit the power of adversary. Indeed, the same effect of executing a contract among dishonest participants can be obtained by the adversary using the deposit rules, only.

Stipulation. To perform a stipulation, turning a contract advertising into an active contract, a few steps are needed. For instance, consider $G = A : ! 1\mathbb{B} @ x | A : ? 1\mathbb{B} @ y | A : \text{secret } a$, and let C be an arbitrary contract, only involving A . Assuming A honest, our semantics gives the following steps:

$$\begin{aligned} & \langle A, 1\mathbb{B} \rangle_x | \langle A, 1\mathbb{B} \rangle_y | \{G\}C \\ & \rightarrow \langle A, 1\mathbb{B} \rangle_x | \langle A, 1\mathbb{B} \rangle_y | \{G\}C | \{A : a\#N\} | A[\# \triangleright \{G\}C] = \Gamma \\ & \rightarrow \Gamma | A[x \triangleright \{G\}C] \\ & \rightarrow \langle A, 1\mathbb{B} \rangle_y | \langle C, 1\mathbb{B} \rangle_z | \{A : a\#N\} \end{aligned}$$

The rules require all participants to commit to their secrets (and to their lengths). Above, this step adds $\{A : a\#N\} | A[\# \triangleright \{G\}C]$ to the configuration, where N is the committed length. After that, all participants must authorize to spend their persistent deposits. Above, this step adds $A[x \triangleright \{G\}C]$ to the configuration. After all such authorizations have been performed, any participant can spend

the persistent deposits to create the active contract. Above, this step consumes $\langle A, 1\mathbb{B} \rangle_x$ (and all the authorizations) to create $\langle C, 1\mathbb{B} \rangle_z$.

We anticipate that, in the computational setting, committing to a secret a is performed by generating a random nonce s_a and broadcasting its hash $h_a = H(s_a)$. Note that a dishonest A could perform a fake commitment, by broadcasting a value h_a without knowing its preimage. In the symbolic setting, we model this situation by allowing a dishonest A to produce a term $\{A : a\#\perp\}$, where \perp represents an “invalid” hash. In the subsequent steps, A will not be able to reveal the secret, coherently with the fact that, in the computational setting, A is not able to compute a preimage of h_a .

Withdraw. We now exemplify the rules for active contracts. Executing `withdraw` A transfers the whole contract balance to A :

$$\langle \text{withdraw } A + C', v \rangle_x \rightarrow \langle A, v \rangle_y$$

This step terminates the contract, and creates a deposit owned by A , with a fresh name y . Above, `withdraw` A is executed as a branch within a choice: as usual, taking a branch discards the other ones (denoted as C'). Below, C' always denotes the discarded branches.

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

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

The balance of the initial contract, $(v_1 + v_2)\mathbb{B}$, is split between two newly spawned contracts: C_1 , receiving $v_1\mathbb{B}$, and C_2 , receiving $v_2\mathbb{B}$. After this step, the two new contracts are executed independently.

Put. `put` x takes the volatile deposit x within the contract:

$$\langle \text{put } x.C + C', v \rangle_y \mid \langle A, v' \rangle_x \rightarrow \langle C, v + v' \rangle_z$$

This step can only be performed if x is still unspent; otherwise, the `put` prefix is stuck. This action naturally generalizes to multiple volatile deposits \vec{x} (if any of them is spent, the prefix is stuck).

Reveal. `reveal` a if p can be fired when the previously committed secret a has been revealed, and it satisfies the guard p . E.g.:

$$\begin{aligned} & \langle \text{reveal } a \text{ if } |a| = M. C + C', v \rangle_x \mid \{A : a\#N\} \\ \rightarrow & \langle \text{reveal } a \text{ if } |a| = M. C + C', v \rangle_x \mid \{A : a\#N\} \\ \rightarrow & \langle C, v \rangle_y \mid \{A : a\#N\} \quad \text{if } N = M \end{aligned}$$

In the first step, A reveals her secret a . In the second step, any participant can cause the contract to take the `reveal` branch, provided that the length of a is M , as required by the predicate. This action naturally generalizes to the case of multiple secrets \vec{a} (all of the must be revealed), and to the case where `put` and `reveal` are performed atomically, e.g. in an action `put` \vec{x} & `reveal` \vec{a} if p .

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

$$\begin{aligned} & \langle A : \text{withdraw } B + C', v \rangle_x \\ \rightarrow & \langle A : \text{withdraw } B + C', v \rangle_x \mid \{A[x \triangleright A : \text{withdraw } B]\} \rightarrow \langle B, v \rangle_y \end{aligned}$$

In the first step, A authorizes the contract to take the branch `withdraw` B . After that, any participant can fire such branch. When multiple authorizations are required, the branch can be taken only after all of them occur in the configuration.

Time. We now discuss the rules for handling time. These rules describe transitions between timed configurations $\Gamma \mid t$ where t denotes the current time. We always allow time t to advance by a delay $\delta > 0$, through the rule: $\Gamma \mid t \rightarrow \Gamma \mid t + \delta$. We allow a contract branch decorated with `after` t to be taken only when the current time is greater than t . For instance, if $t_0 + \delta \geq t$:

$$\begin{aligned} & \langle \text{after } t : \text{withdraw } B, v \rangle_x \mid t_0 \\ \rightarrow & \langle \text{after } t : \text{withdraw } B, v \rangle_x \mid t_0 + \delta \rightarrow \langle B, v \rangle_y \mid t_0 + \delta \end{aligned}$$

For the branches not guarded by an `after`, we lift transitions from untimed to timed configurations, without making them affect time. Namely, for an untimed transition $\Gamma \rightarrow \Gamma'$, we also have the timed transition $\Gamma \mid t \rightarrow \Gamma' \mid t$. This makes actions instantaneous (similarly to many timed process calculi [50]) reflecting the assumption that participants can always meet deadlines, if they want to.

5 SYMBOLIC STRATEGIES & ADVERSARIES

Symbolic runs. A *symbolic run* R^s is a (possibly infinite) sequence $\Gamma_0 \mid t_0 \xrightarrow{\alpha_0} \Gamma_1 \mid t_1 \xrightarrow{\alpha_1} \dots$ where α_i are the transition labels, Γ_0 contains only deposits, and $t_0 = 0$. If R^s is finite, we write Γ_{R^s} for its last untimed configuration.

Stripping. The strategy of a participant can inspect the whole past run, except for the (lengths of the) unrevealed secrets. The *stripping* of a run censors this information: technically, it replaces each committed secret $\{A : a\#N\}$ with a term $\{A : a\#\perp\}$.

Symbolic participant strategies. A symbolic strategy Σ_A^s is a PPTIME algorithm which allows A to select which action(s) to perform, among those permitted by the BitML semantics. Σ_A^s receives as input a stripped run R_x^s , and outputs a finite set of actions (possibly, time delays) that A wants to perform. The choice among these actions is controlled by the adversary strategy, specified below. We forbid Σ_A^s to output authorizations for participants $B \neq A$. Further, strategies must be *persistent*: if on a run Σ_A^s chooses an action 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. In this way, once Σ_A^s has chosen α (implicitly, giving to the adversary her consent to schedule such action), she cannot change her mind.

Symbolic adversary strategies. The adversary Adv acts on behalf of all the dishonest participants, and controls the scheduling among all participants (including the honest ones). Her symbolic strategy Σ_{Adv}^s takes as input the current run and the sets of moves outputted by the strategies of honest participants. Both the run and the moves are stripped, to prevent the adversary from inferring the lengths of secrets. The output of Σ_{Adv}^s is a single action λ^s (to be appended to the current run), only subject to the following constraints: (i) if λ^s is an authorization by some honest A , then it must be chosen by Σ_A^s ; (ii) if λ^s is a time delay, then all the honest participants must agree. Condition (i) rules out authorization forgeries, while (ii) prevents the adversary from delaying the honest participants. Were this condition dropped, honest participants could be prevented from meeting deadlines: e.g., in the timed commitment contract, the adversary could collude with B by enabling the `after` branch before A had the possibility to reveal her secret.

Symbolic conformance. Strategies are PPTIME algorithms: so, besides the other inputs, they also implicitly take as input a stream r of random bits. Fixing this stream and a set of strategies Σ^s — both for the honest participants and for the adversary — we obtain a unique run, which is made by the sequence of actions chosen by Σ_{Adv}^s when taking as input the outputs of the honest participants’ strategies. We say that this run is *conformant* to Σ^s (and r). When $\Sigma_{Adv}^s \notin \Sigma^s$, we say that R^s conforms to Σ^s (and r) when there exists some Σ_{Adv}^s such that R^s conforms to $\Sigma^s \cup \{\Sigma_{Adv}^s\}$ (and r).

6 COMPUTATIONAL MODEL

In this section we introduce our computational model, which will be the target of the BitML compiler. We start by briefly recapping Bitcoin transactions (referring to [16] for a full formal model).

Transactions. In Bitcoin, *transactions* describe transfers of currency. 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*. For instance, consider the transactions T_1 and T_2 below:

T_1	T_2
in: ...	in: $0 \mapsto (T_1, 0)$
wit: ...	wit: $0 \mapsto sig_{K(A)}$
out: $0 \mapsto (\lambda \zeta. versig_{K(A)}(\zeta), v_0 \mathfrak{B})$	out: $0 \mapsto (\lambda x. H(x) = k, v_0 \mathfrak{B})$
out: $1 \mapsto (\lambda \zeta. versig_{K(B)}(\zeta), v_1 \mathfrak{B})$	absLock: t

The transaction T_1 has two *outputs*: the $v_0 \mathfrak{B}$ in *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* *out*(0) (similarly for the other output). This is the case e.g. of the transaction T_2 above. Its *witness* $sig_{K(A)}$ is the signature of A on T_2 (excluding the *wit* field itself), as required by $T_1.out(0)$.

If T_1 is on the blockchain and its *out*(0) is unspent, A can update the blockchain by appending T_2 . This moves $v_0 \mathfrak{B}^1$ 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 $T_2.absLock$ represents the earliest moment when T_2 can be put on the blockchain. A subsequent transaction can redeem $(v_1 + v_0) \mathfrak{B}$ 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 .

We assume that Bitcoin uses secure cryptographic primitives, i.e. ideal hash functions (which hereafter are modelled according to the random oracle model [21]), and a digital signature scheme which is robust against existential forgery attacks.

Blockchain. A blockchain B is a sequence $(T_0, t_0) \cdots (T_n, t_n)$, where $T_0 \cdots T_n$ are transactions (T_0 is the *coinbase* transaction, i.e. $T_0.in = \perp$, meaning that it does not point to a previous transaction), and $t_0 \cdots t_n$ are timestamps, with $t_i \leq t_j$ for all $i \leq j$. Hereafter, we assume that blockchains are append-only, without forks, and *consistent*, i.e. obtained by appending transactions that respect the Bitcoin protocol (as formalised in [16]). Essentially, these ideal properties are coherent with the results in [17, 34, 40], as long as the adversary does not control a very large mining pool [13, 33].

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 [16]) we neglect these *transaction fees*.

Computational runs. We now introduce the computational counterparts of symbolic runs. These are sequences of *computational labels* λ^c , namely bitstrings 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
T	append a Bitcoin transaction to the blockchain
δ	perform a delay

To compute the hash of a message m , A sends m to the oracle $O \notin \text{Part}$, and waits for the reply $H(m)$. Also these actions are included in the computational labels. Note that reliable message broadcasts can be effectively obtained through the Bitcoin network. We postulate that a computational run begins with a coinbase transaction T_0 , followed by the broadcasts of the public keys of all participants. Each A has two key pairs: K_A for signing messages, and \hat{K}_A for redeeming deposits. For each A (either honest or not), we assume that T_0 has an output redeemable with the private key \hat{K}_A . In this way, each participant starts with some funds (possibly 0), and knows the public keys of the others. By extracting the transactions from a computational run R^c , and assigning their time according to the accumulated delays, we obtain a blockchain, denoted as B_{R^c} .

Stripping. Analogously to the stripping of symbolic runs, we define the *A-stripping* of a computational run R^c as the run obtained by removing from R^c all the messages *not* visible by A , i.e. the messages between O and some other $B \neq A$.

Computational participant strategies. We now introduce the computational counterparts of symbolic strategies. A computational strategy Σ_A^c for A is a PPTIME algorithm which receives as input a (A -stripped) computational run R_*^c , and outputs a finite set of computational labels. The choice among these labels is controlled by the adversary strategy, specified below. We impose a few sanity constraints: (i) we forbid A to impersonate another participant; (ii) if the strategy outputs a transaction T , then T must be a consistent update of the blockchain $B_{R_*^c}$ obtained from the run in input, and all the witnesses of T have already been broadcast in the run; (iii) finally, strategies must be persistent, similarly to the symbolic case. The condition on the witnesses of T in (ii) is not a limitation, since these witnesses become public once T is broadcast on the Bitcoin network. We put this condition since it helps to obtain a sharp correspondence between computational and symbolic runs: indeed, the symbolic counterpart of broadcasting a witness is to give an authorization.

Computational adversary strategies. A computational adversary strategy Σ_{Adv}^c is a PPTIME algorithm taking as input a (Adv -stripped) computational run R_*^c and the moves chosen by each honest participant. The strategy gives as output a single computational label, to be appended to the run. We require Σ_{Adv}^c to obey sanity constraints similar to those of participants strategies, with two differences: (i) Adv can impersonate any other participant (except the oracle O); (ii) Adv can perform a time delay only if all the honest participants agree.

This assumption prevents the adversary from delaying the actions of honest participants, who can therefore always meet their deadlines, without interference by the adversary. Note that, even if the adversary controlled a large portion of the mining power of the

Bitcoin network — so being able to delay the transactions sent by honest participants — honest participants could still protect themselves by setting far enough deadlines in their contracts, making the delay attacks ineffective².

Computational conformance. For each participant A , we indicate with r_A the stream of random bits implicitly used by Σ_A^c . Similarly, we indicate with r_O the stream available to the oracle. Fixing a set of computational strategies Σ^c — both for the honest participants and for the adversary — and a function r from participants to streams, we obtain a unique computational run, made by the sequence of actions chosen by Σ_{Adv}^c when taking as input the outputs of the honest participants’ strategies, and enforcing that the queries to the oracle are always answered. We say that this run is *conformant* to Σ^c and r .

7 COMPILING BITML TO BITCOIN

We now describe how to implement BitML on top of Bitcoin. This is done by *compiling* BitML contract advertisements $\{G\}C$ into a finite set \mathcal{T} of Bitcoin transactions³, signed by all participants (except from those whose authorizations are required after stipulation, via $A : \dots$). In Bitcoin, appending (a subset of) these transactions to the blockchain mimics, in BitML, the semantics of $\{G\}C$ (assuming, as usual, that at least one participant is honest). More precisely, the first transaction in \mathcal{T} to be published (called T_{init}), redeems all the permanent deposits, correspondingly to the stipulation of $\{G\}C$. After that, T_{init} can only be redeemed by one of the transactions in \mathcal{T} which corresponds to the subsequent computation step of $\langle C, v \rangle$ in BitML. This is enforced by requiring, in the output script of T_{init} , suitable signatures by all the participants involved in $\{G\}C$. The same principle is followed for the other transactions in \mathcal{T} : they can be appended to the blockchain only when this corresponds to a computation step in BitML. The only transactions in \mathcal{T} that can be redeemed by transactions *not* in \mathcal{T} are those which correspond to the deposits $\langle A, v \rangle$ obtained by reducing a `withdraw A` contract. Indeed, these transactions can be redeemed by standard transactions signed by A .

We illustrate the compiler through a series of examples, which cover all the primitives of BitML. We refer to Appendix A.5 for a formal definition and further examples.

Before running the compiler, participants generate the following key pairs, and exchange their public parts:

- $K(A)$, used in the compilation of a `withdraw A` contract. We exploit this to guarantee that the deposit obtained after the execution of `withdraw A` can be redeemed with a signature of A (using the private part of $K(A)$).
- $K(D, A)$ (for each subterm D of C), used in the compilation of the subterms of C of the form $D + C'$ to enable the firing of the (initial) action in the D branch. The private part of

²For the sake of simplicity, we do not model such delays. At the price of adding further complexity, we could extend our model to relax this assumption, by allowing the adversary to perform delays under reasonable constraints. As long as the same delays are allowed in both the symbolic and computational models, computational soundness still holds. In such extended model, deadlines in contracts must be chosen so to compensate for these delays.

³An implementation of the BitML compiler is under way, at <https://github.com/bitml-lang/bitml-compiler>. This implementation generates standard Bitcoin transactions by exploiting our BALZaC tool [8]. This is crucial, since the Bitcoin network currently discards non-standard transactions.

$K(D, A)$ is used to compute the witness of the transaction corresponding to the compilation of D .

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

Withdraw. As a first example, we show how to compile:

$$\{G\}C = \{A : ! 1\text{B} @ x \mid B : ! 1\text{B} @ y\} \text{ withdraw } B$$

The transactions obtained from the compiler are the following, where T_x is the transaction associated to A ’s deposit $\langle A, 1\text{B} \rangle_x$ and T_y is the transaction associated to B ’s deposit $\langle B, 1\text{B} \rangle_y$:

T_{init}	T_B
in: $0 \mapsto T_x, 1 \mapsto T_y$	in: T_{init}
wit: $0 \mapsto \text{sig}_{K(A)}, 1 \mapsto \text{sig}_{K(B)}$	wit: $\text{sig}_{K(\text{withdraw } B, \{A, B\})}$
out: $(\lambda \zeta. \text{versig}_{K(\text{withdraw } B, \{A, B\})}(\zeta), 2\text{B})$	out: $(\lambda \zeta. \text{versig}_{K(B)}(\zeta), 2\text{B})$

In BitML, the stipulation of the contract requires a few steps:

$$\begin{aligned} &\langle A, 1\text{B} \rangle_x \mid \langle B, 1\text{B} \rangle_y \mid \{G\}C \\ \rightarrow^4 &\langle A, 1\text{B} \rangle_x \mid \langle B, 1\text{B} \rangle_y \mid \{G\}C \mid A[x \triangleright \{G\}C] \mid B[y \triangleright \{G\}C] \\ &\quad \mid A[\# \triangleright \{G\}C] \mid B[\# \triangleright \{G\}C] = \Gamma \end{aligned}$$

In Bitcoin, these steps correspond to obtaining the transactions above from the compiler, and exchanging the signatures shown there (signing T_{init} last). The BitML authorization $A[x \triangleright \{G\}C]$ corresponds to the broadcast of A ’s signature in the first witness of T_{init} . The authorization $A[\# \triangleright \{G\}C]$ in this case is immaterial: in general, it corresponds to the commit of A ’s secrets. The authorizations $B[y \triangleright \{G\}C]$ and $B[\# \triangleright \{G\}C]$ are similar, for B . In BitML, the stipulation is then completed with the following step:

$$\Gamma \rightarrow \langle \text{withdraw } B, 2\text{B} \rangle$$

In Bitcoin, this corresponds to appending T_{init} to the blockchain, which redeems T_x and T_y . The last computation step in BitML is:

$$\langle \text{withdraw } B, 2\text{B} \rangle \rightarrow \langle B, 2\text{B} \rangle_z$$

In Bitcoin, this corresponds to appending T_B to the blockchain, which redeems 2B from T_{init} ⁴. After that, 2B are under B ’s control, since the output script of T_B only requires B ’s signature. The unspent transaction T_B corresponds to the BitML deposit $\langle B, 2\text{B} \rangle_z$.

Note that, during the stipulation phase, it is crucial to exchange the signatures on T_{init} after the ones on T_B , otherwise we would lose the correspondence with the BitML semantics. Indeed, were T_{init} signed first, A could refuse to sign T_B , and yet be able to complete the stipulation, by appending T_{init} to the blockchain. Then, B would be prevented to append T_B , i.e. the Bitcoin counterpart of the BitML `withdraw B` action, which would transfer 2B to him.

Authorizations. We exploit the previous example to illustrate the compilation of authorizations. Compiling $A : \text{withdraw } B$ requires only minor changes: in the witness of T_B , the compiler only inserts B ’s signature $\text{sig}_{K(A : \text{withdraw } B, B)}$, while the output script of T_{init} still requires both signatures, so to prevent T_B from being appended

⁴Note that in the example above, an optimized version of our compiler could omit the transaction T_{init} , and transfer the deposit to B directly through T_B (with the in field as in T_{init}). However, this optimization is only possible when `withdraw` does not occur within a choice; further, the semantics of `withdraw` should be revised accordingly to preserve computational soundness. For the sake of simplicity, we prefer to ignore this optimization in the definition of our compiler.

the blockchain before obtaining A 's authorization. In BitML, to perform the `withdraw` action, A must authorize it:

$$\Gamma \rightarrow \Gamma \mid A[A : \text{withdraw } B]$$

In Bitcoin, this step corresponds to broadcasting A 's signature $\text{sig}_{\mathbf{K}(A : \text{withdraw } B, A)}$. After that, the signature can be added to the witness of T_B , which can now be appended to the blockchain.

After. Similarly, compiling after $t : \text{withdraw } B$ requires only a small change to our first example: the transaction T_B now has an `absLock` field set to time t . In this way, T_B can not be appended to the blockchain until time t — coherently to the BitML semantics, where the `withdraw B` action can not be fired until such time.

Split. To illustrate the compilation of the `split` primitive, let:

$$C_s = \text{split } (1\text{\$} \rightarrow \text{withdraw } A \mid 2\text{\$} \rightarrow \text{withdraw } B)$$

Compiling $\{A : ! 1\text{\$} @ x \mid B : ! 1\text{\$} @ y\} C_s$ produces the following transactions, where T_x and T_y correspond to the deposits $\langle A, 2\text{\$} \rangle_x$ and $\langle B, 1\text{\$} \rangle_y$, and $KW^p = \mathbf{K}(\text{withdraw } p, \{A, B\})$ for $p \in \{A, B\}$:

<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_{init}</th></tr> <tr><td style="padding: 2px;">in: $0 \mapsto T_x, 1 \mapsto T_y$</td></tr> <tr><td style="padding: 2px;">wit: $0 \mapsto \text{sig}_{\mathbf{K}(A)}, 1 \mapsto \text{sig}_{\mathbf{K}(B)}$</td></tr> <tr><td style="padding: 2px;">out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(C_s, \{A, B\})}(\zeta), 3\text{\\$})$</td></tr> </table>	T_{init}	in: $0 \mapsto T_x, 1 \mapsto T_y$	wit: $0 \mapsto \text{sig}_{\mathbf{K}(A)}, 1 \mapsto \text{sig}_{\mathbf{K}(B)}$	out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(C_s, \{A, B\})}(\zeta), 3\text{\$})$	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_{split}</th></tr> <tr><td style="padding: 2px;">in: T_{init}</td></tr> <tr><td style="padding: 2px;">wit: $\text{sig}_{\mathbf{K}(C_s, \{A, B\})}$</td></tr> <tr><td style="padding: 2px;">out: $0 \mapsto (\lambda \zeta. \text{versig}_{KW^A}(\zeta), 1\text{\\$})$</td></tr> <tr><td style="padding: 2px;">$1 \mapsto (\lambda \zeta. \text{versig}_{KW^B}(\zeta), 2\text{\\$})$</td></tr> </table>	T_{split}	in: T_{init}	wit: $\text{sig}_{\mathbf{K}(C_s, \{A, B\})}$	out: $0 \mapsto (\lambda \zeta. \text{versig}_{KW^A}(\zeta), 1\text{\$})$	$1 \mapsto (\lambda \zeta. \text{versig}_{KW^B}(\zeta), 2\text{\$})$
T_{init}										
in: $0 \mapsto T_x, 1 \mapsto T_y$										
wit: $0 \mapsto \text{sig}_{\mathbf{K}(A)}, 1 \mapsto \text{sig}_{\mathbf{K}(B)}$										
out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(C_s, \{A, B\})}(\zeta), 3\text{\$})$										
T_{split}										
in: T_{init}										
wit: $\text{sig}_{\mathbf{K}(C_s, \{A, B\})}$										
out: $0 \mapsto (\lambda \zeta. \text{versig}_{KW^A}(\zeta), 1\text{\$})$										
$1 \mapsto (\lambda \zeta. \text{versig}_{KW^B}(\zeta), 2\text{\$})$										
<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_A</th></tr> <tr><td style="padding: 2px;">in: $(T_{split}, 0)$</td></tr> <tr><td style="padding: 2px;">wit: $\text{sig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}$</td></tr> <tr><td style="padding: 2px;">out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(A)}(\zeta), 1\text{\\$})$</td></tr> </table>	T_A	in: $(T_{split}, 0)$	wit: $\text{sig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}$	out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(A)}(\zeta), 1\text{\$})$	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_B</th></tr> <tr><td style="padding: 2px;">in: $(T_{split}, 1)$</td></tr> <tr><td style="padding: 2px;">wit: $\text{sig}_{\mathbf{K}(\text{withdraw } B, \{A, B\})}$</td></tr> <tr><td style="padding: 2px;">out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(B)}(\zeta), 2\text{\\$})$</td></tr> </table>	T_B	in: $(T_{split}, 1)$	wit: $\text{sig}_{\mathbf{K}(\text{withdraw } B, \{A, B\})}$	out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(B)}(\zeta), 2\text{\$})$	
T_A										
in: $(T_{split}, 0)$										
wit: $\text{sig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}$										
out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(A)}(\zeta), 1\text{\$})$										
T_B										
in: $(T_{split}, 1)$										
wit: $\text{sig}_{\mathbf{K}(\text{withdraw } B, \{A, B\})}$										
out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(B)}(\zeta), 2\text{\$})$										

As before, T_{init} gathers A 's and B 's deposits and starts the contract. Then, appending T_{split} to the blockchain splits the contract balance between two different outputs, indexed with 0 and 1. In BitML, this would correspond to the computation step:

$$\langle C_s, 3\text{\$} \rangle \rightarrow \langle \text{withdraw } A, 1\text{\$} \rangle \mid \langle \text{withdraw } B, 2\text{\$} \rangle$$

where the two contracts in the parallel composition can be executed independently (as usual in process calculi). Similarly, the two outputs of T_{split} can be independently redeemed by T_A and T_B . The in field of these transactions specifies, besides the input transaction T_{split} , also the index of the output they want to redeem. For instance, appending T_A corresponds, in BitML, to the step:

$$\langle \text{withdraw } A, 1\text{\$} \rangle \mid \langle \text{withdraw } B, 2\text{\$} \rangle \rightarrow \langle A, 1\text{\$} \rangle \mid \langle \text{withdraw } B, 2\text{\$} \rangle$$

Put. To illustrate the compilation of the `put` primitive, let:

$$C_p = \text{put } x. \text{withdraw } B$$

Compiling $\{A : ? 1\text{\$} @ x \mid A : ! 1\text{\$} @ y \mid B : ! 1\text{\$} @ z\} C_p$ produces the following transactions, where T_x, T_y, T_z are the Bitcoin counterpart of the BitML deposits:

<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_{init}</th></tr> <tr><td style="padding: 2px;">in: $0 \mapsto T_y, 1 \mapsto T_z$</td></tr> <tr><td style="padding: 2px;">wit: $0 \mapsto \text{sig}_{\mathbf{K}(A)}, 1 \mapsto \text{sig}_{\mathbf{K}(B)}$</td></tr> <tr><td style="padding: 2px;">out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(C_p, \{A, B\})}(\zeta), 2\text{\\$})$</td></tr> </table>	T_{init}	in: $0 \mapsto T_y, 1 \mapsto T_z$	wit: $0 \mapsto \text{sig}_{\mathbf{K}(A)}, 1 \mapsto \text{sig}_{\mathbf{K}(B)}$	out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(C_p, \{A, B\})}(\zeta), 2\text{\$})$	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_B</th></tr> <tr><td style="padding: 2px;">in: T_{put}</td></tr> <tr><td style="padding: 2px;">wit: $\text{sig}_{\mathbf{K}(\text{withdraw } B, \{A, B\})}$</td></tr> <tr><td style="padding: 2px;">out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(B)}(\zeta), 3\text{\\$})$</td></tr> </table>	T_B	in: T_{put}	wit: $\text{sig}_{\mathbf{K}(\text{withdraw } B, \{A, B\})}$	out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(B)}(\zeta), 3\text{\$})$
T_{init}									
in: $0 \mapsto T_y, 1 \mapsto T_z$									
wit: $0 \mapsto \text{sig}_{\mathbf{K}(A)}, 1 \mapsto \text{sig}_{\mathbf{K}(B)}$									
out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(C_p, \{A, B\})}(\zeta), 2\text{\$})$									
T_B									
in: T_{put}									
wit: $\text{sig}_{\mathbf{K}(\text{withdraw } B, \{A, B\})}$									
out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(B)}(\zeta), 3\text{\$})$									

T_{put}
in: $0 \mapsto T_{init}, 1 \mapsto T_x$
wit: $0 \mapsto \text{sig}_{\mathbf{K}(C_p, \{A, B\})}, 1 \mapsto \text{sig}_{\mathbf{K}(A)}$
out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(\text{withdraw } B, \{A, B\})}(\zeta), 3\text{\$})$

The transaction T_{init} gathers *only* the persistent deposits. T_{put} has two inputs: T_{init} , which can be redeemed with the signatures of both A and B , and T_x , which can be redeemed with A 's signature only. All these signatures are exchanged before stipulation, hence any participant can append T_{put} to the blockchain — provided that T_x is still unspent. Instead, if T_x has been spent, the contract gets stuck, and the deposit within T_{init} is frozen — coherently with the semantics of BitML, where deposits can be spent arbitrarily.

Reveal / choice. Recall from Section 2 the timed commitment contract $TC = D_1 + D_2$, where $D_1 = \text{reveal } a. \text{withdraw } A$ and $D_2 = \text{after } t. \text{withdraw } B$. The contract precondition is $G = A : ! 1\text{\$} @ x \mid A : \text{secret } a \mid B : ! 0\text{\$} @ y$. Before running the compiler, A generates a random nonce s_a , and broadcasts its hash $h_a = H(s_a)$. Honest participants will choose a nonce of length greater than a public security parameter η . A large enough parameter (e.g., $\eta = 128$) ensures that the other participants cannot infer s_a (assuming H to be preimage resistant), nor its length. Further, A cannot later on reveal a different secret or a different length (assuming collision resistance). After that, both participants can compile $\{G\}TC$, obtaining the following transactions:

<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_{init}</th></tr> <tr><td style="padding: 2px;">in: $0 \mapsto T_x, 1 \mapsto T_y$</td></tr> <tr><td style="padding: 2px;">wit: $0 \mapsto \text{sig}_{\mathbf{K}(A)}, 1 \mapsto \text{sig}_{\mathbf{K}(B)}$</td></tr> <tr><td style="padding: 2px;">out: $(\lambda \zeta \beta. \text{versig}_{\mathbf{K}(D_1, \{A, B\})}(\zeta) \wedge H(\beta) = h_a \wedge \beta \geq \eta$ $\vee \text{versig}_{\mathbf{K}(D_2, \{A, B\})}(\zeta), 1\text{\\$})$</td></tr> </table>	T_{init}	in: $0 \mapsto T_x, 1 \mapsto T_y$	wit: $0 \mapsto \text{sig}_{\mathbf{K}(A)}, 1 \mapsto \text{sig}_{\mathbf{K}(B)}$	out: $(\lambda \zeta \beta. \text{versig}_{\mathbf{K}(D_1, \{A, B\})}(\zeta) \wedge H(\beta) = h_a \wedge \beta \geq \eta$ $\vee \text{versig}_{\mathbf{K}(D_2, \{A, B\})}(\zeta), 1\text{\$})$	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_{reveal}</th></tr> <tr><td style="padding: 2px;">in: T_{init}</td></tr> <tr><td style="padding: 2px;">wit: $\text{sig}_{\mathbf{K}(D_1, \{A, B\})} [s_a]$</td></tr> <tr><td style="padding: 2px;">out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}(\zeta), 1\text{\\$})$</td></tr> </table>	T_{reveal}	in: T_{init}	wit: $\text{sig}_{\mathbf{K}(D_1, \{A, B\})} [s_a]$	out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}(\zeta), 1\text{\$})$	
T_{init}										
in: $0 \mapsto T_x, 1 \mapsto T_y$										
wit: $0 \mapsto \text{sig}_{\mathbf{K}(A)}, 1 \mapsto \text{sig}_{\mathbf{K}(B)}$										
out: $(\lambda \zeta \beta. \text{versig}_{\mathbf{K}(D_1, \{A, B\})}(\zeta) \wedge H(\beta) = h_a \wedge \beta \geq \eta$ $\vee \text{versig}_{\mathbf{K}(D_2, \{A, B\})}(\zeta), 1\text{\$})$										
T_{reveal}										
in: T_{init}										
wit: $\text{sig}_{\mathbf{K}(D_1, \{A, B\})} [s_a]$										
out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}(\zeta), 1\text{\$})$										
<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_A</th></tr> <tr><td style="padding: 2px;">in: T_{reveal}</td></tr> <tr><td style="padding: 2px;">wit: $\text{sig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}$</td></tr> <tr><td style="padding: 2px;">out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(A)}(\zeta), 1\text{\\$})$</td></tr> </table>	T_A	in: T_{reveal}	wit: $\text{sig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}$	out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(A)}(\zeta), 1\text{\$})$	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><th style="text-align: center; padding: 2px;">T_B</th></tr> <tr><td style="padding: 2px;">in: T_{init}</td></tr> <tr><td style="padding: 2px;">wit: $\text{sig}_{\mathbf{K}(D_2, \{A, B\})}$ —</td></tr> <tr><td style="padding: 2px;">out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(B)}(\zeta), 1\text{\\$})$</td></tr> <tr><td style="padding: 2px;">absLock: t</td></tr> </table>	T_B	in: T_{init}	wit: $\text{sig}_{\mathbf{K}(D_2, \{A, B\})}$ —	out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(B)}(\zeta), 1\text{\$})$	absLock: t
T_A										
in: T_{reveal}										
wit: $\text{sig}_{\mathbf{K}(\text{withdraw } A, \{A, B\})}$										
out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(A)}(\zeta), 1\text{\$})$										
T_B										
in: T_{init}										
wit: $\text{sig}_{\mathbf{K}(D_2, \{A, B\})}$ —										
out: $(\lambda \zeta. \text{versig}_{\mathbf{K}(B)}(\zeta), 1\text{\$})$										
absLock: t										

Transaction T_{init} can be redeemed in two ways, according to the two clauses in the disjunction of its output script: (i) with both signatures $\text{sig}_{\mathbf{K}(D_1, A)}$ and $\text{sig}_{\mathbf{K}(D_1, B)}$ (corresponding to the reveal branch), or (ii) with both signatures $\text{sig}_{\mathbf{K}(D_2, A)}$ and $\text{sig}_{\mathbf{K}(D_2, B)}$ (corresponding to the after branch).

In case (i), also the secret value s_a must be provided in the wit field of T_{reveal} . As indicated by the square brackets around s_a in T_{reveal} : such value is *not* provided at compilation time, but added at runtime. Crucially, altering the wit field does not invalidate the signatures $\text{sig}_{\mathbf{K}(D_1, A)}$ and $\text{sig}_{\mathbf{K}(D_1, B)}$ on T_{reveal} (since signatures neglect the wit field), nor the actual identifier of T_{reveal} used in the in field of T_A . This relies on the *SegWit* feature (activated in August 2017), which allows to neglect witnesses in the computation of

transaction identifiers. Revealing the secret and appending T_{reveal} correspond to the following computation steps in BitML (time is omitted, because immaterial in this case):

$$\{A : a\#N\} \mid \langle TC, 1\mathbb{B} \rangle \rightarrow A : a\#N \mid \langle TC, 1\mathbb{B} \rangle \rightarrow \langle \text{withdraw } A, 1\mathbb{B} \rangle$$

After that, anyone can append the transaction T_A to the blockchain to transfer $1\mathbb{B}$ under A 's control. Note that once T_{reveal} is on the blockchain, it will be no longer possible to append T_B , since both transactions want to redeem T_{init} .

In case (ii), the `absLock` guarantees that T_B can be appended to the blockchain only after time t , coherently with the `after` clause in BitML. Indeed, appending T_B (which makes $1\mathbb{B}$ available to B) corresponds to the following step in BitML (where $t' \geq t$):

$$\langle TC, 1\mathbb{B} \rangle \mid t' \rightarrow \langle B, 1\mathbb{B} \rangle \mid t'$$

8 COHERENCE

Our computational soundness result is based on a correspondence between symbolic and computational runs, that we call *coherence*. Intuitively, a symbolic run R^s is coherent with a computational run R^c when each symbolic step in R^s is matched by the computational step corresponding to its implementation in R^c (in symbols, $R^s \sim_r R^c$, where r is the randomness source used by participants). We illustrate the coherence relation following the possible symbolic steps (see Appendix A.6 for the formal details).

Advertisement. Advertising a contract $\{G\}C$ in R^s is performed through the following step:

$$\Gamma \rightarrow \Gamma \mid \{G\}C$$

which requires that all the deposits mentioned in G (either persistent or volatile) occur in Γ . This step corresponds, in R^c , to broadcasting a bitstring which encodes the symbolic term $\{G\}C$. In the bitstring, the deposit names in G are encoded as the identifiers of their corresponding Bitcoin transaction outputs.

Stipulation: committing secrets. After advertisement, committing secrets is done in R^s performing a step such as:

$$\Gamma \mid \{G\}C \rightarrow \Gamma \mid \{G\}C \mid \{A : a\#N\} \mid A[\# \triangleright \{G\}C] = \Gamma_1$$

assuming, for simplicity, that G only requires $A : \text{secret } a$. In R^c , this step corresponds to a broadcast of a message $m(C, h, \vec{k})$ which comprises the encoding of the contract advertisement C , the hash h of the secret of A , and the sequence \vec{k} of all the public keys $K(D, B)$ for any D subterm of C , and any participant B occurring in G . Note that A can obtain \vec{k} by previously exchanging the public keys with other participants in R^c . Message m is then signed using key K_A . We further require that, in R^c , the value h was indeed generated by querying the oracle O with some bitstring of length $N + \eta$. If h is not generated through the oracle, or if the required length is shorter than η , the computational commitment would be coherent, instead, with the symbolic commitment $\{A : a\#\perp\}$, which models a dishonestly chosen secret.

Stipulation: authorizing deposits. The stipulation phase proceeds, in the symbolic model, by providing the authorizations to spend the persistent deposits required by G . This is done in R^s by performing a step such as:

$$\Gamma_1 \rightarrow \Gamma_1 \mid A[x \triangleright \{G\}C] = \Gamma_2$$

assuming, for simplicity, that G only requires $A : !v\mathbb{B} @ x$. In R^c , this step corresponds to the broadcast of a signature on the T_{init} transaction made with A 's key \vec{K}_A . This signature authorizes T_{init} to redeem the transaction output corresponding to the symbolic persistent deposit $\langle A, v\mathbb{B} \rangle_x$. The unsigned T_{init} transaction is obtained from the compiler, using the hashes and the keys in the commitment messages $m(C, h, \vec{k})$ broadcast in R^c by each participant. Note that an honest participant would sign T_{init} only *after* having exchanged with the other participants the signatures for all the other transactions generated by the compiler.

Stipulation: activating the contract. The stipulation phase is finalised by gathering all the required persistent deposits, and producing an active contract with an initial balance equal to their sum. This is done in R^s by performing a step such as:

$$\Gamma_2 \rightarrow \langle C, v\mathbb{B} \rangle \mid \Gamma_0 \quad \text{where } \Gamma = \Gamma_0 \mid \langle A, v\mathbb{B} \rangle_x$$

In R^c , this step corresponds to appending T_{init} to the blockchain.

Contract actions. In the symbolic run, once the active contract $\langle C, v\mathbb{B} \rangle$ is created, it can be executed by performing its actions, causing in R^s steps such as:

$$\langle C, v\mathbb{B} \rangle \rightarrow \langle C', v\mathbb{B} \rangle$$

In the computational run, each of these steps corresponds to appending to the blockchain a transaction. The blockchain initially contains a compiler-generated transaction T_C , with an unspent output (say, at index i) corresponding to C . Performing the computational step consists in appending another compiler-generated transaction $T_{C'}$, which redeems the output i of T_C . In turn, $T_{C'}$ has an output corresponding to C' . The case of `split` steps, where:

$$\langle C, v\mathbb{B} \rangle \rightarrow \langle C_1, v_1\mathbb{B} \rangle \mid \dots \mid \langle C_n, v_n\mathbb{B} \rangle$$

is similar: $T_{C'}$ will now have n outputs, corresponding to the n active contracts produced by the symbolic step.

Some symbolic steps can only be performed under certain conditions. For instance, $C = A : \tau.C'$ (recall that $\tau = \text{reveal } ()$) requires A 's authorization before it can move forward. In the symbolic model, this authorization is provided by the step:

$$\langle C, v\mathbb{B} \rangle_x \rightarrow \langle C, v\mathbb{B} \rangle_x \mid A[x \triangleright C]$$

In the computational setting, T_C can be redeemed with the signatures of *all* the participants. During the stipulation phase, all such signatures of $T_{C'}$ are exchanged, *except* for A 's one. Participant A can give her authorization at runtime by broadcasting such signature, coherently with the symbolic step above.

Another symbolic step which requires the intervention of a given participant is that of firing a `reveal` prefix. Indeed, in a symbolic run, $C = \text{reveal } a.C'$ can only proceed once the secret a has been revealed through a step:

$$\langle C, v \rangle \mid \{A : a\#N\} \rightarrow \langle C, v \rangle \mid A : a\#N$$

In the computational run, this corresponds to broadcasting a preimage of the hash value which was broadcast during the stipulation.

After. Symbolic delays trivially correspond to computational ones.

Deposits. Symbolic steps which manage deposits have an immediate computational counterpart. For instance, joining two deposits $\langle A, v \rangle_x$ and $\langle A, v' \rangle_y$ into a single deposit $\langle A, v + v' \rangle_z$ corresponds to appending to the blockchain a transaction with two inputs T_x and T_y and one output, with value $v + v'$. Its witnesses comprise two signatures with \hat{K}_A , and its output script verifies a signature with the same key. In the symbolic run, before the join action can be performed, A needs to authorize the spending of x and y , using two distinct steps. In the computational run, these steps correspond to two messages, where A broadcasts the two signatures.

Dually, dividing a deposit $\langle A, v + v' \rangle$ into two parts $\langle A, v \rangle$ and $\langle A, v' \rangle$ corresponds, in the computational run, to appending to the blockchain a transaction with a single input and two outputs. Transforming $\langle A, v \rangle_x$ into $\langle B, v \rangle_y$ (i.e., donating a deposit) corresponds, in the computational run, to appending to the blockchain a transaction having as input T_x , as witness a signature with \hat{K}_A , and a single output of value v , redeemable with \hat{K}_B .

Computational broadcasts. As described above, symbolic authorization or reveal steps correspond to computational broadcasts of suitable messages, such as signatures or hash preimages. Computational participants, however, can also broadcast other messages, e.g. for exchanging their public keys during stipulation. Further, dishonest computational participants can broadcast any arbitrary bitstring they can compute in PPTIME. Coherence ignores any computational broadcast which does not correspond to any of the above mentioned symbolic steps. That is, such broadcasts correspond to no symbolic steps. Ignoring these messages does not affect the security of contracts, because the other computational messages (in particular, the appended transactions) are enough to reconstruct, from a computational run, the BitML steps in the symbolic run.

Appending non-compiler-generated transactions. A subtle case in the definition of coherence is when *dishonest* participants append transactions to the blockchain. To illustrate the issue, suppose that a dishonest A owns $v\$, represented in the symbolic run as a term $\langle A, v \rangle_x$, and in the computational run as a transaction T_x redeemable with \hat{K}_A . Since A knows her key \hat{K}_A , she can sign an arbitrary transaction T' which redeems $v\$ from T_x , and append it to the blockchain. Crucially, T' could be a transaction that can never be generated by the BitML compiler. In such case, it is not possible to precisely match this computational step with a symbolic one. To obtain a correspondence, we let the appending of T' to be coherent with the symbolic *destruction* of the deposit $\langle A, v \rangle_x$, which makes it disappear from the symbolic configuration. In subsequent steps, coherence will ignore the descendants of T' in the computational run. While, in principle, this loss of information at the symbolic level could allow for computational attacks without a symbolic counterpart, in fact 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'' unrepresentable at the symbolic level, A can perform symbolically-representable 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 unrepresentable 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.

Note that, unlike A 's deposits, active contracts involving A can not be destroyed by A in the symbolic run. Hence, if in the computational run A can somehow redeem a transaction T which corresponds to an active contract, using a transaction T' which is not symbolically representable, then such computational step is *not* coherent with any symbolic step. This is intended, since in such case A succeeded in an attack which made the active contract deviate from its symbolic behavior.

Deposits and coherence. The following lemma ensures that deposits in a symbolic run R^s have a correspondent transaction output in any computational run R^c coherent with R^s . A similar correspondence also exists for active contracts (see Appendix A.6).

Lemma 1. *Let $R^s \sim_r R^c$. For each deposit $\langle A, v \rangle$ occurring in Γ_{R^s} , there exists a corresponding unspent transaction output in B_{R^c} with value v , redeemable with a signature with key \hat{K}_A .*

9 COMPUTATIONAL SOUNDNESS

To state the correctness of the BitML compiler, we need to describe how to convert any symbolic strategy Σ_A^s to a computational strategy $\Sigma_A^c = \mathfrak{N}(\Sigma_A^s)$, which realizes the symbolic behaviour in the computational model. Here, we provide the key intuition behind the construction (see Appendix A.7 for details). Strategy Σ_A^c receives as input a (stripped) computational run R^c . From this, Σ_A^c can reconstruct a (stripped) symbolic run R^s coherent with R^c . At this point, Σ_A^c runs Σ_A^s on R^s , obtaining a set of symbolic actions Λ^s . This is then converted to a set of corresponding computational moves Λ^c , so that performing any of the computational moves will produce an extension of R^c which is still *coherent* with R^s , possibly extended with one of the symbolic moves in Λ^s . This conversion closely follows the definition of coherence.

Our computational soundness result follows (see Appendix A.8 for its proof). We assume that honest participants have a symbolic strategy, and that their computational strategy is consequently obtained through the mapping \mathfrak{N} sketched above. 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 2 (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 $\Sigma_A^c = \mathfrak{N}(\Sigma_A^s)$ for all $A \in \text{Hon}$, including an arbitrary adversary strategy Σ_{Adv}^c . Fix $k \in \mathbb{N}$. Then, the following set has overwhelming probability:*

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

In the statement above, r represents the randomness used by all participants, including Adv . Formally, r maps each participant to an infinite string of independent and uniformly distributed bits. Once r is fixed, the behavior of the participants is deterministic, resulting in a single run. Hence, the probability that r drives a run R^c satisfying some property p , as done in the statement, can be seen as the probability that a random run R^c , sampled according to the

strategies, satisfies p . The statement considers only computational runs R^c having a *polynomial* length ($\leq \eta^k$) with respect to a security parameter η . This is crucial, since the adversary strategy is run at each computational step, hence in the whole run the adversary can run PPTIME algorithms for $|R^c|$ times. Without a polynomial bound on $|R^c|$, the adversary would be able to run algorithms outside PPTIME before the end of the run, breaking the underlying cryptographic primitives. Overall, the theorem states that if we run the computational strategies Σ^c using r , generating a polynomially-bounded run R^c conformant with Σ^c , then, with overwhelming probability, R^c is coherent with some symbolic run R^s . Further, this run can be obtained running the strategies Σ^s using the randomness $\pi_1(r)$. This stands for the sequence of the even-indexed bits in r (this is a technical artifact of our \mathfrak{N} construction).

Finally, note that Theorem 2 implicitly uses the compiler in two points: the translation \mathfrak{N} (the obtained computational strategies involve the compilation of the contracts used by the symbolic strategies), and the coherence relation \sim .

10 CONCLUSIONS

Our work bridges the gap between the cryptography community, where Bitcoin smart contracts have been investigated first, and the programming languages community. In particular, our computational soundness result 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.

Although BitML can express many of the smart contracts appeared in literature [15], it has some limitations. For instance, it cannot express contingent payments, where a participant A promises to pay B for a value x satisfying a predicate chosen by A (e.g., x is a prime factor of a given large number). Contingent payments can be implemented in Bitcoin similarly to timed commitments: A pays a deposit, which is taken by B after revealing a preimage of $H(x)$ which satisfies the predicate. An off-chain protocol [18] (which exploits zero-knowledge proofs) is used to guarantee that $H(x)$ is indeed the hash of a value x satisfying the predicate (note that, in the Bitcoin scripting language, one can only check trivial predicates, like e.g. equality). BitML could be extended to express contingent payments, by exploiting zero-knowledge proofs similar to those in [18] in the stipulation phase. This would allow our compiler to only generate standard Bitcoin transactions. Another kind of contracts which are not expressible in BitML are those for which one cannot pre-determine at compile time, a *finite* set of transactions, or of signatures, or of execution steps. This is the case, e.g., of crowdfunding contracts [15], where participants invest some money until a given threshold is reached. Here, we do not statically know neither the number of participants, nor their identities, so it is not possible to statically produce (and pre-sign) a set of transactions, as required by BitML. Extending BitML to express this kind of contracts — while preserving our compilation technique — would require suitable extensions of Bitcoin transactions. For instance, recursion could be obtained via extensions similar to those proposed in [49, 53].

REFERENCES

- [1] 2012. Bitcoin wiki - Contracts. <https://en.bitcoin.it/wiki/Contract>. (2012).
- [2] 2012. Bitcoin wiki - Contracts - Using external state. https://en.bitcoin.it/wiki/Contract#Example_4_Using_external_state. (2012).
- [3] 2015. Bitcoin developer guide - Escrow and arbitration. <https://bitcoin.org/en/developer-guide#escrow-and-arbitration>. (2015).
- [4] 2016. Understanding the DAO attack. (June 2016). <http://www.coindesk.com/understanding-dao-hack-journalists/>.
- [5] 2017. IVY. <https://docs.ivy-lang.org/bitcoin/>. (2017).
- [6] 2017. Parity Wallet Security Alert. (July 2017). <https://paritytech.io/blog/security-alert.html>.
- [7] 2017. A Postmortem on the Parity Multi-Sig Library Self-Destruct. (November 2017). <https://goo.gl/Kw3gXi>.
- [8] 2018. BALZaC: Bitcoin Abstract Language, analyZer and Compiler. <https://blockchain.unica.it/balzac/>. (2018).
- [9] Martin Abadi and Phillip Rogaway. 2007. Reconciling Two Views of Cryptography (The Computational Soundness of Formal Encryption). *J. Cryptology* 20, 3 (2007), 395. <https://doi.org/10.1007/s00145-007-0203-0>
- [10] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolic, Sharon Weed Cocco, and Jason Yellick. 2018. Hyperledger Fabric: a distributed operating system for permissioned blockchains. In *EuroSys*. 30:1–30:15. <https://doi.org/10.1145/3190508.3190538>
- [11] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. 2014. Fair Two-Party Computations via Bitcoin Deposits. In *Financial Cryptography Workshops (LNCS)*, Vol. 8438. Springer, 105–121. https://doi.org/10.1007/978-3-662-44774-1_8
- [12] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. 2014. Secure Multiparty Computations on Bitcoin. In *IEEE S & P*. 443–458. <https://doi.org/10.1109/SP.2014.35> First appeared on Cryptology ePrint Archive, <http://eprint.iacr.org/2013/784>.
- [13] Maria Apostolaki, Aviv Zohar, and Laurent Vanbever. 2017. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *IEEE Symp. on Security and Privacy*. 375–392. <https://doi.org/10.1109/SP.2017.29>
- [14] Nicola Atzei, Massimo Bartoletti, and Tiziana Cimoli. 2017. A Survey of Attacks on Ethereum Smart Contracts (SoK). In *Principles of Security and Trust (POST) (LNCS)*, Vol. 10204. Springer, 164–186. https://doi.org/10.1007/978-3-662-54455-6_8
- [15] Nicola Atzei, Massimo Bartoletti, Tiziana Cimoli, Stefano Lande, and Roberto Zunino. 2018. SoK: unraveling Bitcoin smart contracts. In *Principles of Security and Trust (POST) (LNCS)*, Vol. 10804. Springer, 217–242. <https://doi.org/10.1007/978-3-319-89722-6>
- [16] Nicola Atzei, Massimo Bartoletti, Stefano Lande, and Roberto Zunino. 2018. A formal model of Bitcoin transactions. In *Financial Cryptography and Data Security*.
- [17] Christian Badertscher, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. 2017. Bitcoin as a Transaction Ledger: A Composable Treatment. In *CRYPTO*. 324–356. https://doi.org/10.1007/978-3-319-63688-7_11
- [18] Wacław Banasik, Stefan Dziembowski, and Daniel Malinowski. 2016. Efficient Zero-Knowledge Contingent Payments in Cryptocurrencies Without Scripts. In *ESORICS (LNCS)*, Vol. 9879. Springer, 261–280. https://doi.org/10.1007/978-3-319-45741-3_14
- [19] Massimo Bartoletti and Livio Pompianu. 2017. An empirical analysis of smart contracts: platforms, applications, and design patterns. In *Financial Cryptography Workshops (LNCS)*, Vol. 10323. Springer, 494–509.
- [20] Massimo Bartoletti and Roberto Zunino. 2017. Constant-deposit multiparty lotteries on Bitcoin. In *Financial Cryptography Workshops (LNCS)*, Vol. 10323. Springer. <https://doi.org/10.1007/978-3-319-70278-0>
- [21] Mihir Bellare and Phillip Rogaway. 1993. Random Oracles Are Practical: A Paradigm for Designing Efficient Protocols. In *ACM Conference on Computer and Communications Security*. ACM, 62–73. <https://doi.org/10.1145/168588.168596>
- [22] Iddo Bentov and Ranjit Kumaresan. 2014. How to Use Bitcoin to Design Fair Protocols. In *CRYPTO (LNCS)*, Vol. 8617. Springer, 421–439. https://doi.org/10.1007/978-3-662-44381-1_24
- [23] Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cedric Fournet, Anitha Gollamudi, Georges Gonthier, Nadim Kobeissi, Aseem Rastogi, Thomas Sibut-Pinote, Nikhil Swamy, and Santiago Zanella-Beguelin. 2016. Formal Verification of Smart Contracts. In *PLAS*.
- [24] Alex Biryukov, Dmitry Khovratovich, and Sergei Tikhomirov. 2017. Findel: Secure Derivative Contracts for Ethereum. In *Financial Cryptography Workshops (LNCS)*, Vol. 10323. Springer, 453–467. https://doi.org/10.1007/978-3-319-70278-0_28
- [25] BitFury group. 2015. Smart Contracts on Bitcoin Blockchain. (2015). <http://bitfury.com/content/5-white-papers-research/contracts-1.1.1.pdf>.
- [26] Dan Boneh and Moni Naor. 2000. Timed Commitments. In *CRYPTO (LNCS)*, Vol. 1880. Springer, 236–254. <https://doi.org/10.1007/3-540-44598-6>

- [27] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A. Kroll, and Edward W. Felten. 2015. SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In *IEEE S & P*. 104–121. <https://doi.org/10.1109/SP.2015.14>
- [28] Vitalik Buterin. 2013. Ethereum: a next generation smart contract and decentralized application platform. <https://github.com/ethereum/wiki/wiki/White-Paper>. (2013).
- [29] Krishnendu Chatterjee, Amir Kafshdar Goharshady, and Yaron Velner. 2018. Quantitative Analysis of Smart Contracts. In *ESOP*. 739–767. https://doi.org/10.1007/978-3-319-89884-1_26
- [30] Karl Crary and Michael J. Sullivan. 2015. Peer-to-peer affine commitment using Bitcoin. In *ACM Conf. on Programming Language Design and Implementation*. 479–488. <https://doi.org/10.1145/2737924.2737997>
- [31] Christian Decker and Roger Wattenhofer. 2015. A Fast and Scalable Payment Network with Bitcoin Duplex Micropayment Channels. In *Stabilization, Safety, and Security of Distributed Systems (LNCS)*, Vol. 9212. Springer, 3–18. https://doi.org/10.1007/978-3-319-21741-3_1
- [32] Sergi Delgado-Segura, Cristina Pérez-Solà, Guillermo Navarro-Arribas, and Jordi Herrera-Joancomarti. 2017. A fair protocol for data trading based on Bitcoin transactions. *Future Generation Computer Systems* (2017). <https://doi.org/10.1016/j.future.2017.08.021>
- [33] Ittay Eyal and Emin Gün Sirer. 2014. Majority Is Not Enough: Bitcoin Mining Is Vulnerable. In *Financial Cryptography (LNCS)*, Vol. 8437. Springer, 436–454. https://doi.org/10.1007/978-3-662-45472-5_28
- [34] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. 2015. The Bitcoin Backbone Protocol: Analysis and Applications. In *EUROCRYPT (LNCS)*, Vol. 9057. Springer, 281–310. https://doi.org/10.1007/978-3-662-46803-6_10
- [35] David M. Goldschlag, Stuart G. Stubblebine, and Paul F. Syverson. 2010. Temporarily hidden bit commitment and lottery applications. *Int. J. Inf. Sec.* 9, 1 (2010), 33–50. <https://doi.org/10.1007/s10207-009-0094-1>
- [36] Ilya Grishchenko, Matteo Maffei, and Clara Schneidewind. 2018. Foundations and Tools for the Static Analysis of Ethereum Smart Contracts. In *CAV (LNCS)*, Vol. 10981. Springer, 51–78. https://doi.org/10.1007/978-3-319-96145-3_4
- [37] Ilya Grishchenko, Matteo Maffei, and Clara Schneidewind. 2018. A Semantic Framework for the Security Analysis of Ethereum Smart Contracts. In *Principles of Security and Trust (POST) (LNCS)*, Vol. 10804. Springer, 243–269. https://doi.org/10.1007/978-3-319-89722-6_10
- [38] Yoichi Hirai. 2017. Defining the Ethereum Virtual Machine for Interactive Theorem Provers. In *Financial Cryptography Workshops (LNCS)*, Vol. 10323. Springer, 520–535. https://doi.org/10.1007/978-3-319-70278-0_33
- [39] Simon L. Peyton Jones, Jean-Marc Eber, and Julian Seward. 2000. Composing contracts: an adventure in financial engineering, functional pearl. In *International Conference on Functional Programming (ICFP)*. 280–292. <https://doi.org/10.1145/351240.351267>
- [40] Ahmed E. Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. 2016. Hawk: The Blockchain Model of Cryptography and Privacy-Preserving Smart Contracts. In *IEEE Symp. on Security and Privacy*. 839–858. <https://doi.org/10.1109/SP.2016.55>
- [41] Ranjit Kumaresan and Iddo Bentov. 2014. How to Use Bitcoin to Incentivize Correct Computations. In *ACM CCS*. 30–41. <https://doi.org/10.1145/2660267.2660380>
- [42] Ranjit Kumaresan, Tal Moran, and Iddo Bentov. 2015. How to Use Bitcoin to Play Decentralized Poker. In *ACM CCS*. 195–206. <https://doi.org/10.1145/2810103.2813712>
- [43] Eric Lombrozo, Johnson Lau, and Pieter Wuille. 2015. Segregated Witness (Consensus layer). BIP 141, <https://github.com/bitcoin/bips/blob/master/bip-0141.mediawiki>.
- [44] Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. 2016. Making Smart Contracts Smarter. In *ACM CCS*. 254–269. <https://doi.org/10.1145/2976749.2978309>
- [45] Anastasia Mavridou and Aron Laszka. 2018. Designing Secure Ethereum Smart Contracts: A Finite State Machine Based Approach. In *Financial Cryptography and Data Security*.
- [46] Gregory Maxwell. 2016. The first successful Zero-Knowledge Contingent Payment. <https://bitcoincore.org/en/2016/02/26/zero-knowledge-contingent-payments-announcement/>. (2016).
- [47] Andrew Miller and Iddo Bentov. 2017. Zero-Collateral Lotteries in Bitcoin and Ethereum. In *EuroS&P Workshops*. 4–13. <https://doi.org/10.1109/EuroSPW.2017.44>
- [48] Andrew Miller, Iddo Bentov, Ranjit Kumaresan, and Patrick McCorry. 2017. Sprites: Payment Channels that Go Faster than Lightning. *CoRR abs/1702.05812* (2017). [arXiv:1702.05812 http://arxiv.org/abs/1702.05812](http://arxiv.org/abs/1702.05812)
- [49] Malte Möser, Ittay Eyal, and Emin Gün Sirer. 2016. Bitcoin covenants. In *Financial Cryptography Workshops (LNCS)*, Vol. 9604. Springer, 126–141. https://doi.org/10.1007/978-3-662-53357-4_9
- [50] Xavier Nicollin and Joseph Sifakis. 1991. An Overview and Synthesis on Timed Process Algebras. In *CAV*. 376–398. https://doi.org/10.1007/3-540-55179-4_36
- [51] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. 2002. *Isabelle/HOL: a proof assistant for higher-order logic*. Vol. 2283. Springer Science & Business Media.
- [52] Russell O’Connor. 2017. Simplicity: A New Language for Blockchains. In *PLAS*. <http://arxiv.org/abs/1711.03028>
- [53] Russell O’Connor and Marta Piekarska. 2017. Enhancing Bitcoin transactions with covenants. In *Financial Cryptography Workshops (LNCS)*, Vol. 10323. Springer. https://doi.org/10.1007/978-3-319-70278-0_12
- [54] Joseph Poon and Thaddeus Dryja. 2015. The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments. (2015). <https://lightning.network/lightning-network-paper.pdf>
- [55] Ilya Sergey, Amrit Kumar, and Aquinas Hobor. 2018. Scilla: a Smart Contract Intermediate-Level Language. *CoRR abs/1801.00687* (2018).
- [56] Nikhil Swamy, Catalin Hritcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean Karim Zinzindohoue, and Santiago Zanella Béguelin. 2016. Dependent types and multi-monadic effects in F*. In *POPL*. <https://doi.org/10.1145/2837614.2837655>
- [57] Paul F. Syverson. 1998. Weakly Secret Bit Commitment: Applications to Lotteries and Fair Exchange. In *IEEE CSFW*. 2–13. <https://doi.org/10.1109/CSFW.1998.683149>
- [58] Nick Szabo. 1997. Formalizing and Securing Relationships on Public Networks. *First Monday* 2, 9 (1997). <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/548>.
- [59] Petar Tsankov, Andrei Marian Dan, Dana Drachler Cohen, Arthur Gervais, Florian Buenzli, and Martin T. Vechev. 2018. Securify: Practical Security Analysis of Smart Contracts. *CoRR abs/1806.01143* (2018).

A APPENDIX

A.1 Additional BitML examples

We show a few additional examples of BitML contracts. Hereafter, in a term $\text{reveal } \vec{a}.C$, we abbreviate prefixes $\text{reveal } \vec{a}\vec{b}$ if p occurring in C as $\text{reveal } \vec{b}$ if p if some name in \vec{a} occurs in p .

Variable-refund escrow. We propose a variant of the escrow contract in Section 2 where M can issue a *partial* refund of ζv_B to A , and of $(1 - \zeta)v_B$ to B (similarly to [3, 25]). The possible values of ζ are given by a finite set Z in the range $[0, 1]$. We model the new contract as:

$$C = A : \text{withdraw } B + B : \text{withdraw } A + \sum_{\zeta \in Z} M : D_\zeta$$

$$D_\zeta = \text{split} (\zeta v \rightarrow \text{withdraw } A \mid (1 - \zeta)v \rightarrow \text{withdraw } B)$$

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

Intermediated payment. Assume that A wants to send an indirect payment of v_C to C , routing it through an intermediary B who can choose whether to authorize the payment, in this case retaining a fee v_B , 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 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.

Mutual timed commitment. We can also model *mutual* timed commitment as follows, where $t < t'$:

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

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 $2v\text{!}$.

Zero-collateral lottery. We show a two-players lottery which requires no collateral, similarly to [20, 47]. The preconditions just require the 1! bets and the secrets, while the contract is the following, where $d' > d$:

$$\begin{aligned} &\text{reveal } b \text{ if } 0 \leq |b| \leq 1. \left(\begin{aligned} &\text{reveal } a \text{ if } |a| = |b|. \text{withdraw } A \\ &+ \text{reveal } a \text{ if } |a| \neq |b|. \text{withdraw } B \\ &+ \text{after } d' : \text{withdraw } B \end{aligned} \right) \\ &+ \text{after } d : \text{withdraw } A \end{aligned}$$

Here, B must reveal first. If B does not reveal his secret by the deadline d , or the secret has not the expected length, then A can redeem 2! . Otherwise, A in turn must reveal by the deadline d' , or let B redeem 2! . If both A and B reveal, then the winner is determined by comparing the lengths of their secrets. As before, the rational strategy for each player is to choose a secret length 0 or 1, and reveal it. This makes the lottery fair, even in the absence of a collateral.

Rock-Paper-Scissors. Using similar insights, we can craft contracts for other games. For instance, consider *Rock-Paper-Scissors*, a two players hand game where both players choose simultaneously a hand-shape, and the winner is decided along with the following rules: rock beats scissors, scissors beats paper, and paper beats rock. We model the game for two players A and B who bet 1! each, and represent their moves as secrets of length 0 (rock), 1 (paper), and 2 (scissors). We define the following boolean predicate to determine the winner:

$$w(N, M) = (N = 0 \wedge M = 2) \vee (N = 2 \wedge M = 1) \vee (N = 1 \wedge M = 0)$$

The contract preconditions are:

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

while the contract is the following:

$$\begin{aligned} &\text{split} \left(\begin{aligned} &2\text{!} \rightarrow \text{reveal } b \text{ if } 0 \leq |b| \leq 2. \text{withdraw } B \\ &+ \text{after } d : \text{withdraw } A \\ &| 2\text{!} \rightarrow \text{reveal } a \text{ if } 0 \leq |a| \leq 2. \text{withdraw } A \\ &+ \text{after } d : \text{withdraw } B \\ &| 2\text{!} \rightarrow \text{reveal } ab \text{ if } w(|a|, |b|). \text{withdraw } A \\ &+ \text{reveal } ab \text{ if } w(|b|, |a|). \text{withdraw } B \\ &+ \text{reveal } ab \text{ if } |a| = |b|. \text{split}(1\text{!} \rightarrow \text{withdraw } A \\ &| 1\text{!} \rightarrow \text{withdraw } B) \end{aligned} \right) \end{aligned}$$

The contract is split in three parts, each with a balance of 2! : the first two parts allow the players to redeem the collaterals by revealing their secrets in time (similarly to the first version of the

lottery), while the third one computes the winner. The winner is A if $w(|a|, |b|)$, and B if $w(|b|, |a|)$. If a and b have the same length (i.e., they represent the same move), then there is a tie, so the bets are given back to the two players. Notice that if a player chooses a secret of unexpected length, then it may happen that the 2! in the third part of the `split` remain frozen. However, in such case the dishonest player will pay a 2! penalty to the other one. A zero-collateral version of *Rock-Paper-Scissors* can be obtained similarly to the second version of the lottery.

A.2 Supplementary material for Section 4

The semantics of BitML 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*.

Definition 5 (Configurations). We define (untimed) 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 $(\mid, 0)$ 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 6 (LTS of untimed configurations). The LTS of configurations is defined in Figure 3 (deposits) Figure 4 (advertisements and stipulation), and Figure 5 (contracts). The function cv from labels α to sets of names is defined as: $cv(\text{put}(\vec{x}, \vec{a}, y)) = cv(\text{withdraw}(A, v, y)) = cv(\text{split}(y)) = \{y\}$, otherwise $cv(\alpha)$ is empty. We further assume that a transition $\Gamma \rightarrow \Gamma \mid A[\chi]$ is possible only if $A[\chi]$ is not already present in Γ . \diamond

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 $A_i[\vec{x}, i \triangleright y]$ 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 $[\text{DEP-DESTROY}]$.

We discuss a few subtleties in the rules. First, $[\text{C-ADVERTISE}]$ 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

$\Gamma ::=$	untimed configuration	$\Gamma \mid t$	timed configuration ($t \in \mathbb{N}$)
0	empty	$\chi ::=$	authorization to ...
$\mid \{G\}C$	contract advertisement	$\# \triangleright \{G\}C$	commit secrets to stipulate $\{G\}C$
$\mid \langle C, v \rangle_x$	an active contract containing $v\mathfrak{B}$	$\mid x \triangleright \{G\}C$	spend x to stipulate $\{G\}C$
$\mid \langle A, v \rangle_x$	a deposit of $v\mathfrak{B}$ redeemable by A	$\mid x \triangleright D$	take branch D
$\mid A[\chi]$	authorization of A to perform χ	$\mid x, y \triangleright \langle A, v \rangle$	join deposit x with y into a deposit for A
$\mid \{A : a\#N\}$	committed secret of A ($N \in \mathbb{N} \cup \{\perp\}$)	$\mid x \triangleright \langle A, v \rangle, \langle A, v' \rangle$	divide a deposit x in two deposits for A
$\mid A : a\#N$	revealed secret of A ($N \in \mathbb{N}$)	$\mid x \triangleright B$	donate deposit x to B
$\mid \Gamma \mid \Gamma'$	parallel composition	$\mid \vec{x}, i \triangleright y$	destroy i -th deposits in \vec{x} through y

Figure 2: Configurations and authorizations.

for the side condition is that the correctness of our compiler will rely on this assumption.

Rule $[\text{C-AUTHCOMMIT}]$ 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.

Note that the semantics of predicates is well-defined, provided that there exists a unique $A : a\#N$ in Γ for each $|a|$ in the predicate. Our semantics of contracts ensures that is indeed the case.

Definition 7 (LTS of timed configurations). Timed configurations are terms $\Gamma \mid t$, where $t \in \mathbb{N}$ is the global time. The LTS between timed configurations is defined in Figure 6, where labels α are either untimed labels, or delays $\delta \in \mathbb{N}$. \diamond

Example 1 (Timed commitment). We show two computations of the timed commitment contract TC introduced in Section 2, using $G = A : !1\mathfrak{B} @ x \mid A : \text{secret } a \mid B : !0\mathfrak{B} @ y$ as precondition. Let $\Gamma = \langle A, 1\mathfrak{B} \rangle_x \mid \langle B, 0\mathfrak{B} \rangle_y$. Assuming that A is honest, a possible computation where A reveals her secret and then redeems the deposit is the following (here time is immaterial):

$$\Gamma \rightarrow \Gamma \mid \{G\}TC = \Gamma_1 \quad (1)$$

$$\rightarrow \Gamma_1 \mid \{A : a\#N\} \mid A[\# \triangleright \{G\}TC] = \Gamma_2 \quad (2)$$

$$\rightarrow \Gamma_2 \mid B[\# \triangleright \{G\}TC] = \Gamma_3 \quad (3)$$

$$\rightarrow \Gamma_3 \mid A[x \triangleright \{G\}TC] = \Gamma_4 \quad (4)$$

$$\rightarrow \Gamma_4 \mid B[y \triangleright \{G\}TC] \quad (5)$$

$$\rightarrow \langle TC, 1\mathfrak{B} \rangle_{x_1} \mid \{A : a\#N\} = \Gamma' \quad (6)$$

$$\rightarrow \langle TC, 1\mathfrak{B} \rangle_{x_1} \mid A : a\#N \quad (7)$$

$$\rightarrow \langle \text{withdraw } A, 1\mathfrak{B} \rangle_{x_2} \mid A : a\#N \quad (8)$$

$$\rightarrow \langle A, 1\mathfrak{B} \rangle_{x_3} \mid A : a\#N \quad (9)$$

Step (1) advertises $\{G\}TC$. This is possible because both deposits x and y (required by G) are available in Γ . At step (2), A commits to a secret: its length N is a natural, since A is honest. At step (3) also B does his commitment (empty, since G does not require any secrets from B). At steps (4)-(5), A and B give their authorization to stipulate TC , by providing their authorizations to spend the deposits x and y , respectively. At step (6) the contract TC becomes stipulated.

After this step, the bitcoins in the deposits x and y are transferred to the contract. At step (7), A reveals her secret (and consequently, also its length N). After that, the action `reveal` a is performed at step (8), reducing the contract to `withdraw` A , and discarding the `after` branch. Finally, step (9) performs the `withdraw` A action, producing a fresh deposit x_3 with $1\mathfrak{B}$ redeemable by A .

We also show a computation where A does not reveal her secret, and B waits until $t' > t$ to redeem A 's deposit. Starting from the configuration Γ' at time 0, we have the following steps:

$$\Gamma' \mid 0 \rightarrow \Gamma' \mid t' \rightarrow \langle B, 1\mathfrak{B} \rangle_y \mid \{A : a\#N\} \mid t'$$

The first step lets the time pass, by rule $[\text{DELAY}]$. In the second step, B fires the prefix `withdraw` B within the `after`, and in this way he collects $1\mathfrak{B}$. This is obtained by using rule $[\text{C-WITHDRAW}]$ in the premise of $[\text{C-TIMEOUT}]$. \diamond

A.3 Supplementary material for Section 5

If a symbolic run R^s is finite, we write Γ_{R^s} for its last untimed configuration, and δ_{R^s} for the last time. We write $R^s \xrightarrow{\alpha} \dot{R}^s$ when \dot{R}^s extends R^s with the transition $(\Gamma_{R^s}, \delta_{R^s}) \xrightarrow{\alpha} (\Gamma_{\dot{R}^s}, \delta_{\dot{R}^s})$.

Definition 8 (Stripping of symbolic runs). We denote with $\text{strip}(R^s)$ 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

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

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

Lemma 3. *Let $\alpha \neq A : a$, for any A and a . (i) if $R^s \xrightarrow{\alpha} \dot{R}^s$, then $\text{strip}(R^s) \xrightarrow{\alpha} \text{strip}(\dot{R}^s)$; (ii) if $\text{strip}(R^s) \xrightarrow{\alpha} R_*^s$, then $R^s \xrightarrow{\alpha} \dot{R}^s$, for some \dot{R}^s such that $\text{strip}(\dot{R}^s) = \text{strip}(R_*^s)$.*

Definition 10 (Symbolic participant strategies). The symbolic strategy of a participant $A \in \text{Hon}$ is a PPTIME algorithm $\Sigma_A^s(R_*^s, r_A)$,

$$\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} \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{\text{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} \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{\text{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} \text{[DEP-AUTHDONATE]} \quad \frac{\Gamma = A[x \triangleright B] \mid \Gamma' \quad y \text{ fresh}}{\langle A, v \rangle_x \mid \Gamma \xrightarrow{\text{donate}(x,B)} \langle B, v \rangle_y \mid \Gamma'} \text{[DEP-DONATE]} \\
\frac{\vec{x} = x_1 \cdots x_n \quad j \in 1..n \quad y \text{ fresh (except in destroy authorizations for } \vec{x})}{\left(\parallel_{i \in 1..n} \langle A_i, v_i \rangle_{x_i} \right) \mid \Gamma \xrightarrow{A_j:\vec{x},j} \left(\parallel_{i \in 1..n} \langle A_i, v_i \rangle_{x_i} \right) \mid A_j[\vec{x}, j \triangleright y] \mid \Gamma} \text{[DEP-AUTHDESTROY]} \\
\frac{\vec{x} = x_1 \cdots x_n \quad \Gamma = \left(\parallel_{i \in 1..n} A_i[\vec{x}, i \triangleright y] \right) \mid \Gamma'}{\left(\parallel_{i \in 1..n} \langle A_i, v_i \rangle_{x_i} \right) \mid \Gamma \xrightarrow{\text{destroy}(\vec{x})} \Gamma'} \text{[DEP-DESTROY]}
\end{array}$$

Figure 3: Semantics of untimed configurations: rules for deposits.

$$\frac{\begin{array}{l} \{G\}C \text{ contains at least one participant in Hon} \quad \Gamma \text{ contains } \langle A_i, v_i \rangle_{x_i} \text{ for all } A_i : ! v_i \otimes x_i \text{ in } \{G\}C \\ a \text{ fresh, for each } A : \text{secret } a \text{ in } G \quad \Gamma \text{ contains } \langle A_i, v_i \rangle_{x_i} \text{ for all } A_i : ? v_i \otimes x_i \text{ in } \{G\}C \end{array}}{\Gamma \xrightarrow{\text{advertise}(\{G\}C)} \{G\}C \mid \Gamma} \text{[C-ADVERTISE]}$$

$$\frac{\begin{array}{l} a_1 \cdots a_k \text{ secrets of } A \text{ in } G \\ \forall i \in 1..k : \#N : \{A : a_i \# N\} \in \Gamma \quad \forall i \in 1..k : N_i \in \begin{cases} \mathbb{N} & \text{if } A \in \text{Hon} \\ \mathbb{N} \cup \{\perp\} & \text{otherwise} \end{cases} \\ \Delta = \{A : a_1 \# N_1\} \mid \cdots \mid \{A : a_k \# N_k\} \end{array}}{\{G\}C \mid \Gamma \xrightarrow{A:\{G\}C,\Delta} \{G\}C \mid \Gamma \mid \Delta \mid A[\# \triangleright \{G\}C]} \text{[C-AUTHCOMMIT]}$$

$$\frac{\Gamma \text{ contains } B[\# \triangleright \{G\}C] \text{ for all } B \text{ in } G \quad G = A : ! v \otimes x \mid \cdots}{\{G\}C \mid \Gamma \xrightarrow{A:\{G\}C,x} \{G\}C \mid \Gamma \mid A[x \triangleright \{G\}C]} \text{[C-AUTHINT]}$$

$$\frac{G = \left(\parallel_{i \in I} A_i : ! v_i \otimes x_i \right) \mid \left(\parallel_{i \in J} B_i : ? v'_i \otimes y_i \right) \mid \left(\parallel_{i \in K} C_i : \text{secret } a_i \right) \quad x \text{ fresh}}{\{G\}C \mid \Gamma \mid \left(\parallel_{i \in I} \langle A_i, v_i \rangle_{x_i} \right) \mid \left(\parallel_{A \in G} A[\# \triangleright \{G\}C] \right) \mid \left(\parallel_{i \in I} A_i[x_i \triangleright \{G\}C] \right) \xrightarrow{\text{init}(G,C)} \langle C, \sum_{i \in I} v_i \rangle_x \mid \Gamma} \text{[C-INT]}$$

Figure 4: Semantics of untimed configurations: rules for advertisement and stipulation.

$A, B, \dots \in \text{Part}$	Participants	Γ, Δ, \dots	Configuration (untimed)	r	Random generator
x, y, \dots	Deposit names	$\Gamma \mid t$	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	$R_s^s, \hat{R}_s^s, \dots$	Stripped symb. runs	R^c, \hat{R}^c, \dots	Computational runs
G, G', \dots	Preconditions	λ^s	Symbolic label	$R_s^c, \hat{R}_s^c, \dots$	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	$\mathbb{B}_{\text{adv}}(\cdot)$	Compiler	Σ_A^c	Computational strategy
				$\mathfrak{N}(\cdot)$	Strategy mapping

Table 1: Summary of notation.

$$\begin{array}{c}
\frac{\vec{v} = v_1 \cdots v_k \quad \vec{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 } \vec{v} \rightarrow \vec{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}{\langle A : a \# N \rangle \xrightarrow{A:a} A : a \# N} \text{[C-AUTHREV]} \quad \frac{\vec{x} = x_1 \cdots x_m \quad \Gamma = \|_{i=1}^m \langle A_i, v_i \rangle_{x_i} \quad z \text{ fresh} \quad \vec{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 } \vec{x} \ \& \ \text{reveal } \vec{a} \ \text{if } p. C, v \rangle_y \mid \Gamma \mid \Delta \mid \Gamma' \xrightarrow{\text{put}(\vec{x}, \vec{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 A[x \triangleright D] \mid \Gamma} \text{[C-AUTHCONTROL]} \\
\frac{D \equiv A_1 : \cdots : A_k : \text{after } t_1 : \cdots : \text{after } t_m : D' \quad D' \neq A : \cdots \quad D' \neq \text{after } t : \cdots \quad \langle D', v \rangle_x \mid \Gamma \xrightarrow{\alpha} \Gamma' \quad x \in cv(\alpha)}{\langle D + C, v \rangle_x \mid \|_{i=1}^k A_i[x \triangleright D] \mid \Gamma \xrightarrow{\alpha} \Gamma'} \text{[C-CONTROL]} \\
\llbracket \text{true} \rrbracket_\Gamma = \text{true} \quad \llbracket p_1 \wedge p_2 \rrbracket_\Gamma = \llbracket p_1 \rrbracket_\Gamma \text{ and } \llbracket p_2 \rrbracket_\Gamma \quad \llbracket \neg p \rrbracket_\Gamma = \text{not } \llbracket p \rrbracket_\Gamma \quad \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 \quad \llbracket |a| \rrbracket_\Gamma = N \text{ if } \Gamma \text{ contains } A : a \# N \quad \llbracket E_1 \bullet E_2 \rrbracket_\Gamma = \llbracket E_1 \rrbracket_\Gamma \bullet \llbracket E_2 \rrbracket_\Gamma \quad (\bullet \in \{+, -\})
\end{array}$$

Figure 5: Semantics of untimed configurations: rules for contracts.

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

Figure 6: Semantics of timed configurations.

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 $B : \cdots \in \Sigma_A^s(\text{strip}(R^s), r_A)$, 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.

Definition 11 (Symbolic adversary strategies). A symbolic adversary strategy is a PPTIME algorithm $\Sigma_{\text{Adv}}^s(R_*^s, r_{\text{Adv}}, \vec{\Lambda}^s)$, taking as input a stripped symbolic run R_*^s , a random sequence r_{Adv} , and a list $\vec{\Lambda}^s = \Lambda_1^s \cdots \Lambda_k^s$ 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 $\Sigma_{\text{Adv}}^s(\text{strip}(R^s), r_{\text{Adv}}, \vec{\Lambda}^s) = \lambda^s$, one of the following holds:

- (1) $\lambda^s = (A_i, j)$, $\Lambda_i^s = \alpha_1 \cdots \alpha_m$, and $\alpha_j = A_i : \cdots$;

- (2) $\lambda^s = \alpha$, $\alpha \neq A : \cdots$ for any A , $\alpha \notin \mathbb{N}$, and $R^s \xrightarrow{\alpha}$;
- (3) $\lambda^s = B : \cdots$, $\lambda^s \neq B : a$ 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 $\text{strip}(R^s)$, we have $\Sigma_{\text{Adv}}^s(R_*^s, r_{\text{Adv}}, \epsilon) = B : \{G\}C, \Delta$ where $\{B : a \# N\}$ in Δ , for some $N \neq \perp$ and $\{G\}C$.

If $\Sigma_{\text{Adv}}^s(\text{strip}(R^s), r_{\text{Adv}}, \vec{\Lambda}^s) = B : \{G\}C, \Delta$ for $B \notin \text{Hon}$, we ask that $\Sigma_{\text{Adv}}^s(\text{strip}(R^s), r_{\text{Adv}}, \epsilon) = B : \{G\}C, \Delta$. \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 $\vec{\Lambda}^s$ generated at the point R_*^s . We consequently require that the N chosen by Σ_{Adv}^s does not depend on $\vec{\Lambda}^s$. This restriction does *not* limit the power of Adv , who can first perform a sequence of actions λ^s depending on $\vec{\Lambda}^s$, appending them to R^s , and then choose N depending on such actions λ^s .

Lemma 4. Let $R_*^s = \text{strip}(R^s)$ for some symbolic run R^s . Also, let $\text{Hon} = \{A_1, \dots, A_k\}$, and let $\vec{\Lambda}^s = \Lambda_1^s \cdots \Lambda_k^s$, where $\Lambda_i^s = \Sigma_{A_i}^s(R_*^s, r_{A_i})$ for $i \in 1..k$. If $\Sigma_{\text{Adv}}^s(R_*^s, r, \vec{\Lambda}^s) = \lambda^s$, then:

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

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) $R^s = \Gamma \mid 0$, 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 $\vec{\Lambda}_*^s = \text{strip}(\Lambda_1^s \cdots \Lambda_k^s)$, if $\lambda^s = \Sigma_{Adv}^s(\text{strip}(\hat{R}^s), r_{Adv}, \vec{\Lambda}_*^s)$ 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 $(\Sigma^s \Sigma_{Adv}^s, r)$. \diamond

A.4 Supplementary material for Section 6

Key pairs. We associate each A with two key pairs: 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 $K_{B_i}(r_{Adv})$, $\hat{K}_{B_i}(r_{Adv})$ for the key pairs generated from r_{Adv} . Given any key pair K , we denote with K^P its public part, and with K^S its private part.

Security parameter. Our model uses PPTIME algorithms w.r.t. a security parameter $\eta \in \mathbb{N}$. We extend the randomness source also to \mathbf{O} , assuming $r_{\mathbf{O}} \in \{0, 1\}^\omega$ to be defined. For simplicity, we assume that all communications (except those involving \mathbf{O}) are (reliable) broadcasts. The security parameter η is such that: (i) $H(m)$ comprises η random bits. (ii) the i -th key pair is generated using the i -th 2η bits from r_{Adv} .

Definition 13 (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:

$$\begin{aligned} T_0 \cdots A_i \rightarrow * : (K_{A_i}^P(r_{A_i}), \hat{K}_{A_i}^P(r_{A_i})) \cdots & \quad (\forall A_i \in \text{Hon}) \\ \cdots B_j \rightarrow * : (K_{B_j}^P(r_{Adv}), \hat{K}_{B_j}^P(r_{Adv})) \cdots & \quad (\forall B_j \notin \text{Hon}) \end{aligned}$$

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 . \diamond

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.

Definition 14 (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:

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

Borrowing from [16], we write $\mathbf{B} \triangleright (\mathbf{T}, t)$ when (\mathbf{T}, t) consistently updates \mathbf{B} .

Definition 15 (Blockchain of a computational run). For each R^c , we define the blockchain \mathbf{B}_{R^c} inductively as follows:

$$\mathbf{B}_{T_0} = (T_0, 0) \quad \mathbf{B}_{R^c \lambda^c} = \begin{cases} \mathbf{B}_{R^c}(\mathbf{T}, \delta_{R^c}) & \text{if } \lambda^c = \mathbf{T} \\ \mathbf{B}_{R^c} & \text{otherwise} \end{cases} \quad \diamond$$

Definition 16 (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) $\lambda^c = \mathbf{A} \rightarrow * : m$ or $\lambda^c = \mathbf{A} \rightarrow \mathbf{O} : m$, for some m ;
- (2) $\lambda^c = \mathbf{T}$, where $\mathbf{B}_{R_*^c} \triangleright (\mathbf{T}, \delta_{R_*^c})$, where in R_*^c we can find (in this order): (i) all the inputs of \mathbf{T} ; (ii) a message $\mathbf{B} \rightarrow * : \mathbf{T}$, for some \mathbf{B} ; (iii) for each witness w in \mathbf{T} , a message $\mathbf{B} \rightarrow * : w$, for some \mathbf{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:

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

Persistency ensures that, if A at a certain point wants to perform some λ^c (sending m or putting \mathbf{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 \mathbf{T} , one has to make both \mathbf{T} and its witnesses public. This models the fact that, before \mathbf{T} actually appears on the blockchain, it has to be broadcast on the Bitcoin network, potentially enabling an adversary to know \mathbf{T} beforehand.

Definition 17 (Computational adversary strategies). A computational adversary strategy is a PPTIME algorithm $\Sigma_{Adv}^c(R_*^c, r_{Adv}, \vec{\Lambda}^c)$, taking as input a (Adv -stripped) computational run R_*^c , a random sequence r_{Adv} , and a list $\vec{\Lambda}^c$ 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 $\Sigma_{Adv}^c(R_*^c, r_{Adv}, \vec{\Lambda}^c) = \lambda^c$, one of the following holds:

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

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

Further, in both Items 3 and 4, given n the number of distinct queries to 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_{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 O , modelling the hash functionality as in the random oracle model. We let queries have higher priority than other actions.

Lemma 5. *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) .*

A.5 Supplementary material for Section 7

We compile a contract advertisement $\{G\}C$ into a sequence of standard Bitcoin transactions. 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:

- $K(A)$, for each $A \in \text{PartG}$;
- $K(D, A)$, 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 $K(D, \mathcal{P})$ the set of key pairs $\{K(D, A_1), \dots, K(D, A_n)\}$.

Definition 19 (BitML compiler). The function $\mathbb{B}_{\text{adv}}(\cdot)$ from contract advertisements to sequences of Bitcoin transactions is defined by the rules in Figure 7. \diamond

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 2. We compile the timed commitment $\{G\}TC$ of Section 2, 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 8, where:

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

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 8. Doing this, A commits to her secret, whose hash h_a occurs in the output script $T_{\text{init}}.\text{out}$. 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

A.6 Supplementary material for Section 8

Hereafter, we assume that at least one of the participants in $\{G\}C$ has a secret a in G . Since rule $[\text{C-ADVERTISE}]$ 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.

$G = (\ _{i \in I} A_i : ? v_i \otimes x_i) \mid (\ _{i \in J} B_i : ! v'_i \otimes y_i) \mid (\ _{i \in K} C_i : \text{secret } a_i)$ $C = \sum_{i=1}^m D_i \quad v = \sum_{i \in J} v'_i$ $e_i = \mathbb{B}_{\text{out}}(D_i) \ (\forall i \in 1..m) \quad \vec{x} = \uplus_{i=1}^m \text{fv}(e_i)$ $\mathcal{T}_i = \mathbb{B}_D(D_i, D_i, T_{\text{init}}, 0, v, \text{PartG}, 0) \ (\forall i \in 1..m)$	<table border="1"> <tr><th>T_{init}</th></tr> <tr><td>in: $i \mapsto \text{txout}(y_i) \ (\forall i \in J)$</td></tr> <tr><td>wit: \perp</td></tr> <tr><td>out: $(\lambda \vec{x}. \bigvee_{i=1}^m e_i, v)$</td></tr> </table>	T_{init}	in: $i \mapsto \text{txout}(y_i) \ (\forall i \in J)$	wit: \perp	out: $(\lambda \vec{x}. \bigvee_{i=1}^m e_i, v)$
T_{init}					
in: $i \mapsto \text{txout}(y_i) \ (\forall i \in J)$					
wit: \perp					
out: $(\lambda \vec{x}. \bigvee_{i=1}^m e_i, v)$					

$$\mathbb{B}_{\text{adv}}(\{G\}C) = T_{\text{init}} \mathcal{T}_1 \cdots \mathcal{T}_m$$

$C = \sum_{i=1}^m D_i \quad I = \{z_1, \dots, z_k\}$ $e_i = \mathbb{B}_{\text{out}}(D_i) \ (\forall i \in 1..m)$ $\vec{x} = \uplus_{i=1}^m \text{fv}(e_i)$ $\mathcal{T}_i = \mathbb{B}_D(D_i, D_i, T_C, 0, v, \text{PartG}, 0) \ (\forall i \in 1..m)$	<table border="1"> <tr><th>T_C</th></tr> <tr><td>in: $0 \mapsto (T, o), i \mapsto \text{txout}(z_i) \ (\forall i \in 1..k)$</td></tr> <tr><td>wit: $0 \mapsto \text{sig}_{\mathbb{K}(D_p, \mathcal{P})}, i \mapsto \text{sig}_{\mathbb{K}(\text{part}(z_i))} \ (\forall i \in 1..k)$</td></tr> <tr><td>out: $(\lambda \vec{x}. \bigvee_{i=1}^m e_i, v)$</td></tr> <tr><td>absLock: t</td></tr> </table>	T_C	in: $0 \mapsto (T, o), i \mapsto \text{txout}(z_i) \ (\forall i \in 1..k)$	wit: $0 \mapsto \text{sig}_{\mathbb{K}(D_p, \mathcal{P})}, i \mapsto \text{sig}_{\mathbb{K}(\text{part}(z_i))} \ (\forall i \in 1..k)$	out: $(\lambda \vec{x}. \bigvee_{i=1}^m e_i, v)$	absLock: t
T_C						
in: $0 \mapsto (T, o), i \mapsto \text{txout}(z_i) \ (\forall i \in 1..k)$						
wit: $0 \mapsto \text{sig}_{\mathbb{K}(D_p, \mathcal{P})}, i \mapsto \text{sig}_{\mathbb{K}(\text{part}(z_i))} \ (\forall i \in 1..k)$						
out: $(\lambda \vec{x}. \bigvee_{i=1}^m e_i, v)$						
absLock: t						

$$\mathbb{B}_C(C, D_p, T, o, v, I, \mathcal{P}, t) = T_C \mathcal{T}_1 \cdots \mathcal{T}_m$$

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

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

$$D = \text{withdraw } A$$

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

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

$$D = A : D'$$

$$\mathbb{B}_D(D, D_p, T, o, v, \mathcal{P}, t) = \mathbb{B}_D(D', D_p, T, o, v, \mathcal{P} \setminus \{A\}, t)$$

$$D = \text{after } t' : D'$$

$$\mathbb{B}_D(D, D_p, T, o, v, \mathcal{P}, t) = \mathbb{B}_D(D', D_p, T, o, v, \mathcal{P}, \max\{t, t'\})$$

$\vec{C} = C_1 \cdots C_n \quad C_i = \sum_{j=1}^{k_i} D_{i,j} \ (\forall i \in 1..n)$ $\vec{v} = v_1 \cdots v_n \quad e_{i,j} = \mathbb{B}_{\text{out}}(D_{i,j}) \ (\forall i \in 1..n, j \in 1..k_i)$ $\vec{x}_i = \uplus_{j=1}^{k_i} \text{fv}(e_{i,j}) \ (\forall i \in 1..n)$ $\mathcal{T}_{i,j} = \mathbb{B}_D(D_{i,j}, D_{i,j}, T_C, i-1, v_i, \text{PartG}, 0) \ (\forall i \in 1..n, j \in 1..k_i)$	<table border="1"> <tr><th>T_C</th></tr> <tr><td>in: (T, o)</td></tr> <tr><td>wit: $\text{sig}_{\mathbb{K}(D_p, \mathcal{P})}$</td></tr> <tr><td>out: $i-1 \mapsto (\lambda \vec{x}_i. \bigvee_{j=1}^{k_i} e_{i,j}, v_i) \ (\forall i \in 1..n)$</td></tr> <tr><td>absLock: t</td></tr> </table>	T_C	in: (T, o)	wit: $\text{sig}_{\mathbb{K}(D_p, \mathcal{P})}$	out: $i-1 \mapsto (\lambda \vec{x}_i. \bigvee_{j=1}^{k_i} e_{i,j}, v_i) \ (\forall i \in 1..n)$	absLock: t
T_C						
in: (T, o)						
wit: $\text{sig}_{\mathbb{K}(D_p, \mathcal{P})}$						
out: $i-1 \mapsto (\lambda \vec{x}_i. \bigvee_{j=1}^{k_i} e_{i,j}, v_i) \ (\forall i \in 1..n)$						
absLock: t						

$$\mathbb{B}_{\text{par}}(\vec{C}, D_p, T, o, \vec{v}, \mathcal{P}, t) = T_C(\mathcal{T}_{i,j})_{i \in 1..n, j \in 1..k_i}$$

$$\mathbb{B}(\text{true}) = \text{true} \quad \mathbb{B}(p_1 \wedge p_2) = \mathbb{B}(p_1) \wedge \mathbb{B}(p_2) \quad \mathbb{B}(\neg p) = \neg \mathbb{B}(p) \quad \mathbb{B}(e_1 \circ e_2) = \mathbb{B}(e_1) \circ \mathbb{B}(e_2)$$

$$\mathbb{B}(N) = N \quad \mathbb{B}(|a|) = |a| - \eta \quad \mathbb{B}(e_1 \bullet e_2) = \mathbb{B}(e_1) \bullet \mathbb{B}(e_2)$$

Figure 7: Compiling contracts to Bitcoin transactions.

<table border="1"> <tr><th>T_{init}</th></tr> <tr><td>in: $0 \mapsto (T_A, 0), 1 \mapsto (T_B, 0)$</td></tr> <tr><td>wit: \perp</td></tr> <tr><td>out: $(\lambda \zeta_A \zeta_B b. e_1 \vee e_2, v)$</td></tr> </table>	T_{init}	in: $0 \mapsto (T_A, 0), 1 \mapsto (T_B, 0)$	wit: \perp	out: $(\lambda \zeta_A \zeta_B b. e_1 \vee e_2, v)$	<table border="1"> <tr><th>T'</th></tr> <tr><td>in: $(T_{\text{init}}, 0)$</td></tr> <tr><td>wit: $\text{sig}_{\mathbb{K}(D_1, \{A, B\})}$</td></tr> <tr><td>out: $(\lambda \zeta_A \zeta_B b. e', v)$</td></tr> </table>	T'	in: $(T_{\text{init}}, 0)$	wit: $\text{sig}_{\mathbb{K}(D_1, \{A, B\})}$	out: $(\lambda \zeta_A \zeta_B b. e', v)$	<table border="1"> <tr><th>T'_A</th></tr> <tr><td>in: $(T', 0)$</td></tr> <tr><td>wit: $\text{sig}_{\mathbb{K}(\text{withdraw } A, \{A, B\})}$</td></tr> <tr><td>out: $(\lambda \zeta. \text{versig}_{\mathbb{K}(A)}(\zeta), v)$</td></tr> </table>	T'_A	in: $(T', 0)$	wit: $\text{sig}_{\mathbb{K}(\text{withdraw } A, \{A, B\})}$	out: $(\lambda \zeta. \text{versig}_{\mathbb{K}(A)}(\zeta), v)$	<table border="1"> <tr><th>T'_B</th></tr> <tr><td>in: $(T_{\text{init}}, 0)$</td></tr> <tr><td>wit: $\text{sig}_{\mathbb{K}(D_2, \{A, B\})}$</td></tr> <tr><td>out: $(\lambda \zeta. \text{versig}_{\mathbb{K}(B)}(\zeta), v)$</td></tr> <tr><td>absLock: t</td></tr> </table>	T'_B	in: $(T_{\text{init}}, 0)$	wit: $\text{sig}_{\mathbb{K}(D_2, \{A, B\})}$	out: $(\lambda \zeta. \text{versig}_{\mathbb{K}(B)}(\zeta), v)$	absLock: t
T_{init}																				
in: $0 \mapsto (T_A, 0), 1 \mapsto (T_B, 0)$																				
wit: \perp																				
out: $(\lambda \zeta_A \zeta_B b. e_1 \vee e_2, v)$																				
T'																				
in: $(T_{\text{init}}, 0)$																				
wit: $\text{sig}_{\mathbb{K}(D_1, \{A, B\})}$																				
out: $(\lambda \zeta_A \zeta_B b. e', v)$																				
T'_A																				
in: $(T', 0)$																				
wit: $\text{sig}_{\mathbb{K}(\text{withdraw } A, \{A, B\})}$																				
out: $(\lambda \zeta. \text{versig}_{\mathbb{K}(A)}(\zeta), v)$																				
T'_B																				
in: $(T_{\text{init}}, 0)$																				
wit: $\text{sig}_{\mathbb{K}(D_2, \{A, B\})}$																				
out: $(\lambda \zeta. \text{versig}_{\mathbb{K}(B)}(\zeta), v)$																				
absLock: t																				

Figure 8: Transactions obtained by compiling the timed commitment contract.

Definition 20 (Coherence). We inductively define the relation $\text{coher}(R^s, R^c, r, \text{txout}, \text{sechash}, \kappa)$, 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: $\text{coher}(R^s, R^c, r, \text{txout}, \text{sechash}, \kappa)$ holds if all the following conditions hold: (i) $R^s = \Gamma_0 \mid 0$, with Γ_0 initial; (ii) $R^c = T_0 \dots$ initial; (iii) all the public keys in R^c are generated from r , according to Definition 13; (iv) txout maps exactly the x of $\langle A, v \rangle_x$ in Γ_0 to an output in T_0 of value $v\vec{b}$, and spendable with $\hat{K}_A(r_A)$; (v) $\text{dom sechash} = \emptyset$; (vi) $\text{dom } \kappa = \emptyset$.

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

- (1) $\alpha = \text{advertise}(\{G\}C, \mathcal{L}^c = A \rightarrow * : C)$, where C is obtained by encoding $\{G\}C$ as a bitstring, representing each x in it as the transaction output $\text{txout}'(x)$. Further, $\text{txout}' = \text{txout}$, $\text{sechash}' = \text{sechash}$, 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 $B, \mathcal{L}^c = B \rightarrow * : (C, \vec{h}, \vec{k})$ (signed by A), where \vec{h} is a sequence comprising a bitstring h_i with $|h_i| = \eta$ for each secret a_i in Δ , and \vec{k} is a sequence of keys, as the one produced by the stipulation protocol. We require that \mathcal{L}^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 $\text{sechash}'(b)$ for any $b \in \text{dom}(\text{sechash}')$. (v) $\text{txout} = \text{txout}'$. (vi) sechash extends $\text{sechash}'$ so that for each secret a_i we have $\text{sechash}(a_i) = h_i$. (vii) If $A \in \text{Hon}$, we define κ by extending κ' according to \vec{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) $\mathcal{L}^c = B \rightarrow * : m$ for some B , where m is the signature of the transaction T_{init} of $\mathbb{B}_{\text{adv}}(\{G\}C)$ relatively to the input x with $\hat{K}_A(r_A)$. The parameters of the compiler are set as follows: **part**, **PartG** and **val** are inferred from G , we let $\text{txout} = \text{txout}'$, $\text{sechash} = \text{sechash}'$, and $\mathbf{K}(B) = \hat{K}_B^p(r_B)$, $\mathbf{K}(D, B) = \kappa'(\{G\}C, D, B)$ for each B , and D subterm of C . (ii) For some B , we have $B \rightarrow * : T_{\text{init}}$ occurring in \dot{R}^c . (iii) \mathcal{L}^c is the first occurrence in \dot{R}^c of a broadcast of m after T_{init} . (iv) $\text{txout} = \text{txout}'$, $\text{sechash} = \text{sechash}'$, 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) $\mathcal{L}^c = T_{\text{init}}$ where T_{init} is the first transaction

in $\mathbb{B}_{\text{adv}}(\{G\}C)$. The needed compiler parameters are obtained as in Item 3. (iii) $\text{sechash} = \text{sechash}'$, $\kappa = \kappa'$, 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) $\mathcal{L}^c = B \rightarrow * : m$, where m is a signature with key $\kappa'(\{G\}C, D, A)$ of the first transaction T in $\mathbb{B}_D(D, D, T', o, v, \text{PartG}, 0)$, where $(T', o) = \text{txout}'(x)$. The compiler parameters are obtained as in Item 3. (iv) $\text{txout} = \text{txout}'$, $\text{sechash} = \text{sechash}'$, and $\kappa = \kappa'$. (v) \dot{R}^c contains $B \rightarrow * : T$ for some B , and m is the first signature of T in $\dot{R}^c \mathcal{L}^c$ after the first broadcast of T .
- (6) $\alpha = \text{put}(\vec{x}, \vec{a}, y)$, where: (i) $\vec{x} = x_1 \dots x_k$. (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 $D = \dots : \text{put} \dots \text{reveal} \dots C'$. 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 advertisement, for some G and C'' . (iv) $\mathcal{L}^c = T$ where T is the first transaction of $\mathbb{B}_C(C', D, T', o, v', \vec{x}, \text{PartG}, t)$, where $(T', o) = \text{txout}'(y)$. The compiler parameters are obtained as in Item 3. (v) txout extends txout' so that y' is mapped to $(T, 0)$, $\text{sechash} = \text{sechash}'$, and $\kappa = \kappa'$.
- (7) $\alpha = A : a$, where: (i) $\mathcal{L}^c = B \rightarrow * : m$ from some B with $|m| \geq \eta$. (ii) $\dot{R}^c = \dots (B \rightarrow O : m)(O \rightarrow B : \text{sechash}'(a)) \dots$, for some B . (iii) $\text{txout} = \text{txout}'$, $\text{sechash} = \text{sechash}'$ 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 $m' = (C, \vec{h}, \vec{k})$. (v) \mathcal{L}^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 $\langle C_0, v_0 \rangle_{x_0} \mid \dots \mid \langle C_k, v_k \rangle_{x_k}$ where $D = \dots : \text{split } \vec{v} \rightarrow \vec{C}$ and $\vec{C} = C_0 \dots C_k$. 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) $\mathcal{L}^c = T$ where T is the first transaction of $\mathbb{B}_{\text{par}}(\vec{C}, D, T', o, \text{PartG}, t)$ where $(T', o) = \text{txout}'(y)$. The compiler parameters are obtained as for Item 3. (iv) txout extends txout' mapping each x_i to (T, i) , $\text{sechash} = \text{sechash}'$, 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 $D = \dots : \text{withdraw } A$. (ii) In \dot{R}^s , we find that $\langle D + C, v \rangle_y$ has $\{G\}C'$ as its ancestor advertisement. (iii) $\mathcal{L}^c = T$ where T is the first transaction of $\mathbb{B}_D(D, D, T', o, v, \text{PartG}, 0)$ where $(T', o) = \text{txout}'(y)$. The compiler parameters are obtained as for Item 3. (iv) txout extends txout' mapping x to $(T, 0)$, $\text{sechash} = \text{sechash}'$, 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 $\text{txout}'(x)$ and $\text{txout}'(x')$, and a single output of value $v + v'$ redeemable with $\hat{K}_A(r_A)$. (iii) $\mathcal{L}^c = B \rightarrow * : m'$ for some B, m' , where m' is the signature of T with $\hat{K}_A(r_A)$. (iv) \mathcal{L}^c is the first broadcast of m' in \dot{R}^c after the first broadcast of T . (v) $\text{txout} = \text{txout}'$, $\text{sechash} = \text{sechash}'$, 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 $\text{txout}'(x)$ and $\text{txout}'(x')$, and having one output of value $v + v'$ redeemable with $\hat{K}_A(r_A)$. (iii) txout extends txout' mapping y to $(T, 0)$, $\text{sechash} = \text{sechash}'$, 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) $\alpha = A : \vec{y}, j$, where: (i) $\vec{y} = y_1 \cdots y_k$. (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 $\text{txout}'(y_i)$ 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 $\hat{K}_A(r_A)$, corresponding to the j -th input. (v) λ^c is the first broadcast of m in \dot{R}^c after the first broadcast of 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) $\text{txout} = \text{txout}'$, $\text{sechash} = \text{sechash}'$, and $\kappa = \kappa'$.
- (17) $\alpha = \text{destroy}(\vec{x})$, where: (i) $\vec{x} = x_1 \cdots x_k$. (ii) In \dot{R}^s , α consumes $\langle A_i, v_i \rangle_{x_i}$ to obtain 0. (iii) $\lambda^c = T$ from some T having as inputs $\text{txout}'(x_1), \dots, \text{txout}'(x_k)$, 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) $\text{txout} = \text{txout}'$, $\text{sechash} = \text{sechash}'$, and $\kappa = \kappa'$.
- (18) $\alpha = \delta = \lambda^c$, and $\text{txout} = \text{txout}'$, $\text{sechash} = \text{sechash}'$, and $\kappa = \kappa'$.

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

- (1) $\lambda^c = T$ where no input of T belongs to ran txout .
- (2) $\lambda^c = A \rightarrow O : m$ or $\lambda^c = O \rightarrow A : m$, for some A, m .
- (3) $\lambda^c = A \rightarrow * : m$, 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 $\text{coher}(R^s, R^c, r, \text{txout}, \text{sechash}, \kappa)$ for some $\text{txout}, \text{sechash}$, and κ . \diamond

The following lemma is the active contracts analogous of Lemma 1. Both results are proved by induction on the definition of coherence.

Lemma 6. *Let $\text{coher}(R^s, R^c, r, \text{txout}, \text{sechash}, \kappa)$. For each active contract $\langle C, v \rangle_x$ occurring in Γ_{R^s} , there exists a corresponding unspent transaction output (T, o) in \mathcal{B}_{R^c} with value v . Further, T is generated by the invoking the compiler as $\mathbb{B}_C(C, D_p, T', o', v, I, \mathcal{P}, t)$ for some values of $D_p, T', o', I, \mathcal{P}, t$, or as $\mathbb{B}_{\text{par}}(\vec{C}, D_p, T', o', \vec{v}, \mathcal{P}, t)$ for some values of $\vec{C}, D_p, T', o', \vec{v}, \mathcal{P}, t$ such that $C = \vec{C}_{o+1}$ and $v = \vec{v}_{o+1}$ using parameters $\text{txout}, \text{sechash}, \kappa$.*

PROOF. By induction on the definition of coherence. The correspondence is witnessed by the txout mapping. \square

A.7 From symbolic to computational strategies

Our computational soundness result is based on a mapping from 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 4. In the computational model, this is performed by following a *stipulation protocol* (Definition 21), 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 bit-strings, using transaction outputs (T, o) instead of deposit names x . In the stipulation protocol, we assume as given this representation 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 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 $\mathbb{B}_{\text{adv}}(\{G\}C)$, and finally she puts the generated T_{init} transaction on the ledger.

Definition 21 (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 C is the following.

- (1) A decodes C , constructing a symbolic contract advertisement $\{G\}C$; in doing this, A chooses distinct symbolic names for all the transaction outputs in 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 $B \in \text{PartG} \setminus \{A\}$. 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 \vec{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 $A : \text{secret } a$ in G . Then, A computes the hashes $\vec{h} = h_1 \cdots h_k$ of secret nonces (by querying O), and broadcasts all these hashes as a single message $m(C, \vec{h}, \vec{k})$. Dually, A receives the hashes \vec{h}' from the other participants. When doing this, A defines sechash using the first (correctly signed) $m(C, \vec{h}', \vec{k})$ in R_*^c which has no duplicate hashes, and has no hashes already occurring (signed) in R_*^c .
- (6) A computes $\mathbb{B}_{\text{adv}}(\{G\}C)$, 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.

Remark on replay attacks Item 5 of Definition 21 makes any honest participant reject any hash which has been already used in a previous commitment. This is crucial to avoid replay attacks during the commitment of the hashes: otherwise, it would be easy for an adversary to win the *OddsEvens* game in Section 2. The attack would proceed as follows. The adversary plays the role of **A**. After **B** has committed to secret s_b (by broadcasting $H(s_b)$), **A** commits to the same hash. Later on, after **B** has revealed s_b , **A** reveals $s_a = s_b$. In this way, **A** would ensure that the two secrets are identical, hence have the same length, which causes **A** to win the game.

Under the assumption that hashes are modelled according to the random oracle model, checking against duplicates in the stipulation protocol avoids replay attacks like the one above. Instead, under some weaker assumptions on the hash function, e.g. collision resistance, a replay attack⁵ would not be prevented by Item 5. For instance, let H' be a collision resistant hash function, and let:

$$H(X) = X_0 \parallel H'(X_1 \cdots X_n) \quad (\text{where } X = X_0 \cdots X_n)$$

In this way, also the hash function H is collision resistant. Let $\text{flip}(X)$ be the function which toggles the first bit of a bit string X . Our hash function H , by definition, satisfies:

$$\begin{aligned} H(\text{flip}(X)) &= H(\text{flip}(X_0) \parallel X_1 \cdots X_n) \\ &= \text{flip}(X_0) \parallel H'(X_1 \cdots X_n) \\ &= \text{flip}(X_0 \parallel H'(X_1 \cdots X_n)) \\ &= \text{flip}(H(X)) \end{aligned}$$

In this way, after receiving $H(s_b)$, the adversary can compute and broadcast $\text{flip}(H(s_b))$. Since $\text{flip}(H(s_b))$ is different from $H(s_b)$ (in the first bit), it is accepted as a commitment of **A** by our stipulation protocol. Further, after s_b is revealed, **A** can compute $\text{flip}(s_b)$, and reveal it as a preimage of $\text{flip}(H(s_b)) = H(\text{flip}(s_b))$. Since s_b and $\text{flip}(s_b)$ have the same length, the adversary **A** wins the game.

Parsing computational runs 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 20. 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 .

⁵We would like to thank Andrew Miller for pointing out this attack, which would be possible when weakening the assumptions about the hash function.

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 16 and 17, 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 $[\text{DEP-DESTROY}]$ 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 $[\text{C-AUTHCOMMIT}]$ 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 22 (From symbolic to computational strategies). Let Σ_A^s be a symbolic strategy, with $A \in \text{Hon}$. We define $\mathfrak{N}(\Sigma_A^s) = \Sigma_A^c$ 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 21), 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

A.8 Supplementary material for Section 9

Proof of Theorem 2 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 \hat{R}^c of R^c having a corresponding \hat{R}^s which is coherent (to \hat{R}^c) and conforming (to the computational strategies). We have that $R^c = \hat{R}^c \lambda^c \cdots$. We now show that either $\hat{R}^c \lambda^c$ has a corresponding symbolic run $\hat{R}^s \hat{R}^s$ which is coherent and conforming to the symbolic strategies (contradicting the maximality of \hat{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, \hat{R}^s must be either empty, or contain a single symbolic action.

We proceed by cases on λ^c :

- (1) $\lambda^c = B \rightarrow * : m$. Then, coherence must hold for some \hat{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 \hat{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 (\check{R}^s empty) the run $\check{R}^s \check{R}^s = \check{R}^s$ 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 A . Since $\Sigma_A^c = \mathfrak{N}(\Sigma_A^s)$, it follows that, at some time in the past, Σ_A^s enabled the authorization or reveal. By persistency, it is also enabled at the end of \check{R}^s , hence Σ_{Adv}^s can choose such action, and achieve conformance.

- (2) If $\lambda^c = T$, we consider the following subcases according to the inputs of T :
- (a) If no input of T belongs to $\text{ran } txout$, then coherence and conformance are achieved taking \check{R}^s to be empty.
 - (b) Otherwise, if at least one of the inputs of T belongs to $\text{ran } txout$, then we look in \check{R}^s for all the deposits and active contracts corresponding to such inputs. By definition of computational strategy, we must find in \check{R}^c a (first) broadcast $B \rightarrow * : T$ followed by a (first) broadcast $B \rightarrow * : m$ for all witnesses m of T . By the coherence of \check{R}^c , in \check{R}^s the messages $B \rightarrow * : m$ correspond to suitable authorization/reveal moves for each of the (counterparts of the) inputs of T . We consider the following subcases:
 - (i) If all the inputs are deposits, then we let \check{R}^s perform the symbolic move corresponding to T (e.g., init or join). Note that if T can not be represented symbolically, we can choose \check{R}^s to perform a destroy . Such moves are feasible symbolically since we already have their authorizations. Such \check{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 T' of T must correspond to an active contract. This must be originated from an advertisement, which has to involve at least one honest participant A , by definition of the symbolic semantics (rule $[\text{C-ADVERTISE}]$). However, by construction, our compiler makes T' require the signatures of all the participants, hence including A . Since such signature must occur as a witness in T , the adversary Adv must have forged it (with negligible probability), or must have obtained it from A . In the latter case, \check{R}^s contains an authorization for the symbolic move corresponding to λ^c . By choosing \check{R}^s accordingly, we obtain a coherent and conforming run.
- (3) Finally, if $\lambda^c = \delta$, we simply choose \check{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 \mathfrak{N} , this must also be the case in Λ^s , resulting in conformance. \square