Renegotiation and recursion in Bitcoin contracts

Massimo Bartoletti[1], Maurizio Murgia[2], and Roberto Zunino[2]

[1] University of Cagliari, Italy
[2] University of Trento, Italy

Abstract. BitML is a process calculus to express smart contracts that can be run on Bitcoin. One of its current limitations is that, once a contract has been stipulated, the participants cannot renegotiate its terms: this prevents expressing common financial contracts, where funds have to be added by participants at run-time. In this paper, we extend BitML with a new primitive for contract renegotiation. At the same time, the new primitive can be used to write recursive contracts, which was not possible in the original BitML. We show that, despite the increased expressiveness, it is still possible to execute BitML on standard Bitcoin, preserving the security guarantees of BitML.

1 Introduction

Smart contracts — computer protocols that regulate the exchange of assets in trustless environments — have become popular with the growth of interest in blockchain technologies. Mainstream blockchain platforms like Ethereum, Libra, and Cardano, feature expressive high-level languages for programming smart contracts. This flexibility has a drawback in that it may open the door to attacks that steal or tamper with the assets controlled by vulnerable contracts [4,21].

An alternative approach, pursued first by Bitcoin and more recently also by Algorand, is to sacrifice the expressiveness of smart contracts to reduce the attack surface. For instance, Bitcoin has a minimal language for transaction redeem scripts, containing only a limited set of logic, arithmetic, and cryptographic operations. Despite the limited expressiveness of these scripts, it is possible to encode a variety of smart contracts (like gambling games, escrow services, crowdfunding systems, etc.) by suitably chaining transactions [1,2,3,5,8,9,12,17,18,19,20,22]. The common trait of these works is that they render contracts as cryptographic protocols, where participants can exchange/sign messages, read the blockchain, and append transactions. Verifying the correctness of these protocols is hard, since it requires to reason in a computational model, where participants can manipulate arbitrary bitstrings, only being constrained to use PPTIME algorithms.

Departing from this approach, BitML [10] allows to write Bitcoin contracts in a high-level, process-algebraic language. BitML features a compiler that translates contracts into sets of standard Bitcoin transactions, and a sound and complete verification technique of relevant trace properties [11]. The computational soundness of the compiler guarantees that the execution of the compiled contract is coherent with the semantics of the source BitML specification, even in

the presence of adversaries. Although BitML can express many of the Bitcoin contracts presented in the literature [6], it is not "Bitcoin-complete", i.e. there exist contracts that can be executed on Bitcoin, but are not expressible in BitML.

For instance, consider a zero-coupon bond [16], where an investor A pays $1\mbox{Ƀ}$ upfront to a bank B, and receives back $2\mbox{Ƀ}$ after a maturity date (say, year 2030). We can express this contract in BitML as follows. First, as a precondition to the stipulation of the contract, we require both A and B to provide a deposit: A's deposit is $1\mbox{Ƀ}$, while B's deposit is $2\mbox{Ƀ}$. In BitML, we write this precondition as:

$$\mathsf{A} \colon 1\mbox{Ƀ} \, @ \, x_1 \mid \mathsf{B} \colon 2\mbox{Ƀ} \, @ \, x_2$$

where x_1 and x_2 are the identifiers of transactions containing the required amount of bitcoins ($\mbox{Ƀ}$). Under this precondition, we can specify the zero-coupon bond contract *ZCB* as follows:

$$ZCB \;=\; \texttt{split} \left(1\mbox{Ƀ} \to \texttt{withdraw } \mathsf{B} \;\mid\; 2\mbox{Ƀ} \to \texttt{after } 2030\colon\texttt{withdraw } \mathsf{A} \right)$$

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

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

In this paper we propose an extension of BitML that overcomes this issue. The idea is to allow the contract participants to *renegotiate* it after stipulation, in a controlled way. Renegotiation makes it possible to inject in the contract new funds, that were not specified in the original precondition. We can use this feature to solve the issue with the *ZCB* contract. The new precondition is $\mathsf{A} \colon 1\mbox{Ƀ} \, @ \, x_1$, i.e. we only require A's deposit. The revised contract is:

$$ZCB2 \;=\; \texttt{split} \left(1\mbox{Ƀ} \to \texttt{withdraw } \mathsf{B} \;\mid\; 0\mbox{Ƀ} \to * \colon \texttt{rngt } \mathtt{X} \langle \rangle \right)$$
$$\mathtt{X} \langle \rangle \;=\; \{ \mathsf{B} \colon 2\mbox{Ƀ} \, @ \, d \} \, \texttt{after } 2030\colon\texttt{withdraw } \mathsf{A}$$

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

The revised contract $ZCB2$ solves the problem of ZCB, in that it no longer freezes 2฿ for the whole duration of the bond: the bank could choose to renegotiate the contract, paying 2฿, just before the maturity date. This flexibility comes at a cost, since A loses the guarantee to eventually receive 2฿. To address this issue we need to add, as in the real world, an external mechanism. More specifically, we assume an insurance company I that, for an annual premium of $p\text{฿}$ paid by the bank, covers a face amount of $f\text{฿}$ (with $2 > f > 10p$):

$$\mathsf{A}\colon 1\text{฿} \, @ \, x_1 \mid \mathsf{B}\colon p\text{฿} \, @ \, x_2 \mid \mathsf{I}\colon f\text{฿} \, @ \, x_3$$

We revise the bond contract as follows:

$$
\begin{aligned}
ZCB3 \;=\; & \mathtt{split}\,\bigl(1\text{฿} \to \mathtt{withdraw}\ \mathsf{B} \\
& \qquad\quad \mid p\text{฿} \to \mathtt{withdraw}\ \mathsf{I} \\
& \qquad\quad \mid f\text{฿} \to *:\mathtt{rngt}\ \mathtt{X}\langle 1\rangle + \mathtt{after}\,2021\colon \mathtt{withdraw}\ \mathsf{A}\bigr) \\
\mathtt{X}\langle n \in 1..9\rangle \;=\; & \{\mathsf{B}\colon p\text{฿} \, @ \, d\} \\
& \mathtt{split}\,\bigl(p\text{฿} \to \mathtt{withdraw}\ \mathsf{I} \\
& \qquad\quad \mid f\text{฿} \to *:\mathtt{rngt}\ \mathtt{X}\langle n+1\rangle + \mathtt{after}\,(2021+n)\colon \mathtt{withdraw}\ \mathsf{A}\bigr) \\
\mathtt{X}\langle 10\rangle \;=\; & \{\mathsf{B}\colon 2\text{฿} \, @ \, d\} \\
& \mathtt{split}\,\bigl(f\text{฿} \to \mathtt{withdraw}\ \mathsf{I} \\
& \qquad\quad \mid 2\text{฿} \to \mathtt{after}\,2030\colon \mathtt{withdraw}\ \mathsf{A}\bigr)
\end{aligned}
$$

The contract starts by transferring 1฿ to the bank, and the first year of the premium to the insurer. The remaining $f\text{฿}$ are transferred to the renegotiated contract $\mathtt{X}\langle 1\rangle$, or, if the renegotiation is not completed by 2021, to the investor. The contracts $\mathtt{X}\langle n\rangle$, for $n \in 1..9$, allow the insurer to receive the annual premium until 2030: if the bank does not renegotiate the contract for the following year (paying the corresponding premium), then the investor can redeem the face amount of $f\text{฿}$. Finally, the contract $\mathtt{X}\langle 10\rangle$ can be triggered if the bank deposits the 2฿: when this happens, the face amount is given back to the insurer, and the investor can redeem 2฿ after the maturity date.

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

Contributions. We summarise our main contributions as follows:

- We extend BitML with the renegotiation primitive $*$: `rngt X⟨⟩`, suitably adapting the language syntax and semantics. The new primitive increases the expressiveness of BitML: besides allowing participants to provide new deposits and secrets at run-time, it also allows for *unbounded* recursion.
- We extend the BitML compiler to the new primitive, making it possible to execute renegotiations on Bitcoin. We accordingly extend the computational soundness result in [10], guaranteeing that the BitML semantics is coherent with the actual Bitcoin executions, also in the presence of adversaries.
- We exploit renegotiation to design a new gambling game where two players repeatedly flip coins, and whoever wins twice in a row takes the pot (a form of unbounded recursion). We prove the game to be fair.
- We introduce alternative renegotiation primitives, which allow participants to choose some parameters (e.g. the amounts to be deposited) at renegotiation time, and to change the set of participants involved in the renegotiated contract. We show that both primitives can be executed on Bitcoin *as is*. We also introduce a primitive that, at the price of minor Bitcoin extensions, supports *non-consensual* renegotiations, which are automatically triggered by the contract without requiring the participants' agreement.

Because of space constraints, we relegate part of the technicalities to Appendix A.

2 BitML with renegotiation and recursion

We start by formalising contract preconditions. We use $A, B, \ldots$ to range over participants. We assume a set of *deposit names* $x, y, \ldots$, a set of deposit variables $d, e, \ldots$, and a set of *secret names* $a, b, \ldots$. We use $\chi, \chi', \ldots$ to range over deposit names and variables, and v, v' to range over non-negative values.

Definition 1 (Contract precondition). *Contract preconditions have the following syntax (the deposits χ in a contract precondition G must be distinct):*

$$
\begin{array}{llll}
G ::= & A\!:\!v \,@\, \chi & \textit{deposit of } v \text{\fontencoding{U}\selectfont B} \textit{ put by } A \\
& | \ A\!:\!\texttt{secret}\, a & \textit{secret committed by } A \\
& | \ G \,|\, G & \textit{composition} & \diamond
\end{array}
$$

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

To define contracts, we assume a set of recursion variables, ranged over by $X, Y, \ldots$, and a language of *static expressions* $\mathcal{E}, \mathcal{E}', \ldots$, formed by integer constants k, integer variables $\alpha, \beta, \ldots$, and the usual arithmetic operators. We omit to define the syntax and semantics of static expressions, since they are standard. We assume that a closed static expression evaluates to a 32-bit value. We use the bold notation for sequences, e.g. $\boldsymbol{x}$ denotes a finite sequence of deposit names.

$$C ::= \sum_{i \in I} D_i \qquad \text{contract}$$

$D ::= \qquad\qquad\qquad$ guarded contract

$\quad$ `reveal` $\boldsymbol{a}$ `if` $p.\, C \qquad$ reveal secrets (if p is true)

$\quad |$ `withdraw` A $\qquad\quad$ transfer the balance to A

$\quad |$ `split` $\boldsymbol{v} \to \boldsymbol{C} \qquad\;$ split the balance

$\quad |$ A $: D \qquad\qquad\quad$ wait for A's authorization

$\quad |$ `after` $\mathcal{E} : D \qquad\quad$ wait until time $\mathcal{E}$

$\quad |\; * :$ `rngt` X$\langle\boldsymbol{\mathcal{E}}\rangle \qquad$ renegotiate the contract

$p ::= true \qquad$ truth

$\quad |\; p \wedge p \qquad$ conjunction

$\quad |\; \neg p \qquad$ negation

$\quad |\; E = E \qquad$ equality

$\quad |\; E < E \qquad$ less than

$E ::= \mathcal{E} \qquad$ static expression

$\quad |\; a \qquad$ secret

$\quad |\; E + E \qquad$ addition

$\quad |\; E - E \qquad$ subtraction

Fig. 1: Syntax of BitML contracts.

Definition 2 (Contract). *Contracts are terms with the syntax in Figure 1, where: (i) each recursion variable* X *has a unique defining equation* X$(\boldsymbol{\alpha}) = \{G\}C$*; (ii) renegotiations* $* :$ `rngt` X$\langle\boldsymbol{\mathcal{E}}\rangle$ *have the correct number of arguments; (iii) the names* $\boldsymbol{a}$ *in* `reveal` $\boldsymbol{a}$ `if` p *are distinct, and they include those occurring in* p*; (iv) in a prefix* `split` $\boldsymbol{v} \to \boldsymbol{C}$*, the sequences* $\boldsymbol{v}$ *and* $\boldsymbol{C}$ *have the same length. We denote with* 0 *the empty sum. We assume that the order of decorations is immaterial, e.g.,* `after` $\mathcal{E} :$ A $:$ B $: D$ *is equivalent to* B $:$ A $:$ `after` $\mathcal{E} : D$. $\qquad\qquad\diamond$

A contract C is a choice among guarded contracts D_i. A guarded contract `reveal` $\boldsymbol{a}$ `if` $p.\, C'$ continues as C' once all the secrets $\boldsymbol{a}$ have been revealed and satisfy the predicate p. The guarded contract `split` $(v_1 \to C_1 \mid \cdots \mid v_n \to C_n)$ divides the contract into n contracts C_i, each one with balance v_i. The sum of the v_i must coincide with the current balance. The action `withdraw` A transfers the whole balance to A. 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 $: D$, requiring to wait for A's authorization, and the decoration `after` $\mathcal{E} : D$, requiring to wait until the time specified by the static expression $\mathcal{E}$. The action $* :$ `rngt` X$\langle\boldsymbol{\mathcal{E}}\rangle$ allows the participants involved in the contract to renegotiate it. Intuitively, if X$(\boldsymbol{\alpha}) = \{G\}C$, then the contract continues as $C\{\boldsymbol{\mathcal{E}}/\boldsymbol{\alpha}\}$ if all the participants give their authorization, and satisfy the precondition G.

Definition 3 (Contract advertisement). *A contract advertisement is a term* $\{G\}C$ *such that: (i) each secret name in* C *occurs in* G*; (ii)* G *requires a deposit from each* A *in* $\{G\}C$*; (iii) each* $* :$ `rngt` X$\langle\boldsymbol{\mathcal{E}}\rangle$ *in* C *refers to a defining equation* X$(\boldsymbol{\alpha}) = \{G'\}C'$ *where the participants in* G' *are the same as those in* G. $\qquad\diamond$

The second condition is used to guarantee that the contract is stipulated only if *all* the involved participants give their authorizations. The last condition is only used to simplify the technical development. We outline in Section 5 how to relax it, by allowing renegotiations to exclude some participants, or to include new ones, which were not among those who originally stipulated the contract.

We now extend the reduction semantics of BitML [10], by focussing on the new renegotiation primitive. Because of space limitations, here we just provide the underlying intuition, relegating the full formalisation to Appendix A. We start by defining the configurations of the semantics.

Definition 4 (Configuration). *Configurations have the following syntax:*

$$
\begin{array}{lll}
\Gamma ::= & 0 & \textit{empty} \\
& | \ \{G\}^x C & \textit{contract advertisement (name x is optional)} \\
& | \ \langle C, v \rangle_x & \textit{active contract containing $v\text{B}$} \\
& | \ \langle \mathsf{A}, v \rangle_x & \textit{deposit of $v\text{B}$ redeemable by A} \\
& | \ \mathsf{A}[\chi] & \textit{authorization of A to perform action χ} \\
& | \ \{\mathsf{A} : a \# N\} & \textit{committed secret of A ($N \in \mathbb{N} \cup \{\bot\}$)} \\
& | \ \mathsf{A} : a \# N & \textit{revealed secret of A ($N \in \mathbb{N}$)} \\
& | \ \mathsf{A} : d \leftarrow x & \textit{A's deposit variable d assigned to deposit name x} \\
& | \ \Gamma \mid \Gamma' & \textit{parallel composition}
\end{array}
$$

We denote with $\Gamma \mid t$ a timed *configuration, where $t \in \mathbb{N}$ is a global time.* ◇

We illustrate configurations and their semantics through a series of examples.

Deposits. A deposit $\langle \mathsf{A}, v \rangle_x$ can be subject to several operations, like e.g. split into two smaller deposits, join with another deposit, transfer to another participant, or destroy. In all cases A must authorise the operation. For instance, to authorize the join of two deposits, A can perform the following step:

$$\langle \mathsf{A}, v \rangle_x \mid \langle \mathsf{A}, v' \rangle_y \ \rightarrow \ \langle \mathsf{A}, v \rangle_x \mid \langle \mathsf{A}, v' \rangle_y \mid \mathsf{A}[x, y \triangleright \langle \mathsf{A}, v + v' \rangle]$$

where $x, y \triangleright \langle \mathsf{A}, v + v' \rangle$ is the authorization of A to spend x. After A also provides the dual authorization to spend y, anyone can actually join the deposits:

$$\langle \mathsf{A}, v \rangle_x \mid \langle \mathsf{A}, v' \rangle_y \mid \mathsf{A}[x, y \triangleright \langle \mathsf{A}, v + v' \rangle] \mid \mathsf{A}[y, x \triangleright \langle \mathsf{A}, v + v' \rangle] \ \rightarrow \ \langle \mathsf{A}, v + v' \rangle_z$$

Advertisement. Any participant can broadcast a new contract advertisement $\{G\}C$, provided that all the deposits mentioned in G exist in the current configuration, and that the names of the secrets declared in G are fresh.

Stipulation. To stipulate an advertised contract $\{G\}C$, all the participants mentioned in it must fulfill the preconditions, and authorise the stipulation. For instance, let $G = \mathsf{A}{:}1 \, @ \, x \mid \mathsf{B}{:}1 \, @ \, y \mid \mathsf{A} : \mathtt{secret} \, a$, and let C be an arbitrary contract involving only A and B. The stipulation starts from a configuration containing the advertisement and the participants' deposits:

$$\Gamma \ = \ \{G\}C \mid \langle \mathsf{A}, 1 \rangle_x \mid \langle \mathsf{B}, 1 \rangle_y$$

At this point the participants must commit to their secrets (in this case, only A has a secret). This is rendered as a sequence of steps:

$$\Gamma \ \rightarrow^* \ \Gamma \mid \{\mathsf{A} : a \# N\} \mid \mathsf{A}[\# \triangleright \{G\}C] \mid \mathsf{B}[\# \triangleright \{G\}C] \ = \ \Gamma'$$

where $\{A : a\#N\}$ represents the fact that A has committed to the secret N, while $A[\# \triangleright \{G\}C]$ and $B[\# \triangleright \{G\}C]$ represent ending the commitment phase (these steps might seem redundant, but they are useful to obtain a step-by-step correspondence between BitML executions and Bitcoin executions).

After that, A and B must perform an additional sequence of steps to authorize the transfer of their deposits x, y to the contract:

$$\Gamma' \to^* \Gamma' \mid A[x \triangleright \{G\}C] \mid B[y \triangleright \{G\}C] = \Gamma''$$

where $A[x \triangleright \{G\}C]$ and $B[y \triangleright \{G\}C]$ are the authorizations to spend x and y.

At this point all the needed authorizations have been given, so the advertisement can be turned into an active contract. This step consumes the deposits and all the authorizations, and creates an active contract, with a fresh name z:

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

Renegotiation. We illustrate the steps to renegotiate $X\langle \alpha \rangle = \{G\}C$, where $G = A\!:\!1\,@\,d \mid B\!:\!1\,@\,e \mid A : \mathtt{secret}\,a$, and C is an arbitrary contract involving only A and B, and possibly containing the integer variable α in static expressions. Here, G requires A and B to spend two 1$\overset{\!}{\mathtt{B}}$ deposits, and A to commit to a secret. Unlike in the case of contract stipulation above, deposits names are unknown before renegotiation, so we use the deposit variables d, e to refer to them.

Consider a configuration $\langle * : \mathtt{rngt}\ X\langle k \rangle + C'', v \rangle_x \mid \Gamma$, where C'' contains the branches alternative to the renegotiation. A possible execution of the action $* : \mathtt{rngt}\ X\langle k \rangle$ starts as follows:

$$\langle * : \mathtt{rngt}\ X\langle k \rangle + C'', v \rangle_x \mid \Gamma \to \langle * : \mathtt{rngt}\ X\langle k \rangle + C'', v \rangle_x \mid \{G'\}^x C' \mid \Gamma = \Gamma'$$

where the advertisement $\{G'\}^x C'$ is obtained by transforming $\{G\}C$ as follows: (i) variables d, e are renamed into fresh ones d', e', and similarly the secret name a into a', (ii) the static expressions in C are evaluated, assuming $\alpha = k$, and replaced with their results. The superscript x in the advertisement is used to record that, when the renegotiation is concluded, the contract x must be reduced.

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

$$\Gamma' \to^* \Gamma' \mid A : d' \leftarrow y \mid \{A : a'\#N\} \mid A[\# \triangleright \{G'\}^x C']$$
$$\mid B : e' \leftarrow z \mid B[\# \triangleright \{G'\}^x C'] \qquad\qquad = \Gamma''$$

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

$$\Gamma'' \to^* \Gamma'' \mid A[y \triangleright \{G'\}^x C'] \mid A[x \triangleright \{G'\}^x C']$$
$$\mid B[z \triangleright \{G'\}^x C'] \mid B[x \triangleright \{G'\}^x C'] = \Gamma'''$$

Finally, the new contract is stipulated. This closes the old contract, and transfers its balance to the newly generated one, with a fresh name x':

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

Note that the branches in C'' are discarded only in the last step above, where we complete the renegotiation. Before this step, it would have been possible to take one of the branches in C''', aborting the renegotiation.

Withdraw. Executing `withdraw A` transfers the whole contract balance to A:

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

After the execution, the alternative branch C' is discarded, and a fresh deposit of v for A is created. Note that the active contract x is terminated.

Split. The `split` primitive divides the contract balance in n parts, each one controlled by its own contract. For instance, if $n = 2$:

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

After this step, the new spawned contracts C_1 and C_2 are executed concurrently.

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

$$\langle (\mathtt{reveal}\ ab\ \mathtt{if}\ a = b.\ C) + C', v \rangle_x \mid \Gamma \rightarrow \langle C, v \rangle_y \mid \Gamma$$

The terms $\mathsf{A} : a\#N$ and $\mathsf{B} : b\#N$ represent the fact that the secrets a and b have been revealed. Crucially, only the participant who performed the commitment can add the corresponding term to the configuration.

Authorizations. A branch decorated by $\mathsf{A} : \cdots$ can be taken only if the participant A has provided her authorization. For instance:

$$\langle \mathsf{A} : \mathtt{withdraw}\ \mathsf{B} + C', v \rangle_x \mid \mathsf{A}[x \rhd \mathsf{A} : \mathtt{withdraw}\ \mathsf{B}] \rightarrow \langle \mathsf{B}, v \rangle_y$$

The leftmost configuration contains the term $\mathsf{A}[x \rhd \mathsf{A} : \mathtt{withdraw}\ \mathsf{B}]$, which represents A's authorization to take the branch `withdraw B`. This enables the step to be taken. When multiple authorizations are required, the branch can be taken only after all of them occur in the configuration.

Time constraints. We represent time in configurations as $\Gamma \mid t$, where Γ is the untimed part of the configuration and t is the current time. We always allow the time to advance through the rule $\Gamma \mid t \rightarrow \Gamma \mid t + \delta$, for all $\delta > 0$. A branch decorated with `after d` can be taken only if time d has passed. For instance:

$$\langle \mathtt{after}\ d : \mathtt{withdraw}\ \mathsf{B}, v \rangle_x \mid t \rightarrow \langle \mathsf{B}, v \rangle_y \mid t \qquad \text{if } t \geq d$$

For the branches not guarded by `after`, we lift transitions from untimed to timed configurations: namely, for an untimed transition $\Gamma \rightarrow \Gamma'$, we also have the timed transition $\Gamma \mid t \rightarrow \Gamma' \mid t$. This reflects the assumption that participants can always meet deadlines, if they want to.

3 Executing BitML on Bitcoin

To execute a BitML contract, participants first compile it to a set of Bitcoin transactions, and then append these transactions to the blockchain, each following their own strategy. Participants' strategies can involve other actions besides appending transactions, like e.g. broadcasting signatures on given transactions (which corresponds, in BitML, to add an authorization to the configuration), revealing secrets, and waiting some time (see Definition 16 in [10]). The coherence between the BitML semantics and the execution on Bitcoin is guaranteed by a step-by-step correspondence between the transitions of the BitML semantics and the actions performed by participants on the Bitcoin network.

In this section we illustrate the compiler and the execution protocol through a couple of examples, focussing on the new renegotiation primitive. The needed background on Bitcoin will be introduced along with these examples. We relegate the formal definition of the compilation rules to Appendix A.2.

Zero-coupon bond. Recall the ZCB contract from Section 1:

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

The precondition $A : 1 @ x_1 \mid B : 2 @ x_2$ requires A to deposit 1B in the contract, and B to deposit 2B. In Bitcoin, this precondition corresponds to requiring two unspent transactions redeemable by A and B, and containing the required amounts. We represent these transactions as follows, using the notation in [7]:

T_{x_1}
in: $\cdots$
wit: $\cdots$
out: $(\lambda x.\mathsf{versig}_{\mathbf{K}(A)}(x), 1\text{B})$

T_{x_2}
in: $\cdots$
wit: $\cdots$
out: $(\lambda x.\mathsf{versig}_{\mathbf{K}(B)}(x), 2\text{B})$

The transaction T_{x_1} is a record with three fields (T_{x_2} is similar). The in field points to one or more previous transactions in the blockchain. The field out is a pair, whose first element is a boolean predicate (with parameter x), and the second element, 1B, is the amount that a subsequent transaction satisfying the predicate can redeem from T_{x_1}. Here, the predicate $\mathsf{versig}_{\mathbf{K}(A)}(x)$ is true when x is a signature of A on the redeeming transaction (i.e., one having T_{x_1} as in).

The contract ZCB is compiled into the transactions in Figure 2. The first one that can be appended to the blockchain is T_{init}. This requires a few conditions to be met: (i) T_{x_1} and T_{x_2} are *unspent* on the blockchain, i.e. no other transactions spend them; (ii) the amount specified in the out field of T_{init} does not exceed the sum of the amounts in T_{x_1} and T_{x_2}; (iii) the predicates in the out fields of T_{x_1} and T_{x_2} are true, after replacing the formal parameters with the signatures $sig_{\mathbf{K}(A)}$ and $sig_{\mathbf{K}(B)}$, contained in the wit field of T_{init}. The contract ZCB becomes stipulated once T_{init} is on the blockchain.

After that, the split action can be performed by either A or B, by redeeming T_{init} with T_{split}. This transaction uses $\mathbf{K}(ZCB, \{A, B\})$, a set of two key pairs, each one owned by each participant. These keys are only used in this step, to ensure that no transaction but T_{split} can redeem T_{init}.

Fig. 2: Transactions obtained by compiling the *ZCB* contract.

The transaction T_{split} creates two unspent outputs (indexed by 0 and 1), corresponding to the two parallel components of the `split`, each with its own balance. These outputs can be redeemed independently, by different transactions. The output at index 0 can only be redeemed by T_B (note that T_B's in field refers to the output 0 of T_{split}), transferring $1\dot{\mathsf{B}}$ to B. No other redemption is possible, since such output requires a signature with a specific key set, i.e. $\mathbf{K}(\texttt{withdraw B}, \{A, B\})$, which is not used for any other purpose. Further, the output of T_B can be redeemed with B's key, without A's one. Similarly, the output 1 of T_{split} can be redeemed by T_A, which in turns transfers $2\dot{\mathsf{B}}$ to A. The `absLock` field in T_A ensures that this may only happen after time 2030.

Zero-coupon bond with renegotiation. Compiling *ZCB2* yields the transactions:

Once these three transactions are on the blockchain, the only enabled action in the corresponding BitML contract is $* : \texttt{rngt X}\langle\rangle$, which asks $2\dot{\mathsf{B}}$ from B as a precondition. At the Bitcoin level, satisfying this precondition requires B to broadcast the identifier of a transaction T_y holding $2\dot{\mathsf{B}}$ and redeemable by himself. In BitML, this corresponds to choosing the deposit name y for the deposit variable d. Then, participants compile the contract advertisement $\{B : 2 @ d\}C$, where $C = \texttt{after 2030:withdraw A}$, after replacing d with y. The compiler produces the following transactions:

The renegotiation succeeds once $\mathsf{T}_{init}^{\mathtt{X}\langle\rangle}$ is on the blockchain. After that, any participant can perform the `withdraw A`, by appending T_A to the blockchain.

As shown by this example, the compiler handles renegotiation as follows:

– at stipulation time, it does not produce transactions for any $* : \mathtt{rngt}\, \mathtt{X}\langle\rangle$;
– at renegotiation time, the participants broadcast the identifiers of their new deposits, and the commitments of their new secrets. Static expressions are then evaluated, and replaced by their value. Finally, the new contract is compiled as usual, with the exception that the new initial transaction has an extra input, which transfers the balance of the caller contract to the callee (in the *ZCB2* example, this extra input is $(\mathsf{T}_{split}, 1)$ within $\mathsf{T}_{init}^{\mathtt{X}\langle\rangle}$).

Computational soundness The main result of [10] is computational soundness, which ensures that each execution trace at the Bitcoin level has a corresponding one in the semantics of BitML. This was achieved by formalizing the semantics of Bitcoin using a computational model, where participants can exchange bitstrings as messages, and append transactions to the blockchain. Then, a coherence relation was defined to relate symbolic runs to computational ones, essentially matching symbolic moves with their implementation in Bitcoin.

Our extension of BitML with renegotiation still enjoys computational soundness. The argument is similar, and requires extending the coherence relation to the new primitive. In particular, the reduction:

$$\{G\}^x C \mid \Gamma \;\to\; \{G\}^x C \mid \Gamma \mid \parallel_i \{\mathsf{A} : a_i \# N_i\} \mid \parallel_j \mathsf{A} : d_j \leftarrow x_j \mid \mathsf{A}[\# \rhd \{G\}^x C]$$

corresponds, in Bitcoin, to A broadcasting a message which contains the hashes of her secrets and the transaction identifiers that she wishes to use as deposits.

Instead, the reduction:

$$\{G\}^x C \mid \Gamma \;\to\; \{G\}^x C \mid \Gamma \mid \mathsf{A}[x \rhd \{G\}^x C]$$

corresponds to A signing all the transactions obtained by compiling the new contract, and broadcasting the signatures. A participant signs T_{init} only after receiving the signatures of the other transactions from all the other participants.

Computational soundness requires that each contract involves at least one participant, say A, who follows the Bitcoin implementation of BitML. In particular, A follows the stipulation and renegotiation protocols correctly, i.e. signing nothing but the protocol messages, and signing T_{init} last. We also make the usual assumptions on computational adversaries: they can only run PPTIME algorithms, and they can break the underlying cryptography with negligible probability, only. Consequently, we only consider computational runs of polynomial length (with respect to the security parameter). This is because in longer runs the adversary would be able to break the cryptography by brute force.

Below, we provide an intuitive statement of computational soundness. The formal statement is in Appendix A.3.

Theorem 1 (Computational soundness). *Under the hypotheses above, each Bitcoin-level computational run has a corresponding coherent BitML run, with overwhelming probability.*

$$CFG = \quad \texttt{reveal } b \texttt{ if } 0 \le b \le 1.\Big($$
$$\texttt{reveal } ab \texttt{ if } a = b.\ (* : \texttt{rngt } \mathtt{X_A}\langle 1 \rangle\ +\ \texttt{after } 3 : \mathit{Split_A})$$
$$+\ \texttt{reveal } ab \texttt{ if } a \ne b.\ (* : \texttt{rngt } \mathtt{X_B}\langle 1 \rangle\ +\ \texttt{after } 3 : \mathit{Split_B})$$
$$+\ \texttt{after } 2 : \texttt{withdraw } \mathsf{B}\Big)$$
$$+\ \texttt{after } 1 : \texttt{withdraw } \mathsf{A}$$

$$\mathtt{X_A}\langle n \rangle = \{\mathsf{A}:\texttt{secret } a \mid \mathsf{B}:\texttt{secret } b\}$$
$$\texttt{reveal } b \texttt{ if } 0 \le b \le 1.\Big($$
$$\texttt{reveal } ab \texttt{ if } a = b.\ \texttt{withdraw } \mathsf{A}$$
$$+\ \texttt{reveal } ab \texttt{ if } a \ne b.\ (* : \texttt{rngt } \mathtt{X_B}\langle n+1 \rangle\ +\ \texttt{after } (3n+3) : \mathit{Split_B})$$
$$+\ \texttt{after } (3n+2) : \texttt{withdraw } \mathsf{B}\Big)$$
$$+\ \texttt{after } (3n+1) : \texttt{withdraw } \mathsf{A}$$

$$\mathtt{X_B}\langle n \rangle = \{\mathsf{A}:\texttt{secret } a \mid \mathsf{B}:\texttt{secret } b\}$$
$$\texttt{reveal } b \texttt{ if } 0 \le b \le 1.\Big($$
$$\texttt{reveal } ab \texttt{ if } a = b.\ (* : \texttt{rngt } \mathtt{X_A}\langle n+1 \rangle\ +\ \texttt{after } (3n+3) : \mathit{Split_A})$$
$$+\ \texttt{reveal } ab \texttt{ if } a \ne b.\ \texttt{withdraw } \mathsf{B}$$
$$+\ \texttt{after } (3n+2) : \texttt{withdraw } \mathsf{B}\Big)$$
$$+\ \texttt{after } (3n+1) : \texttt{withdraw } \mathsf{A}$$

$$\mathit{Split_A} = \texttt{split } (4 \to \texttt{withdraw } \mathsf{A} \mid 2 \to \texttt{withdraw } \mathsf{B})$$

$$\mathit{Split_B} = \texttt{split } (4 \to \texttt{withdraw } \mathsf{B} \mid 2 \to \texttt{withdraw } \mathsf{A})$$

Fig. 3: A recursive coin flipping game.

4 A fair recursive coin flipping game

To illustrate recursion in our extended BitML, we introduce a simple game where two players repeatedly flip coins, and the one who wins two consecutive flips takes the pot. The precondition requires each player to deposit 3 ฿ and choose a secret:

$$\mathsf{A}:3\,@\,x \mid \mathsf{A}:\texttt{secret } a \mid \mathsf{B}:3\,@\,y \mid \mathsf{B}:\texttt{secret } b$$

The contract CFG (Figure 3) asks B to reveal his secret first: if B waits too much, A can withdraw the contract funds after time 1. Then, it is A's turn to reveal (before time 2, otherwise B can withdraw the funds). The current flip winner is A if the secrets of A and B are equal, otherwise it is B. At this point, the contract can be renegotiated as $\mathtt{X_A}\langle 1 \rangle$ or $\mathtt{X_B}\langle 1 \rangle$, depending on the flip winner (the parameter 1 represents the round). If players do not agree on the renegotiation, then the funds are split fairly, according to the current expected win.

The contract $\mathtt{X_A}\langle n \rangle$ requires A and B to generate fresh secrets for the n-th turn. If A wins again, she can withdraw the pot, otherwise the contract can be renegotiated as $\mathtt{X_B}\langle n+1 \rangle$. If the players do not agree on the renegotiation, the pot is split fairly between them. The contract $\mathtt{X_B}$ is similar.

The following theorem states that our coin flipping game is fair. Fairness ensures that the expected payoff of a *rational* player is always non-negative, notwithstanding the behaviour of the other player. Rational players must choose random secrets in $\{0, 1\}$. Indeed, non uniformly distributed secrets can make

the adversary bias the coin flip in her favour. Further, choosing a secret different from 0 or 1 would decrease the player payoff. Indeed, B would be prevented from revealing his secrets (by the predicate in the `reveal b`), and so A could win after the timeout. If A chooses a secret different from 0 or 1, she makes B win the round (since B wins when the secrets are different). Rationality also requires to reveal secrets in time (before the alternative `after` branch is enabled), and to take the *Split* branch if restipulation does not occur in time. This ensures that, when renegotiation happens, there is still time to reveal the round secrets. Indeed, a late renegotiation could enable the other player to win by timeout.

Theorem 2. *The expected payoff of a rational player is always non-negative.*

Proof (Sketch). First, we consider the case where renegotiation always happens. A rational player wins each coin flip with probability $1/2$, at least: so, the probability of winning the whole game is also $1/2$, at least. In the general case, the renegotiation at the end of each round may fail. When this happens, the rational player takes the *Split* branch, distributing the pot according to the expected payoff in the *current* game state, thus ensuring the fairness of the whole game. The player who won the last coin flip is expected to win $p\,₿$, with $p = 1/2 \cdot 6 + 1/2 \cdot (1/2 \cdot p + 1/2 \cdot 0)$, giving $p = 4$. Accordingly, the *Split* contracts transfer $4\,₿$ to the winner of the last flip and $(6-4)\,₿ = 2\,₿$ to the other player.

5 More expressive renegotiation primitives

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

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

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

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

installments payment plan as $\mathtt{IPP}\langle 1\rangle$, with the following defining equations:

$$\mathtt{IPP}\langle\alpha < 5\rangle = \{\mathsf{A}\!:\!1\,@\,d\}\big(\mathtt{split}\ 1 \to \mathtt{withdraw}\ \mathsf{B}\ |\ 0 \to *:\mathtt{rngt}\ \mathtt{IPP}\langle\alpha+1\rangle\big)$$

$$\mathtt{IPP}\langle 5\rangle = \{\mathsf{A}\!:\!1\,@\,d\}\big(\mathtt{split}\ 1 \to \mathtt{withdraw}\ \mathsf{B}\ |\ 0 \to *:\mathtt{rngt}\ \mathtt{Y}\langle\mathsf{B}:k,\mathsf{B}:v\rangle\big)$$

$$\mathtt{Y}\langle\alpha \neq 1,\beta\rangle = \{\mathsf{A}\!:\!\beta\,@\,d\}\big(\mathtt{split}\ \beta \to \mathtt{withdraw}\ \mathsf{B}\ |\ 0 \to *:\mathtt{rngt}\ \mathtt{Y}\langle\alpha-1,\beta\rangle\big)$$

$$\mathtt{Y}\langle 1,\beta\rangle = \{\mathsf{A}\!:\!\beta\,@\,d\}\ \mathtt{withdraw}\ \mathsf{B}$$

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

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

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

$$\mathsf{A}:\mathtt{withdraw}\ \mathsf{B}\ +\ \mathsf{B}:\mathtt{withdraw}\ \mathsf{A}\ +\ \mathsf{A}:\mathsf{M}:\mathtt{rngt}\ \mathtt{Refd}_\mathsf{A}\langle\rangle\ +\ \mathsf{B}:\mathsf{M}:\mathtt{rngt}\ \mathtt{Refd}_\mathsf{B}\langle\rangle$$

$$\mathtt{Refd}_\mathsf{P} = \{\mathsf{P}\!:\!0.1\,@\,d\}\ \mathtt{split}\ \big(0.1 \to \mathtt{withdraw}\ \mathsf{M}\ |\ 1 \to \mathtt{withdraw}\ \mathsf{P}\big)$$

where $\mathsf{A}:\mathsf{M}:\mathtt{rngt}\ \mathtt{Refd}_\mathsf{A}\langle\rangle$ means that only A and M need to agree in order for the contract $\mathtt{Refd}_\mathsf{A}$ to be executed, resolving the dispute. In this case it is crucial that the renegotiation is possible even without the agreement between A and B. Indeed, if M decides to refund A (by authorizing $\mathtt{Refd}_\mathsf{A}$), it is not to be expected that also B agrees. Similarly to the one discussed before, also this extension can be implemented on-top of Bitcoin. The computational soundness property is preserved, under the assumption that at least one participant in any renegotiation is *honest*, i.e. it follows the renegotiation protocol. Crucially, if a renegotiation only involves dishonest participants, the renegotiated contract could be anything, not necessarily that prescribed in the original contract.

Non-consensual renegotiation. In the variants of $*:\mathtt{rngt}$ discussed before, renegotiation requires one or more participants to agree. Hence, each use of $*:\mathtt{rngt}$ must include suitable alternative branches, to be fired in case the renegotiation fails. In certain scenarios, we may want to renegotiate the contract without the participants having to agree. To this purpose, we can introduce a new primitive

`call X⟨⟩`, which continues as `X⟨⟩` without requiring anyone to agree. For simplicity, we assume the defining equations of this primitive of the form $X(\boldsymbol{\alpha}) = \{v\}C$, where v represents the amount of ฿ added to the contract, by anyone.

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

$$C = \{A\!:\!1 \, @ \, x \mid B\!:\!2 \, @ \, y\}(\texttt{call } X_A\langle 2\rangle + \texttt{after } 1\!:\!\texttt{withdraw } B)$$
$$X_A\langle n\rangle = \{2\}(\texttt{call } X_B\langle n+1\rangle + \texttt{after } n\!:\!\texttt{withdraw } A)$$
$$X_B\langle n\rangle = \{2\}(\texttt{call } X_A\langle n+1\rangle + \texttt{after } n\!:\!\texttt{withdraw } B)$$

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

Implementing the `call` primitive on top of Bitcoin seems unfeasible: even if it were possible to use complex off-chain multiparty computation protocols [15], doing so might be impractical. Rather, we would like to extend Bitcoin as much as needed for the new primitive. In our implementation of BitML, we compile contracts to sets of transactions and make participants sign them. In standard BitML this is doable since, at stipulation time, we can finitely over-approximate the reducts of the original contract. Recursion can make this set infinite, e.g. $X_A\langle 2\rangle, X_A\langle 3\rangle, \ldots$, hence impossible to compile and sign statically. A way to cope with this is to extend Bitcoin with *malleable* signatures which only cover the part of the transaction not affected by the parameter n in $X_B\langle n\rangle$. Further, signatures must not cover the in fields of transactions, since they change as recursion unfolds. In this way, the same signature can be reused for each call.

Adding malleability provides flexibility, but poses some risks. For instance, instead of redeeming the transaction corresponding to $X_A\langle n\rangle$ with the transaction of $X_B\langle n+1\rangle$ one could instead use the transaction of $X_B\langle n+100\rangle$, since the two transactions have the same signature. To overcome this problem, we could add a new opcode to allow the output script of $X_B\langle n\rangle$ to access the parameter in the redeeming transaction, so to verify that it is indeed $n+1$ as intended. Similarly, to check that we have 2฿ more in the new transaction, an opcode could provide the value of the new output. The same goal could be achieved by adapting the techniques used in [23,24] to realize *covenants*.

6   Conclusions

We have investigated linguistic primitives to renegotiate BitML contracts, and their implementation on standard Bitcoin. By relaxing this constraint, e.g. assuming the extended UTXO model of [13], we could devise more expressive

primitives. Note that the existing verification techniques for BitML [11], which are based on model-checking the whole state space, cannot be directly applied in the presence of recursion. Type-based approaches, like e.g. the one in [14], could be devised for our setting.

References

1. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Fair two-party computations via Bitcoin deposits. In: Financial Cryptography Workshops. LNCS, vol. 8438, pp. 105–121. Springer (2014). https://doi.org/10.1007/978-3-662-44774-1_8

2. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multiparty computations on Bitcoin. In: IEEE S & P. pp. 443–458 (2014). https://doi.org/10.1109/SP.2014.35, first appeared on Cryptology ePrint Archive, http://eprint.iacr.org/2013/784

3. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multiparty computations on Bitcoin. Commun. ACM **59**(4), 76–84 (2016). https://doi.org/10.1145/2896386

4. Atzei, N., Bartoletti, M., Cimoli, T.: A survey of attacks on Ethereum smart contracts (SoK). In: Principles of Security and Trust (POST). LNCS, vol. 10204, pp. 164–186. Springer (2017). https://doi.org/10.1007/978-3-662-54455-6_8

5. Atzei, N., Bartoletti, M., Cimoli, T., Lande, S., Zunino, R.: SoK: unraveling Bitcoin smart contracts. In: POST. LNCS, vol. 10804, pp. 217–242. Springer (2018). https://doi.org/10.1007/978-3-319-89722-6

6. Atzei, N., Bartoletti, M., Lande, S., Yoshida, N., Zunino, R.: Developing secure Bitcoin contracts with BitML. In: ESEC/FSE (2019). https://doi.org/https://doi.org/10.1145/3338906.3341173

7. Atzei, N., Bartoletti, M., Lande, S., Zunino, R.: A formal model of Bitcoin transactions. In: Financial Cryptography and Data Security. LNCS, vol. 10957. Springer (2018). https://doi.org/10.1007/978-3-662-58387-6

8. Banasik, W., Dziembowski, S., Malinowski, D.: Efficient zero-knowledge contingent payments in cryptocurrencies without scripts. In: ESORICS. LNCS, vol. 9879, pp. 261–280. Springer (2016). https://doi.org/10.1007/978-3-319-45741-3_14

9. Bartoletti, M., Zunino, R.: Constant-deposit multiparty lotteries on Bitcoin. In: Financial Cryptography Workshops. LNCS, vol. 10323. Springer (2017). https://doi.org/10.1007/978-3-319-70278-0

10. Bartoletti, M., Zunino, R.: BitML: a calculus for Bitcoin smart contracts. In: ACM CCS (2018). https://doi.org/10.1145/3243734.3243795

11. Bartoletti, M., Zunino, R.: Verifying liquidity of bitcoin contracts. In: POST. LNCS, vol. 11426. Springer (2019)

12. Bentov, I., Kumaresan, R.: How to use Bitcoin to design fair protocols. In: CRYPTO. LNCS, vol. 8617, pp. 421–439. Springer (2014). https://doi.org/10.1007/978-3-662-44381-1_24

13. Chakravarty, M.M., Chapman, J., MacKenzie, K., Melkonian, O., Jones, M.P., Wadler, P.: The extended UTXO model. In: Workshop on Trusted Smart Contracts (2020)

14. Das, A., Balzer, S., Hoffmann, J., Pfenning, F.: Resource-aware session types for digital contracts. CoRR **abs/1902.06056** (2019)

15. Gudgeon, L., Moreno-Sanchez, P., Roos, S., McCorry, P., Gervais, A.: Sok: Off the chain transactions. IACR Cryptology ePrint Archive **2019**, 360 (2019)
16. Jones, S.L.P., Eber, J., Seward, J.: Composing contracts: an adventure in financial engineering, functional pearl. In: International Conference on Functional Programming (ICFP). pp. 280–292 (2000). https://doi.org/10.1145/351240.351267
17. Kumaresan, R., Bentov, I.: How to use Bitcoin to incentivize correct computations. In: ACM CCS. pp. 30–41 (2014). https://doi.org/10.1145/2660267.2660380
18. Kumaresan, R., Bentov, I.: Amortizing secure computation with penalties. In: ACM CCS. pp. 418–429 (2016). https://doi.org/10.1145/2976749.2978424
19. Kumaresan, R., Moran, T., Bentov, I.: How to use Bitcoin to play decentralized poker. In: ACM CCS. pp. 195–206 (2015). https://doi.org/10.1145/2810103.2813712
20. Kumaresan, R., Vaikuntanathan, V., Vasudevan, P.N.: Improvements to secure computation with penalties. In: ACM CCS. pp. 406–417 (2016). https://doi.org/10.1145/2976749.2978421
21. Luu, L., Chu, D.H., Olickel, H., Saxena, P., Hobor, A.: Making smart contracts smarter. In: ACM CCS. pp. 254–269 (2016). https://doi.org/10.1145/2976749.2978309
22. Miller, A., Bentov, I.: Zero-collateral lotteries in Bitcoin and Ethereum. In: EuroS&P Workshops. pp. 4–13 (2017). https://doi.org/10.1109/EuroSPW.2017.44
23. Möser, M., Eyal, I., Sirer, E.G.: Bitcoin covenants. In: Financial Cryptography Workshops. LNCS, vol. 9604, pp. 126–141. Springer (2016). https://doi.org/10.1007/978-3-662-53357-4_9
24. O'Connor, R., Piekarska, M.: Enhancing Bitcoin transactions with covenants. In: Financial Cryptography Workshops. LNCS, vol. 10323. Springer (2017). https://doi.org/10.1007/978-3-319-70278-0_12

A Supplementary material

A.1 Semantics of BitML

We denote with Hon the non-empty set of the *honest* participants. We use notation $\{G\}^{\circ}C$ to refer to a contract advertisement of the form $\{G\}C$, or $\{G\}^x C$ where x is immaterial. $[\![\mathcal{E}]\!]$ is the evaluation of static expression $\mathcal{E}$, defined as expected. Note that the evaluation function is undefined if $\mathcal{E}$ contains free variables. We overload $[\![-]\!]$ to contracts and contract advertisements: $[\![C]\!]$ is the contract obtained by substituting all the occurring static expressions with their valuation, and $[\![\{G\}C]\!] = \{G\}[\![C]\!]$. We write $\equiv_\alpha$ for equivalence of contract advertisement up-to α-conversion of secret and deposit names. We write $\{G\}C \equiv X\langle\mathcal{E}\rangle$ if: (i) $X(\boldsymbol{\alpha}) = \{G'\}C'$; (ii) $[\![\{G'\}C'\{[\![\boldsymbol{\mathcal{E}}]\!]/\boldsymbol{\alpha}\}]\!] \equiv_\alpha \{G\}C$.

The semantic rules for advertisement and stipulation are in Figure 4. The rules for actions and for deposits are unchanged (see [10]).

A.2 Compiler

Figure 5 shows the new rules to be added to the ones in [10] for compiling renegotiations.

Compiling an advertisement used for renegotiation $\{G\}^z C$ is similar to compiling a regular advertisement $\{G\}C$. The main difference is that, in the renegotiation case, we also have to redeem the transaction output relative to the parent contract z, denoted with $\mathbf{txout}(z)$. Therefore, the generated transaction T_{init} includes $\mathbf{txout}(z)$ as an additional input. Further, the value v of T_{init} consists of the amounts of the deposits and the balance of the parent contract z, denoted with $\mathbf{val}(z)$. The compiler also handles time-constrained renegotiations as in $\mathtt{after}\, t : * : \mathtt{rngt}\, X\langle\mathcal{E}\rangle$ by using t in the absLock field of T_{init}, so to prevent the renegotiation to be completed earlier than t.

Note that compiling $* : \mathtt{rngt}\, X\langle\mathcal{E}\rangle$ generates *no transactions*. This is because the transactions for $X\langle\mathcal{E}\rangle$ do not have to be generated at stipulation time, but only at renegotiation time.

A.3 Computational soundness

Below, $\aleph(\Sigma_\mathsf{A}^s)$ denotes the computational strategy obtained by following the symbolic strategy Σ_A^s of honest participant A, as defined in [10]. This is extended to deal with renegotiation using the same protocol as stipulation, with minor differences as explained in Section 3.

Theorem 3 (Computational soundness). *Let $\boldsymbol{\Sigma}^s$ be a set of symbolic strategies for all $\mathsf{A} \in$ Hon. Let $\boldsymbol{\Sigma}^c$ be a set of computational strategies such that $\Sigma_\mathsf{A}^c = \aleph(\Sigma_\mathsf{A}^s)$ for all $\mathsf{A} \in$ Hon, including an arbitrary adversary strategy Σ_{Adv}^c. Fix $k \in \mathbb{N}$. Then, the following set has overwhelming probability:*

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

$$\dfrac{\Gamma \text{ contains } \langle \mathsf{A}_i, v_i\rangle_{x_i} \text{ for all } \mathsf{A}_i\!:\!v_i\,@\,x_i \text{ in } \{G\}C \qquad \mathtt{secrets}(G) \text{ fresh} \qquad \mathtt{part}(\{G\}C)\cap\mathsf{Hon}\neq\emptyset}{\Gamma \xrightarrow{\;advertise(\{G\}C)\;} \{G\}C \mid \Gamma}\;[\text{C-Adv}]$$

$$\dfrac{\begin{array}{ll} \{G\}C' \equiv \mathtt{X}\langle\boldsymbol{\mathcal{E}}\rangle & \text{all names in } G \text{ fresh}\\ \mathtt{secrets}(G) \text{ fresh} & \{G'\}C'' \not\equiv \mathtt{X}\langle\boldsymbol{\mathcal{E}}\rangle, \text{for all } \{G'\}^x C'' \in \Gamma\\ d \text{ fresh, for each } \mathsf{A}\!:\!v\,@\,d \text{ in } G & \Gamma \text{ contains } \langle\mathsf{A}_i,v_i\rangle_{x_i} \text{ for all } \mathsf{A}_i\!:\!v_i\,@\,x_i \text{ in } G \end{array}}{\langle *:\mathtt{rngt}\ \mathtt{X}\langle\boldsymbol{\mathcal{E}}\rangle + C, v\rangle_x \mid \Gamma \xrightarrow{\;advertise(\{G\}^x C')\;} \langle *:\mathtt{rngt}\ \mathtt{X}\langle\boldsymbol{\mathcal{E}}\rangle + C, v\rangle_x \mid \{G\}^x C' \mid \Gamma}\;[\text{C-Rngt}]$$

$$\dfrac{\begin{array}{ll} a_1\cdots a_k \text{ secrets of } \mathsf{A} \text{ in } G & \forall i\in 1..k: \nexists N: \{\mathsf{A}:a_i\#N\}\in\Gamma\\ d_1\cdots d_h \text{ deposit variables of } \mathsf{A} \text{ in } G & \\ \forall i\in 1..h: \nexists x:(\mathsf{A}:d_i\leftarrow x)\in\Gamma & \forall i\in 1..k: N_i\in \begin{cases}\mathbb{N} & \text{if } \mathsf{A}\in\mathsf{Hon}\\ \mathbb{N}\cup\{\bot\} & \text{otherwise}\end{cases}\\ \mathsf{A}\in\mathsf{Hon} \implies \begin{cases}\Gamma \text{ contains } \langle\mathsf{A},v_i\rangle_{x_i}, \text{ for all } i\in 1..h, \text{ and}\\ x_i\neq x_j \text{ for all } i\neq j\in 1..h, \text{ and}\\ x_i\neq x \text{ for all } i\in 1..h, x \text{ such that } \exists v: \mathsf{A}\!:\!v\,@\,x\in G\end{cases}\\ \Delta = \{\mathsf{A}:a_1\#N_1\}\mid\cdots\mid\{\mathsf{A}:a_k\#N_k\}\mid\mathsf{A}:d_1\leftarrow x_1\mid\ldots\mid\mathsf{A}:d_h\leftarrow x_h \end{array}}{\{G\}^\circ C \mid \Gamma \xrightarrow{\;\mathsf{A}:\{G\}^\circ C,\Delta\;} \{G\}^\circ C \mid \Gamma \mid \Delta \mid \mathsf{A}[\#\rhd\{G\}^\circ C]}\;[\text{C-AuthCommit}]$$

$$\dfrac{\Gamma \text{ contains } \mathsf{B}[\#\rhd\{G\}^\circ C] \text{ for all } \mathsf{B} \text{ in } G \qquad G = \mathsf{A}\!:\!v\,@\,\chi\mid\cdots \qquad \Gamma\vdash\chi = x}{\{G\}^\circ C \mid \Gamma \xrightarrow{\;\mathsf{A}:\{G\}^\circ C,x\;} \{G\}^\circ C \mid \Gamma \mid \mathsf{A}[x\rhd\{G\}^\circ C]}\;[\text{C-AuthInitDep}]$$

$$\dfrac{\Gamma \text{ contains } \mathsf{B}[\#\rhd\{G\}^x C] \text{ for all } \mathsf{B} \text{ in } G}{\{G\}^x C \mid \Gamma \xrightarrow{\;\mathsf{A}:\{G\}C,x\;} \{G\}^x C \mid \Gamma \mid \mathsf{A}[x\rhd\{G\}^x C]}\;[\text{C-AuthInitContr}]$$

$$\dfrac{\begin{array}{c} \{G\}C \equiv \mathtt{X}\langle\boldsymbol{\mathcal{E}}\rangle \quad y \text{ fresh} \quad G = \big(\,\|_{i\in I}\mathsf{A}_i\!:\!v_i\,@\,x_i\big)\mid\big(\,\|_{i\in J}\mathsf{B}_i\!:\!v_i'\,@\,d_i\big)\mid\big(\,\|_{i\in K}\mathsf{C}_i\!:\!\mathtt{secret}\ a_i\big)\\ \Delta = \big(\,\|_{i\in I}\langle\mathsf{A}_i,v_i\rangle_{x_i}\big)\mid\big(\,\|_{i\in J}\langle\mathsf{B}_i,v_i'\rangle_{x_i'}\big)\mid\big(\,\|_{i\in J}\mathsf{B}_i:d_i\leftarrow x_i'\big)\\ \mid\big(\,\|_{\mathsf{A}\in G}\mathsf{A}[\#\rhd\{G\}^x C]\mid\mathsf{A}[x\rhd\{G\}^x C]\big)\mid\big(\,\|_{i\in I}\mathsf{A}_i[x_i\rhd\{G\}^x C]\big)\mid\big(\,\|_{i\in J}\mathsf{B}_i[x_i'\rhd\{G\}^x C]\big) \end{array}}{\langle *:\mathtt{rngt}\ \mathtt{X}\langle\boldsymbol{\mathcal{E}}\rangle + C', v\rangle_x \mid \{G\}^x C \mid \Gamma \mid \Delta \xrightarrow{\;init(x,G,C)\;} \langle C,\sum_{i\in I}v_i + \sum_{i\in J}v_i' + v\rangle_y \mid \Gamma}\;[\text{C-Init}]$$

$$\overline{\Gamma\vdash x = x} \qquad \overline{\Gamma\mid\mathsf{A}:d\leftarrow x\vdash d = x} \qquad cv(init(x,G,C)) = \{x\}$$

Fig. 4: Semantics of advertisement and stipulation.

$$\dfrac{\begin{array}{ll} G = \big(\,\|_{i\in J}\mathsf{A}_i\!:\!v_i'\,@\,y_i\big)\mid\big(\,\|_{i\in K}\mathsf{B}_i\!:\!\mathtt{secret}\ a_i\big) & \\ C = \sum_{i=1}^m D_i & v = \mathbf{val}(z) + \sum_{i\in J}v_i'\\ e_i = \overset{\smile}{\mathsf{B}}_{\mathbf{out}}(D_i)\ (\forall i\in 1..m) & \boldsymbol{x} = \biguplus_{i=1}^m \mathrm{fv}(e_i)\\ \mathcal{T}_i = \overset{\smile}{\mathsf{B}}_{\mathbf{D}}(D_i,D_i,\mathsf{T}_{init},0,v,\mathsf{PartG},0)\ (\forall i\in 1..m) & \end{array}}{\overset{\smile}{\mathsf{B}}_{\mathbf{adv}}(\{G\}^z C,t) = \mathsf{T}_{init}\mathcal{T}_1\cdots\mathcal{T}_m}$$

T_{init}
in: $\begin{array}{l}0\mapsto\mathbf{txout}(z)\\ i+1\mapsto\mathbf{txout}(y_i)\ (\forall i\in J)\end{array}$
wit: $\bot$
out: $\big(\lambda\boldsymbol{x}.\,\bigvee_{i=1}^m e_i,\ v\big)$
absLock: t

$$\dfrac{D = *:\mathtt{rngt}\ \mathtt{X}\langle\boldsymbol{\mathcal{E}}\rangle}{\overset{\smile}{\mathsf{B}}_{\mathbf{D}}(D,D_p,\mathsf{T},o,v,\mathcal{P},t) = \epsilon \quad (\text{empty sequence})}$$

Fig. 5: Two compilation rules (see [10] for the others).