

Formal investigation of the Extended UTxO model

Laying the foundations for the formal verification of smart contracts

ORESTIS MELKONIAN, Utrecht University, The Netherlands

This report serves as the proposal of my MSc thesis, supervised by Wouter Swierstra from Utrecht University and Manuel Chakravarty from IOHK.

1 INTRODUCTION

Blockchain technology has opened a whole array of interesting new applications (e.g. secure multi-party computation[Andrychowicz et al. 2014], fair protocol design fair[Bentov and Kumaresan 2014], zero-knowledge proof systems[Goldreich et al. 1991]). Nonetheless, reasoning about the behaviour of such systems is an exceptionally hard task, mainly due to their distributed nature. Moreover, the fiscal nature of the majority of these applications requires a much higher degree of rigorosity compared to conventional IT applications, hence the need for a more formal account of their behaviour.

The advent of smart contracts (programs that run on the blockchain itself) gave rise to another source of vulnerabilities. One primary example of such a vulnerability caused by the use of smart contracts is the DAO attack¹, where a security flaw on the model of Ethereum’s scripting language led to the exploitation of a venture capital fund worth 150 million dollars at the time. The solution was to create a hard fork of the Ethereum blockchain, clearly going against the decentralized spirit of cryptocurrencies. Since these (possibly Turing-complete) programs often deal with transactions of significant funds, it is of utmost importance that one can reason and ideally provide formal proofs about their behaviour in a concurrent/distributed setting.

Research Question. The aim of this thesis is to provide a mechanized formal model of an abstract distributed ledger equipped with smart contracts, in which one can begin to formally investigate the expressiveness of the extended UTxO model. Moreover, we hope to lay a foundation for a formal comparison with account-based models used in Ethereum. Put concisely, the research question posed is:

*How much expressiveness do we gain by extending the UTxO model?
Is it as expressive as the account-based model used in Ethereum?*

Overview. Section 2 reviews some basic definitions related to blockchain technology and introduces important literature, which will be the main subject of study throughout the development of our reasoning framework. Section 3 describes the technology we will use to formally reason about the problem at hand and some key design decisions we set upfront. Section 4 presents the progress made thus far in terms of (mechanized) formal verification, as well as problems we have encountered and also expect along the way. Section 5 discusses next steps for the remainder of the thesis, as well as a rough estimate on when these milestones will be completed.

¹[https://en.wikipedia.org/wiki/The_DAO_\(organization\)](https://en.wikipedia.org/wiki/The_DAO_(organization))

2 BACKGROUND

2.1 Distributed Ledger Technology: Blockchain

Cryptocurrencies rely on distributed ledgers, where there is no central authority managing the accounts and keeping track of the history of transactions.

One particular instance of distributed ledgers are blockchain systems, where transactions are bundled together in blocks, which are linearly connected with hashes and distributed to all peers. The blockchain system, along with a consensus protocol deciding on which competing fork of the chain is to be included, maintains an immutable distributed ledger (i.e. the history of transactions).

Validity of the transactions is tightly coupled with a consensus protocol, which makes sure peers in the network only validate well-behaved and truthful transactions and are, moreover, properly incentivized to do so.

The absence of a single central authority that has control over all assets of the participants allows for shared control of the evolution of data (in this case transactions) and generally leads to more robust and fair management of assets.

While cryptocurrencies are the major application of blockchain systems, one could easily use them for any kind of valuable asset, or even as general distributed databases.

2.2 Smart Contracts

Most blockchain systems come equipped with a scripting language, where one can write *smart contracts* that dictate how a transaction operates. A smart contract could, for instance, pose restrictions on who can redeem the output funds of a transaction.

One could view smart contracts as a replacement of legal frameworks, providing the means to conduct contractual relationships completely algorithmically.

While previous work on writing financial contracts [Peyton Jones et al. 2000] suggests it is fairly straightforward to write such programs embedded in a general-purpose language (in this case Haskell) and to reason about them with *equational reasoning*, it is restricted in the centralized setting and, therefore, does not suffice for our needs.

Things become much more complicated when we move to the distributed setting of a blockchain [Bhargavan et al. 2016; Sergey et al. 2018; Setzer 2018]. Hence, there is a growing need for methods and tools that will enable tractable and precise reasoning about such systems.

Numerous scripting languages have appeared recently [Seijas et al. 2016], spanning a wide spectrum of expressiveness and complexity. While language design can impose restrictions on what a language can express, most of these restrictions are inherited from the accounting model to which the underlying system adheres.

In the next section, we will discuss the two main forms of accounting models:

- (1) **UTxO-based:** stateless models based on *unspent transaction outputs*
- (2) **Account-based:** stateful models that explicitly model interaction between *user accounts*

2.3 UTxO-based: Bitcoin

The primary example of a UTxO-based blockchain is Bitcoin [Nakamoto 2008]. Its blockchain is a linear sequence of *blocks* of transactions, starting from the initial *genesis* block. Essentially, the blockchain acts as a public log of all transactions that have taken place, where each transaction refers to outputs of previous transactions, except for the initial *coinbase* transaction of each block. Coinbase transactions have no inputs, create new currency and reward the miner of that block with a fixed amount. Bitcoin also provides a cryptographic protocol to make sure no adversary can tamper with the transactional history, e.g. by making the creation of new blocks computationally hard and invalidating the “truthful” chain statistically impossible.

Formal investigation of the Extended UTxO model

A crucial aspect of Bitcoin's design is that there are no explicit addresses included in the transactions. Rather, transaction outputs are actually program scripts, which allow someone to claim the funds by giving the proper inputs. Thus, although there are no explicit user accounts in transactions, the effective available funds of a user are all the *unspent transaction outputs* (UTxO) that he can claim (e.g. by providing a digital signature).

2.3.1 SCRIPT. In order to write such scripts in the outputs of a transaction, Bitcoin provides a low-level, Forth-like, stack-based scripting language, called SCRIPT. SCRIPT is intentionally not Turing-complete (e.g. it does not provide looping structures), in order to have more predictable behaviour. Moreover, only a very restricted set of “template” programs are considered standard, i.e. allowed to be relayed from node to node.

SCRIPT Notation. Programs in script are a linear sequence of either data values (e.g. numbers, hashes) or built-in operations (distinguished by their `OP_` prefix).

The stack is initially considered empty and we start reading inputs from left to right. When we encounter a data item, we simply push it to the top of the stack. On encountering an operation, we pop the necessary number of arguments from the stack, apply the operation and push the result back. The evaluation function $\llbracket _ \rrbracket$ executes the given program and returns the final result at the top of the stack. For instance, adding two numbers looks like this:

$$\llbracket [1\ 2\ \text{OP_ADD}] \rrbracket = 3$$

P2PKH. The most frequent example of a ‘standard’ program in SCRIPT is the *pay-to-pubkey-hash* (P2PKH) type of scripts. Given a hash of a public key `<pub#>`, a P2PKH output carries the following script:

`OP_DUP OP_HASH <pub#> OP_EQ OP_CHECKSIG`

where `OP_DUP` duplicates the top element of the stack, `OP_HASH` replaces the top element with its hash, `OP_EQ` checks that the top two elements are equal, `OP_CHECKSIG` verifies that the top two elements are a valid pair of a digital signature of the transaction data and a public key hash.

The full script will be run when the output is claimed (i.e. used as input in a future transaction) and consists of the P2PKH script, preceded by the digital signature of the transaction by its owner and a hash of his public key. Given a digital signature `<sig>` and a public key hash `<pub>`, a transaction is valid when the execution of the script below evaluates to True.

`<sig> <pub> OP_DUP OP_HASH <pub#> OP_EQ OP_CHECKSIG`

To clarify, assume a scenario where Alice want to pay Bob ₿ 10. Bob provides Alice with the cryptographic hash of his public key (`<pub#>`) and Alice can submit a transaction of ₿ 10 with the following output script:

`OP_DUP OP_HASH <pub#> OP_EQ OP_CHECKSIG`

After that, Bob can submit another transaction that uses this output by providing the digital signature of the transaction `<sig>` (signed with his private key) and his public key `<pub>`. It is easy to see that the resulting script evaluates to True.

P2SH. A more complicated script type is *pay-to-script-hash* (P2SH), where output scripts simply authenticate against a hash of a *redeemer* script `<red#>`:

`OP_HASH <red#> OP_EQ`

A redeemer script $\langle \text{red} \rangle$ resides in an input which uses the corresponding output. The following two conditions must hold for the transaction to go through:

- (1) $\llbracket \langle \text{red} \rangle \rrbracket = \text{True}$
- (2) $\llbracket \langle \text{red} \rangle \text{ OP_HASH } \langle \text{red}\# \rangle \text{ OP_EQ} \rrbracket = \text{True}$

Therefore, in this case the script residing in the output is simpler, but inputs can also contain arbitrary redeemer scripts (as long as they are of a standard “template”).

In this thesis, we will model scripts in a much more general, mathematical sense, so we will eschew from any further investigation of properties particular to `SCRIPT`.

2.3.2 The BitML Calculus. Although Bitcoin is the most widely used blockchain to date, many aspects of it are poorly documented. In general, there is a scarcity of formal models, most of which are either introductory or exploratory.

One of the most involved and mature previous work on formalizing the operation of Bitcoin is the Bitcoin Modelling Language (BitML) [Bartoletti and Zunino 2018]. First, an idealistic *process calculus* that models Bitcoin contracts is introduced, along with a detailed small-step reduction semantics that models how contracts interact and its non-determinism accounts for the various outcomes.

The semantics consist of transitions between *configurations*, abstracting away all the cryptographic machinery and implementation details of Bitcoin. Consequently, such operational semantics allow one to reason about the concurrent behaviour of the contracts in a *symbolic* setting.

The authors then provide a compiler from BitML contracts to ‘standard’ Bitcoin transactions, proven correct via a correspondence between the symbolic model and the computational model operating on the Bitcoin blockchain. We will return for a more formal treatment of BitML in Section 4.2.

2.3.3 Extended UTxO. In this work, we will consider the version of the UTxO model used by IOHK’s Cardano blockchain². In contrast to Bitcoin’s *proof-of-work* consensus protocol [Nakamoto 2008], Cardano’s *Ouroboros* protocol [Kiayias et al. 2017] is *proof-of-stake*. This, however, does not concern our study of the abstract accounting model, thus we refrain from formally modelling and comparing different consensus techniques.

The actual extension we care about is the inclusion of *data scripts* in transaction outputs, which essentially provide the validation script in the corresponding input with additional information of an arbitrary type.

This extension of the UTxO model has already been implemented³, but only informally documented⁴. The reason to extend the UTxO model with data scripts is to bring more expressive power to UTxO-based blockchains, hopefully bringing it on par with Ethereum’s account-based scripting model (see Section 2.4).

However, there is no formal argument to support this claim, and it is the goal of this thesis to provide the first formal investigation of the expressiveness introduced by this extension.

2.4 Account-based: Ethereum

On the other side of the spectrum, lies the second biggest cryptocurrency today, Ethereum [Buterin et al. 2014]. In contrast to UTxO-based systems, Ethereum has a built-in notion of user addresses and operates on a stateful accounting model. It goes even further to distinguish *human accounts* (controlled by a public-private key pair) from *contract accounts* (controlled by some EVM code).

²www.cardano.org

³<https://github.com/input-output-hk/plutus/tree/master/wallet-api/src/Ledger>

⁴<https://github.com/input-output-hk/plutus/blob/master/docs/extended-utxo/README.md>

Formal investigation of the Extended UTxO model

This added expressiveness is also reflected in the quasi-Turing-complete low-level stack-based bytecode language in which contract code is written, namely the *Ethereum Virtual Machine* (EVM). EVM is mostly designed as a target, to which other high-level user-friendly languages will compile to.

Solidity. The most widely adopted language that targets the EVM is *Solidity*, whose high-level object-oriented design makes writing common contract use-cases (e.g. crowdfunding campaigns, auctions) rather straightforward.

One of *Solidity*'s most distinguishing features is the concept of a contract's *gas*; a limit to the amount of computational steps a contract can perform. At the time of the creation of a transaction, its owner specifies a certain amount of gas the contract can consume and pays a transaction fee proportional to it. In case of complete depletion (i.e. all gas has been consumed before the contract finishes its execution), all global state changes are reverted as if the contract had never been run. This is a necessary ingredient for smart contract languages that provide arbitrary looping behaviour, since non-termination of the validation phase is certainly undesirable.

If time permits, we will initially provide a formal justification of *Solidity* and proceed to formally compare the extended UTxO model against it. Since *Solidity* is a fully-fledged programming language with lots of features (e.g. static typing, inheritance, libraries, user-defined types), it makes sense to restrict our formal study to a compact subset of *Solidity* that is easy to reason about. This is the approach also taken in Featherweight Java [Igarashi et al. 2001]; a subset of Java that omits complex features such as reflection, in favour of easier behavioural reasoning and a more formal investigation of its semantics. In the same vein, we will try to introduce a lightweight version of *Solidity*, which we will refer to as *Featherweight Solidity*.

3 METHODOLOGY

3.1 Scope

At this point, we have to stress the fact that we are not aiming for a formalization of a fully-fledged blockchain system with all its bells and whistles, but rather focus on the underlying accounting model. Therefore, we will omit details concerning cryptographic operations and aspects of the actual implementation of such a system. Instead, we will work on an abstract layer that postulates the well-behavedness of these subcomponents, which will hopefully lend itself to more tractable reasoning and give us a clear overview of the essence of the problem.

Restricting the scope of our attempt is also motivated from the fact that individual components such as cryptographic protocols are orthogonal to the functionality we study here. This lack of tight cohesion between the components of the system allows one to safely factor each one out and formalize it independently.

It is important to note that this is not always the case for every domain. A prominent example of this are operating systems, which consist of intricately linked subcomponents (e.g. drivers, memory modules), thus making it impossible to trivially divide the overall proof into small independent ones. In order to overcome this complexity burden, one has to invent novel ways of modular proof mechanization, as exemplified by *CertiKOS* [Chen et al. 2016], a formally verified concurrent OS.

3.2 Proof Mechanization

Fortunately, the sub-components of the system we are examining are not no interdependent, thus lending themselves to separate treatment. Nonetheless, the complexity of the sub-system we care about is still high and requires rigorous investigation. Therefore, we choose to conduct our formal

study in a mechanized manner, i.e. using a proof assistant along the way and formalizing all results in Type Theory. Proof mechanization will allow us to discover edge cases and increase the confidence of the model under investigation.

3.3 Agda

As our proof development vehicle, we choose Agda [Norell 2008], a dependently-typed total functional language similar to Haskell [Hudak et al. 1992].

Agda embraces the *Curry-Howard correspondence*, which states that types are isomorphic to statements in (intuitionistic) logic and their programs correspond to the proofs of these statements [Martin-Löf and Sambin 1984]. Through its unicode-based *mixfix* notational system, one can easily translate a mathematical theorem into a valid Agda type. Moreover, programs and proofs share the same structure, e.g. induction in the proof manifests itself as recursion in the program.

While Agda is not ideal for large software development, its flexible notation and elegant design is suitable for rapid prototyping of new ideas and exploratory purposes. We do not expect to hit such problems, since we will stay on a fairly abstract level which postulates cryptographic operations and other implementation details.

Limitation. The main limitation of Agda lies in its lack of a proper proof automation system. While there has been work on providing Agda with such capabilities [Kokke and Swierstra 2015], it requires moving to a meta-programming mindset which would be an additional programming hindrance.

A reasonable alternative would be to use Coq [Barras et al. 1997], which provides a pragmatic scripting language for programming *tactics*, i.e. programs that work on proof contexts and can generate new sub-goals. This approach to proof mechanization has, however, been criticized for widening the gap between informal proofs and programs written in a proof assistant. This clearly goes against the aforementioned principle of *proofs-as-programs*.

3.4 The IOHK approach

At this point, we would like to mention the specific approach taken by IOHK⁵. In contrast to numerous other companies currently creating cryptocurrencies, its main focus is on provably correct protocols with a strong focus on peer-reviewing and robust implementations, rather than fast delivery of results. This is evidenced by the choice of programming languages (Agda/Coq/Haskell/Scala) – all functional programming languages with rich type systems – and the use of *property-based testing* [Claessen and Hughes 2011] for the production code.

IOHK’s distinct feature is that it advocates a more rigorous development pipeline; ideas are initially worked on paper by pure academics, which create fertile ground for starting formal verification in Agda/Coq for more confident results, which result in a prototype/reference implementation in Haskell, which informs the production code-base (also written in Haskell) on the properties that should be tested.

Since this thesis is done in close collaboration with IOHK, it is situated on the second step of aforementioned pipeline; while there has been work on writing papers about the extended UTxO model along with the actual implementation in Haskell, there is still no complete and mechanized account of its properties.

3.5 Functional Programming Principles

One last important manifestation of the functional programming principles behind IOHK is the choice of a UTxO-based cryptocurrency itself.

⁵<https://iohk.io/>

Formal investigation of the Extended UTxO model

On the one hand, one can view a UTxO ledger as a dataflow diagram, whose nodes are the submitted transactions and edges represent links between transaction inputs and outputs. On the other hand, account-based ledgers rely on a global state and transaction have a much more complicated specification.

The key point here is that UTxO-based transaction are just pure mathematical functions, which are much more straightforward to model and reason about. Coming back to the principles of functional programming, one could contrast this with the difference between functional and imperative programs. One can use *equational reasoning* for functional programs, due to their *referential transparency*, while this is not possible for imperative programs that contain side-effectful commands. Therefore, we hope that these principles will be reflected in the proof process itself; one would reason about purely functional UTxO-based ledgers in a compositional manner.

4 PRELIMINARY RESULTS

This section gives an overview of the progress made so far in the on-going Agda formalization of the two main subjects of study, namely the Extended UTxO model and the BitML calculus. For the sake of brevity, we refrain from showing the full Agda code along with the complete proofs, but rather provide the most important datatypes and formalized results and explain crucial design choices we made along the way. Furthermore, we will omit notational burden imposed by technicalities particular to Agda, such as *universe polymorphism* and *proof irrelevance*.

4.1 Formal Model I: Extended UTxO

We now set out to model the accounting model of a UTxO-based ledger. We will provide a inherently-typed model of transactions and ledgers; this gives rise to a notion of *weakening* of available addresses, which we formalize. Moreover, we showcase the reasoning abilities of our model by giving an example of a correct-by-construction ledger. All code is publicly available on Github⁶.

We start with the basic types, keeping them abstract since we do not care about details arising from the encoding in an actual implementation:

postulate

Address : Set

Value : Set

$\mathbb{B} : \mathbb{N} \rightarrow \textit{Value}$

We assume there are types representing addresses and bitcoin values, but also require the ability to construct a value out of a natural number. In the examples that follow, we assume the simplest representation, where both types are the natural numbers.

There is also the notion of the *state* of a ledger, which will be provided to transaction scripts and allow them to have stateful behaviour for more complicated schemes (e.g. imposing time constraints).

record *State* : Set **where**

field *height* : $\mathbb{N}$

$\vdots$

The state components have not been finalized yet, but can easily be extended later when we

⁶<https://github.com/omelkonian/formal-utxo>

actually investigate examples with expressive scripts that make use of state information, such as the current length of the ledger (*height*).

As mentioned previously, we will not dive into the verification of the cryptological components of the model, hence we postulate an *irreversible* hashing function which, given any value of any type, gives back an address (i.e. a natural number) and is furthermore injective (i.e. it is highly unlikely for two different values to have the same hash).

postulate

```

$$\begin{aligned} \_ \# &: \forall \{A : \text{Set}\} \rightarrow A \rightarrow \text{Address} \\ \# \text{-injective} &: \forall \{x\ y : A\} \rightarrow x \# \equiv y \# \rightarrow x \equiv y \end{aligned}$$

```

4.1.1 Transactions. In order to model transactions that are part of a distributed ledger, we need to first define transaction *inputs* and *outputs*.

```
record TxOutputRef : Set where
  constructor _@_
  field id      : Address
         index : ℕ

record TxInput { R D : Set } : Set where
  field outputRef : TxOutputRef
         redeemer  : State → R
         validator : State → Value → R → D → Bool
```

Output references consist of the address that a transaction hashes to, as well as the index in this transaction's list of outputs. *Transaction inputs* refer to some previous output in the ledger, but also contain two types of scripts. The *redeemer* provides evidence of authorization to spend the output. The *validator* then checks whether this is so, having access to the current state of the ledger, the bitcoin output and data provided by the redeemer and the *data script* (residing in outputs). It is also noteworthy that we immediately model scripts by their *denotational semantics*, omitting unnecessary details relating to concrete syntax, lexing and parsing.

Transaction outputs send a bitcoin amount to a particular address, which either corresponds to a public key hash of a blockchain participant (P2PKH) or a hash of a next transaction's script (P2SH). Here, we opt to embrace the *inherently-typed* philosophy of Agda and model available addresses as module parameters. That is, we package the following definitions in a module with such a parameter, as shown below:

```
module UTxO (addresses : List Address) where
  record TxOutput { D : Set } : Set where
    field value      : Value
           address   : Index addresses
           dataScript : State → D

  record Tx : Set where
    field inputs  : Set⟨ TxInput ⟩
           outputs : List TxOutput
```


Formal investigation of the Extended UTxO model

$$\begin{aligned} \text{forge} & : \text{Value} \\ \text{fee} & : \text{Value} \\ \text{Ledger} & : \text{Set} \\ \text{Ledger} & = \text{List Tx} \end{aligned}$$

Transaction outputs consist of a bitcoin amount and the address (out of the available ones) this amount is sent to, as well as the data script, which provides extra information to the aforementioned validator and allows for more expressive schemes. Investigating exactly the extent of this expressiveness is one of the main goals of this thesis.

For a transaction to be submitted, one has to check that each input can actually spend the output it refers to. At this point of interaction, one must combine all scripts, as shown below:

$$\begin{aligned} \text{runValidation} & : (i : \text{TxInput}) \rightarrow (o : \text{TxOutput}) \rightarrow D \ i \equiv D \ o \rightarrow \text{State} \rightarrow \text{Bool} \\ \text{runValidation } i \ o \ \text{refl } st & = \text{validator } i \ st \ (\text{value } o) \ (\text{redeemer } i \ st) \ (\text{dataScript } o \ st) \end{aligned}$$

Note that the intermediate types carried by the respective input and output must align, evidenced by the equality proof that is required as an argument.

4.1.2 Unspent Transaction Outputs. With the basic modelling of a ledger and its transaction in place, it is fairly straightforward to inductively define the calculation of a ledger's unspent transaction outputs:

$$\begin{aligned} \text{unspentOutputs} & : \text{Ledger} \rightarrow \text{Set} \langle \text{TxOutputRef} \rangle \\ \text{unspentOutputs } [] & = \emptyset \\ \text{unspentOutputs } (tx :: txs) & = (\text{unspentOutputs } txs \setminus \text{spentOutputsTx } tx) \cup \text{unspentOutputsTx } tx \\ \text{where} \\ \text{spentOutputsTx}, \text{unspentOutputsTx} & : \text{Tx} \rightarrow \text{Set} \langle \text{TxOutputRef} \rangle \\ \text{spentOutputsTx} & = (\text{outputRef } \langle \$ \rangle _) \circ \text{inputs} \\ \text{unspentOutputsTx } tx & = ((tx \#) \ @ _) \langle \$ \rangle (\text{indices } (\text{outputs } tx)) \end{aligned}$$

4.1.3 Validity of Transactions. In order to submit a transaction, one has to make sure it is valid with respect to the current ledger. We model validity as a record indexed by the transaction to be submitted and the current ledger:

$$\begin{aligned} \text{record } \text{IsValidTx } (tx : \text{Tx}) \ (l : \text{Ledger}) & : \text{Set} \text{ where} \\ \text{field} \\ \text{validTxRefs} & : \\ \forall i \rightarrow i \in \text{inputs } tx \rightarrow \\ \text{Any } (\lambda t \rightarrow t \# \equiv \text{id } (\text{outputRef } i)) \ l \\ \\ \text{validOutputIndices} & : \\ \forall i \rightarrow (i \in : i \in \text{inputs } tx) \rightarrow \\ \text{index } (\text{outputRef } i) < \\ \text{length } (\text{outputs } (\text{lookupTx } l \ (\text{outputRef } i) \ (\text{validTxRefs } i \ i \in))) \end{aligned}$$

```

validOutputRefs :
   $\forall i \rightarrow i \in \text{inputs } tx \rightarrow$ 
     $\text{outputRef } i \in \text{unspentOutputs } l$ 

validDataScriptTypes :
   $\forall i \rightarrow (i \in :i \in \text{inputs } tx) \rightarrow$ 
     $D \ i \equiv D \ (\text{lookupOutput } l \ (\text{outputRef } i) \ (\text{validTxRefs } i \ i \in) \ (\text{validOutputIndices } i \ i \in))$ 

```

```

preservesValues :
   $\text{forge } tx + \text{sum} \ (\text{mapWith } \in (\text{inputs } tx) \ \lambda \ \{i\} \ i \in \rightarrow$ 
     $\text{lookupValue } l \ i \ (\text{validTxRefs } i \ i \in) \ (\text{validOutputIndices } i \ i \in))$ 
     $\equiv$ 
     $\text{fee } tx + \text{sum} \ (\text{value } <\$> \ \text{outputs } tx)$ 

noDoubleSpending :
   $\text{noDuplicates} \ (\text{outputRef } <\$> \ \text{inputs } tx)$ 

allInputsValidate :
   $\forall i \rightarrow (i \in :i \in \text{inputs } tx) \rightarrow$ 
    let  $\text{out} : \text{TxOutput}$ 
       $\text{out} = \text{lookupOutput } l \ (\text{outputRef } i) \ (\text{validTxRefs } i \ i \in) \ (\text{validOutputIndices } i \ i \in)$ 
    in  $\forall (st : \text{State}) \rightarrow$ 
       $T \ (\text{runValidation } i \ \text{out} \ (\text{validDataScriptTypes } i \ i \in) \ st)$ 

validateValidHashes :
   $\forall i \rightarrow (i \in :i \in \text{inputs } tx) \rightarrow$ 
    let  $\text{out} : \text{TxOutput}$ 
       $\text{out} = \text{lookupOutput } l \ (\text{outputRef } i) \ (\text{validTxRefs } i \ i \in) \ (\text{validOutputIndices } i \ i \in)$ 
    in  $\text{toN} \ (\text{address } \text{out}) \equiv (\text{validator } i) \ \#$ 

```

The first four conditions make sure the transaction references and types are well-formed, namely that inputs refer to actual transactions (*validTxRefs*, *validOutputIndices*) which are unspent so far (*validOutputRefs*), but also that intermediate types used in interacting inputs and outputs align (*validDataScriptTypes*).

The last four validation conditions are more interesting, as they ascertain the validity of the submitted transaction, namely that the bitcoin values sum up properly (*preservesValues*), no output is spent twice (*noDoubleSpending*), validation succeeds for each input-output pair (*allInputsValidate*) and outputs hash to the hash of their corresponding validator script (*validateValidHashes*).

The definitions of lookup functions are omitted, as they are uninteresting. The only important design choice is that, instead of modelling lookups as partial functions (i.e. returning *Maybe*), they require a membership proof as an argument moving the responsibility to the caller (as evidenced by their usage in the validity conditions).

Formal investigation of the Extended UTxO model

4.1.4 Weakening Lemma. We have defined everything with respect to a fixed set of available addresses, but it would make sense to be able to include additional addresses without losing the validity of the ledger constructed thus far.

In order to do, we need to first expose the basic datatypes from inside the module, introducing their *primed* version which takes the corresponding module parameter as an index:

```
Ledger' : List Address → Set
Ledger' as = Ledger
  where open import UTxO as
  :
```

We can now precisely define what it means to weaken an address space; one just adds more available addresses without removing any of the pre-existing addresses:

```
weakenTxOutput : Prefix as bs → TxOutput' as → TxOutput' bs
weakenTxOutput pr txOut = txOut { address = inject ≤ addr (prefix-length pr) }
  where open import UTxO bs
```

For simplicity's sake, we allow extension at the end of the address space instead of anywhere in between⁷. Notice also that the only place where weakening takes place are transaction outputs, since all other components do not depend on the available address space.

With the weakening properly defined, we can finally prove the *weakening lemma* for the available address space:

```
weakening : ∀ { as bs : List Address } { tx : Tx' as } { l : Ledger' as }
  → (pr : Prefix as bs)
  → IsValidTx' as tx l
  → IsValidTx' bs (weakenTx pr tx) (weakenLedger pr l)
weakening = ...
```

The weakening lemma states that the validity of a transaction with respect to a ledger is preserved if we choose to weaken the available address space, which we estimate to be useful when we later prove more intricate properties of the extended UTxO model.

4.1.5 Example. To showcase how we can use our model to construct *correct-by-construction* ledgers, let us revisit the example ledger presented in the Chimeric Ledgers paper [Zahmentferner and HK 2018].

Any blockchain can be visually represented as a *directed acyclic graph* (DAG), with transactions as nodes and input-output pairs as edges, as shown in Figure 1.

First, we need to set things up by declaring the list of available addresses and opening our module with this parameter.

⁷Technically, we require *Prefix as bs* instead of the more flexible *as ⊆ bs*.




Fig. 1. Example ledger with six transactions (unspent outputs are coloured in red)

addresses : *List Address*

addresses = 1 :: 2 :: 3 :: []

open import *UTxO* *addresses*

dummyValidator : *State* → *Value* → ℕ → ℕ → *Bool*

dummyValidator = λ _ _ _ _ → *true*

withScripts : *TxOutputRef* → *TxInput*

withScripts *tin* = **record** { *outputRef* = *tin*
 ; *redeemer* = λ _ → 0
 ; *validator* = *dummyValidator*
 }

₿ _ @ _ : *Value* → *Index* *addresses* → *TxOutput*

₿ *v* @ *addr* = **record** { *value* = *v*
 ; *address* = *addr*
 ; *dataScript* = λ _ → 0
 }

postulate

validator # : ∀ { *i* : *Index* *addresses* } → *toN* *i* ≡ *dummyValidator* #

Note that, since we will not utilize the expressive power of scripts in this example, we also provide convenient short cuts for defining inputs and outputs with dummy default scripts. Furthermore, we postulate that the addresses are actually the hashes of validators scripts, corresponding to the P2SH scheme in Bitcoin.

We can then proceed to define the individual transactions⁸ depicted in Figure 1:

⁸ The first sub-index of each variable refers to the order the transaction are submitted, while the second sub-index refers to which output of the given transaction we select.

Formal investigation of the Extended UTxO model

```

 $t_1, t_2, t_3, t_4, t_5, t_6 : Tx$ 
 $t_1 = \text{record } \{ \text{inputs} = []$ 
     $;\text{ outputs} = [\text{\$ } 1000 @ 0]$ 
     $;\text{ forge} = \text{\$ } 1000$ 
     $;\text{ fee} = \text{\$ } 0$ 
 $\}$ 

 $t_2 = \text{record } \{ \text{inputs} = [\text{withScripts } t_{10}]$ 
     $;\text{ outputs} = \text{\$ } 800 @ 1 :: \text{\$ } 200 @ 0 :: []$ 
     $;\text{ forge} = \text{\$ } 0$ 
     $;\text{ fee} = \text{\$ } 0$ 
 $\}$ 

 $t_3 = \text{record } \{ \text{inputs} = [\text{withScripts } t_{21}]$ 
     $;\text{ outputs} = [\text{\$ } 199 @ 2]$ 
     $;\text{ forge} = \text{\$ } 0$ 
     $;\text{ fee} = \text{\$ } 1$ 
 $\}$ 

 $t_4 = \text{record } \{ \text{inputs} = [\text{withScripts } t_{30}]$ 
     $;\text{ outputs} = [\text{\$ } 207 @ 1]$ 
     $;\text{ forge} = \text{\$ } 10$ 
     $;\text{ fee} = \text{\$ } 2$ 
 $\}$ 

 $t_5 = \text{record } \{ \text{inputs} = \text{withScripts } t_{20} :: \text{withScripts } t_{40} :: []$ 
     $;\text{ outputs} = \text{\$ } 500 @ 1 :: \text{\$ } 500 @ 2 :: []$ 
     $;\text{ forge} = \text{\$ } 0$ 
     $;\text{ fee} = \text{\$ } 7$ 
 $\}$ 

 $t_6 = \text{record } \{ \text{inputs} = \text{withScripts } t_{50} :: \text{withScripts } t_{51} :: []$ 
     $;\text{ outputs} = [\text{\$ } 999 @ 2]$ 
     $;\text{ forge} = \text{\$ } 0$ 
     $;\text{ fee} = \text{\$ } 1$ 
 $\}$ 

```

Finally, we can construct a correct-by-construction ledger, by iteratively submitting each transaction along with the proof that it is valid with respect to the ledger constructed thus far⁹:

⁹ Here, we use a specialized notation of the form $\bullet t_1 \vdash p_1 \oplus t_2 \vdash p_2 \oplus \dots \oplus t_n \vdash p_n$, where each insertion of transaction t_x requires a proof of validity p_x as well. Technically, the $\oplus$ operator has type $(l : \text{Ledger}) \rightarrow (t : Tx) \rightarrow \text{IsValidTx } t \ l \rightarrow \text{Ledger}$.

ex-ledger : *Ledger*
ex-ledger = • $t_1 \vdash$ **record** { *validTxRefs* = $\lambda i ()$
 ; *validOutputIndices* = $\lambda i ()$
 ; *validOutputRefs* = $\lambda i ()$
 ; *validDataScriptTypes* = $\lambda i ()$
 ; *preservesValues* = *refl*
 ; *noDoubleSpending* = *tt*
 ; *allInputsValidate* = $\lambda i ()$
 ; *validateValidHashes* = $\lambda i ()$
 }
 $\oplus t_2 \vdash$ **record** { ... }
 $\oplus t_3 \vdash$ **record** { ... }
 $\oplus t_4 \vdash$ **record** { ... }
 $\oplus t_5 \vdash$ **record** { ... }
 $\oplus t_6 \vdash$ **record** { ... }

The proof validating the submission of the first transaction t_1 is trivially discharged. While the rest of the proofs are quite involved, it is worthy to note that their size/complexity stays constant independently of the ledger length. This is mainly due to the re-usability of proof components, arising from the main functions being inductively defined.

It is now trivial to verify that the only unspent transaction output of our ledger is the output of the last transaction t_6 , as demonstrated below:

utxo : *list* (*unspentOutputs ex-ledger*) $\equiv [t_{60}]$
utxo = *refl*

4.2 Formal Model II: BitML Calculus

Now let us shift our focus to our second subject of study, the BitML calculus for modelling smart contracts. In this subsection we sketch the formalized part of BitML we have covered so far, namely the syntax and small-step semantics of BitML contracts, as well as an example execution of a contract under these semantics. All code is publicly available on Github¹⁰.

First, we begin with some basic definitions that will be used throughout this section:

module *Types* (*Participant* : *Set*) (*Honest* : *List*⁺ *Participant*) **where**

Time : *Set*

Time = $\mathbb{N}$

Value : *Set*

Value = $\mathbb{N}$

¹⁰<https://github.com/omelkonian/formal-bitml>

Formal investigation of the Extended UTxO model

```
record Deposit : Set where  
  constructor _has_  
  field participant : Participant  
        value : Value
```

```
Secret : Set  
Secret = String
```

```
data Arith : List Secret → Set where ...  
N[] : ∀ {s} → Arith s → ℕ  
N[] = ...
```

```
data Predicate : List Secret → Set where ...  
B[] : ∀ {s} → Predicate s → Bool  
B[] = ...
```

Instead of giving a fixed datatype of participants, we parametrise our module with a given *universe* of participants and a non-empty list of honest participants. Representation of time and monetary values is again done using natural numbers, while we model participant secrets as simple strings¹¹. A deposits consists of the participant that owns it and the number of bitcoins it carries. We, furthermore, introduce a simplistic language of logical predicates and arithmetic expressions with the usual constructs (e.g. numerical addition, logical conjunction) and give the usual semantics (predicates on booleans and arithmetic on naturals). A more unusual feature of these expressions is the ability to calculate length of secrets (within arithmetic expressions) and, in order to ensure more type safety later on, all expressions are indexed by the secrets they internally use.

4.2.1 Contracts in BitML. A *contract advertisement* consists of a set of *preconditions*, which require some resources from the involved participants prior to the contract's execution, and a *contract*, which specifies the rules according to which bitcoins are transferred between participants.

Preconditions either require participants to have a deposit of a certain value on their name (volatile or not) or commit to a certain secret. Notice the index of the datatype below, which captures the values of all required deposits:

```
data Precondition : List Value → Set where  
  -- volatile deposit  
  _?_ : Participant → (v : Value) → Precondition [v]  
  -- persistent deposit  
  _!_ : Participant → (v : Value) → Precondition [v]  
  -- committed secret  
  _#_ : Participant → Secret → Precondition []  
  -- conjunction  
  _^_ : Precondition vsl → Precondition vsr → Precondition (vsl ++ vsr)
```

¹¹ Of course, one could provide more realistic types (e.g. words of specific length) to be closer to the implementation, as shown for the UTxO model in Section 4.1.

Moving on to actual contracts, we define them by means of a collection of five types of commands; *put* injects participant deposits and revealed secrets in the remaining contract, *withdraw* transfers the current funds to a participant, *split* distributes the current funds across different individual contracts, *_ : _* requires the authorization from a participant to proceed and *after _ : _* allows further execution of the contract only after some time has passed.

```

data Contract : Value -- the monetary value it carries
                → Values -- the deposits it presumes
                → Set where
    -- collect deposits and secrets
    put _ reveal _ if _ ⇒ _ ⊢ _ :
      (vs : List Value) → (s : Secrets) → Predicate s' → Contract (v + sum vs) vs' → s' ⊆ s
      → Contract v (vs' ++ vs)
    -- transfer the remaining balance to a participant
    withdraw : ∀ {v} → Participant → Contract v []
    -- split the balance across different branches
    split : (cs : List (∃ [v] ∃ [vs] Contract v vs))
      → Contract (sum (proj1 <$> cs)) (concat (proj2 <$> cs))
    -- wait for participant's authorization
    _ : _ : Participant → Contract v vs → Contract v vs
    -- wait until some time passes
    after _ : _ : Time → Contract v vs → Contract v vs

```

There is a lot of type-level manipulation across all constructors, since we need to make sure that indices are calculated properly. For instance, the total value in a contract constructed by the *split* command is the sum of the values carried by each branch. The *put* command¹² additionally requires an explicit proof that the predicate of the *if* part only uses secrets revealed by the same command.

We also introduce an intuitive syntax for declaring the different branches of a *split* command, emphasizing the *linear* nature of the contract's total monetary value:

```

_ ⊖ _ : ∀ {vs : Values} → (v : Value) → Contract v vs → ∃ [v] ∃ [vs] Contract v vs
_ ⊖ _ {vs} v c = v, vs, c

```

Having defined both preconditions and contracts, we arrive at the definition of a contract advertisement:

```

record Advertisement (v : Value) (vsc vsg : List Value) : Set where
  constructor _⟨_⟩ ⊢ _
  field G      : Precondition vs
  C          : Contracts v vs
  valid      : length vsc ≤ length vsg
              × participantsg G ++ participantsc C ⊆ (participant <$> persistentDepositsp G)

```

Notice that in order to construct an advertisement, one has to also provide proof of the contract's

¹² *put* comprises of several components and we will omit those that do not contain any helpful information, e.g. write *put _ ⇒ _* when there are no revealed secrets and the predicate trivially holds.

validity with respect to the given preconditions, namely that all deposit references in the contract are declared in the precondition and each involved participant is required to have a persistent deposit.

To clarify things so far, let us see a simple example of a contract advertisement:

```

open BitML (A | B) [A]+
ex-ad : Advertisement 5 [200] (200 :: 100 :: [])
ex-ad = ⟨ B! 200 ∧ A! 100 ⟩
      split ( 2  $\multimap$  withdraw B
            ⊕ 2  $\multimap$  after 100 : withdraw A
            ⊕ 1  $\multimap$  put [200] ⇒ B : withdraw {201} A ⊢ ...
      )
    ⊢ ...

```

We first need to open our module with a fixed set of participants (in this case *A* and *B*). We then define an advertisement, whose type already says a lot about what is going on; it carries ₧ 5, presumes the existence of at least one deposit of ₧ 200, and requires two deposits of ₧ 200 and ₧ 100.

Looking at the precondition itself, we see that the required deposits will be provided by *B* and *A*, respectively. The contract first splits the bitcoins across three branches: the first one gives ₧ 2 to *B*, the second one gives ₧ 2 to *A* after some time period, while the third one retrieves *B*'s deposit of ₧ 200 and allows *B* to authorise the withdrawal of the remaining funds (currently ₧ 201) from *A*.

We have omitted the proofs that ascertain the well-formedness of the *put* command and the advertisement, as they are straightforward and do not provide any more intuition¹³.

4.2.2 Small-step Semantics. BitML is a *process calculus*, which is geared specifically towards smart contracts. Contrary to most process calculi that provide primitive operators for inter-process communication via message-passing [Hoare 1978], the BitML calculus does not provide such built-in features.

It, instead, provides domain-specific synchronization mechanisms through its *small-step* reduction semantics. These essentially define a *labelled transition system* between *configurations*, where *action* labels are emitted on every transition and represent the required actions of the participants. This symbolic model consists of two layers; the bottom one transitioning between *untimed* configurations and the top one that works on *timed* configurations.

We start with the datatype of actions, which showcases the principal actions required to satisfy an advertisement's preconditions and an action to pick one branch of a collection of contracts (introduced by the choice operator $\oplus$). We have omitted uninteresting actions concerning the manipulation of deposits, such as dividing, joining, donating and destroying them. Since we will often need versions of the types of advertisements/contracts with their indices existentially quantified, we first provide aliases for them.

¹³ In fact, we have defined decidable procedures for all such proofs using the *proof-by-reflection* pattern [Van Der Walt and Swierstra 2012]. These automatically discharge all proof obligations, when there are no variables involved.

AdvertisedContracts : Set

AdvertisedContracts = List (∃[v] ∃[vs^c] ∃[vs^g] Advertisement v vs^c vs^g)

ActiveContracts : Set

ActiveContracts = List (∃[v] ∃[vs] List (Contract v vs))

```
data Action (p : Participant) -- the participant that authorises this action
  :   AdvertisedContracts -- the contract advertisements it requires
    → ActiveContracts      -- the active contracts it requires
    → Values               -- the deposits it requires from this participant
    → List Deposit         -- the deposits it produces
    → Set where

    -- commit secrets to stipulate an advertisement
    #▷ _ : (ad : Advertisement v vsc vsg)
          → Action p [v, vsc, vsg, ad] [] [] []

    -- spend x to stipulate an advertisement
    _▷s_ : (ad : Advertisement v vsc vsg)
          → (i : Index vsg)
          → Action p [v, vsc, vsg, ad] [] [vsg !! i] []

    -- pick a branch
    _▷b_ : (c : List (Contract v vs))
          → (i : Index c)
          → Action p [] [v, vs, c] [] []

    ∴
```

The action datatype is parametrised¹⁴ over the participant who performs it and includes several indices representing the prerequisites the current configuration has to satisfy, in order for the action to be considered valid (e.g. one cannot spend a deposit to stipulate an advertisement that does not exist).

The first index refers to advertisements that appear in the current configuration, the second to contracts that have already been stipulated, the third to deposits owned by the participant currently performing the action and the fourth declares new deposits that will be created by the action (e.g. dividing a deposit would require a single deposit as the third index and produce two other deposits in its fourth index).

Although our indexing scheme might seem a bit heavyweight now, it makes many little details and assumptions explicit, which would bite us later on when we will need to reason about them.

Continuing from our previous example advertisement, let's see an example action where *A* spends the required ₧ 100 to stipulate the example contract¹⁵:

¹⁴ In Agda, datatype parameters are similar to indices, but are not allowed to vary across constructors.

¹⁵ Notice that we have to make all indices of the advertisement explicit in the second index in the action's type signature.

Formal investigation of the Extended UTxO model

$ex-spend : Action\ A\ [5, [200], 200 :: 100 :: [], ex-ad]\ []\ [100]\ []$
 $ex-spend = ex-ad \triangleright^s 1$

Configurations are now built from advertisements, active contracts, deposits, action authorizations and committed/revealed secrets:

```
data Configuration' : --      current      ×      required
                        AdvertisedContracts × AdvertisedContracts
                        → ActiveContracts   × ActiveContracts
                        → List Deposit       × List Deposit
                        → Set where

-- empty
∅ : Configuration' ([], []) ([], []) ([], [])

-- contract advertisement
' : (ad : Advertisement v vsc vsg)
    → Configuration' ([v, vsc, vsg, ad], []) ([], []) ([], [])

-- active contract
⟨_,_⟩c : (c : List (Contract v vs)) → Value
        → Configuration' ([], []) ([v, vs, c], []) ([], [])

-- deposit redeemable by a participant
⟨_,_⟩d : (p : Participant) → (v : Value)
        → Configuration' ([], []) ([], []) ([p has v], [])

-- authorization to perform an action
_[]_ : (p : Participant) → Action p ads cs vs ds
      → Configuration' ([], ads) ([], cs) (ds, ((p has _) <$> vs))

-- committed secret
⟨_:_#_⟩ : Participant → Secret → ℕ ⊔ ⊥
        → Configuration' ([], []) ([], []) ([], [])

-- revealed secret
_:_#_ : Participant → Secret → ℕ
      → Configuration' ([], []) ([], []) ([], [])

-- parallel composition
_ | _ : Configuration' (adsl, radsl) (csl, rcsl) (dsl, rdsl)
      → Configuration' (adsr, radsr) (csr, rcsr) (dsr, rdsr)
      → Configuration' (adsl ++ adsr, radsl ++ (radsr \ adsl))
                        (csl ++ csr, rcsl ++ (rcsr \ csl))
                        ((dsl \ rdsr) ++ dsr, rdsl ++ (rdsr \ dsl))
```

The indices are quite involved, since we need to record both the current advertisements, stipulated contracts and deposits and the required ones for the configuration to become valid. The most interesting case is the parallel composition operator, where the resources provided by the left operand might satisfy some requirements of the right operand. Moreover, consumed deposits

have to be eliminated as there can be no double spending, while the number of advertisements and contracts always grows.

By composing configurations together, we will eventually end up in a *closed* configuration, where all required indices are empty (i.e. the configuration is self-contained):

Configuration : *AdvertisedContracts* $\rightarrow$ *ActiveContracts* $\rightarrow$ *List Deposit* $\rightarrow$ *Set*
Configuration *ads cs ds* = *Configuration'* (*ads*, []) (*cs*, []) (*ds*, [])

We are now ready to declare the inference rules of the bottom layer of our small-step semantics, by defining an inductive datatype modelling the binary step relation between untimed configurations:

data $_ \longrightarrow _ : \text{Configuration } \text{ads } \text{cs } \text{ds} \rightarrow \text{Configuration } \text{ads}' \text{cs}' \text{ds}' \rightarrow \text{Set}$ **where**

DEP-AuthJoin :

$\langle A, v \rangle^d \mid \langle A, v' \rangle^d \mid \Gamma \longrightarrow \langle A, v \rangle^d \mid \langle A, v' \rangle^d \mid A [0 \leftrightarrow 1] \mid \Gamma$

DEP-Join :

$\langle A, v \rangle^d \mid \langle A, v' \rangle^d \mid A [0 \leftrightarrow 1] \mid \Gamma \longrightarrow \langle A, v + v' \rangle^d \mid \Gamma$

C-Advertise : $\forall \{ \Gamma \text{ ad} \}$

$\rightarrow \exists [p \in \text{participants}^g (G \text{ ad})] \ p \in \text{Hon}$

$\rightarrow \Gamma \longrightarrow 'ad \mid \Gamma$

C-AuthCommit : $\forall \{ A \text{ ad } \Gamma \}$

$\rightarrow \text{secrets } (G \text{ ad}) \equiv a_0 \dots a_n$

$\rightarrow (A \in \text{Hon} \rightarrow \forall [i \in 0 \dots n] \ a_i \neq \perp)$

$\rightarrow 'ad \mid \Gamma \longrightarrow 'ad \mid \Gamma \mid \dots \langle A : a_i \# N_i \rangle \dots \mid A [\# \triangleright \text{ad}]$

C-Control : $\forall \{ \Gamma \text{ C } i \text{ D} \}$

$\rightarrow C !! i \equiv A_1 : A_2 : \dots : A : D$

$\rightarrow \langle C, v \rangle^c \mid \dots A_i [C \triangleright^b i] \dots \mid \Gamma \longrightarrow \langle D, v \rangle^c \mid \Gamma$

$\vdots$

There is a total of 18 rules we need to define, but we choose to depict only a representative subset of them. The first pair of rules initially appends the authorisation to merge two deposits to the current configuration (rule *[DEP-AuthJoin]*) and then performs the actual join (rule *[DEP-Join]*). This is a common pattern across all rules, where we first collect authorisations for an action by all involved participants, and then we fire a subsequent rule to perform this action. *[C-Advertise]* advertises a new contract, mandating that at least one of the participants involved in the pre-condition is honest and requiring that all deposits needed for stipulation are available in the surrounding context. *[C-AuthCommit]* allows participants to commit to the secrets required by the contract's pre-condition, but only dishonest ones can commit to the invalid length $\perp$.

Formal investigation of the Extended UTxO model

Lastly, *[C-Control]* allows participants to give their authorization required by a particular branch out of the current choices present in the contract, discarding any time constraints along the way.

It is noteworthy to mention that during the transcriptions of the complete set of rules from the paper [Bartoletti and Zunino 2018] to our dependently-typed setting, we discovered a discrepancy in the *[C-AuthRev]* rule, namely that there was no context Γ . Moreover, in order to later facilitate equational reasoning, we re-factored the *[C-Control]* to not contain the inner step as a hypothesis, but instead immediately inject it in the result operand of the step relation.

The inference rules above have elided any treatment of timely constraints; this is handled by the top layer, whose states are now timed configurations. The only interesting inference rule is the one that handles time decorations of the form *after* $_ : _$, since all other cases are dispatched to the bottom layer (which just ignores timely aspects).

record *Configuration*^t (*ads* : *AdvertisedContracts*) (*cs* : *ActiveContracts*) (*ds* : *Deposits*) : *Set* **where**
constructor $_ @ _$
field *cfg* : *Configuration* *ads cs ds*
time : *Time*

data $_ \longrightarrow_t _ : \text{Configuration}^t \text{ ads cs ds} \rightarrow \text{Configuration}^t \text{ ads' cs' ds'} \rightarrow \text{Set}$ **where**

Action : $\forall \{ \Gamma \Gamma' t \}$

$\rightarrow \Gamma \longrightarrow \Gamma'$

$\rightarrow \Gamma @ t \longrightarrow_t \Gamma' @ t$

Delay : $\forall \{ \Gamma t \delta \}$

$\rightarrow \Gamma @ t \longrightarrow_t \Gamma @ (t + \delta)$

Timeout : $\forall \{ \Gamma \Gamma' t i \text{ contract} \}$

$\rightarrow \text{All } (_ \leq t) (\text{timeDecorations } (\text{contract} !! i))$ *-- all time constraints are satisfied*

$\rightarrow \langle [\text{contract} !! i], v \rangle^c \mid \Gamma \longrightarrow \Gamma'$ *-- resulting state if we pick this branch*

$\rightarrow \langle \langle \text{contract}, v \rangle^c \mid \Gamma \rangle @ t \longrightarrow_t \Gamma' @ t$

Having defined the step relation in this way allows for equational reasoning, a powerful tool for writing complex proofs:

data $_ \twoheadrightarrow _ : \text{Configuration ads cs ds} \rightarrow \text{Configuration ads' cs' ds'} \rightarrow \text{Set}$ **where**

$_ \sqsubset : (M : \text{Configuration ads cs ds}) \rightarrow M \twoheadrightarrow M$

$_ \longrightarrow \langle _ \rangle _ : \forall \{ M N \} (L : \text{Configuration ads cs ds})$

$\rightarrow L \longrightarrow M \rightarrow M \twoheadrightarrow N$

$\rightarrow L \twoheadrightarrow N$

begin $_ : \forall \{ M N \} \rightarrow M \twoheadrightarrow N \rightarrow M \twoheadrightarrow N$

4.2.3 Example. We are finally ready to see a more intuitive example of the *timed-commitment protocol*, where a participant commits to revealing a valid secret a (e.g. "qwerty") to another participant, but loses her deposit of $\text{€ } 1$ if she does not meet a certain deadline t :

$$tc : \text{Advertisement } 1 [] (1 :: 0 :: []) \\ tc = \langle A!1 \wedge A\#a \wedge B!0 \rangle \text{ reveal } [a] \Rightarrow \text{withdraw } A \vdash \dots \oplus \text{after } t : \text{withdraw } B$$

Below is one possible reduction in the bottom layer of our small-step semantics, demonstrating the case where the participant actually meets the deadline:

$$\begin{aligned} tc\text{-semantics} : \langle A, 1 \rangle^d &\rightarrow \langle A, 1 \rangle^d \mid A : a \# 6 \\ tc\text{-semantics} = & \\ \text{begin} & \\ \langle A, 1 \rangle^d & \\ \rightarrow \langle C\text{-Advertise} \rangle & \\ 'tc \mid \langle A, 1 \rangle^d & \\ \rightarrow \langle C\text{-AuthCommit} \rangle & \\ 'tc \mid \langle A, 1 \rangle^d \mid \langle A : a \# 6 \rangle \mid A [\# \triangleright tc] & \\ \rightarrow \langle C\text{-AuthInit} \rangle & \\ 'tc \mid \langle A, 1 \rangle^d \mid \langle A : a \# 6 \rangle \mid A [\# \triangleright tc] \mid A [tc \triangleright^s 0] & \\ \rightarrow \langle C\text{-Init} \rangle & \\ \langle tc, 1 \rangle^c \mid \langle A : a \# \text{inj}_1 6 \rangle & \\ \rightarrow \langle C\text{-AuthRev} \rangle & \\ \langle tc, 1 \rangle^c \mid A : a \# 6 & \\ \rightarrow \langle C\text{-Control} \rangle & \\ \langle [\text{reveal } [a] \Rightarrow \text{withdraw } A \vdash \dots], 1 \rangle^c \mid A : a \# 6 & \\ \rightarrow \langle C\text{-PutRev} \rangle & \\ \langle [\text{withdraw } A], 1 \rangle^c \mid A : a \# 6 & \\ \rightarrow \langle C\text{-Withdraw} \rangle & \\ \langle A, 1 \rangle^d \mid A : a \# 6 & \end{aligned}$$

□

At first, A holds a deposit of $\text{€ } 1$, as required by the contract's precondition. Then, the contract is advertised and the participants slowly provide the corresponding prerequisites (i.e. A commits to a secret via $[C\text{-AuthCommit}]$ and spends the required deposit via $[C\text{-AuthInit}]$, while B does not do anything). After all pre-conditions have been satisfied, the contract is stipulated (rule $[C\text{-Init}]$) and the secret is successfully revealed (rule $[C\text{-AuthRev}]$). Finally, the first branch is picked (rule $[C\text{-Control}]$) and A retrieves her deposit back (rules $[C\text{-PutRev}]$ and $[C\text{-Withdraw}]$).

4.3 Reasoning Modulo Permutation

In the definitions above, we have assumed that $(_|_, \emptyset)$ forms a commutative monoid, which allowed us to always present the required sub-configuration individually on the far left of a composite configuration. While such definitions enjoy a striking similarity to the ones appearing in the original paper [Bartoletti and Zunino 2018] (and should always be preferred in an informal textual setting), this approach does not suffice for a mechanized account of the model. After all, this explicit treatment of all intuitive assumptions/details is what makes our approach robust and will lead to a deeper understanding of how these systems behave. To overcome this intricacy, we

Formal investigation of the Extended UTxO model

introduce an *equivalence relation* on configurations, which holds when they are just permutations of one another:

$_ \approx _ : \text{Configuration ads cs ds} \rightarrow \text{Configuration ads cs ds} \rightarrow \text{Set}$
 $c \approx c' = \text{cfgToList } c \rightsquigarrow \text{cfgToList } c'$

where

open import *Data.List.Permutation* **using** ($_ \rightsquigarrow _$)
 $\text{cfgToList} : \text{Configuration}' p_1 p_2 p_3 \rightarrow \text{List } (\exists[p_1] \exists[p_2] \exists[p_3] \text{Configuration}' p_1 p_2 p_3)$
 $\text{cfgToList } \emptyset = []$
 $\text{cfgToList } (l \mid r) = \text{cfgToList } l ++ \text{cfgToList } r$
 $\text{cfgToList } \{p_1\} \{p_2\} \{p_3\} c = [p_1, p_2, p_3, c]$

Given this reordering mechanism, we now need to generalise all our inference rules to implicitly reorder the current and next configuration of the step relation. We achieve this by introducing a new variable for each of the operands of the resulting step relations, replacing the operands with these variables and requiring that they are re-orderings of the previous configurations, as shown in the following generalisation of the *[DEP-AuthJoin]* rule¹⁶:

DEP-AuthJoin :

$$\frac{\begin{array}{l} \Gamma' \approx \langle A, v \rangle^d \mid \langle A, v' \rangle^d \mid \Gamma \quad \in \text{Configuration ads cs } (A \text{ has } v :: A \text{ has } v' :: ds) \\ \rightarrow \Gamma'' \approx \langle A, v \rangle^d \mid \langle A, v' \rangle^d \mid A[0 \leftrightarrow 1] \mid \Gamma \quad \in \text{Configuration ads cs } (A \text{ has } (v + v') :: ds) \end{array}}{\rightarrow \Gamma' \rightarrow \Gamma''}$$

Unfortunately, we now have more proof obligations of the re-ordering relation lying around, which makes reasoning about our semantics rather tedious. We are currently investigating different techniques to model such reasoning up to equivalence:

- *Quotient types* [Altenkirch et al. 2011] allow equipping a type with an equivalence relation. If we assume the axiom that two elements of the underlying type are *propositionally* equal when they are equivalent, we could discharge our current proof burden trivially by reflexivity. Unfortunately, while one can easily define *setoids* in Agda, there is not enough support from the underlying type system to make reasoning about such an equivalence as easy as with built-in equality.
- Going a step further into more advanced notions of equality, we arrive at *homotopy type theory* [hom 2013], which tries to bridge the gap between reasoning about isomorphic objects in informal pen-paper proofs and the way we achieve this in mechanized formal methods. Again, realizing practical systems with such an enriched theory is a topic of current research [Cohen et al. 2016] and no mature implementation exists yet, so we cannot integrate it with our current development in any pragmatic way.
- The crucial problems we have encountered so far are attributed to the non-deterministic nature of BitML, which is actually inherent in any process calculus. Building upon this idea, we plan to take a step back and investigate different reasoning techniques for a minimal process calculus. Once we have an approach that is more suitable, we will incorporate it in our full-blown BitML calculus.

¹⁶ In fact, it is not necessary to reorder both ends for the step relation; at least one would be adequate.

5 PLANNING

In this section, I describe possible next steps I plan to investigate during the remainder of my thesis. It is impossible to accurately predict what will be achieved in the following five months and there will definitely be some surprises along the way, but I hope it will give realistic expectations of the final results of my thesis.

5.1 Extended UTxO: Multi-currency

Many major blockchain systems today support the creation of secondary cryptocurrencies, which are independent of the main currency. In Bitcoin, for instance, *colored coins* allow transactions to assign additional meaning to their outputs (e.g. each coin could correspond to a real-world asset, such as company shares) [Rosenfeld 2012].

This approach, however, has the disadvantage of larger transactions and less efficient processing. One could instead bake the multi-currency feature into the base system, mitigating the need for larger transactions and slow processing. Building on the abstract UTxO model, there are current research efforts on a general framework that provides mechanisms to establish and enforce monetary policies for multiple currencies [Zahnentferner 2019].

Fortunately, the extensions proposed by the multi-currency are orthogonal to the functionality I have currently formalized. In order to achieve this, one has to generalize the *Value* datatype to account for multiple currencies. Hence, I plan to integrate this with my current formal development of the extended UTxO model and, by doing so, provide the first formalization of a UTxO ledger that supports multiple cryptocurrencies.

5.2 BitML: Towards Completeness

Continuing my work on the formalization of the BitML paper [Bartoletti and Zunino 2018], there is still a lot of theoretical results to be covered:

- While I currently have the symbolic model in place, there is still no formalization of *symbolic strategies*, where one can reason about different adversary strategies and prove that certain scenarios are impossible.
- Another import task is to define the computational model; a counterpart of the symbolic model augmented with pragmatic computational properties to more closely resemble the low-level details of Bitcoin.
- When both symbolic and computational strategies have been formalized, I will be able to finally prove the correctness of the BitML compiler, which translates high-level BitML contracts to low-level standard Bitcoin transactions. The symbolic model concerns the input of the compiler, while the computational one concerns the output. This endeavour will involve implementing the actual translation and proving *coherence* between the symbolic and the computational model. Proving coherence essentially requires providing a (weak) *simulation* between the two models; each step in the symbolic part is matched by (multiple) steps in the computational one.

5.3 UTxO-BitML Integration

So far I have worked separately on the two models under study, but it would be interesting to see whether these can be intertwined in some way. This would possibly involve a translation from BitML contracts to contracts modelled in our extended UTxO models, along with corresponding meta-theoretical properties (e.g. validity of UTxO transactions correspond to another notion of validity of BitML contracts).

Moreover, and it would be beneficial to review the different modelling techniques used across both models, identifying their key strengths and witnesses. With this in mind, I could refactor crucial parts of each model for the sake of elegance, clarity and ease of reasoning.

5.4 Plutus Integration

In my current formalization of the extended UTxO model, scripts are immediately modelled by their denotations (i.e. pure mathematical functions). This is not accurate, however, since scripts are actually pieces of program text. However, there is current development by James Chapman of IOHK to formalize the meta-theory of Plutus, Cardano’s scripting language¹⁷.

Since we mostly care about Plutus as a scripting language, it would be possible to replace the denotations with actual Plutus Core source code and utilize the formalized meta-theory to acquire the denotational semantics when needed.

5.5 Featherweight Solidity

One of the posed research questions concerns the expressiveness of the extended UTxO model with respect to Ethereum-like account-based ledgers.

In order to investigate this in a formal manner, one has to initially model a reasonable subset of Solidity, so a next step would be to model *Featherweight Solidity*, taking inspiration from the approach taken in the formalization of Java using *Featherweight Java* [Igarashi et al. 2001]. Fortunately, I will not have to start from scratch, since there have been recent endeavours in F* to analyse and verify Ethereum smart contracts, which already describe a simplified model of Solidity [Bhargavan et al. 2016].

As a next step, one could try out different example contracts in Solidity and check whether they can be transcribed to contracts appropriate for an extended UTxO ledger.

5.6 Proof Automation

Last but not least, our current dependently-typed approach to formalizing our models has led to a significant proof burden, as evidenced by the complicated type signatures presented throughout this proposal. This certainly makes the reasoning process quite tedious and time consuming, so a reasonable task would be to implement automatic proof-search procedures using Agda meta-programming [Kokke and Swierstra 2015].

5.7 Timetable

I have assembled a detailed timetable in Figure 2, positioning the aforementioned tasks across the whole timespan of my thesis.

I expect to complete the tasks that are more tightly coupled with my current development, i.e. incorporating multi-currency features and the formalized Plutus meta-theory into the extended UTxO model, completing the formalization of the BitML paper and implementing proof-search automation to facilitate easier reasoning. The merging of these two subjects of study (i.e. UTxO and BitML) is somewhat unclear at this stage, but I hope to at least provide a proof-of-concept translation, even if this comes without significant meta-theoretic results like coherence.

Other tasks, such as a mature model of Featherweight Solidity and formal results comparing it to UTxO, are sadly outside the scope of this thesis due to time constraints. Nonetheless, I will strive for a prototype model with lots of examples and hope my work will lay the foundations to further investigate these topics in future research.

¹⁷<https://github.com/input-output-hk/plutus-metatheory>




Fig. 2. My workplan.

REFERENCES

2013. Homotopy Type Theory: Univalent Foundations of Mathematics. *CoRR* abs/1308.0729 (2013). arXiv:1308.0729 <http://arxiv.org/abs/1308.0729>
- Thorsten Altenkirch, Thomas Anberrée, and Nuo Li. 2011. Definable quotients in type theory. *Draft paper* (2011), 48–49.
- Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. 2014. Secure multiparty computations on bitcoin. In *Security and Privacy (SP), 2014 IEEE Symposium on*. IEEE, 443–458.
- Bruno Barras, Samuel Boutin, Cristina Cornes, Judicaël Courant, Jean-Christophe Filliatre, Eduardo Gimenez, Hugo Herbelin, Gerard Huet, Cesar Munoz, Chetan Murthy, et al. 1997. *The Coq proof assistant reference manual: Version 6.1*. Ph.D. Dissertation. Inria.
- Massimo Bartoletti and Roberto Zunino. 2018. *BitML: a calculus for Bitcoin smart contracts*. Technical Report. Cryptology ePrint Archive, Report 2018/122.
- Iddo Bentov and Ranjit Kumaresan. 2014. How to use bitcoin to design fair protocols. In *International Cryptology Conference*. Springer, 421–439.
- Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cedric Fournet, A Gollamudi, G Gonthier, N Kobeissi, A Rastogi, T Sibut-Pinote, N Swamy, and S Zanella-Béguelin. 2016. Short paper: Formal verification of smart contracts. In *Proceedings of the 11th ACM Workshop on Programming Languages and Analysis for Security (PLAS), in conjunction with ACM CCS*. 91–96.
- Vitalik Buterin et al. 2014. A next-generation smart contract and decentralized application platform. *white paper* (2014).
- Hao Chen, Xiongnan Newman Wu, Zhong Shao, Joshua Lockerman, and Ronghui Gu. 2016. Toward compositional verification of interruptible OS kernels and device drivers. In *ACM SIGPLAN Notices*, Vol. 51. ACM, 431–447.
- Koen Claessen and John Hughes. 2011. QuickCheck: a lightweight tool for random testing of Haskell programs. *Acm sigplan notices* 46, 4 (2011), 53–64.
- Cyril Cohen, Thierry Coquand, Simon Huber, and Anders Mörtberg. 2016. Cubical Type Theory: a constructive interpretation of the univalence axiom. *CoRR* abs/1611.02108 (2016). arXiv:1611.02108 <http://arxiv.org/abs/1611.02108>
- Oded Goldreich, Silvio Micali, and Avi Wigderson. 1991. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM (JACM)* 38, 3 (1991), 690–728.
- Charles Antony Richard Hoare. 1978. Communicating sequential processes. In *The origin of concurrent programming*. Springer, 413–443.

Formal investigation of the Extended UTxO model

- Paul Hudak, Simon Peyton Jones, Philip Wadler, Brian Boutel, Jon Fairbairn, Joseph Fasel, María M Guzmán, Kevin Hammond, John Hughes, Thomas Johnsson, et al. 1992. Report on the programming language Haskell: a non-strict, purely functional language version 1.2. *ACM SigPlan notices* 27, 5 (1992), 1–164.
- Atsushi Igarashi, Benjamin C Pierce, and Philip Wadler. 2001. Featherweight Java: a minimal core calculus for Java and GJ. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 23, 3 (2001), 396–450.
- Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. 2017. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual International Cryptology Conference*. Springer, 357–388.
- Wen Kokke and Wouter Swierstra. 2015. Auto in agda. In *International Conference on Mathematics of Program Construction*. Springer, 276–301.
- Per Martin-Löf and Giovanni Sambin. 1984. *Intuitionistic type theory*. Vol. 9. Bibliopolis Naples.
- Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. (2008).
- Ulf Norell. 2008. Dependently typed programming in Agda. In *International School on Advanced Functional Programming*. Springer, 230–266.
- Simon Peyton Jones, Jean-Marc Eber, and Julian Seward. 2000. Composing contracts: an adventure in financial engineering. *ACM SIG-PLAN Notices* 35, 9 (2000), 280–292.
- Meni Rosenfeld. 2012. Overview of colored coins. *White paper, bitcoil. co. il* 41 (2012).
- Pablo Lamela Seijas, Simon J Thompson, and Darryl McAdams. 2016. Scripting smart contracts for distributed ledger technology. *IACR Cryptology ePrint Archive* 2016 (2016), 1156.
- Ilya Sergey, Amrit Kumar, and Aquinas Hobor. 2018. Scilla: a smart contract intermediate-level language. *arXiv preprint arXiv:1801.00687* (2018).
- Anton Setzer. 2018. Modelling Bitcoin in Agda. *arXiv preprint arXiv:1804.06398* (2018).
- Paul Van Der Walt and Wouter Swierstra. 2012. Engineering proof by reflection in Agda. In *Symposium on Implementation and Application of Functional Languages*. Springer, 157–173.
- Joachim Zahnentferner. 2019. Multi-Currency Ledgers. (2019), To Appear.
- Joachim Zahnentferner and Input Output HK. 2018. *Chimeric ledgers: Translating and unifying utxo-based and account-based cryptocurrencies*. Technical Report. Cryptology ePrint Archive, Report 2018/262, 2018. <https://eprint.iacr.org>