

1 Formal specification of the Cardano blockchain 2 ledger, mechanized in Agda

3 Andre Knispel ✉ 
Input Output, Global

4 James Chapman ✉ 
Input Output, Global

5 Joosep Jääger ✉
Input Output, Global

6 Ulf Norell ✉
7 QuviQ, Sweden

Orestis Melkonian ✉ 
Input Output, Global

Alasdair Hill ✉
Input Output, Global

William DeMeo ✉ 
Input Output, Global

8 — Abstract —

9 Blockchain systems comprise critical software that handle substantial monetary funds, rendering
10 them excellent candidates for *formal verification*. One of their core components is the underlying
11 ledger that does all the accounting: keeping track of transactions and their validity, etc.

12 Unfortunately, previous theoretical studies are typically confined to an idealized setting, while
13 specifications for real implementations are scarce; either the functionality is directly implemented
14 without a proper specification, or at best an informal specification is written on paper.

15 The present work expands beyond prior meta-theoretical investigations of the EUTxO model to
16 encompass the full scale of the Cardano blockchain: our formal specification describes a hierarchy of
17 modular transitions that covers all the intricacies of a realistic blockchain, such as fully expressive
18 smart contracts and decentralized governance.

19 It is mechanized in a proof assistant, thus enjoys a higher standard of rigor: type-checking prevents
20 minor oversights that were frequent in previous informal approaches; key meta-theoretical properties
21 can now be formally proven; it is an *executable* specification against which the implementation in
22 production is being tested for conformance; and it provides firm foundations for smart contract
23 verification.

24 Apart from a safety net to keep us in check, the formalization also provides a guideline for the
25 ledger design: one informs the other in a symbiotic way, especially in the case of state-of-the-art
26 features like decentralized governance, which is an emerging sub-field of blockchain research that
27 however mandates a more exploratory approach.

28 All the results presented in this paper have been mechanized in the Agda proof assistant and
29 are publicly available. In fact, this document is itself a literate Agda script and all rendered code
30 has been successfully type-checked.

31 **2012 ACM Subject Classification** Theory of computation → Type theory; Theory of computation
32 → Logic and verification; Theory of computation → Program specifications

33 **Keywords and phrases** blockchain, distributed ledgers, UTxO, Cardano, formal verification, Agda

34 **Digital Object Identifier** [10.4230/OASICS.FMBC.2024.7](https://doi.org/10.4230/OASICS.FMBC.2024.7)

35 **1** Introduction

36 This paper gives a high-level overview of the Cardano ledger specification in the Agda proof
37 assistant, which is one of three core pieces of the Cardano blockchain:

- 38 ■ **Networking:** deals with sending messages across the internet.
- 39 ■ **Consensus:** establishes a common order of valid blocks.
- 40 ■ **Ledger:** decides whether a sequence of blocks is valid.

41 Such *separation of concerns* is crucial to enable a rigidly formal study of each individual
42 component; the ledger is based on the *Extended UTxO* model (EUTxO), an extension of



© Andre Knispel;
licensed under Creative Commons License CC-BY 4.0
5th International Workshop on Formal Methods for Blockchains (FMBC 2024).
Editors: Bruno Bernardo and Diego Marmosoler; Article No. 7; pp. 7:1–7:19



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

43 Bitcoin’s model of unspent transaction outputs [19] – in contrast to Ethereum’s account-based
 44 model [8] – to accommodate fully expressive *smart contracts* that run on the blockchain.
 45 Luckily for us, EUTxO enjoys a well-studied meta-theory [9, 10] that is also mechanized
 46 in Agda, albeit in a much simpler setting where a single ledger feature is considered at a
 47 time, but not how multiple concurrent features interact. We take this to the next level by
 48 scaling up these prior theoretical results to match the complexity of the real world: the
 49 Cardano blockchain being one of the top ten cryptocurrencies today by market capitalization,
 50 it handles gigabytes of transactions that transfer hundred of millions US dollars, while
 51 simultaneously supporting all these features plus many more that have not been formally
 52 studied before.

53 We are happy to report that the formalization overhead has proven minuscule compared
 54 to the development effort of the actual implementation, measured either by lines of code (~10
 55 thousand lines of Agda formalization *versus* ~200 thousand of Haskell implementation) or
 56 by number of man hours put in so far (only a couple of full-time formal methods engineers
 57 *versus* tens of production developers). The result is a *mechanized* document that leaves little
 58 room for error, additionally proves crucial invariants of the overall system, *e.g.*, that the
 59 global value carried by the system stays constant, formally stated in Section 4. It doubles as
 60 an executable reference implementation that we can utilize in production for conformance
 61 testing. All of our work, much like this paper, is mechanized in Agda and is publicly available:

62 <https://github.com/IntersectMBO/formal-ledger-specifications>

63 **Scope.** Cardano’s evolution proceeds in *eras*, each introducing a new vital feature to the
 64 previous ones. While we would ideally want to provide a multitude of formal artifacts, each
 65 describing a single era in full detail, the specification formalized here is that of the **Voltaire**
 66 era that introduces *decentralized governance* as described in the Cardano Improvement
 67 Proposal (CIP) 1694.¹ This stems from the fact that the design of the blockchain happens in
 68 tandem with the formal specification; one informs the other in an intricate, non-linear fashion.
 69 Thus arises a pragmatic need to think of the process as an act of balance between keeping
 70 up with the *past*, *i.e.*, going back to previous eras and incrementally incorporating their
 71 features, and co-evolving with the current design of the *future* ledger capabilities. Therefore,
 72 we set aside details of the previous **Byron**, **Shelley**, and **Alonzo** eras while at the same
 73 time missing orthogonal features related to smart contracts brought in the **Babbage** era.

74 **Transitions as relations.** The ledger can itself be conceptually divided into multiple
 75 sub-components, each described by a transition between states that only contains the relevant
 76 parts of the overarching ledger state and possibly some internal auxiliary information that is
 77 discarded at the outer layer. These transitions are not independent, but form a hierarchy
 78 of “state machines” where some higher-level transition might demand successful transition
 79 of a sub-component down the dependency graph as one of its premises. Eventually, these
 80 cascading transitions all get combined to dictate the top-level transition that handles an
 81 individual block of transactions submitted to the blockchain.

82 Formally, we formulate such (labeled) transitions as relations X between the environment
 83 Γ inherited from a higher layer, an initial state s , a signal b that acts as user input, and a
 84 final state s' :

$$85 \quad \Gamma \vdash s \xrightarrow[X]{b} s' \quad \begin{array}{c|c} \textit{Environments} & \textit{States} \\ \textit{(Signals)} & \end{array} \\ \hline \textit{Possible transitions}$$

¹ <https://github.com/cardano-foundation/CIPs/blob/17771640/CIP-1694/README.md>

86 We will henceforth present such transitions as shown on the right; a *tritych* defining
 87 environments and possibly signals (top left), states (top right), and the rules that *inductively*
 88 define the transition (bottom).

89 1.1 Agda preliminaries

90 In Agda, the aforementioned ledger transitions are modeled as *inductive families* of type:

91 $_ \vdash _ \rightarrow (_ _)_ : Env \rightarrow State \rightarrow Signal \rightarrow State \rightarrow Type$

92 **Reflexive transitive closure.** We will often need to apply a transition repeatedly until
 93 we arrive at a final state, which corresponds to the standard mathematical construction of
 94 taking the relation's *reflexive transitive closure*:

95 `data $_ \vdash _ \rightarrow (_ _)*_ : Env \rightarrow State \rightarrow List Signal \rightarrow State \rightarrow Type$ where`

<p>96 <code>base :</code></p> $\frac{}{\Gamma \vdash s \rightarrow (\ []) * s}$	<p style="color: green;">step :</p> <ul style="list-style-type: none"> • $\Gamma \vdash s \rightarrow (b \quad) \ s'$ • $\Gamma \vdash s' \rightarrow (bs \quad) * s''$ $\frac{}{\Gamma \vdash s \rightarrow (b :: bs) * s''}$
--	--

98 **Finite sets & maps.** One particular trait we inherited from previous pen-and-paper
 99 iterations of the ledger specification is a heavy use of set theory, which goes against Agda's
 100 foundations in Type Theory, both technically and in a philosophical sense. To remedy this,
 101 we have developed an in-house library for conducting *Axiomatic Set Theory* within the type-
 102 theoretic setting of Agda [18]; we stay in its finite fragment for this application. Crucially, the
 103 type of sets is entirely *abstract*: there is no way to utilize proof-by-computation (*e.g.*, as one
 104 would do when modeling sets as lists of distinct elements), so that all proofs eventually resort
 105 to the axioms and the library's implementation details stay irrelevant. At the same time,
 106 when extracting executable code the library provides a properly executable implementation—
 107 the abstraction layer only exists at compile-time. Implementing this abstraction layer helped
 108 us greatly reduce code complexity and size over a previous list-based approach. In fact, it is
 109 highly encouraged to provide *multiple* implementations without affecting the formalization
 110 and the validity of the established proofs therein.

111 Equipped with the axioms provided by the library, *e.g.*, the ability to construct power
 112 sets \mathbb{P} , it is remarkably easy to define common set-theoretic concepts like set inclusion and
 113 extensional equality of sets (left), as well as re-purpose sets of key-value pairs to model *finite*
 114 *maps*² by imposing uniqueness of keys (right):

<p>115 $_ \subseteq _ : \{A : Type\} \rightarrow \mathbb{P} A \rightarrow \mathbb{P} A \rightarrow Type$</p> $X \subseteq Y = \forall \{x\} \rightarrow x \in X \rightarrow x \in Y$	<p>$_ \rightarrow _ : Type \rightarrow Type \rightarrow Type$</p> $A \rightarrow B = \exists \lambda (\mathfrak{R} : \mathbb{P} (A \times B)) \rightarrow$ $\forall \{a \ b \ b'\} \rightarrow (a , b) \in \mathfrak{R} \rightarrow (a , b') \in \mathfrak{R} \rightarrow b \equiv b'$
<p>$_ \approx _ : \{A : Type\} \rightarrow \mathbb{P} A \rightarrow \mathbb{P} A \rightarrow Type$</p> $X \approx Y = X \subseteq Y \times Y \subseteq X$	

² It is natural to think of maps as partial functions, but unrestricted Agda functions would not do here.

116 **2 Fundamental entities**117 **2.1 Cryptographic primitives**

118 There are two types of credentials that can be used on Cardano: `VKey` and script credentials.
 119 `VKey` credentials use a public key signing scheme (Ed25519) for verification. Some serialized
 120 (`Ser`) data can be signed, and `isSigned` is the property that a public `VKey` signed some data
 121 with a given signature (`Sig`). There are also other cryptographic primitives in the Cardano
 122 ledger, for example KES and VRF used in the consensus layer, but we omit those here.

123 Script credentials correspond to a hash of a script that has to be executed by the ledger
 124 as part of transaction validation. There are two different types of scripts, native and Plutus,
 125 but the details of this are not relevant for the rest of this paper.

126 `VKey Sig Ser : Type` `isSigned : VKey → Ser → Sig → Type`

127 In the specification, all definitions that require these primitives must accept these as
 128 additional arguments. To streamline this process, these definitions are bundled into a record
 129 and, using Agda's module system, are quantified only once per file. We are using this pattern
 130 many times, either to introduce additional abstraction barriers or to effectively provide
 131 foreign functions within a safe environment. Additionally, particularly fundamental interfaces
 132 like the one presented above are sometimes re-bundled transitively into larger records, which
 133 further streamlines the interface. This is in stark contrast to the Haskell implementation,
 134 which often needs to repeat tens of type class constraints on many functions in a module.

135 **2.2 Addresses**

136 There are various types of addresses used for storing funds in the UTxO set, which all contain
 137 a payment `Credential` and optionally a staking `Credential`. `Addr` is the union of all of those
 138 types. A `Credential` is a hash of a public key or script, types for which are kept abstract. The
 139 most common type of address is a `BaseAddr` which must include a staking `Credential`.

140 There is also a special type of address (not included in `Addr`) without a payment credential,
 141 called a reward address. It is not used for interacting with the UTxO set, but instead used
 142 to refer to reward accounts [31].

143 `Credential = KeyHash ⊔ ScriptHash`

<p>144 <code>record BaseAddr : Type where</code> <code> pay : Credential</code> <code> stake : Credential</code></p>	<p><code>record RwdAddr : Type where</code> <code> stake : Credential</code></p>
--	--

145 `Addr = BaseAddr ⊔ ...`

146 **2.3 Base types**

147 The basic units of currency and time are `Coin`, `Slot` and `Epoch`, which we treat as natural
 148 numbers, while an implementation might use isomorphic but more complicated types (for
 149 example to represent the beginning of time in a special way).

150 `Coin = Slot = Epoch = ℕ`

151 A `Coin` is the smallest unit of currency, a `Slot` is the smallest unit of time (corresponding to 1
 152 second in the main chain), and an `Epoch` is a fixed number of slots (corresponding to 5 days

153 in the main chain). Every slot, a stake pool has a random chance to be able to mint a block,
 154 and one block every five slots is expected [13].

155 **3 Advancing the blockchain**

156 **3.1 Protocol parameters**

157 We start with adjustable protocol parameters. In contrast to constants such as the length of
 158 an `Epoch`, these parameters can be changed while the system is running via the governance
 159 mechanism. They can affect various features of the system, such as minimum fees, maximum
 160 and minimum sizes of certain components, and more.

161 The full specification contains well over 20 parameters, while we only list a few. The
 162 maximum sizes should be self-explanatory, while `a` and `b` are the coefficients of a polynomial
 163 used in the calculation of the minimum fee for transactions (*c.f.*, function `minfee` in
 164 Appendix B).

```
165 record PParams : Type where
166   maxBlockSize maxTxSize a b : ℕ
```

167 **3.2 Extending the blockchain block-by-block**

168 `CHAIN` is the main state machine describing the ledger. Since it is not invoked from any
 169 other state machine, it does not have an environment. It invokes two other state machines,
 170 `NEWEPOCH` and `LEDGER*`, where the former detects if the new block `b` is in a new epoch.
 171 In that case, `NEWEPOCH` takes care of various bookkeeping tasks, such as counting votes for
 172 the governance system and updating stake distributions for consensus. For a basic version
 173 that detects whether a new epoch has been entered, see Appendix C. The potentially updated
 174 state is then given to `LEDGER*`, which is the reflexive-transitive closure of `LEDGER` and
 175 applies all the transactions in the block in sequence. Finally, `CHAIN` updates `ChainState` with
 176 the resulting states.

177 There is a key property about `NEWEPOCH`, which is that it never gets stuck, i.e. that
 178 for all states, environments and signals it always transitions to a new state. This property is
 179 proven in our development.

```
180 record Block : Type where
   ts : List Tx
   slot : Slot
```

```
record NewEpochState : Type where
   lastEpoch : Epoch
   acnt       : Acnt
   ls        : LState
   es        : EnactState
   fut       : RatifyState
```

```
record ChainState : Type where
   newEpochState : NewEpochState
```

181 `CHAIN` :

```
182 • mkNewEpochEnv s ⊢ s.newEpochState → ( epoch slot ,NEWEPOCH ) nes
183 • [ slot ⊗ constitution .proj₁ .proj₂ ⊗ pparams .proj₁ ⊗ es ] ⊢ nes .ls → ( ts ,LEDGER* ) ls'
184 _____
185 _ ⊢ s → ( b ,CHAIN ) updateChainState s nes
```

186 **3.3 Extending the ledger transaction-by-transaction**

187 Transaction processing is broken down into three separate parts: accounting & witnessing
 188 (UTXOW), application of certificates (CERT) and processing of governance votes & proposals
 189 (GOV).

<pre> 189 record LEnv : Type where slot : Slot ppolicy : Maybe ScriptHash pparams : PParams enactState : EnactState </pre>	<pre> record LState : Type where utxoSt : UTxOState govSt : GovState certState : CertState </pre>
---	--

```

191 LEDGER :
192 • mkUTxOEnv  $\Gamma \vdash \text{utxoSt} \rightarrow (\text{tx}, \text{UTXOW}) \text{ utxoSt}'$ 
193 •  $\llbracket \text{epoch slot} \otimes \text{pparams} \otimes \text{txvote} \otimes \text{txwdrls} \rrbracket \vdash \text{certState} \rightarrow (\text{txcerts}, \text{CERT}*) \text{ certState}'$ 
194 •  $\llbracket \text{txid} \otimes \text{epoch slot} \otimes \text{pparams} \otimes \text{enactState} \rrbracket \vdash \text{govSt} \rightarrow (\text{txgov txb}, \text{GOV}*) \text{ govSt}'$ 
195
196  $\Gamma \vdash s \rightarrow (\text{tx}, \text{LEDGER}) \llbracket \text{utxoSt}' \otimes \text{govSt}' \otimes \text{certState}' \rrbracket$ 

```

197 (The notation $\llbracket \dots \otimes \dots \rrbracket$ constructs records of any type by giving their fields in order.)

198 **4 UTxO**199 **4.1 Witnessing**

200 Transaction witnessing checks that all required signatures are present and all required scripts
 201 accept the validity of the given transaction. *witsKeyHashes* and *witsScriptHashes* is the set
 202 of hashes of keys/scripts included in the transaction.

```

203 UTXOW-inductive :
204 • witsVKeyNeeded ppolicy utxo txb  $\subseteq$  witsKeyHashes
205 • scriptsNeeded ppolicy utxo txb  $\equiv$  witsScriptHashes
206 •  $\forall [ (vk, \sigma) \in \text{vkSigs} ] \text{isSigned vk (txidBytes txid)} \sigma$ 
207 •  $\forall [ s \in \text{scriptsP1} ] \text{validP1Script witsKeyHashes txvldt s}$ 
208 •  $\Gamma \vdash s \rightarrow (\text{tx}, \text{UTXO}) s'$ 
209
210  $\Gamma \vdash s \rightarrow (\text{tx}, \text{UTXOW}) s'$ 

```

211 **4.2 Accounting**

212 Accounting is handled by the UTXO state machine. The preconditions for UTXO-inductive
 213 ensure various properties or prevent attacks. For example, if txins was allowed to be empty,
 214 one could make a transaction that only spends from reward accounts. This does not require a
 215 specific hash to be present in the transaction body, so such a transaction could be repeatable in
 216 certain scenarios. The equation between produced and consumed ensures that the transaction
 217 is properly balanced. For details on some of these functions, see Appendix B.

<pre> 218 record UTXOEnv : Type where slot : Slot pparams : PParams Deposits = DepositPurpose → Coin </pre>	<pre> record UTXOState : Type where utxo : UTXO deposits : Deposits fees donations : Coin </pre>
<pre> 219 UTXO-inductive : 220 • txins ≠ ∅ 221 • txins ⊆ dom utxo 222 • minfee pp tx ≤ txfee 223 • txsize ≤ maxTxSize pp 224 • consumed pp s txb ≡ produced pp s txb 225 • coin mint ≡ 0 </pre> <hr/> <pre> 226 227 Γ ⊢ s →(tx ,UTXO) [(utxo txins) ∪ outs txb ⊗ updateDeposits pp txb deposits ⊗ fees + txfee ⊗ donations + txdonation] </pre>	

228 ► **Property 4.1 (Value preservation).**

229 *Let $getCoin$ be the sum of all coins contained within a $UTXOState$. Then, for all $\Gamma \in UTXOEnv$,*
 230 *$s, s' \in UTXOState$ and $tx \in Tx$, if $tx.body.txid \notin \text{map } proj_1 (\text{dom } (s . UTXOState.utxo))$ and Γ*
 231 *$\vdash s \rightarrow(tx ,UTXO) s'$ then $getCoin s \equiv getCoin s'$.*

232 Note that this is one of the most important properties of a UTXO-based ledger, as
 233 evidenced by its central place in EUTxO's meta-theory [9, 10].

234 5 Decentralized Governance

235 5.1 Entities and actions

236 The governance framework has three bodies of governance, the constitutional committee,
 237 delegated representatives and stake pool operators, corresponding to the roles **CC**, **DRep**
 238 and **SPO**. Proposals relevant to the governance system come in the form of Governance
 239 Actions. They are identified by their **GovActionID**, which consists of the **Txid** belonging to
 240 the transaction that proposed it and the index within that transaction (a transaction can
 241 propose multiple governance actions at once).

```

242 GovActionID = TxId × ℕ
243 data GovRole : Type where
244   CC DRep SPO : GovRole
245 data GovAction : Type where
246   NoConfidence      : GovAction
247   NewCommittee      : Credential → Epoch → ℙ Credential → ℚ → GovAction
248   NewConstitution   : DocHash → Maybe ScriptHash → GovAction
249   TriggerHF         : ProtVer → GovAction
250   ChangePParams    : PParamsUpdate → GovAction
251   TreasuryWdrl     : (RwdAddr → Coin) → GovAction
252   Info              : GovAction

```

253 For the meaning of these individual actions, see [12].

254 **5.2 Votes and proposals**

255 Before a `Vote` can be cast it must be packaged together with further information, such as
 256 who is voting and for which governance action. This information is combined in the `GovVote`
 257 record. To propose a governance action, a `GovProposal` needs to be submitted. Beside the
 258 proposed action, it requires a deposit, which will be returned to `returnAddr`.

<pre>259 data Vote : Type where yes no abstain : Vote</pre>	<pre>record GovVote : Type where gid : GovActionID role : GovRole credential : Credential vote : Vote</pre>	<pre>record GovProposal : Type where action : GovAction deposit : Coin returnAddr : RwdAddr</pre>
---	--	---

260 **5.3 Enactment**

261 Enactment of a governance action is carried out via the `ENACT` state machine. We just show
 262 two example rules for this state machine—there is one corresponding to each constructor of
 263 `GovAction`. For an explanation of the hash protection scheme, see Appendix A.

<pre>264 record EnactEnv : Type where gid : GovActionID treasury : Coin epoch : Epoch</pre>	<pre>record EnactState : Type where cc : HashProtected (Maybe ((Credential → Epoch) × ℚ)) constitution : HashProtected (DocHash × Maybe ScriptHash) pv : HashProtected ProtVer pparams : HashProtected PParams withdrawals : RwdAddr → Coin</pre>
---	---

265 `Enact-NewConst :`

266
$$\frac{}{\llbracket gid \otimes t \otimes e \rrbracket \vdash s \rightarrow (\text{NewConstitution } dh \ sh \ , \text{ENACT}) \text{ record } s \{ \text{constitution} = (dh \ , \ sh) \ , \ gid \}}$$

269 `Enact-Wdrl :`

270
$$\text{let } newWdrls = s .\text{withdrawals} \cup wdrl \text{ in } \sum [x \leftarrow newWdrls] x \leq t$$

272
$$\llbracket gid \otimes t \otimes e \rrbracket \vdash s \rightarrow (\text{TreasuryWdrl } wdrl \ , \text{ENACT}) \text{ record } s \{ \text{withdrawals} = newWdrls \}$$

273 (The `record` keyword indicates a record update, i.e. we take the existing `EnactState` and
 274 update one of its fields.)

275 **5.4 Voting and Proposing**

276 The order of proposals is maintained by keeping governance actions in a list—this acts as a
 277 tie breaker when multiple competing actions might be able to be ratified at the same time.

<pre> record GovActionState : Type where votes : (GovRole × Credential) → Vote returnAddr : RwdAddr expiresIn : Epoch action : GovAction prevAction : NeedsHash action GovState = List (GovActionID × GovActionState) </pre>	<pre> record GovEnv : Type where txid : TxId epoch : Epoch pparams : PParams enactState : EnactState </pre>
---	--

279 **GOV-Vote :**

- 280 • $(aid, ast) \in \text{fromList } s$
- 281 • $\text{canVote pparams (action ast) role}$

282
283 $(\Gamma, k) \vdash s \rightarrow (\text{sig}, \text{GOV}) \text{ addVote } s \text{ aid role cred } v$

285 **GOV-Propose :**

- 286 • $\text{actionWellFormed } a \equiv \text{true}$
- 287 • $d \equiv \text{govActionDeposit}$

288
289 $(\Gamma, k) \vdash s \rightarrow (\text{inj}_2 \text{ prop}, \text{GOV}) \text{ addAction } s (\text{govActionLifetime} + \text{epoch}) (\text{txid}, k) \text{ addr } a \text{ prev}$

290 5.5 Ratification

291 Governance actions are *ratified* through on-chain voting actions. Different kinds of governance
292 actions have different ratification requirements but always involve at least *two of the three*
293 governance bodies. The voting power of the **DRep** and **SPO** roles is proportional to the stake
294 delegated to them, while the constitutional committee has individually elected members
295 where each member has the same voting power.

296 Some actions take priority over others and, when enacted, delay all further ratification to
297 the next epoch boundary. This allows a changed government to reevaluate existing proposals.

<pre> record RatifyEnv : Type where stakeDistrs : StakeDistrs currentEpoch : Epoch dreps : Credential → Epoch </pre>	<pre> record RatifyState : Type where es : EnactState removed : ℙ (GovActionID × GovActionState) delay : Bool </pre>
--	---

299 **RATIFY-Accept :**

- 300 • $\text{accepted } \Gamma \text{ es } st$
- 301 • $\neg \text{delayed action prevAction es } d$
- 302 • $\llbracket a \text{ .proj}_1 \otimes \text{treasury} \otimes \text{currentEpoch} \rrbracket \vdash \text{es} \rightarrow (\text{action}, \text{ENACT}) \text{ es}'$

303
304 $\Gamma \vdash \llbracket \text{es} \otimes \text{removed} \quad \otimes d \quad \rrbracket \rightarrow (a, \text{RATIFY})$
305 $\llbracket \text{es}' \otimes \{ a \} \cup \text{removed} \otimes \text{delayingAction action} \rrbracket$

307 **RATIFY-Reject :**

- 308 • $\neg \text{accepted } \Gamma \text{ es } st$
- 309 • $\text{expired currentEpoch } st$

7:10 Formal specification of the Cardano blockchain ledger, mechanized in Agda

```

310
311   Γ ⊢ [ [ es ⊗ removed ⊗ d ] → ( a ,RATIFY ) [ [ es ⊗ { a } ∪ removed ⊗ d ] ]
312
313   RATIFY-Continue :
314     ( • → accepted Γ es st • → expired currentEpoch st )
315     ⊔ ( • accepted Γ es st
316         • ( delayed action prevAction es d
317             ⊔ ( ∀ es' → ¬ [ [ a .proj₁ ⊗ treasury ⊗ currentEpoch ] ⊢ es → ( action ,ENACT ) es' ) ) )
318
319   Γ ⊢ [ [ es ⊗ removed ⊗ d ] → ( a ,RATIFY ) [ [ es ⊗ removed ⊗ d ] ]

```

320 The main new ingredients for the rules of the `RATIFY` state machine are:

- 321 ■ `accepted`, which is the property that there are sufficient votes from the required bodies to pass this action;
- 322 ■ `delayed`, which expresses whether an action is delayed;
- 323 ■ `expired`, which becomes true a certain number of epochs after the action has been proposed.

324 The three `RATIFY` rules correspond to the cases where an action can be ratified and enacted (in which case it is), or it is expired and can be removed, or, otherwise it will be kept around for the future. This means that all governance actions eventually either get accepted and enacted via `RATIFY-Accept` or rejected via `RATIFY-Reject`. It is not possible to remove actions by voting against them, one has to wait for the action to expire.

6 Transactions

330 A transaction is made up of a transaction body and a collection of witnesses.

```

332   Ix TxId : Type
333   TxIn  = TxId × Ix
334   TxOut = Addr × Value × Maybe DataHash
335   UTxO  = TxIn → TxOut

```

```

record TxBody : Type where
  txins  : ℙ TxIn
  txouts : Ix → TxOut
  txfee  : Coin
  txvote : List GovVote
  txprop : List GovProposal
  txsize : ℕ
  txid   : TxId

record TxWitnesses : Type where
  vkSigs : VKey → Sig
  scripts : ℙ Script

record Tx : Type where
  body  : TxBody
  wits  : TxWitnesses

```

337 Some key ingredients in the transaction body are:

- 338 ■ A set of transaction inputs (`txins`), each of which identifies an output from a previous transaction. A transaction input (`TxIn`) consists of a transaction ID and an index to uniquely identify the output.
- 339 ■ An indexed collection of transaction outputs (`txouts`). A transaction output (`TxOut`) is an address paired with a multi-asset `Value` (see [10]).
- 340 ■ A transaction fee (`txfee`), whose value will be added to the fee pot.

- 344 ■ The size (`txsize`) and the hash (`txid`) of the serialized form of the transaction that was
 345 included in the block. Cardano’s serialization is not canonical, so any information that is
 346 necessary but lost during deserialisation must be preserved by attaching it to the data
 347 like this.

348 7 Compiling to a Haskell implementation & Conformance testing

349 In order to deliver on our promise that the specification is also *executable*, there is still some
 350 work to be done given that all transitions have been formulated as relations.

351 This is precisely the reason we also manually prove that each and every transition of the
 352 previous sections is indeed *computational*:

```
353 record Computational (⟦_⟧_ → (⟦_, X ⟧)_ : C → S → Sig → S → Type) : Type where
354   compute          : C → S → Sig → Maybe S
355   compute-correct : compute Γ s b ≡ just s' ⇔ Γ ⊢ s → (⟦ b , X ⟧) s'
```

356 The definition above captures what it means for a (small-step) relation to be accurately
 357 computed by a function `compute`, which given as input an environment, source state, and
 358 signal, outputs the resulting state or an error for invalid transitions. Most importantly, such
 359 a function must be *sound* and *complete*: it does not return output states that are not related,
 360 and, *vice versa*, all related states are successfully returned. An alternative interpretation is
 361 that this rules out *non-determinism* across all ledger transitions, *i.e.*, there cannot be two
 362 distinct states arising from the same inputs.

363 There is one last obstacle that hinders execution: we have leveraged Agda’s module
 364 system³ to parameterize our specification over some abstract types and functions that we
 365 assume as given, *e.g.*, the cryptographic primitives. As a final step, we instantiate these
 366 parameters with concrete definitions, either by manually providing them within Agda, or
 367 deferring to the Haskell *foreign function interface* to reuse existing Haskell ones that have no
 368 Agda counterpart.

369 Equipped with a fully concrete specification and the `Computational` proofs for each relation,
 370 it is finally possible to generate executable Haskell code using Agda’s MAlonzo compilation
 371 backend.⁴ The resulting Haskell library is then deployed as part of the automated testing
 372 setup for the Cardano ledger in production, so as to ensure the developers have faithfully
 373 implemented the specification. This is made possible by virtue of the implementation
 374 mirroring the specification’s structure to define transitions, which one can then test by
 375 randomly generating environments/states/signals, and executing both state machines on
 376 these same random inputs to compare the final results for *conformance*.

377 One small caveat remains though: production code might use different data structures,
 378 mainly for reasons of *performance*, which are not isomorphic to those used in the specification
 379 and might require non-trivial translation functions and notions of equality to perform
 380 the aforementioned tests. In the future, we plan to also formalize these more efficient
 381 representations in Agda and prove that soundness is preserved regardless.

382 8 Related Work

383 **EUTxO.** The approach we followed is a natural evolution of prior meta-theoretical results

³ <https://agda.readthedocs.io/en/v2.6.4/language/module-system.html#parameterised-modules>

⁴ <https://agda.readthedocs.io/en/v2.6.4/tools/compilers.html#ghc-backend>

384 on the EUTxO model [9, 10], but now employed at a much larger scale to cover all the
 385 features of a realistic ledger: epochs, protocol parameters, decentralized governance, *etc.*

386 All this complexity does not come for free though: one has to be economical about
 387 which properties to prove of the resulting system, and this might entail limiting oneself
 388 to mechanizing just the core properties, such as global value preservation as we saw with
 389 Property 4.1, otherwise the whole effort can quickly become practically infeasible to maintain
 390 from a software-engineering perspective.

391 **Formal Methods, generally.** The overarching methodology—formally specifying the
 392 system under design—is by no means particular to the blockchain space. A principal success
 393 story in the wider computing world nowadays is definitely the *WebAssembly* language,
 394 an alternative to Javascript to act as a compilation target for web applications with
 395 performance and security in mind [16], which was designed in tandem with a formalization
 396 of its semantics [29].

397 Apart from keeping programming language designers honest by making sure no edge
 398 cases are overlooked, it allows the language to evolve in a much more robust fashion: every
 399 future extension has to pass through a rigorous process which eventually involves extending
 400 the formalization itself.

401 While the WebAssembly line of work [29, 30] provided much inspiration for us, we believe
 402 our approach to be even more radical by mitigating the need for informal processes altogether:
 403 the formalization *is* the specification!

404 **Formal Methods, specifically for blockchain.** The work presented here fits well within
 405 Cardano’s vision for *agile formal methods* [17], which strikes a good balance between a fully
 406 certified implementation (too much effort, too few resources) and an informal, under-specified
 407 product (quicker, easier, but far less trustworthy). Instead of demanding the impossible by
 408 extracting the actual production from the formalization itself, we find the sweet spot lies in
 409 the middle: extracting a *reference implementation* in Haskell and using *conformance testing*
 410 to ensure the system in production behaves as it should (*c.f.*, Section 7).

411 Apart from our work, there are very few mechanized results on UTXO-based blockchains
 412 (modeled after Bitcoin [19]), and all of them invariably are formulated on a idealized
 413 setting [26, 1, 9, 10], abstracting away the complexity that ensues when multiple features
 414 interact. Thus, the mechanized specification presented here for the Cardano ledger is the
 415 first of its kind, and we hope this sets a higher standard for subsequent work and pushes
 416 forward a more formal agenda for blockchain research in the future.

417 Although not directly comparable to our use case, account-based blockchains (modeled
 418 after Ethereum [8]) fair better in this respect, with plenty of formal method tools available,
 419 ranging from model checking [15, 28] to full-blown formal verification [11, 7, 23]. Notable
 420 blockchains that spearhead progress in this direction include Tezos [5, 6, 14], Zilliqa and its
 421 Scilla smart-contract language [25, 24], and Concordium [3, 21, 2, 27, 20]. The main difference
 422 with our work lies in *readability*, partly due to the choice of tool (Agda being notorious for
 423 its beautiful renderings but lack of proper support for practical “big” proofs that arise in
 424 large scale software verification projects, where tactic-based proof assistants like Coq [4]
 425 and Isabelle [22] are more common), and the point where mechanization is placed within
 426 the development pipeline: most aforementioned work builds upon informal pen-and-paper
 427 documents and some of its aspects are only mechanized *a-posteriori*. Having said that,
 428 the fundamental split stems from a completely different *target audience*; our formalization
 429 is meant to be read by researchers, formal methods engineers, compiler engineers, and
 430 developers alike. In contrast, the majority of the aforementioned work is primarily targeted
 431 at a select team of experts which complement other (informal) documentation and software.

432 **9 Conclusion**

433 We have outlined the mechanized specification of the EUTxO-based ledger rules of the
434 Cardano blockchain, by taking a *bird's-eye view* of the hierarchy of transitions handling
435 different sub-components in a modular way.

436 Although space limitations preclude us from exhaustively fleshing out all the gory details
437 of our formalization, we hope to have conveyed the general *design principles* that will be
438 helpful to others when attempting to mechanize something of this kind and at this scale.
439 In the little space we could afford for more thorough details, we made a conscious choice
440 of putting emphasis on the most novel aspect of the current era of the Cardano blockchain:
441 *decentralized governance*. There, the introduction of the notions of voting, ratification, and
442 enactment complicate the ledger rules of previous eras—albeit in a fairly orthogonal way,
443 which we found particularly satisfying.

444 A mechanized formal artifact of this kind is rigid enough to eliminate any ambiguity
445 that would often arise in pen-and-paper specifications, all the while sustaining a readable
446 document that is accessible to a wide audience and allows for varied uses.

447 By virtue of conducting our work within a proof assistant based on *constructive* logic,
448 our result extends beyond a purely theoretical exercise to an *executable* resource that can be
449 leveraged as a *reference implementation*, against which a system-in-production can be tested
450 for conformance.

451 Last but not least, it is evident that developing a ledger on these foundations opens up
452 a plethora of opportunities for further formalization work, *e.g.*, instantiating the abstract
453 notion of scripts with actual *Plutus* scripts brings us close to enabling practical smart
454 contract verification where developers write their programs immediately in Agda, prove
455 properties about their behavior, and then extract Plutus code they can deploy to the actual
456 Cardano blockchain. All these point to bright prospects for formal methods in UTxO-based
457 blockchains, which we are excited to explore in the future and hope that others do as well.

458 — References

- 459 1 Fahad F. Alhabardi, Arnold Beckmann, Bogdan Lazar, and Anton Setzer. Verification of
460 Bitcoin Script in Agda using weakest preconditions for access control. In Henning Basold,
461 Jesper Cockx, and Silvia Ghilezan, editors, *27th International Conference on Types for Proofs
462 and Programs, TYPES 2021, June 14-18, 2021, Leiden, The Netherlands (Virtual Conference)*,
463 volume 239 of *LIPICs*, pages 1:1–1:25. Schloss Dagstuhl - Leibniz-Zentrum für Informatik,
464 2021. doi:10.4230/LIPICs.TYPES.2021.1.
- 465 2 Danil Annenkov, Mikkel Milo, Jakob Botsch Nielsen, and Bas Spitters. Extracting smart
466 contracts tested and verified in Coq. In Catalin Hritcu and Andrei Popescu, editors, *CPP '21:
467 10th ACM SIGPLAN International Conference on Certified Programs and Proofs, Virtual Event,
468 Denmark, January 17-19, 2021*, pages 105–121. ACM, 2021. doi:10.1145/3437992.3439934.
- 469 3 Danil Annenkov, Jakob Botsch Nielsen, and Bas Spitters. Concert: a smart contract
470 certification framework in Coq. In Jasmin Blanchette and Catalin Hritcu, editors, *Proceedings
471 of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs ,
472 CPP 2020, New Orleans, LA, USA, January 20-21, 2020*, pages 215–228. ACM, 2020.
473 doi:10.1145/3372885.3373829.
- 474 4 Bruno Barras, Samuel Boutin, Cristina Cornes, Judicaël Courant, Jean-Christophe Filliatre,
475 Eduardo Gimenez, Hugo Herbelin, Gerard Huet, Cesar Munoz, Chetan Murthy, et al. *The
476 Coq proof assistant reference manual: Version 6.1*. PhD thesis, Inria, 1997.
- 477 5 Bruno Bernardo, Raphaël Cauderlier, Guillaume Claret, Arvid Jakobsson, Basile Pesin, and
478 Julien Tesson. Making tezos smart contracts more reliable with Coq. In Tiziana Margaria
479 and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods, Verification and
480 Validation: Applications - 9th International Symposium on Leveraging Applications of Formal
481 Methods, ISoLA 2020, Rhodes, Greece, October 20-30, 2020, Proceedings, Part III*, volume
482 12478 of *Lecture Notes in Computer Science*, pages 60–72. Springer, 2020. doi:10.1007/
483 978-3-030-61467-6_5.
- 484 6 Bruno Bernardo, Raphaël Cauderlier, Zhenlei Hu, Basile Pesin, and Julien Tesson. Mi-cho-coq,
485 a framework for certifying Tezos smart contracts. In Emil Sekerinski, Nelma Moreira, José N.
486 Oliveira, Daniel Ratiu, Riccardo Guidotti, Marie Farrell, Matt Luckcuck, Diego Marmosler,
487 José Creissac Campos, Troy Astarte, Laure Gonnord, Antonio Cerone, Luis Couto, Brijesh
488 Dongol, Martin Kutrib, Pedro Monteiro, and David Delmas, editors, *Formal Methods. FM
489 2019 International Workshops - Porto, Portugal, October 7-11, 2019, Revised Selected Papers,
490 Part I*, volume 12232 of *Lecture Notes in Computer Science*, pages 368–379. Springer, 2019.
491 doi:10.1007/978-3-030-54994-7_28.
- 492 7 Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Anitha Gollamudi, Georges
493 Gonthier, Nadim Kobeissi, Natalia Kulatova, Aseem Rastogi, Thomas Sibut-Pinote, Nikhil
494 Swamy, et al. Formal verification of smart contracts: Short paper. In *Proceedings of the 2016
495 ACM Workshop on Programming Languages and Analysis for Security*, pages 91–96, 2016.
- 496 8 Vitalik Buterin. A next-generation smart contract and decentralized application platform
497 (white paper). [https://ethereum.org/content/whitepaper/whitepaper-pdf/Ethereum_](https://ethereum.org/content/whitepaper/whitepaper-pdf/Ethereum_Whitepaper_-_Buterin_2014.pdf)
498 [Whitepaper_-_Buterin_2014.pdf](https://ethereum.org/content/whitepaper/whitepaper-pdf/Ethereum_Whitepaper_-_Buterin_2014.pdf), 2014.
- 499 9 Manuel M. T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian,
500 Michael Peyton Jones, and Philip Wadler. The Extended UTXO model. In Matthew Bernhard,
501 Andrea Bracciali, L. Jean Camp, Shin'ichiro Matsuo, Alana Maurushat, Peter B. Rønne, and
502 Massimiliano Sala, editors, *Financial Cryptography and Data Security - FC 2020 International
503 Workshops, AsiaUSEC, CoDeFi, VOTING, and WTSC, Kota Kinabalu, Malaysia, February
504 14, 2020, Revised Selected Papers*, volume 12063 of *Lecture Notes in Computer Science*, pages
505 525–539. Springer, 2020. doi:10.1007/978-3-030-54455-3_37.
- 506 10 Manuel M. T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian, Jann
507 Müller, Michael Peyton Jones, Polina Vinogradova, and Philip Wadler. Native custom tokens
508 in the Extended UTXO model. In Tiziana Margaria and Bernhard Steffen, editors, *Leveraging
509 Applications of Formal Methods, Verification and Validation: Applications - 9th International*

- 510 *Symposium on Leveraging Applications of Formal Methods, ISoLA 2020, Rhodes, Greece,*
 511 *October 20-30, 2020, Proceedings, Part III*, volume 12478 of *Lecture Notes in Computer*
 512 *Science*, pages 89–111. Springer, 2020. doi:10.1007/978-3-030-61467-6_7.
- 513 11 Xiaohong Chen, Daejun Park, and Grigore Roşu. A language-independent approach to smart
 514 contract verification. In *International Symposium on Leveraging Applications of Formal*
 515 *Methods*, pages 405–413. Springer, 2018.
- 516 12 Jared Corduan, Matthias Benkort, Kevin Hammond, Charles Hoskinson, Andre Knispel, and
 517 Samuel Leathers. A first step towards on-chain decentralized governance. [https://cips.](https://cips.cardano.org/cip/CIP-1694)
 518 [cardano.org/cip/CIP-1694](https://cips.cardano.org/cip/CIP-1694), 2023.
- 519 13 Bernardo Machado David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. Ouroboros
 520 Praos: An adaptively-secure, semi-synchronous proof-of-stake protocol. *IACR Cryptology*
 521 *ePrint Archive*, 2017:573, 2017.
- 522 14 Christopher Goes. Compiling Quantitative Type Theory to Michelson for compile-time
 523 verification and run-time efficiency in juvix. In Tiziana Margaria and Bernhard Steffen, editors,
 524 *Leveraging Applications of Formal Methods, Verification and Validation: Applications - 9th*
 525 *International Symposium on Leveraging Applications of Formal Methods, ISoLA 2020, Rhodes,*
 526 *Greece, October 20-30, 2020, Proceedings, Part III*, volume 12478 of *Lecture Notes in Computer*
 527 *Science*, pages 146–160. Springer, 2020. doi:10.1007/978-3-030-61467-6_10.
- 528 15 Neville Grech, Michael Kong, Anton Jurisevic, Lexi Brent, Bernhard Scholz, and Yannis
 529 Smaragdakis. Madmax: Surviving out-of-gas conditions in Ethereum smart contracts.
 530 *Proceedings of the ACM on Programming Languages*, 2(OOPSLA):1–27, 2018.
- 531 16 Andreas Haas, Andreas Rossberg, Derek L. Schuff, Ben L. Titzer, Michael Holman, Dan
 532 Gohman, Luke Wagner, Alon Zakai, and J. F. Bastien. Bringing the web up to speed with
 533 WebAssembly. In Albert Cohen and Martin T. Vechev, editors, *Proceedings of the 38th*
 534 *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI*
 535 *2017, Barcelona, Spain, June 18-23, 2017*, pages 185–200. ACM, 2017. doi:10.1145/3062341.
 536 3062363.
- 537 17 Philipp Kant, Kevin Hammond, Duncan Coutts, James Chapman, Nicholas Clarke, Jared
 538 Corduan, Neil Davies, Javier Díaz, Matthias Güdemann, Wolfgang Jeltsch, Marcin Szamotulski,
 539 and Polina Vinogradova. Flexible formality: Practical experience with agile formal methods.
 540 In Aleksander Byrski and John Hughes, editors, *Trends in Functional Programming - 21st*
 541 *International Symposium, TFP 2020, Krakow, Poland, February 13-14, 2020, Revised Selected*
 542 *Papers*, volume 12222 of *Lecture Notes in Computer Science*, pages 94–120. Springer, 2020.
 543 doi:10.1007/978-3-030-57761-2_5.
- 544 18 Andre Knispel. Constructive zf-style set theory in type theory. unpublished, 2023. URL:
 545 <https://whatisrt.github.io/papers/ZF-style-set-theory-in-type-theory.pdf>.
- 546 19 S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. [https://bitcoin.org/en/](https://bitcoin.org/en/bitcoin-paper)
 547 [bitcoin-paper](https://bitcoin.org/en/bitcoin-paper), October 2008.
- 548 20 Eske Hoy Nielsen, Danil Annenkov, and Bas Spitters. Formalising decentralised exchanges in
 549 Coq. In Robbert Krebbers, Dmitriy Traytel, Brigitte Pientka, and Steve Zdancewic, editors,
 550 *Proceedings of the 12th ACM SIGPLAN International Conference on Certified Programs and*
 551 *Proofs, CPP 2023, Boston, MA, USA, January 16-17, 2023*, pages 290–302. ACM, 2023.
 552 doi:10.1145/3573105.3575685.
- 553 21 Jakob Botsch Nielsen and Bas Spitters. Smart contract interactions in Coq. In Emil Sekerinski,
 554 Nelma Moreira, José N. Oliveira, Daniel Ratiu, Riccardo Guidotti, Marie Farrell, Matt
 555 Luckcuck, Diego Marmosoler, José Creissac Campos, Troy Astarte, Laure Gonnord, Antonio
 556 Cerone, Luis Couto, Brijesh Dongol, Martin Kutrib, Pedro Monteiro, and David Delmas,
 557 editors, *Formal Methods. FM 2019 International Workshops - Porto, Portugal, October 7-11,*
 558 *2019, Revised Selected Papers, Part I*, volume 12232 of *Lecture Notes in Computer Science*,
 559 pages 380–391. Springer, 2019. doi:10.1007/978-3-030-54994-7_29.
- 560 22 Tobias Nipkow, Lawrence C Paulson, and Markus Wenzel. *Isabelle/HOL: a proof assistant for*
 561 *higher-order logic*, volume 2283. Springer Science & Business Media, 2002.

- 562 23 George Pirlea and Ilya Sergey. Mechanising blockchain consensus. In June Andronick and
563 Amy P. Felty, editors, *Proceedings of the 7th ACM SIGPLAN International Conference on*
564 *Certified Programs and Proofs, CPP 2018, Los Angeles, CA, USA, January 8-9, 2018*, pages
565 78–90. ACM, 2018. doi:10.1145/3167086.
- 566 24 Ilya Sergey, Amrit Kumar, and Aquinas Hobor. Temporal properties of smart contracts. In
567 Tiziana Margaria and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods,*
568 *Verification and Validation. Industrial Practice - 8th International Symposium, ISoLA 2018,*
569 *Limassol, Cyprus, November 5-9, 2018, Proceedings, Part IV*, volume 11247 of *Lecture Notes*
570 *in Computer Science*, pages 323–338. Springer, 2018. doi:10.1007/978-3-030-03427-6_25.
- 571 25 Ilya Sergey, Vaivaswatha Nagaraj, Jacob Johannsen, Amrit Kumar, Anton Trunov, and Ken
572 Chan Guan Hao. Safer smart contract programming with Scilla. *Proc. ACM Program. Lang.*,
573 3(OOPSLA):185:1–185:30, 2019. doi:10.1145/3360611.
- 574 26 Anton Setzer. Modelling Bitcoin in Agda. *CoRR*, abs/1804.06398, 2018. URL: <http://arxiv.org/abs/1804.06398>,
575 [arXiv:1804.06398](https://arxiv.org/abs/1804.06398).
- 576 27 Søren Eller Thomsen and Bas Spitters. Formalizing Nakamoto-style proof of stake. In *34th*
577 *IEEE Computer Security Foundations Symposium, CSF 2021, Dubrovnik, Croatia, June 21-25,*
578 *2021*, pages 1–15. IEEE, 2021. doi:10.1109/CSF51468.2021.00042.
- 579 28 Petar Tsankov. Security analysis of smart contracts in Datalog. In *International Symposium*
580 *on Leveraging Applications of Formal Methods*, pages 316–322. Springer, 2018.
- 581 29 Conrad Watt. Mechanising and verifying the WebAssembly specification. In June Andronick
582 and Amy P. Felty, editors, *Proceedings of the 7th ACM SIGPLAN International Conference on*
583 *Certified Programs and Proofs, CPP 2018, Los Angeles, CA, USA, January 8-9, 2018*, pages
584 53–65. ACM, 2018. doi:10.1145/3167082.
- 585 30 Conrad Watt, Maja Trela, Peter Lammich, and Florian Märkl. Wasmref-isabelle: A verified
586 monadic interpreter and industrial fuzzing oracle for WebAssembly. *Proc. ACM Program.*
587 *Lang.*, 7(PLDI):100–123, 2023. doi:10.1145/3591224.
- 588 31 Joachim Zahentferner. Chimeric ledgers: Translating and unifying UTXO-based and account-
589 based cryptocurrencies. *Cryptology ePrint Archive, Report 2018/262*, 2018. URL: <https://eprint.iacr.org/2018/262>.

A Governance helper calculations

The design of the hash protection mechanism is elaborated here. The issue at hand is that different actions of the same type may override each other, and they allow for partial modifications to the state. So if arbitrary actions were allowed to be applied, the system may end up in a particular state that was never intended and voted for.

In the original design of the governance system, the fix for this issue was to allow only a single governance action of each type to be enacted per epoch. This restriction is a potentially severe limitation and may open the door to some types of attacks.

The final design instead requires some types of governance actions to reference the ID of the parent they are building on, similar to a Merkle tree. Then, in a single epoch the system can take arbitrarily many steps down that tree, and since IDs are unforgeable, the system is only ever in a state that was publically known prior to voting.

There are two governance actions where this mechanism is not required, because they either commute naturally or they do not actually affect the state. For these it is more convenient to not enforce dependencies.

```

606 NeedsHash : GovAction → Type
607 NeedsHash NoConfidence      = GovActionID
608 NeedsHash (NewCommittee _ _ _) = GovActionID
609 NeedsHash (NewConstitution _ _) = GovActionID
610 NeedsHash (TriggerHF _)       = GovActionID
611 NeedsHash (ChangePParams _)   = GovActionID
612 NeedsHash (TreasuryWdrl _)    = T
613 NeedsHash Info                = T

```

```

614
615 HashProtected : Type → Type
616 HashProtected A = A × GovActionID

```

The two functions adjusting the state in GOV are `addVote` and `addAction`.

- `addVote` inserts (and potentially overrides) a vote made for a particular governance action by a credential in a role.
- `addAction` adds a new proposed action at the end of a given `GovState`, properly initializing all the required fields.

```

623 addVote : GovState → GovActionID → GovRole → Credential → Vote → GovState
624 addVote s aid r kh v = map modifyVotes s
625   where modifyVotes = λ (gid , s') → gid , record s'
626         { votes = if gid ≡ aid then insert (votes s') (r , kh) v else votes s' }
627
628 addAction : GovState
629           → Epoch → GovActionID → RwdAddr → (a : GovAction) → NeedsHash a
630           → GovState
631 addAction s e aid addr a prev = s :: (aid , record
632   { votes = ∅ ; returnAddr = addr ; expiresIn = e ; action = a ; prevAction = prev })

```

634 **B** UTxO

635 Some of the functions used to define the `UTxO` and `UTxOW` state machines are defined here;
 636 `inject` is the function takes a `Coin` and turns it into a multi-asset `Value` [10].

```
637 outs : TxBody → UTxO
638 outs tx = mapKeys (tx .txid ,_) (tx .txouts)
639
640 minfee : PParams → Tx → Coin
641 minfee pp tx = pp .a * tx .body .txsize + pp .b
```

```
642
643 consumed : PParams → UTxOState → TxBody → Value
644 consumed pp st txb
645   = balance (st .utxo | txb .txins)
646   + txb .mint
647   + inject (depositRefunds pp st txb)
```

```
648
649 produced : PParams → UTxOState → TxBody → Value
650 produced pp st txb
651   = balance (outs txb)
652   + inject (txb .txfee)
653   + inject (newDeposits pp st txb)
654   + inject (txb .txdonation)
```

```
655
656 credsNeeded : Maybe ScriptHash → UTxO → TxBody → ℙ (ScriptPurpose × Credential)
657 credsNeeded p utxo txb
658   = map (λ (i , o) → (Spend i , payCred (proj1 o))) ((utxo | txins) )
659   ∪ map (λ a → (Rwrd a , RwdAddr.stake a)) (dom $ txwdrls .proj1)
660   ∪ map (λ c → (Cert c , cwitness c)) (fromList txcerts)
661   ∪ map (λ x → (Mint x , inj2 x)) (policies mint)
662   ∪ map (λ v → (Vote v , GovVote.credential v)) (fromList txvote)
663   ∪ (if p then (λ {sh} → map (λ p → (Propose p , inj2 sh)) (fromList txprop))
664       else ∅)
665   where open TxBody txb
```

```
666
667 witsVKeyNeeded : Maybe ScriptHash → UTxO → TxBody → ℙ KeyHash
668 witsVKeyNeeded sh = mapPartial isInj1 o2 map proj2 o2 credsNeeded sh
```

```
669
670 scriptsNeeded : Maybe ScriptHash → UTxO → TxBody → ℙ ScriptHash
671 scriptsNeeded sh = mapPartial isInj2 o2 map proj2 o2 credsNeeded sh
```

672

673 **C** Advancing epochs

674 The `NEWEPOCH` state machine is responsible for detecting epoch changes: either the epoch
 675 remains unchanged (`NEWEPOCH-Not-New`), or the immediately next epoch is reached and
 676 the state is updated subject to some ratification requirements (`NEWEPOCH-New`).

```

677 NEWEPOCH-New :
678 •  $e \equiv \text{succ lastEpoch}$ 
679 •  $\text{record } \{ \text{currentEpoch} = e ; \text{treasury} = \text{treasury} ; \text{GState } g\text{State} ; \text{NewEpochEnv } \Gamma \}$ 
680    $\vdash \llbracket es \otimes \emptyset \otimes \text{false} \rrbracket \rightarrow \langle \text{govSt}' , \text{RATIFY*} \rangle \text{ fut}'$ 
681   

---


682    $\Gamma \vdash nes \rightarrow \langle e , \text{NEWEPOCH} \rangle \llbracket e \otimes \text{acct}' \otimes \text{ls}' \otimes es \otimes \text{fut}' \rrbracket$ 
683
684 NEWEPOCH-Not-New :
685  $e \not\equiv \text{succ lastEpoch}$ 
686   

---


687    $\Gamma \vdash nes \rightarrow \langle e , \text{NEWEPOCH} \rangle nes$ 
688

```