

A Formal Specification of the Cardano Ledger

(git revision 1.1-486-g301fede)

Jared Corduan
jared.corduan@iohk.io

Polina Vinogradova
polina.vinogradova@iohk.io

Matthias GÜdemann
matthias.gudemann@iohk.io

February 20, 2019

Abstract

This document defines the rules for extending a ledger with transactions. The transactions will affect both UTxO and stake delegation. It is intended to serve as the specification for random generators of transactions which adhere to the rules presented here.

List of Contributors

Nicolás Arqueros, Nicholas Clarke, Duncan Coutts, Ruslan Dudin, Sebastien Guillemot, Vincent Hanquez, Ru Horlick, Michael Hueschen, Philipp Kant, Jean-Christophe Mincke, Damian Nadales, Nicolas Di Prima.

Contents

1	Introduction	4
2	Notation	5
3	Cryptographic primitives	6
4	Addresses	7
5	Protocol Parameters	8
6	Transactions	10
7	UTxO	12
7.1	UTxO Transitions	12
7.2	Deposits and Refunds	17
7.3	Witnesses	20
8	Delegation	22
8.1	Delegation Definitions	22
8.2	Delegation Transitions	24
8.3	Delegation Rules	26
8.4	Stake Pool Rules	28
8.5	Delegation and Pool Combined Rules	30
9	Ledger State Transition	33

10 Rewards and the Epoch Boundary	35
10.1 Overview of the Reward Calculation	35
10.2 Helper Functions	35
10.3 Stake Distribution Calculation	36
10.4 Snapshot Transition	39
10.5 Pool Reaping Transition	41
10.6 Protocol Parameter Update Transition	43
10.7 Complete Epoch Boundary Transition	45
10.8 Rewards Distribution Calculation	47
10.9 Reward Transition	52
11 Properties	56
11.1 Validity of a Ledger State	56
11.2 Ledger Properties	56
11.3 Ledger State Properties for Delegation Transitions	58
11.4 Ledger State Properties for Staking Pool Transitions	59
11.5 Properties of Numerical Calculations	60
12 Non-Integral Calculations	62
12.1 Types of Non-Integral Calculations	62
12.2 Implementation of Non-Integer Calculations	62
12.2.1 Function Simplification	62
12.2.2 Properties of Function Approximation	63
A Proofs	64
References	65

List of Figures

1	Non-standard map operators	5
2	Cryptographic definitions	6
3	Definitions used in Addresses	7
4	Definitions used in Protocol Parameters	9
5	Definitions used in the UTxO transition system	11
6	Functions used in UTxO rules	13
7	UTxO transition-system types	14
8	UTxO inference rules	16
9	Functions used in Deposits - Refunds	18
10	Functions used in Deposits - Decay	19
11	Functions used in witness rule	20
12	UTxO with witness transition-system types	20
13	UTxO with witnesses inference rules	21
14	Delegation Definitions	23
15	Delegation Transitions	25
16	Delegation Inference Rules	27
17	Pool Inference Rule	29
18	Delegation and Pool Combined Transition Type	30
19	Delegation and Pool Combined Transition Rules	31
20	Delegation sequence transition type	31
21	Delegation sequence rules	32
22	Ledger transition-system types	33

23	Ledger inference rule	34
24	Helper Functions used in Rewards and Epoch Boundary	36
25	Epoch definitions	37
26	Stake Distribution Function	38
27	Snapshot transition-system types	39
28	Snapshot Inference Rule	40
29	Pool Reap Transition	41
30	Pool Reap Inference Rule	42
31	New Proto Param transition-system types	43
32	New Proto Param Inference Rule	44
33	Epoch transition-system types	45
34	Epoch Inference Rule	46
35	Functions used in the Reward Calculation	48
36	Functions used in the Reward Splitting	49
37	The Reward Calculation	51
38	Rewards transition-system types	52
39	Preservation of Value	53
40	Rewards inference rules	55
41	Definitions and Functions for Valid Ledger State	56
42	Definitions and Functions for Stake Delegation in Ledger States	58

1 Introduction

This document is a formal specification of the functionality of the ledger on the blockchain. The blockchain layer of the protocol and the interaction between the ledger and the blockchain layer is presented in a separate document, see [Formal Methods Team \(TODO\)](#). The details of the background and the larger context for the design decisions formalized in this document are presented in [Kant et al. \(2018\)](#)

In this work, we present important properties any implementation of the ledger must have. Specifically, we model the following aspects of the functionality of the ledger on the blockchain:

Preservation of value Every coin in the system is accounted for, and the total amount is unchanged by every transaction and epoch change. In particular, every coin is accounted for by one of the following categories:

- Circulation (UTxO)
- Deposit pot
- Fee pot
- Reserves (monetary expansion)
- Rewards (account addresses)
- Reward pot (undistributed)
- Treasury

Witnesses Authentication of parts of the transaction data by means of cryptographic entities (such as signatures and private keys) contained in these transactions.

Delegation Validity of delegation certificates, which delegate block-signing rights.

Stake Staking rights associated to an address.

While the blockchain protocol is a reactive system driven by the arrival of blocks causing updates to the ledger, the formal description is a collection of rules which is a static description of what a *valid ledger* is. The specifics of the semantics we use to define and apply the rules we present in this document are explained in detail in [Formal Methods Team \(2018\)](#). A valid ledger state can only be reached by applying a sequence of inference rules, and any valid ledger state is reachable by applying some sequence of these rules.

The structure of the rules we give here is such that their application is deterministic. That is, given a specific initial state and relevant environmental constants, there is no ambiguity about which rule should be applied at any given time (i.e. which state transition is allowed to take place). This is an important property which reflects the reality of the implementation — the blockchain evolves in a particular way given some user activity and the passage of time, and its behaviour is never unexpected.

2 Notation

The transition system is explained in [Formal Methods Team \(2018\)](#).

Powerset Given a set X , $\mathbb{P} X$ is the set of all the subsets of X .

Sequences Given a set X , X^* is the set of sequences having elements taken from X . The empty sequence is denoted by ϵ , and given a sequence Λ , $\Lambda;x$ is the sequence that results from appending $x \in X$ to Λ .

Functions $A \rightarrow B$ denotes a **total function** from A to B . Given a function f we write $f a$ for the application of f to argument a .

Inverse Image Given a function $f : A \rightarrow B$ and $b \in B$, we write $f^{-1} b$ for the **inverse image** of f at b , which is defined by $\{a \mid f a = b\}$.

Maps and partial functions $A \mapsto B$ denotes a **partial function** from A to B , which can be seen as a map (dictionary) with keys in A and values in B . Given a map $m \in A \mapsto B$, notation $a \mapsto b \in m$ is equivalent to $m a = b$.

Map Operations Figure 1 describes some non-standard map operations.

Relations A relation on $A \times B$ is a subset of $A \times B$. Both maps and functions can be thought of as relations. A function $f : A \rightarrow B$ is a relation consisting of pairs $(a, f(a))$ such that $a \in A$. A map $m : A \mapsto B$ is a relation consisting of pairs (a, b) such that $a \mapsto b \in m$. Given a relation R on $A \times B$, we define the inverse relation R^{-1} to be all pairs (b, a) such that $(a, b) \in R$. Similarly, given a function $f : A \rightarrow B$ we define inverse relation f^{-1} to consist of all pairs $(f(a), a)$. Finally, given two relations $R \subseteq A \times B$ and $S \subseteq B \times C$, we define the composition $R \circ S$ to be all pairs (a, c) such that $(a, b) \in R$ and $(b, c) \in S$ for some $b \in B$.

In Figure 1, we specify the notation we use in the rest of the document.

$set \triangleleft map = \{k \mapsto v \mid k \mapsto v \in map, k \in set\}$	domain restriction
$set \not\leftarrow map = \{k \mapsto v \mid k \mapsto v \in map, k \notin set\}$	domain exclusion
$map \triangleright set = \{k \mapsto v \mid k \mapsto v \in map, v \in set\}$	range restriction
$map \not\triangleright set = \{k \mapsto v \mid k \mapsto v \in map, v \notin set\}$	range exclusion
$A \triangle B = (A \setminus B) \cup (B \setminus A)$	symmetric difference
$M \cup_{\rightarrow} N = (\text{dom } N \not\leftarrow M) \cup N$	union override right
$M \cup_{\leftarrow} N = M \cup (\text{dom } M \not\leftarrow N)$	union override left
$M \cup_{+} N = (M \triangle N) \cup \{k \mapsto v_1 + v_2 \mid k \mapsto v_1 \in M \wedge k \mapsto v_2 \in N\}$	union override plus (for monoidal values)

Figure 1: Non-standard map operators

3 Cryptographic primitives

figure 2 introduces the cryptographic abstractions used in this document. we begin by listing the abstract types, which are meant to represent the corresponding concepts in cryptography. Only the functionality explicitly stated in the figures below is assumed within the scope of this paper. That is, their exact implementation remains open to interpretation, and we do not rely on any additional properties derived from the study or implementation of public key cryptography outside this work. The types and rules we give here are needed in order to guarantee certain security properties of the delegation process, which we discuss later.

The cryptographic concepts required for the formal definition of witnessing include public-private key pairs, one-way functions and signatures. The constraint we introduce states that a signature of some data signed with a (private) key is only correct whenever we can verify it using the corresponding public key.

Besides basic cryptographic abstractions, we also make use of some abstract data storage properties in this document in order to build necessary definitions and make judgement calls about them.

Abstract data types in this paper are essentially placeholders with names indicating the data types they are meant to represent in an implementation. Derived types are made up of data structures (i.e. products, lists, finite maps, etc.) built from abstract types. The underlying structure of a data type is implementation-dependent, and furthermore, the way the data is stored on physical storage can vary as well.

Serialization is a physical manifestation of data on a given storage device. In this document, the properties and rules we state involving serialization are assumed to hold true independently of the storage medium and style of data organization chosen for an implementation.

<i>Abstract types</i>		
$sk \in SKey$		private signing key
$vk \in VKey$		public verifying key
$hk \in HashKey$		hash of a key
$\sigma \in Sig$		signature
$d \in Data$		data
<i>Derived types</i>		
$(sk, vk) \in KeyPair$		signing-verifying key pairs
<i>Abstract functions</i>		
$hashKey \in VKey \rightarrow HashKey$		hashKey function
$verify \in \mathbb{P} (VKey \times Data \times Sig)$		verification relation
$sign \in SKey \rightarrow Data \rightarrow Sig$		signing function
<i>Constraints</i>		
$\forall (sk, vk) \in KeyPair, d \in Data, \sigma \in Sig \cdot sign\ sk\ d = \sigma \implies verify\ vk\ d\ \sigma$		
<i>Notation for serialized and verified data</i>		
$\llbracket x \rrbracket$		serialised representation of x
$\mathcal{V}_{vk} \llbracket d \rrbracket_{\sigma} = verify\ vk\ d\ \sigma$		shorthand notation for verify

Figure 2: Cryptographic definitions

4 Addresses

Addresses are described in section 4.2 of the delegation design document [Kant et al. \(2018\)](#). The types needed for the addresses are defined in Figure 3. There are three types of UTxO addresses:

- Base addresses, $\text{Addr}_{\text{base}}$, containing the hash of a payment key and the hash of a staking key,
- Pointer addresses, Addr_{ptr} , containing the hash of a payment key and a pointer to a stake key registration certificate,
- Enterprise addresses, $\text{Addr}_{\text{enterprise}}$, containing only the hash of a payment key (and which have no staking rights).

Together, these three address types make up the Addr type, which will be used in transaction outputs in Section 7.

Note that for security, privacy, and usability reasons, the staking (delegating) key pair associated with an address should be different from its payment key pair. Before the stake key is registered and delegated to an existing stake pool, the payment key can be used for transactions, though it will not receive rewards from staking. Once a stake key is registered, the shorter pointer addresses can be generated.

Finally, there is an account style address Addr_{rwd} which contains the hash of a staking key. These account addresses will only be used for receiving rewards from the proof of stake leader election. Appendix A of [Kant et al. \(2018\)](#) explains this design choice. The mechanism for transferring rewards from these accounts will be explained in Section 7, and follows [Zahentferner \(2018\)](#).

<i>Abstract types</i>			
	$slot \in \text{Slot}$	absolute slot	
	$ix \in \text{Ix}$	index	
<i>Derived types</i>			
$(s, t, c) \in \text{Ptr}$	=	$\text{Slot} \times \text{Ix} \times \text{Ix}$	certificate pointer
$addr \in \text{Addr}_{\text{base}}$	=	$\text{HashKey}_{\text{pay}} \times \text{HashKey}_{\text{stake}}$	base address
$addr \in \text{Addr}_{\text{ptr}}$	=	$\text{HashKey}_{\text{pay}} \times \text{Ptr}$	pointer address
$addr \in \text{Addr}_{\text{enterprise}}$	=	$\text{HashKey}_{\text{pay}}$	enterprise address
$addr \in \text{Addr}$	=	$\text{Addr}_{\text{base}} \uplus \text{Addr}_{\text{ptr}} \uplus \text{Addr}_{\text{enterprise}}$	output address
$acct \in \text{Addr}_{\text{rwd}}$	=	$\text{HashKey}_{\text{stake}}$	reward account
<i>Accessor Functions</i>			
$\text{paymentHK} \in \text{Addr} \rightarrow \text{HashKey}_{\text{pay}}$			hash of payment key from addr
$\text{stakeHK}_b \in \text{Addr}_{\text{base}} \rightarrow \text{HashKey}_{\text{stake}}$			hash of stake key from base addr
$\text{stakeHK}_r \in \text{Addr}_{\text{rwd}} \rightarrow \text{HashKey}_{\text{stake}}$			hash of stake key from reward account
$\text{addrPtr} \in \text{Addr}_{\text{ptr}} \rightarrow \text{Ptr}$			pointer from pointer addr
<i>Constructor Functions</i>			
$\text{addr}_{\text{rwd}} \in \text{HashKey}_{\text{stake}} \rightarrow \text{Addr}_{\text{rwd}}$			construct a reward account

Figure 3: Definitions used in Addresses

5 Protocol Parameters

The rules for the ledger depend on several parameters and are contained in the PParams type defined in Figure 4.

The type Coin is defined as an alias for the integers. Negative values will not be allowed in UTXO outputs or reward accounts, and \mathbb{Z} is only chosen over \mathbb{N} for its additive inverses.

The minfee function calculates the minimum fee that must be paid by a transaction. This value depends on the protocol parameters and the size of the transaction.

Two time related types are introduced, Epoch and Duration. A Duration is the difference between two slots, as given by $-_s$.

One global constant is defined, SlotsPerEpoch, representing the number of slots in an epoch. As a global constant, this value can only be changed by updating the software.

Lastly, there are two functions, epoch and firstSlot for converting between epochs and slots.

Abstract types

$fparams \in \text{FeeParams}$ min fee parameters
 $dur \in \text{Duration}$ difference between slots
 $epoch \in \text{Epoch}$ epoch

Derived types

$coin \in \text{Coin} = \mathbb{Z}$ unit of value

Protocol Parameters

$\text{PParams} = \left(\begin{array}{ll} fparams \in \text{FeeParams} & \text{min fee parameters} \\ keyDeposit \in \text{Coin} & \text{stake key deposit} \\ keyMinRefund \in [0, 1] & \text{stake key min refund} \\ keyDecayRate \in [0, \infty) & \text{stake key decay rate} \\ poolDeposit \in \text{Coin} & \text{stake pool deposit} \\ poolMinRefund \in [0, 1] & \text{stake pool min refund} \\ poolDecayRate \in [0, \infty) & \text{stake pool decay rate} \\ movingAvgWeight \in [0, 1] & \text{moving average weight} \\ movingAvgExp \in (0, \infty) & \text{moving average exponent} \\ E_{max} \in \text{Epoch} & \text{epoch bound on pool retirement} \\ n_{opt} \in \mathbb{N}^+ & \text{desired number of pools} \\ a_0 \in (0, \infty) & \text{pool influence} \\ \tau \in [0, 1] & \text{treasury expansion} \\ \rho \in [0, 1] & \text{monetary expansion} \end{array} \right)$

Accessor Functions

$fparams, keyDeposit, keyMinRefund, keyDecayRate, poolDeposit, poolMinRefund, poolDecayRate,$
 $movingAvgWeight, movingAvgExp, emax, nopt, influence, \tau, \rho$

Abstract Functions

$\text{minfee} \in \text{PParams} \rightarrow \text{Tx} \rightarrow \text{Coin}$ minimum fee calculation
 $(-_s) \in \text{Slot} \rightarrow \text{Slot} \rightarrow \text{Duration}$ duration between slots

Global Constants

$\text{SlotsPerEpoch} \in \mathbb{N}$ slots per epoch

Derived Functions

$\text{epoch} \in \text{Slot} \rightarrow \text{Epoch}$ epoch of a slot
 $\text{epoch slot} = \text{slot} \text{ div } \text{SlotsPerEpoch}$

$\text{firstSlot} \in \text{Epoch} \rightarrow \text{Slot}$ first slot of an epoch
 $\text{firstSlot } e = e \cdot \text{SlotsPerEpoch}$

Figure 4: Definitions used in Protocol Parameters

6 Transactions

Transactions are defined in Figure 5. A transaction body, `TxBody`, is made up of six pieces:

- A set of transaction inputs. The `TxIn` derived type identifies an output from a previous transaction. It consists of a transaction id and an index to uniquely identify the output.
- An indexed collection of transaction outputs. The `TxOut` type is an address paired with a coin value.
- A list of certificates, which will be explained in detail in Section 8.
- A transaction fee. This value will be added to the fee pot and eventually handed out as stake rewards.
- A time to live. A transaction will be deemed invalid if processed after this slot.
- A mapping of reward account withdrawals. The type `Wdrl` is a finite map that maps a reward address to the coin value to be withdrawn. The coin value must be equal to the full value contained in the account. Explicitly stating these values ensures that error messages can be precise about why a transaction is invalid.

A transaction, `Tx`, is a transaction body together with:

- A collection of witnesses, represented as a finite map from payment verification keys to signatures.

Additionally, the `UTxO` type will be used by the ledger state to store all the unspent transaction outputs. It is a finite map from transaction inputs to transaction outputs that are available to be spent.

Finally, `txid` computes the transaction id of a given transaction. This function must produce a unique id for each unique transaction.

<i>Abstract types</i>			
		$txid \in TxId$	transaction id
<i>Derived types</i>			
$(txid, ix) \in TxIn$	=	$TxId \times Ix$	transaction input
$(addr, c) \in TxOut$	=	$Addr \times Coin$	transaction output
$utxo \in UTxO$	=	$TxIn \mapsto TxOut$	unspent tx outputs
$wdr1 \in Wdr1$	=	$Addr_{rwd} \mapsto Coin$	reward withdrawal
<i>Transaction Types</i>			
$txbody \in TxBody$	=	$\mathbb{P} TxIn \times (Ix \mapsto TxOut) \times DCert^* \times Coin \times Slot \times Wdr1$	
$tx \in Tx$	=	$TxBody \times (VKey \mapsto Sig)$	
<i>Accessor Functions</i>			
$txins \in Tx \rightarrow \mathbb{P} TxIn$			transaction inputs
$txouts \in Tx \rightarrow (Ix \mapsto TxOut)$			transaction outputs
$txcerts \in Tx \rightarrow DCert^*$			delegation certificates
$txfee \in Tx \rightarrow Coin$			transaction fee
$txttl \in Tx \rightarrow Slot$			time to live
$txwdr1s \in Tx \rightarrow Wdr1$			withdrawals
$txbody \in Tx \rightarrow TxBody$			transaction body
$txwits \in Tx \rightarrow (VKey \mapsto Sig)$			witnesses
<i>Abstract Functions</i>			
		$txid \in Tx \rightarrow TxId$	compute transaction id

Figure 5: Definitions used in the UTxO transition system

7 UTxO

A key constraint that must always be satisfied as a result and precondition of a valid ledger state transition is called the *general accounting property*, or the *preservation of value* condition. Every piece of software that is a part of the implementation of the Cardano cryptocurrency must function in such a way as to not result in a violation of this rule. If this condition is not satisfied, it is an indicator of incorrect accounting, potentially due to malicious disruption or a bug.

The preservation of value is expressed as an equality that uses values in the ledger state and the environment, as well as the values in the body of the signal transaction. We have defined the rules of the delegation protocol in a way that should consistently satisfy the preservation of value. In the future, we hope to give a formally-verified proof that every *valid* ledger state satisfies this property.

In this section, we discuss the relevant accounting that needs to be done as a result of processing a transaction, i.e. the deposits for all certificates, transaction fees, transaction withdrawals, and refunds for individual deregistration, so that we may keep track of whether the preservation of value is satisfied. Stake pool retirement refunds are not triggered by a transaction (but rather, happen at the epoch boundary), and are therefore not considered in our state change rules invoked due to a signal transaction.

Note, that when a transaction is issued by a wallet to be applied to the ledger state (i.e. processed), we define the rules in this section in such a way that it is impossible to apply only some parts of a transaction (e.g. only certain certificates). Every part of the transaction must be valid and it must be live, otherwise it is ignored entirely. It is the wallet's responsibility to inform the user that a transaction failed to be processed.

7.1 UTxO Transitions

Figure 6 defines functions needed for the UTxO transition system.

- The function `outs` creates unspent outputs generated by a transaction, so that they can be added to the ledger state. For each output in the transaction, `outs` maps the transaction id and output index to the output.
- The `ubalance` function calculates sum total of all the coin in a given UTxO.
- The `wbalance` function calculates sum total of all the withdrawals in a transaction.
- The calculation `consumed` gives the value consumed by the transaction `tx` in the context of the protocol parameters, the current UTxO on the ledger, and the registered stake keys. This calculation is a sum of all coin in the inputs of `tx`, reward withdrawals, and stake key deposit refunds. Some of the definitions used in this function will be defined in Section 7.2. In particular, `keyRefunds` is defined in Figure 9 and `StakeKeys` is defined in Figure 14.
- The calculation `produced` gives the value produced by the transaction `tx` in the context of the protocol parameters and the registered stake pools. This calculation is a sum of all coin in the outputs of `tx`, the transaction fee, and all needed deposits. Some of the definitions used in this function will be defined in Section 7.2. In particular, `deposits` is defined in Figure 9 and `StakePools` is defined in Figure 14.

The preservation of value property holds for a transaction, for a given ledger state, exactly when the results of `consumed` equal the results of `produced`. Moreover, when the property holds, value is only moved between transaction outputs, the reward accounts, the fee pot, and the deposit pot.

Note that the `consumed` function takes the registered stake pools (`stools`) as a parameter only in order to determine which pool registration certificates are new (and thus require a deposit) and which ones are updates. Registration will be discussed more in Section 8.

$\text{outs} \in \text{Tx} \rightarrow \text{UTxO}$ $\text{outs } tx = \{(txid \ tx, ix) \mapsto txout \mid ix \mapsto txout \in txouts \ tx\}$	tx outputs as UTxO
$\text{ubalance} \in \text{UTxO} \rightarrow \text{Coin}$ $\text{ubalance } utxo = \sum_{(_ \mapsto (_, c)) \in utxo} c$	UTxO balance
$\text{wbalance} \in \text{Wdrl} \rightarrow \text{Coin}$ $\text{wbalance } ws = \sum_{_ \mapsto c \in ws} c$	withdrawal balance
$\text{consumed} \in \text{PParams} \rightarrow \text{UTxO} \rightarrow \text{StakeKeys} \rightarrow \text{Wdrl} \rightarrow \text{Tx} \rightarrow \text{Coin}$ $\text{consumed } pp \ utxo \ stkeys \ rewards \ tx =$ $\text{ubalance } (txins \ tx \triangleleft \ utxo) + \text{wbalance } (txwdrls \ tx)$ $+ \text{keyRefunds } pp \ stkeys \ tx$	value consumed
$\text{produced} \in \text{PParams} \rightarrow \text{StakePools} \rightarrow \text{Tx} \rightarrow \text{Coin}$ $\text{produced } pp \ stpools \ tx =$ $\text{ubalance } (\text{outs } tx) + \text{txfee } tx + \text{deposits } pp \ stpools \ (dcerts \ tx)$	value produced

Figure 6: Functions used in UTxO rules

The types for the UTxO transition are given in Figure 7. The environment, $UTxOEnv$, consists of:

- The current slot.
- The protocol parameters.
- The registered stake keys (which will be explained in Section 8, Figure 14).
- The registered stake pools (also explained in Section 8, Figure 14).

The current slot and the registrations are need for the refund calculations described in Section 7.2. The state needed for the UTxO transition, $UTxOState$, consists of:

- The current UTxO.
- The deposit pot.
- The fee pot.

The signal for the UTxO transition is a transaction.

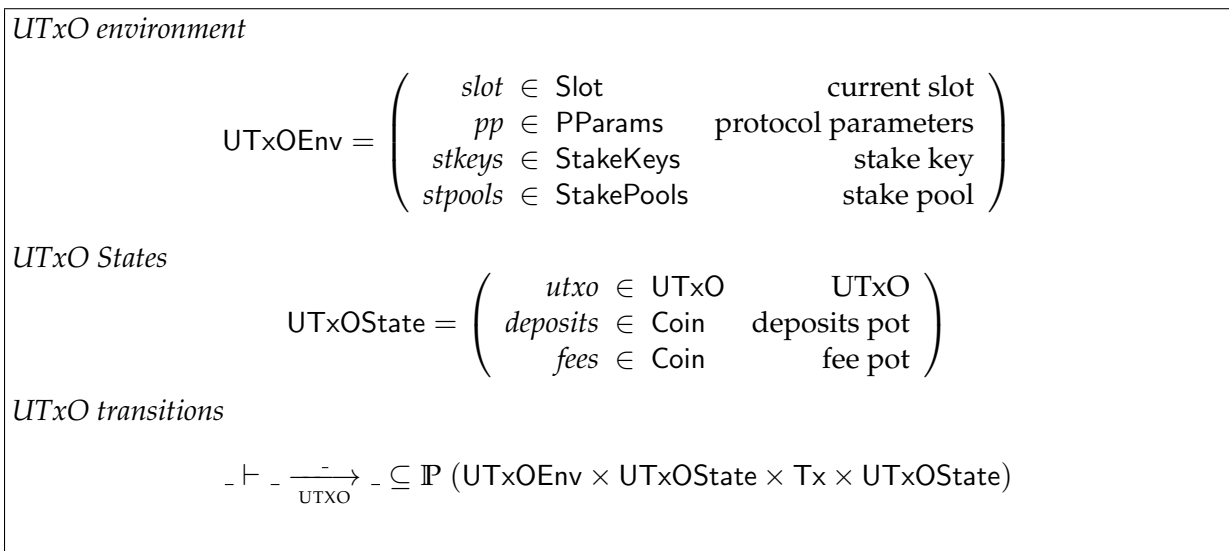


Figure 7: UTxO transition-system types

The UTxO transition system is given in Figure 8. Rule 1 specifies the conditions under which a transaction can be applied to a particular $UTxOState$ in environment $UTxOEnv$:

The transition contains the following predicates:

- The transaction is live (the current slot is less than its time to live).
- The transaction has at least one input. The global uniqueness of transaction inputs prevents replay attacks. By requiring that all transactions spend at least one input, the entire transaction is safe from such attacks. A delegation certificate by itself, for example, does not have this property.
- The fee paid by the transaction has to be greater than or equal to the minimum fee, which is based on the size of the transaction. A user or wallet might choose to create a fee larger than necessary in exchange for a faster processing time.
- Each input spent in the transaction must be in the set of unspent outputs.

- The *preservation of value* property must hold. In other words, the amount of value produced by the transaction must be the same as the amount consumed.
- The coin value of each new output must be non-negative.

If all the predicates are satisfied, the state is updated as follows:

- Update the UTxO:
 - Remove from the UTxO all the $(txin, txout)$ pairs associated with the *txins*'s in the *inputs* list of *tx*.
 - Add all the *outputs* of *tx* to the UTxO, associated with the txid *tx*
- Add all new deposits to the deposit pot and subtract all refunds.
- Add the transaction fee to the fee pot. Additionally, for any refund returned by this transaction, add the amount of the original deposit which has decayed to the fee pot. The amount decayed will depend on the time to live of the transaction and will be explained further in Section 7.2.

The accounting for the reward withdrawals is not done in this transition system. The rewards are tracked with the delegation state and will be removed in the final delegation transition, see 11.

Note here that output entries for both the deposit refunds and the rewards withdrawals must be included in the body of the transaction carrying the deregistration certificates (requesting these refunds) and the reward requests. It is the job of the wallet to calculate the value of these refunds and withdrawals, and generate the correct outputs to include in the outputs list of *tx* such that applying this transaction results in a valid ledger update adding correct amounts of coin to the right addresses.

The approach of including refunds and rewards directly in the *outputs* gives great flexibility to the management of the coin value obtained from these accounts, i.e. it can be directed to any address. However, it means there is no direct link between the *wdrls* requests (similarly, the key deregistration certificate addresses and refund amounts) and the *outputs*. We verify that the included outputs are correct and authorized through the preservation of value condition and witnessing the transaction. The combination of the preservation of value and witnessing, described in Section 7.3, assures that the ledger state is updated correctly.

The main difference, however, in how rewards and refunds work is that refunds come from a *deposits* pool, which is a single coin value indicating the total decayed amount of all the deposits ever made, while rewards come from individual accounts where a reward is accumulated to a specific address.

Note that the *refunded* and *decayed* values added together give what the full, non-decayed refund for all the key deregistration certificates in *tx* would be, and this total value is always removed from the *deposits* amount on the ledger. The *refunded* amount is returned to the certificate author, and the *decayed* amount is transferred over to *fees* (this allows the ledger to adhere to the preservation of value).

Note also that the reason only the decayed value of requested refunds from *this epoch* is transferred to fees is that at the epoch boundary, the total decayed value for the whole epoch for both the individual and pool deposits is transferred into the fees (independent of refund requests).

$$\begin{array}{c}
\text{txttl } tx \geq \text{slot} \quad \text{txins } tx \neq \emptyset \quad \text{minfee } pp \text{ tx} \leq \text{txfee } tx \quad \text{txins } tx \subseteq \text{dom } utxo \\
\text{consumed } pp \text{ utxo } \text{stkeys } \text{rewards } tx = \text{produced } pp \text{ stpools } tx \\
\\
\forall (- \mapsto (-, c)) \in \text{txouts } tx, c \geq 0 \\
\\
\text{refunded} = \text{keyRefunds } pp \text{ stkeys } tx \\
\text{decayed} = \text{decayedTx } pp \text{ stkeys } tx \\
\text{depositChange} = (\text{deposits } pp \text{ stpools } \text{dcerts } tx) - (\text{refunded} + \text{decayed}) \\
\hline
\text{UTxO-inductive} \\
\text{slot} \\
pp \\
stkeys \\
stools \\
\vdash \left(\begin{array}{c} utxo \\ deposits \\ fees \end{array} \right) \xrightarrow[\text{UTxO}]{tx} \left(\begin{array}{c} (\text{txins } tx \not\subseteq utxo) \cup \text{outs } tx \\ deposits + \text{depositChange} \\ fees + \text{txfee } tx + \text{decayed} \end{array} \right) \\
(1)
\end{array}$$

Figure 8: UTxO inference rules

7.2 Deposits and Refunds

Deposits are described in appendix B.2 of the delegation design document [Kant et al. \(2018\)](#). These deposit functions were used above in the UTxO transition, [7.1](#). Deposits are used for stake key registration certificate and pool registration certificates, which will be explained in [Section 8](#). In particular, the function `cwitness`, which gets the certificate witness from a certificate, will be defined later. [Figure 9](#) defines the deposit and refund functions.

- The function `deposits` returns the total deposits that have to be made by a transaction. This calculation is based on the protocol parameters. Specifically, for a given transaction, it sums up the values of the stake key deposits and the stake pool deposits. Those certificates which are updates of stake pool parameters of already registered pool keys should not (and are, in fact, not allowed to) make a deposit.
- The function `refund` calculates the deposit refund with an exponential decay.
- The function `keyRefund`, calculates the refund for an individual stake key registration deposit, based on the slot when it was created and the slot passed to the function. The creation slot should always exist in the map `stkeys` passed to the function, and this would be a good property to prove about the transition system.
- The function `keyRefunds`, in turn, uses `keyRefund` to calculate the total value to be refunded to all individual key deregistration certificate authors in a transaction.

It is important to note here that instead of the *current* slot number, the time to live of *tx* is passed to the `certRefunds` function within the summation in `keyRefunds`. The reason for this is that the refunds for any key deregistration certificates are, in fact, included in the *tx* itself — meaning that the coin value of the refund must be explicitly specified in the outputs of the transaction. So, the value of the included refund must be calculated before this transaction is ever processed, and be the same *no matter when the tx is actually processed* in order to allow the system to continue to satisfy the general accounting property.

It is impossible to predict the exact slot number in which *tx* will be processed, but it will be some time before slot number `txttl tx`. So, this is the slot number value used in both the calculation to generate the refund coin value in the outputs of *tx* and in the general accounting property equation.

Note also that `keyRefunds` calculates the total individual refunds for a transaction based on *current* protocol parameters. This means that any deposits made prior to a change will be different from their corresponding (decayed) refunds in the case of key deregistration after a change in protocol parameters. Constants may only change at the epoch boundary, and ensuring there are always sufficient funds for all refunds in the *deposits* pool is part of the protocol constant change transition, described in [Section 10](#).

The protocol parameters are not expected to change often, and using the current ones for the calculation is a deliberate simplification choice, which does not introduce any inconsistencies into the system rules or properties. In particular, the general accounting property is not violated.

[Figure 10](#) defines the decays functions.

- The function `decayedKey` calculates how much of a stake key deposit has decayed since the last epoch. Again, this is done using the time to live of the transaction (and not the current slot, as explained above). At the epoch boundaries, decayed portions of deposits are moved to the reward pot, so between epochs we need only account for what has decayed since the last epoch. The value is calculated by subtracting the refund calculation based at the epoch boundary from the refund calculation based at the time to live of the transaction.

- The function `decayedTx` calculates the total decayed deposits associated with all the refunds in a given transaction. This function was used earlier in the UTxO transition in Figure 8.

Recall that the stake pool retirement refunds are issued not when a certificate scheduling the retirement is processed, but at the epoch boundary for which the retirement is scheduled. The decayed value over the full previous epoch is also accounted for at the boundary change. For details of this accounting, see Section [Section 10](#).

`deposits` \in `PParams` \rightarrow `StakePools` \rightarrow `DCert*` \rightarrow `Coin` total deposits for transaction

`deposits` pp $stools$ $certs =$

$$\sum_{c \in certs \cap DCert_{regkey}} (\text{keyDeposit } pp) + \sum_{\substack{c \in certs \cap DCert_{regpool} \\ (cwitness\ c) \notin stools}} (\text{poolDeposit } pp)$$

`refund` \in `Coin` \rightarrow $[0, 1]$ \rightarrow $(0, \infty)$ \rightarrow `Duration` \rightarrow `Coin`

refund calculation

$$\text{refund } d_{val} \ d_{min} \ \lambda \ \delta = \left[d_{val} \cdot \left(d_{min} + (1 - d_{min}) \cdot e^{-\lambda \cdot \delta} \right) \right]$$

`keyRefund` \in `Coin` \rightarrow $[0, 1]$ \rightarrow $(0, \infty)$ \rightarrow

`StakeKeys` \rightarrow `Slot` \rightarrow `DCert_{deregkey}` \rightarrow `Coin`

key refund for a certificate

`keyRefund` $d_{val} \ d_{min} \ \lambda \ stkeys \ slot \ c =$

$$\begin{cases} 0 & \text{if } cwitness\ c \notin \text{dom } stkeys \\ \text{refund } d_{val} \ d_{min} \ \lambda \ \delta & \text{otherwise} \end{cases}$$

where $\delta = slot -_s (stkeys (cwitness\ c))$

`keyRefunds` \in `PParams` \rightarrow `StakeKeys` \rightarrow `Tx` \rightarrow `Coin`

key refunds for a transaction

`keyRefunds` $pp \ stkeys \ tx =$

$$\sum_{\substack{c \in dcerts\ tx \\ c \in DCert_{deregkey}}} \text{keyRefund } d_{val} \ d_{min} \ \lambda \ stkeys (txttl\ tx) \ c$$

where

$$d_{val} = \text{keyDeposit } pp$$

$$d_{min} = \text{keyMinRefund } pp$$

$$\lambda = \text{keyDecayRate } pp$$

Figure 9: Functions used in Deposits - Refunds

$\text{decayedKey} \in \text{PParams} \rightarrow \text{StakeKeys} \rightarrow \text{Slot} \rightarrow \text{DCert}_{\text{deregkey}} \rightarrow \text{Coin}$ decayed since epoch
 $\text{decayedKey } pp \text{ stkeys } cslot \ c =$

$$\begin{cases} 0 & \text{if } cwitness \ c \notin \text{dom } stkeys \\ epochRefund - currentRefund & \text{otherwise} \end{cases}$$

where

$$\begin{aligned} created &= stkeys \ (cwitness \ c) \\ start &= \max \ (\text{firstSlot } epoch \ cslot) \ created \\ epochRefund &= \text{keyRefund } d_{val} \ d_{min} \ \lambda \ stkeys \ start \ c \\ currentRefund &= \text{keyRefund } d_{val} \ d_{min} \ \lambda \ stkeys \ cslot \ c \\ d_{val} &= \text{keyDeposit } pp \\ d_{min} &= \text{keyMinRefund } pp \\ \lambda &= \text{keyDecayRate } pp \end{aligned}$$

$\text{decayedTx} \in \text{PParams} \rightarrow \text{StakeKeys} \rightarrow \text{Tx} \rightarrow \text{Coin}$ decayed deposit portions
 $\text{decayedTx } pp \text{ stkeys } tx =$

$$\sum_{\substack{c \in \text{dcerts } tx \\ c \in \text{DCert}_{\text{deregkey}}}} \text{decayedKey } pp \text{ stkeys } (txttl \ tx) \ c$$

Figure 10: Functions used in Deposits - Decay

7.3 Witnesses

The purpose of witnessing is make sure that the intended action is authorized by the holder of the signing key, providing replay protection as a consequence. Replay prevention is an inherent property of UTxO type accounting since transaction IDs are unique, and we require all transaction to consume at least one input.

A transaction is witnessed by a signature and a verification key corresponding to this signature. The witnesses, together with the transaction body, form a full transaction. Every witness in a transaction signs the transaction body. Moreover, the witnesses are represented as finite maps from verification keys to signatures, so that any key that is required to sign a transaction only provides a single witness. This means that, for example, transaction which includes a delegation certificate and a reward withdrawal corresponding to the same stake key still only includes one signature.

Figure 11 defines the function which gathers all the (hashes of) verification keys needed to witness a given transaction. This consists of:

- payment keys for outputs being spent
- stake keys for reward withdrawals
- stake keys for delegation certificates (all five types)
- stake keys for the pool owners in a pool registration certificate

$$\begin{aligned}
 & \text{witsNeeded} \in \text{UTxO} \rightarrow \text{Tx} \rightarrow \mathbb{P} \text{HashKey} && \text{hashkeys for needed witnesses} \\
 & \text{witsNeeded } utxo \ tx = \\
 & \quad \{ \text{paymentHK } a \mid i \mapsto (a, _) \in utxo, i \in \text{txins } tx \} \cup \\
 & \quad \{ \text{stakeHK}_r \ a \mid a \mapsto _ \in \text{txwdrls } tx \} \cup \\
 & \quad \{ \text{cwitness } c \mid c \in \text{txcerts } tx \} \cup \\
 & \quad \bigcup_{\substack{c \in \text{txcerts } tx \\ c \in \text{DCert}_{\text{regpool}}}} \text{poolOwners } c
 \end{aligned}$$

Figure 11: Functions used in witness rule

The UTxOW transition system adds witnessing to the previous UTxO transition system. Figure 12 defines the type for this transition.

Figure 13 defines UTxOW transition. It has two predicates:

- Every signature in the transaction is a valid signature of the transaction body.
- The set of (hashes of) verification keys given by the transaction is exactly the set of needed (hashes of) verification keys.

If the predicates are satisfied, the state is transitioned by the UTxO transition rule.

$$\begin{aligned}
 & \text{UTxO with witness transitions} \\
 & \quad _ \vdash _ \xrightarrow{\text{UTxOW}} _ \subseteq \mathbb{P} (\text{UTxOEnv} \times \text{UTxOState} \times \text{Tx} \times \text{UTxOState})
 \end{aligned}$$

Figure 12: UTxO with witness transition-system types

$$\begin{array}{c}
(utxo, _, _) = utxoSt \\
\forall vk \mapsto \sigma \in txwits\ tx, \mathcal{V}_{vk}[\![txbody\ tx]\!]_{\sigma} \\
witsNeeded\ utxo\ tx = \{hashKey\ vk \mid vk \in dom(txwits\ tx)\} \\
utxoEnv \vdash utxoSt \xrightarrow[UTXO]{tx} utxoSt' \\
\text{UTxO-wit} \text{-----} \\
utxoEnv \vdash utxoSt \xrightarrow[UTXOW]{tx} \mathbf{utxoSt'}
\end{array} \tag{2}$$

Figure 13: UTxO with witnesses inference rules

8 Delegation

We briefly describe the motivation and context for delegation. The full context is contained in [Kant et al. \(2018\)](#).

Stake is said to be *active* in the blockchain protocol when it is eligible for participation in the leader election. In order for stake to become active, the associated verification stake key must be registered and its staking rights must be delegated to an active stake pool. Individuals who wish to participate in the protocol can register themselves as a stake pool.

Stake keys are registered (deregistered) through the use of registration (deregistration) certificates. Registered stake keys are delegated through the use of delegation certificates. Finally, stake pools are registered (retired) through the use of registration (retirement) certificates.

Stake pool retirement is handled a bit differently than stake key deregistration. Stake keys are considered inactive as soon as a deregistration certificate is applied to the ledger state. Stake pool retirement certificates, however, specify the epoch in which it will retire.

Delegation requires the following to be tracked by the ledger state: the registered stake keys, the delegation map from registered stake keys to stake pools, pointers associated with stake keys, the registered stake pools, and upcoming stake pool retirements. Additionally, the blockchain protocol rewards eligible stake, and so we must also include a mapping from active stake keys to rewards.

8.1 Delegation Definitions

In [Figure 14](#) we give the delegation primitives. Here we introduce the following primitive datatypes used in delegation:

- $\text{DCert}_{\text{regkey}}$: a stake key registration certificate.
- $\text{DCert}_{\text{deregkey}}$: a stake key de-registration certificate.
- $\text{DCert}_{\text{delegate}}$: a stake key delegation certificate.
- $\text{DCert}_{\text{regpool}}$: a stake pool registration certificate.
- $\text{DCert}_{\text{retirepool}}$: a stake key retirement certificate.
- DCert : any one of of the five certificate types above.

The following derived types are introduced:

- StakeKeys represents registered stake keys, and is represented by a finite map from hashkeys to slot when it was registered.
- StakePools represents registered stake pools, and has the same type as StakeKeys .
- PoolParam represents the parameters found in a stake pool registration certificate that must be tracked:
 - the pool owners.
 - the pool cost.
 - the pool margin.
 - the pool pledge.
 - the pool reward account.

The idea of pool owners is explained in Section 4.4.4 of [Kant et al. \(2018\)](#). The pool cost and margin indicate how much more of the rewards pool leaders get than the members. The pool pledge is explained in Section 5.1 of [Kant et al. \(2018\)](#). The pool reward account is where all pool rewards go.

Accessor functions for certificates and pool parameters are also defined, but only the *cwitness* accessor function needs explanation. It does the following:

- For a $\text{DCert}_{\text{regkey}}$ certificate, *cwitness* returns the hashkey of the key being registered.
- For a $\text{DCert}_{\text{deregkey}}$ certificate, *cwitness* returns the hashkey of the key being de-registered.
- For a $\text{DCert}_{\text{delegate}}$ certificate, *cwitness* returns the hashkey of the key that is delegating (and not the key to which the stake in being delegated to).
- For a $\text{DCert}_{\text{regpool}}$ certificate, *cwitness* returns the hashkey of the key of the pool operator.
- For a $\text{DCert}_{\text{retirepool}}$ certificate, *cwitness* returns the hashkey of the key of the pool operator.

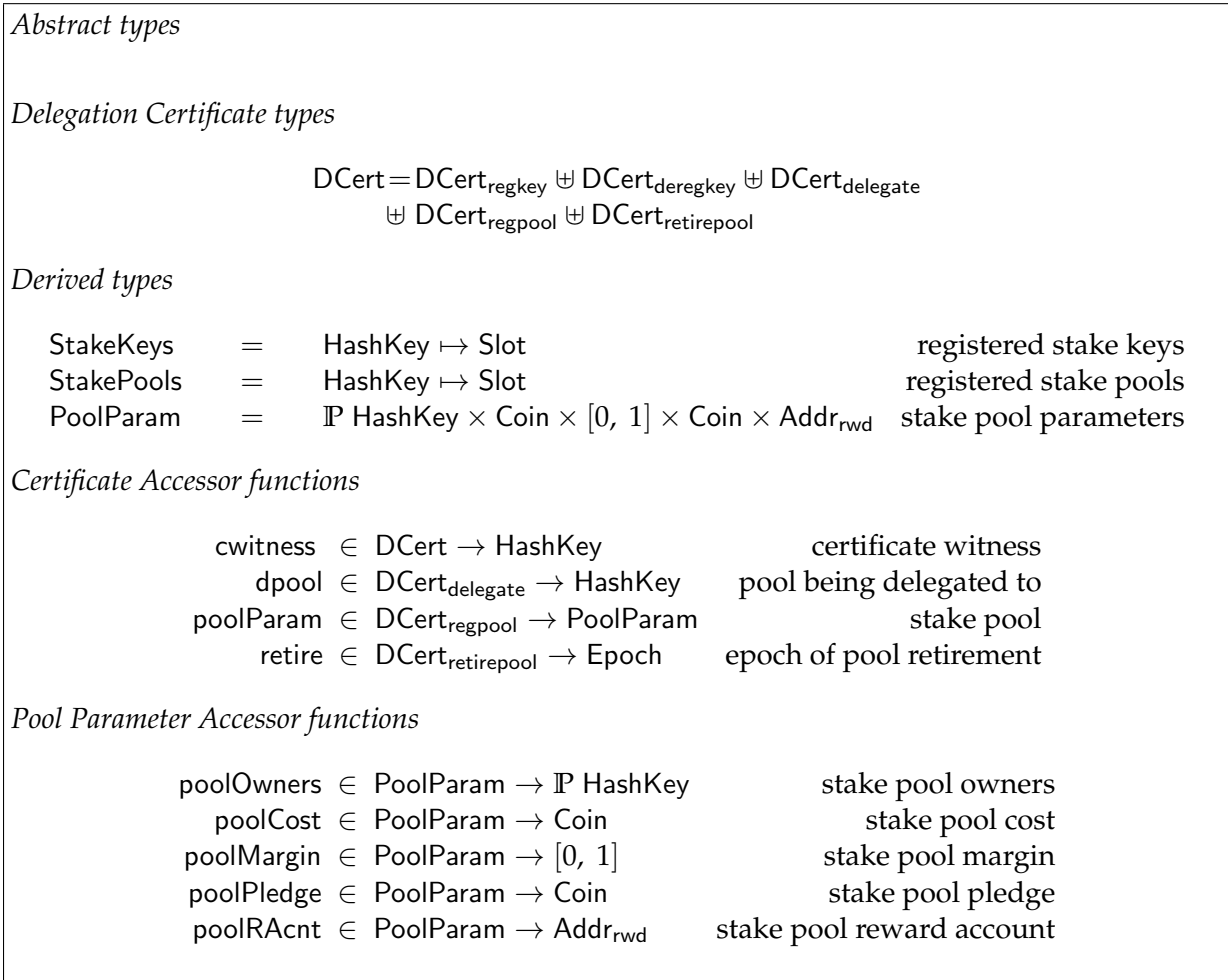


Figure 14: Delegation Definitions

8.2 Delegation Transitions

In [Figure 15](#) we give the delegation and stake pool state transition types. We define two separate parts of the ledger state.

- DState keeps track of the delegation state, consisting of:
 - *stkeys* tracks the registered stake keys. It consists of a finite mapping from hashkeys to the slot of the registration.
 - *rewards* stores the rewards accumulated by stake keys. These are represented by a finite map from reward addresses to the accumulated rewards.
 - *delegations* stores the delegation relation, mapping stake keys to the pool to which is delegates.
 - *ptrs* maps stake keys to the position of the registration certificate in the blockchain. This is needed to lookup the stake hashkey of a pointer address.
- PState keeps track of the stake pool information:
 - *stools* tracks the registered stake pools. It consists of a finite mapping from hashkeys to the slot of the registration.
 - *poolParams* tracks the parameters associated with each stake pool, such as their costs and margin.
 - *retiring* tracks stake pool retirements, using a map from hashkeys to the epoch in which it will retire.
 - *avgs* stores the latest value of the pool's performance moving average. This value quantifies the desirability of delegating to a given pool, based on past performance.

The environment for the state transition for DState contains the current slot number and the index for the current certificate pointer. The environment for the state transition for PState contains the current slot number and the protocol parameters.

Delegation States

$$\begin{aligned}
 \text{DState} &= \left(\begin{array}{ll}
 \text{stkeys} \in \text{StakeKeys} & \text{registered stake keys} \\
 \text{rewards} \in \text{Addr}_{\text{rwd}} \mapsto \text{Coin} & \text{rewards} \\
 \text{delegations} \in \text{HashKey}_{\text{stake}} \mapsto \text{HashKey}_{\text{pool}} & \text{delegations} \\
 \text{ptrs} \in \text{Ptr} \mapsto \text{HashKey} & \text{pointer to hashkey}
 \end{array} \right) \\
 \\
 \text{PState} &= \left(\begin{array}{ll}
 \text{stpools} \in \text{StakePools} & \text{registered pools to creation time} \\
 \text{poolParams} \in \text{HashKey}_{\text{pool}} \mapsto \text{PoolParam} & \text{registered pools to pool parameters} \\
 \text{retiring} \in \text{HashKey}_{\text{pool}} \mapsto \text{Epoch} & \text{retiring stake pools} \\
 \text{avgs} \in \text{HashKey}_{\text{pool}} \mapsto \mathbb{R}^{\geq 0} & \text{performance moving average}
 \end{array} \right)
 \end{aligned}$$

Delegation Environment

$$\text{DEnv} = \left(\begin{array}{ll}
 \text{slot} \in \text{Slot} & \text{slot} \\
 \text{ptr} \in \text{Ptr} & \text{certificate pointer}
 \end{array} \right)$$

Pool Environment

$$\text{PEnv} = \left(\begin{array}{ll}
 \text{slot} \in \text{Slot} & \text{slot} \\
 \text{pp} \in \text{PParams} & \text{protocol parameters}
 \end{array} \right)$$

Delegation Transitions

$$_ \vdash _ \xrightarrow{\text{DELEG}} _ \in \mathbb{P} (\text{DEnv} \times \text{DState} \times \text{DCert} \times \text{DState})$$

$$_ \vdash _ \xrightarrow{\text{POOL}} _ \in \mathbb{P} (\text{PEnv} \times \text{PState} \times \text{DCert} \times \text{PState})$$

Figure 15: Delegation Transitions

8.3 Delegation Rules

The rules for registering and delegating stake keys are given in [Figure 16](#). Note that section 5.2 of [Kant et al. \(2018\)](#) describes how a wallet would help a user choose a stake pool, though these concerns are independent of the ledger rules.

- Stake key registration is handled by [Equation \(3\)](#), since it contains the precondition that the certificate has type $DCert_{regkey}$. All the equations in DELEG and POOL follow this same pattern of matching on certificate type.

There is also a precondition on registration that the hashkey associated with the certificate witness of the certificate is not already found in the current list of stake keys.

Registration causes the following state transformation:

- The key is added to the set of registered stake keys.
 - A reward account is created for this key, with a starting balance of zero.
 - The certificate pointer is mapped to the new stake key.
- Stake key deregistration is handled by [Equation \(4\)](#). There is a precondition that the key has been registered, and that the reward balance is zero. Deregistration causes the following state transformation:
 - The key is removed from the collection of registered keys.
 - The reward account is removed.
 - The key is removed from the delegation relation.
 - The certificate pointer is removed.
- Stake key delegation is handled by [Equation \(5\)](#). There is a precondition that the key has been registered. Delegation causes the following state transformation:
 - The delegation relation is updated so that stake stake key is delegated to the given stake pool. The use of union override here allows us to use the same rule to perform both an initial delegation and an update to an existing delegation.

$$\begin{array}{c}
\text{Deleg-Reg} \\
\hline
\begin{array}{c}
c \in \text{DCert}_{\text{regkey}} \quad hk = \text{cwitness } c \quad hk \notin \text{dom } stkeys \\
slot \\
ptr \vdash \left(\begin{array}{c} stkeys \\ rewards \\ delegations \\ ptrs \end{array} \right) \xrightarrow{\text{DELEG } c} \left(\begin{array}{c} stkeys \cup \{hk \mapsto slot\} \\ rewards \cup \{\text{addr}_{\text{rwd}} hk \mapsto 0\} \\ delegations \\ ptrs \cup \{ptr \mapsto hk\} \end{array} \right)
\end{array}
\end{array} \quad (3)$$

$$\begin{array}{c}
\text{Deleg-Dereg} \\
\hline
\begin{array}{c}
c \in \text{DCert}_{\text{deregkey}} \quad hk = \text{cwitness } c \\
hk \in \text{dom } stkeys \quad hk \mapsto 0 \in \text{rewards} \\
slot \\
ptr \vdash \left(\begin{array}{c} stkeys \\ rewards \\ delegations \\ ptrs \end{array} \right) \xrightarrow{\text{DELEG } c} \left(\begin{array}{c} \{hk\} \triangleleft stkeys \\ \{\text{addr}_{\text{rwd}} hk\} \triangleleft rewards \\ \{hk\} \triangleleft delegations \\ \{ptr\} \triangleleft ptrs \end{array} \right)
\end{array}
\end{array} \quad (4)$$

$$\begin{array}{c}
\text{Deleg-Deleg} \\
\hline
\begin{array}{c}
c \in \text{DCert}_{\text{delegate}} \quad hk = \text{cwitness } c \quad hk \in \text{dom } stkeys \\
slot \\
ptr \vdash \left(\begin{array}{c} stkeys \\ rewards \\ delegations \\ ptrs \end{array} \right) \xrightarrow{\text{DELEG } c} \left(\begin{array}{c} stkeys \\ rewards \\ delegations \sqcup \{hk \mapsto \text{dpool } c\} \\ ptrs \end{array} \right)
\end{array}
\end{array} \quad (5)$$

Figure 16: Delegation Inference Rules

8.4 Stake Pool Rules

The rules for updating the part of the ledger state defining the current stake pools are given in [Figure 17](#). The calculation of stake distribution is described in [Section 10.3](#).

In the pool rules, the stake pool is identified with the hashkey of the pool operator. For each rule, again, we first check that a given certificate c is of the correct type.

- Stake pool registration is handled by [Equation \(6\)](#). It is required that the pool not be currently registered. Registration causes the following state transformation:
 - The key is added to the set of registered stake pools.
 - The pool’s parameters are stored.
- Stake pool parameter updates are handled by [Equation \(6\)](#). This rule, which also matches on the certificate type `DCertRegPool`, is distinguished from [Equation \(6\)](#) by the requirement that the pool be registered. This rule also ends stake pool retirements. Reregistration causes the following state transformation:
 - The pool’s parameters are updated.
 - The pool is removed from the collection of retiring pools.
 - Note that *stools* is **not** updated. The registration creation slot does not change.
- Stake pool retirements are handled by [Equation \(8\)](#). Given a slot number *slot*, the application of this rule requires that the planned retirement epoch e stated in the certificate is in the future, i.e. after *epoch*, the epoch of the current slot number in this context, as well as that it is less than E_{\max} epochs after the current one. It is also required that the pool be registered. Note that imposing the E_{\max} constraint on the system is not strictly necessary. However, forcing stake pools to announce their retirement a shorter time in advance will curb the growth of the *retiring* list in the ledger state.

The pools scheduled for retirement must be removed from the *retiring* state variable at the end of the epoch they are scheduled to retire in. This non-signaled transition (triggered, instead, directly by a change of current slot number in the environment), along with all other transitions that take place at the epoch boundary, are described in [Section 10](#).

Reregistration causes the following state transformation:

- The pool is marked to retire on the given epoch. If it was previously retiring, the retirement epoch is now updated.

$$\begin{array}{c}
\text{Pool-Reg} \\
\hline
\begin{array}{c}
c \in \text{DCert}_{\text{regpool}} \quad hk = \text{cWitness } c \quad hk \notin \text{dom } stpools \\
\text{slot} \\
pp \vdash \left(\begin{array}{c} stpools \\ poolParams \\ retiring \\ avgs \end{array} \right) \xrightarrow{\text{POOL } c} \left(\begin{array}{c} stpools \cup \{hk \mapsto \text{slot}\} \\ poolParams \cup \{hk \mapsto \text{poolParam } c\} \\ retiring \\ avgs \end{array} \right)
\end{array} \\
(6)
\end{array}$$

$$\begin{array}{c}
\text{Pool-reReg} \\
\hline
\begin{array}{c}
c \in \text{DCert}_{\text{regpool}} \quad hk = \text{cWitness } c \quad hk \in \text{dom } stpools \\
\text{slot} \\
pp \vdash \left(\begin{array}{c} stpools \\ poolParams \\ retiring \\ avgs \end{array} \right) \xrightarrow{\text{POOL } c} \left(\begin{array}{c} stpools \\ poolParams \cup \{hk \mapsto \text{poolParam } c\} \\ \{hk\} \cup \cancel{retiring} \\ avgs \end{array} \right)
\end{array} \\
(7)
\end{array}$$

$$\begin{array}{c}
\text{Pool-Retire} \\
\hline
\begin{array}{c}
c \in \text{DCert}_{\text{retirepool}} \quad hk = \text{cWitness } c \quad hk \in \text{dom } stpools \\
e = \text{retire } c \quad \text{epoch} = \text{epoch } slot \quad \text{epoch} < e < \text{epoch} + (\text{emax } pp) \\
\text{slot} \\
pp \vdash \left(\begin{array}{c} stpools \\ poolParams \\ retiring \\ avgs \end{array} \right) \xrightarrow{\text{POOL } c} \left(\begin{array}{c} stpools \\ poolParams \\ retiring \cup \{hk \mapsto e\} \\ avgs \end{array} \right)
\end{array} \\
(8)
\end{array}$$

Figure 17: Pool Inference Rule

8.5 Delegation and Pool Combined Rules

We now combine the delegation and pool transition systems. Figure 18 gives the state, environment, and transition type for the combined transition.

Delegation and Pool Combined Environment

$$\text{DPEnv} = \left(\begin{array}{ll} \text{slot} \in \text{Slot} & \text{slot} \\ \text{ptr} \in \text{Ptr} & \text{certificate pointer} \\ \text{pp} \in \text{PParams} & \text{protocol parameters} \end{array} \right)$$

Delegation and Pool Combined State

$$\text{DPState} = \left(\begin{array}{ll} \text{dstate} \in \text{DState} & \text{delegation state} \\ \text{pstate} \in \text{PState} & \text{pool state} \end{array} \right)$$

Delegation and Pool Combined Transition

$$_ \vdash _ \xrightarrow[\text{DELPL}]{_} _ \in \mathbb{P} (\text{DPEnv} \times \text{DPState} \times \text{DCert} \times \text{DPState})$$

Figure 18: Delegation and Pool Combined Transition Type

Figure 19, gives the rules for the combined transition. Note that for any given certificate, at most one of the two rules (Equation (9) and Equation (10)) will be successful, since the pool certificates are disjoint from the delegation certificates.

Delegation and Pool Combined Rules

$$\text{Delpl-Del} \frac{\text{slot} \quad \text{ptr} \vdash \text{dstate} \xrightarrow[\text{DELEG}]{c} \text{dstate}'}{\left(\begin{array}{c} \text{slot} \\ \text{ptr} \\ \text{pp} \end{array} \right) \vdash \left(\begin{array}{c} \text{dstate} \\ \text{pstate} \end{array} \right) \xrightarrow[\text{DELPL}]{c} \left(\begin{array}{c} \text{dstate}' \\ \text{pstate} \end{array} \right)} \quad (9)$$

$$\text{Delpl-Pool} \frac{\text{slot} \quad \text{pp} \vdash \text{pstate} \xrightarrow[\text{POOL}]{c} \text{pstate}'}{\left(\begin{array}{c} \text{slot} \\ \text{ptr} \\ \text{pp} \end{array} \right) \vdash \left(\begin{array}{c} \text{dstate} \\ \text{pstate} \end{array} \right) \xrightarrow[\text{DELPL}]{c} \left(\begin{array}{c} \text{dstate} \\ \text{pstate}' \end{array} \right)} \quad (10)$$

Figure 19: Delegation and Pool Combined Transition Rules

We now describe a transition system that processes the list of certificates inside a transaction. It is defined recursively from the transition system in Figure 19 above.

Figure 20 defines the types for the delegation certificate sequence transition.

Certificate Sequence Environment

$$\text{DPSEnv} = \left(\begin{array}{cc} \text{slot} \in \text{Slot} & \text{slot} \\ \text{txIx} \in \text{Ix} & \text{transaction index} \\ \text{pp} \in \text{PParams} & \text{protocol parameters} \end{array} \right)$$

$$_ \vdash _ \xrightarrow[\text{DELEGS}]{_} _ \in \mathbb{P} (\text{DPSEnv} \times \text{DPState} \times \text{DCert}^* \times \text{DPState})$$

Figure 20: Delegation sequence transition type

Figure 21 defines the transition system recursively. This definition guarantees that a certificate list (and therefore, the transaction carrying it) cannot be processed unless every certificate in it is valid. For example, if a transaction is carrying a certificate that schedules a pool retirement in a past epoch, the whole transaction will be invalid.

- The base case, when the list is empty, is captured by Equation (11). In the base case, we address one final accounting detail not yet covered by the UTxO transition, namely setting the reward account balance to zero for any account that made a withdrawal. There is therefore a precondition that all withdrawals are correct, where correct means that there is a reward account for each stake key, and that the balance matches that of the reward being withdrawn. The base case triggers the following state transformation:
 - Reward accounts are set to zero for each corresponding withdrawal.
- The inductive case, when the list is non-empty, is captured by Equation (12). It constructs a certificate pointer given the current slot and transaction index, calls DELPL on the next certificate in the list, and inductively calls DELEGS on the rest of the list. The inductive case triggers the following state transformation:

- The delegation and pool states are (inductively) updated by the results of DELEGS, which is then updated according to DELPL.

$$\begin{array}{c}
 \text{Seq-delg-base} \frac{
 \begin{array}{l}
 wdrls = txwdrls \ tx \quad wdrls \subseteq rewards \\
 rewards' = rewards \sqcup \{(w, 0) \mid w \in \text{dom } wdrls\}
 \end{array}
 }{
 \left(\begin{array}{c} slot \\ txIx \\ pp \end{array} \right) \vdash \left(\begin{array}{c} stkeys \\ rewards \\ delegations \\ ptrs \\ stpools \\ poolParams \\ retiring \\ avgs \end{array} \right) \xrightarrow[\text{DELEGS}]{\epsilon} \left(\begin{array}{c} stkeys \\ \mathbf{rewards'} \\ delegations \\ ptrs \\ stpools \\ poolParams \\ retiring \\ avgs \end{array} \right)
 }
 \quad (11)
 \end{array}$$

$$\begin{array}{c}
 c \in \text{DCert}_{\text{delegate}} \Rightarrow \text{dpool } c \in \text{dom } stpools \\
 ptr = (slot, txIx, \text{len } \Gamma - 1) \\
 \left(\begin{array}{c} slot \\ txIx \\ pp \end{array} \right) \vdash dpstate \xrightarrow[\text{DELEGS}]{\Gamma} dpstate' \\
 \left(\begin{array}{c} slot \\ ptr \\ pp \end{array} \right) \vdash dpstate' \xrightarrow[\text{DELPL}]{c} dpstate'' \\
 \text{Seq-delg-ind} \frac{
 \left(\begin{array}{c} slot \\ ptr \\ pp \end{array} \right) \vdash dpstate' \xrightarrow[\text{DELPL}]{c} dpstate''
 }{
 \left(\begin{array}{c} slot \\ txIx \\ pp \end{array} \right) \vdash dpstate \xrightarrow[\text{DELEGS}]{\Gamma; c} \mathbf{dpstate''}
 }
 \quad (12)
 \end{array}$$

Figure 21: Delegation sequence rules

9 Ledger State Transition

The entire state transformation of the ledger state caused by a valid transaction can now be given as the combination of the UTxO transition and the delegation transitions.

Figure 22 defines the types for this transition. The environment for this rule is consists of:

- The current slot.
- The transaction index within the current block.
- The protocol parameters.

The ledger state consists of:

- The UTxO state.
- The delegation and pool states.

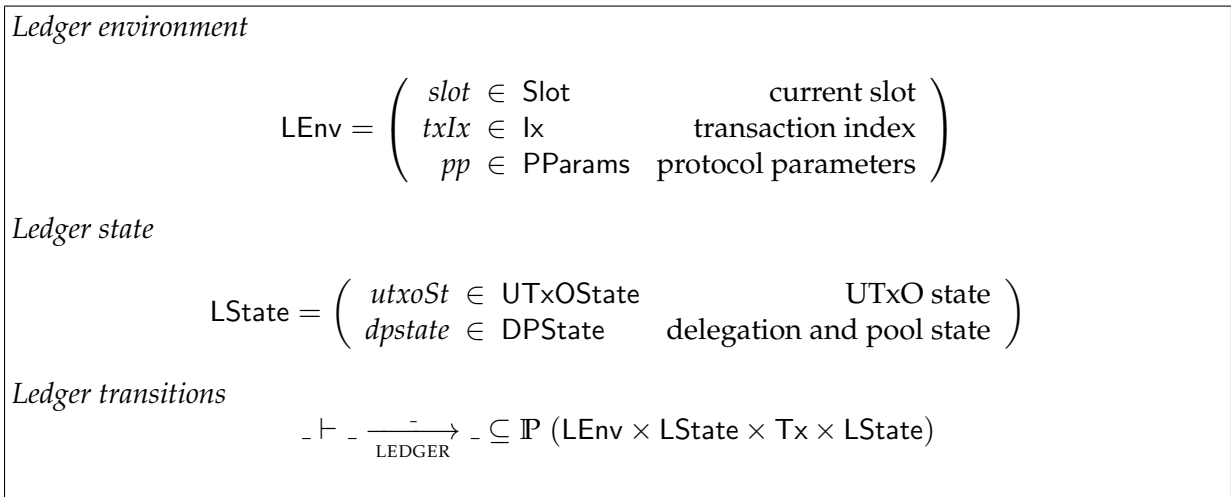


Figure 22: Ledger transition-system types

Figure 22 defines the ledger state transition. It has a single rule, which first calls the UTXOW transition, and then calls the DELEGS transition.

$$\begin{array}{c}
(dstate, pstate) = dpstate \\
(stkeys, -, -, -) = dstate \\
(-, -, stpools, -) = pstate \\
\\
\left(\begin{array}{c} slot \\ pp \\ stkeys \\ stpools \end{array} \right) \vdash utxoSt \xrightarrow[\text{UTXOW}]{tx} utxoSt' \\
\\
\left(\begin{array}{c} slot \\ txIx \\ pp \end{array} \right) \vdash dpstate \xrightarrow[\text{DELEGS}]{tx} dpstate' \\
\hline
\text{ledger} \left(\begin{array}{c} slot \\ txIx \\ pp \end{array} \right) \vdash \left(\begin{array}{c} utxoSt \\ dpstate \end{array} \right) \xrightarrow[\text{LEDGER}]{tx} \left(\begin{array}{c} utxoSt' \\ dpstate' \end{array} \right)
\end{array} \tag{13}$$

Figure 23: Ledger inference rule

10 Rewards and the Epoch Boundary

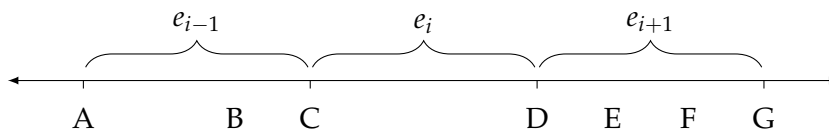
This chapter introduces two main transition systems. Neither transition is triggered by a transaction, and in fact have no signal.

The first one, defined in Section 10.7, involves calculations that occur at the epoch boundary. This includes taking stake distribution snapshots (Sections 10.2 and 10.4), retiring stake pools (Section 10.5), and performing protocol updates (Section 10.6). The second transition, defined in Sections 10.8 and 10.9, distributes the leader election rewards.

10.1 Overview of the Reward Calculation

The rewards for a given epoch e_i involve the two epochs surrounding it. In particular, the stake distribution will come from the previous epoch, and the rewards will be calculated in the following epoch. More concretely:

- (A) A stake distribution snapshot is taken at the beginning of epoch e_{i-1} .
- (B) The randomness for leader election is fixed during epoch e_{i-1} .
- (C) Epoch e_i begins.
- (D) Epoch e_i ends. A snapshot is taken of the stake pool performance during epoch e_i . A snapshot is also taken of the fee pot and the decayed deposit values.
- (E) The snapshots from (D) are stable and the reward calculation can begin.
- (F) Rewards are given out.



We must therefore store the last three stake distributions. The mnemonic “mark, set, go” will be used to keep track of the snapshots, where the label “mark” refers to the most recent snapshot, and “go” refers to the snapshot that is ready to be used in the reward calculation. In the above diagram, the snapshot taken at (A) is labeled “mark” during epoch e_{i-1} , “set” during epoch e_i , and “go” during epoch e_{i+1} . At (G) the snapshot taken at (A) is no longer needed and will be discarded.

The two main transition systems in this section are:

- The transition system named EPOCH, which is defined in Section 10.7, covering what happens at the epoch boundary, such as at (A), (C), (D), and (G).
- The transition named REWARD, which is defined in Section 10.9, covering the reward calculation which happens between (E) and (F).

10.2 Helper Functions

Figure 24 defines four helper functions needed throughout the rest of the section.

- The function obligation calculates the the minimal amount of coin needed to pay out all deposit refunds, as of the current slot.

- The function `poolRefunds` is used to calculate the total refunds that must be distributed for stake pools scheduled to retire. Note that this calculation takes a slot number corresponding to the epoch boundary slot when the calculation is performed. The returned map maps pool operator hashkeys to the refunds, which will ultimately be returned to the registered reward account.
- The function `poolStake` filters the stake distribution to one stake pool.
- The function `updateAvg`s calculates the new performance moving averages.

Total possible refunds

$\text{obligation} \in \text{PParams} \rightarrow \text{StakeKeys} \rightarrow \text{StakePools} \rightarrow \text{Slot} \rightarrow \text{Coin}$

$\text{obligation } pp \text{ } stkeys \text{ } stpools \text{ } cslot =$

$$\sum_{(_ \mapsto s) \in stkeys} \text{refund } d_{\text{val}} \text{ } d_{\text{min}} \text{ } \lambda_d \text{ } (cslot -_s s) + \sum_{(_ \mapsto s) \in stpools} \text{refund } p_{\text{val}} \text{ } p_{\text{min}} \text{ } \lambda_p \text{ } (cslot -_s s)$$

where $d_{\text{val}}, d_{\text{min}}, \lambda_d = \text{keyDeposit } pp, \text{keyMinRefund } pp, \text{keyDecayRate } pp$
 $p_{\text{val}}, p_{\text{min}}, \lambda_p = \text{poolDeposit } pp, \text{poolMinRefund } pp, \text{poolDecayRate } pp$

Pool refunds

$\text{poolRefunds} \in \text{PParams} \rightarrow (\text{HashKey}_{\text{pool}} \mapsto \text{Epoch}) \rightarrow \text{Slot} \rightarrow (\text{HashKey}_{\text{pool}} \mapsto \text{Coin})$

$\text{poolRefunds } pp \text{ } retiring \text{ } cslot = \{hk \mapsto \text{refund } p_{\text{val}} \text{ } p_{\text{min}} \text{ } \lambda \text{ } (cslot -_s (\text{slot } e)) \mid hk \mapsto e \in retiring\}$

where $p_{\text{val}}, p_{\text{min}}, \lambda_p = \text{poolDeposit } pp, \text{poolMinRefund } pp, \text{poolDecayRate } pp$

Filter Stake to one Pool

$\text{poolStake} \in \text{HashKey}_{\text{pool}} \rightarrow (\text{HashKey}_{\text{stake}} \mapsto \text{HashKey}_{\text{pool}}) \rightarrow \text{Stake} \rightarrow \text{Stake}$

$\text{poolStake } hk \text{ } delegs \text{ } stake = \text{dom } (delegs \triangleleft \{hk\}) \triangleright stake$

Update Moving Averages

$\text{updateAvg} \in \text{PParams} \rightarrow \text{Avg} \rightarrow \text{BlocksMade} \rightarrow (\text{HashKey}_{\text{stake}} \mapsto \text{HashKey}_{\text{pool}}) \rightarrow \text{Stake} \rightarrow \text{Avg}$

$\text{updateAvg} \text{ } pp \text{ } avgs \text{ } blocks \text{ } delegs \text{ } stake =$

$\{hk \mapsto \text{movingAvg } pp \text{ } hk \text{ } n \text{ } (\text{expected } hk) \text{ } avgs \mid hk \mapsto n \in blocks\}$

where

$$tot = \sum_{_ \mapsto c \in stake} c$$

$$\text{expected } hk = \left(\sum_{_ \mapsto c \in (\text{poolStake } hk \text{ } delegs \text{ } stake)} c \right) \cdot \text{SlotsPerEpoch} / tot$$

Figure 24: Helper Functions used in Rewards and Epoch Boundary

10.3 Stake Distribution Calculation

This section defines the stake distribution calculations. Figure 25 introduces three new derived types:

- BlocksMade represents the number of blocks each stake pool produced during an epoch.
- Stake represents the amount of stake (in Coin) controlled by each stake pool.
- Avgs represents the performance moving averages of the stake pools.

Derived types

$blocks \in \text{BlocksMade}$	=	$\text{HashKey}_{pool} \mapsto \mathbb{N}$	blocks made by stake pools
$stake \in \text{Stake}$	=	$\text{HashKey}_{stake} \mapsto \text{Coin}$	stake
$avgs \in \text{Avgs}$	=	$\text{HashKey}_{pool} \mapsto \mathbb{R}^{\geq 0}$	performance moving averages

Figure 25: Epoch definitions

The stake distribution calculation is given in Figure 26.

- aggregate_+ takes a relation on $A \times B$, where B is any monoid, and returns a map from each $a \in A$ to the sum of all $b \in B$ such that $(a, b) \in B$.
- stakeDistr uses the aggregate_+ function and several relations to compute the stake distribution, mapping each hashkey to the total coin under its control. Keys that are not both registered and delegated are filtered out. The relation passed to aggregate_+ is made up of:
 - stakeHK_b^{-1} , relating hashkeys to (base) addresses
 - $(\text{addrPtr} \circ \text{ptr})^{-1}$, relating hashkeys to (pointer) addresses
 - range utxo , relating addresses to coins
 - $\text{stakeHK}_r^{-1} \circ \text{rewards}$, relating (reward) addresses to coins

The notation for relations is explained in Section 2.

Aggregation (for a monoid B)

$$\text{aggregate}_+ \in \mathbb{P} (A \times B) \rightarrow (A \mapsto B)$$

$$\text{aggregate}_+ R = \left\{ a \mapsto \sum_{(a,b) \in R} b \mid a \in \text{dom } R \right\}$$

Stake Distribution (using functions and maps as relations)

$\text{stakeDistr} \in \text{UT} \times \text{O} \rightarrow \text{DState} \rightarrow \text{PState} \rightarrow \text{Stake}$

$\text{stakeDistr } utxo \ dstate \ pstate = (\text{dom } activeDelegs) \triangleleft (\text{aggregate}_+ \ stakeRelation)$

where

$(stkeys, rewards, delegations, ptrs) = dstate$

$(stpools, _ _ _) = pstate$

$stakeRelation = \left(\left(\text{stakeHK}_b^{-1} \cup (\text{addrPtr} \circ ptr)^{-1} \right) \circ (\text{range } utxo) \right) \cup \left(\text{stakeHK}_r^{-1} \circ rewards \right)$

$activeDelegs = (\text{dom } stkeys) \triangleleft delegations \triangleright (\text{dom } stpools)$

Figure 26: Stake Distribution Function

10.4 Snapshot Transition

The state transition types for stake distribution snapshots are given in Figure 27. The type `Snapshots` contains the information needing to be saved on the epoch boundary:

- $pstake_{mark}$, $pstake_{set}$, and $pstake_{go}$ are the three stake distribution snapshots (paired with the corresponding delegation map), as explained in Section 10.1.
- $poolsSS$ stores the pool parameters from the epoch boundary.
- $blocksSS$ stores the performance of the completed epoch.
- $feeSS$ stores the fees and decayed deposit amounts at the epoch boundary.

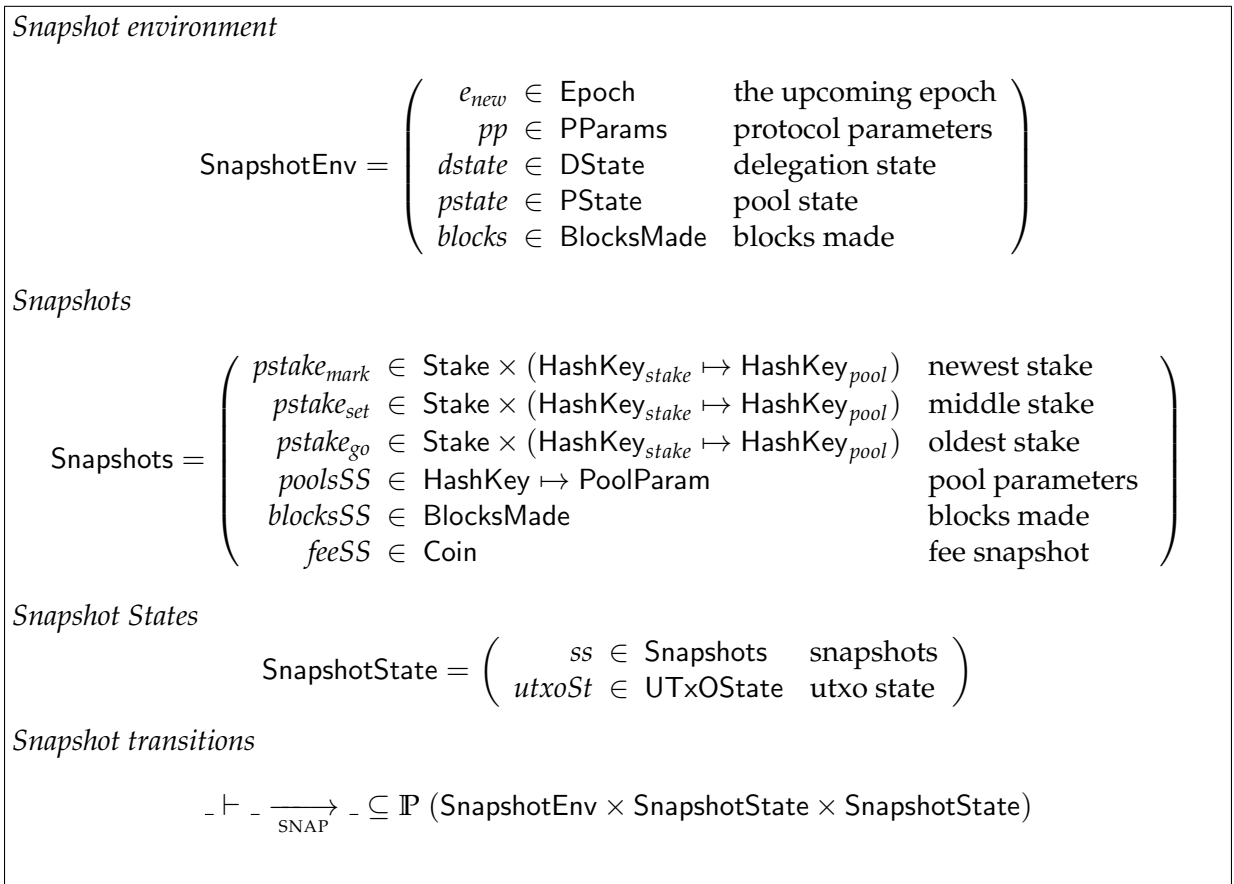


Figure 27: Snapshot transition-system types

The snapshot transition rule is given in Figure 28. This transition has no preconditions and results in the following state change:

- The oldest snapshot is replaced with the penultimate one.
- The penultimate snapshot is replaced with the newest one.
- The newest snapshot is replaced with one just calculated.
- The pool parameters are stored.
- The pool performance is stored.
- The fees and decayed deposits are stored in $feeSS$. Note that this value will not change between epochs, unlike the $fees$ and $deposits$ values in the `UTxO` state.

- In the UTxO state, the decayed deposit amounts are moved from the deposit pot to the fee pool. Note that in the reward transition (Section 10.9), the value $feeSS$ will be removed from the fee pot in the UTxO state. The decay is calculated based on *the first slot in the upcoming epoch*.

$$\begin{array}{l}
 (utxo, deposits, fees) = utxoSt \\
 (stkeys, _, delegations, _) = dstate \\
 (stpools, poolParams, _, _) = pstate \\
 stake = stakeDistr \ utxo \ dstate \ pstate \\
 slot = firstSlot \ e_{new} \\
 oblg = obligation \ pp \ stkeys \ stpools \ slot \\
 decayed = deposits - oblg
 \end{array}
 \quad (14)$$

$$\begin{array}{l}
 \text{Snapshot} \\
 e_{new} \\
 pp \\
 dstate \\
 pstate \\
 blocks
 \end{array}
 \vdash
 \begin{array}{c}
 \left(\begin{array}{c}
 pstake_{mark} \\
 pstake_{set} \\
 pstake_{go} \\
 poolsSS \\
 blocksSS \\
 feeSS \\
 \\
 utxo \\
 deposits \\
 fees
 \end{array} \right)
 \xrightarrow{\text{SNAP}}
 \begin{array}{c}
 \left(\begin{array}{c}
 (stake, delegations) \\
 \\
 pstake_{mark} \\
 pstake_{set} \\
 poolParams \\
 blocks \\
 fees + decayed \\
 \\
 utxo \\
 oblg \\
 fees + decayed
 \end{array} \right)
 \end{array}$$

Figure 28: Snapshot Inference Rule

10.5 Pool Reaping Transition

Figure 29 defines the types for the pool reap transition, which is responsible for removing pools slated for retirement in the given epoch.

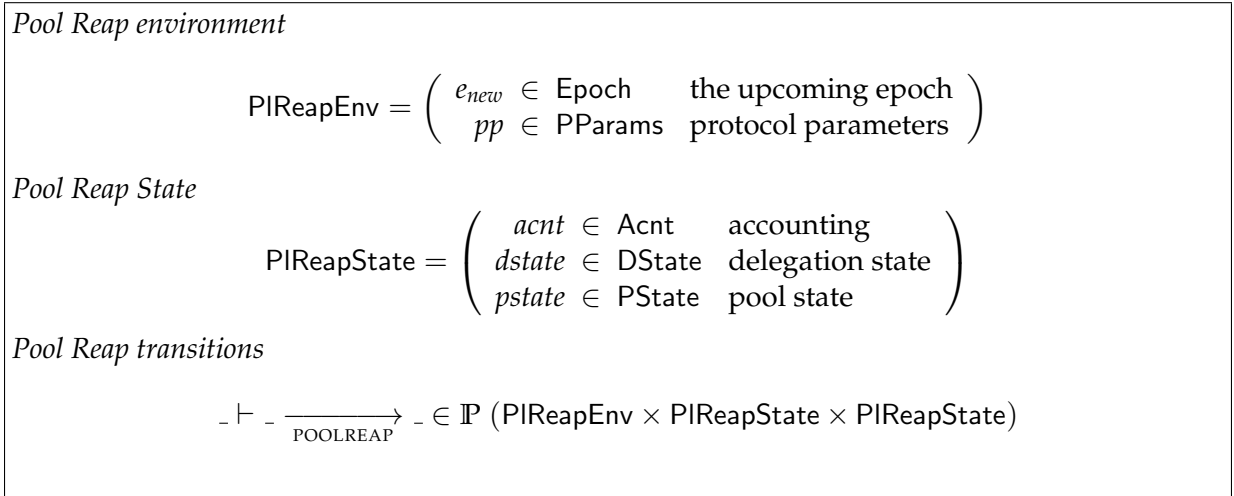


Figure 29: Pool Reap Transition

The pool-reap transition rule is given in Figure 30. This transition has no preconditions and results in the following state change:

- For each retiring pool, the refund for the pool registration deposit is added to the pool's registered reward account, provided the reward account is still registered.
- The sum of all the refunds attached to unregistered reward accounts are added to the treasury.
- Any delegation to a retiring pool is removed.
- Each retiring pool is removed from all four maps in the pool state.

$$\begin{aligned}
& \text{retired} = \text{retiring}^{-1} e_{\text{new}} \\
& \text{pr} = \text{poolRefunds } pp \text{ retiring } (\text{firstSlot } e_{\text{new}}) \\
& \text{rewardAcnts} = \{hk \mapsto \text{poolRAcnt } pool \mid hk \mapsto pool \in \text{retired} \triangleleft \text{poolParams}\} \\
& \text{refunds} = \left\{ a \mapsto c \mid \begin{array}{l} hk \mapsto c \in pr, \\ hk \mapsto a \in \text{rewardAcnts}, \\ a \in \text{dom rewards} \end{array} \right\} \\
& \text{unclaimed} = \sum_{\substack{hk \mapsto c \in pr \\ hk \mapsto a \in \text{rewardAcnts}, \\ a \notin \text{dom rewards}}} c
\end{aligned}$$

Pool-Reap

$$\begin{array}{c}
e_{\text{new}} \\
pp \vdash
\end{array}
\left(\begin{array}{c}
\text{treasury} \\
\text{reserves} \\
\text{rewardPot} \\
\\
\text{stkeys} \\
\text{rewards} \\
\text{delegations} \\
\text{ptrs} \\
\\
\text{stpools} \\
\text{poolParams} \\
\text{retiring} \\
\text{avgs}
\end{array} \right)
\begin{array}{c}
\longrightarrow \\
\text{POOLREAP}
\end{array}
\left(\begin{array}{c}
\text{treasury} + \text{unclaimed} \\
\text{reserves} \\
\text{rewardPot} \\
\\
\text{stkeys} \\
\text{rewards} \cup_+ \text{refunds} \\
\text{delegations} \triangleright \text{retired} \\
\text{ptrs} \\
\\
\text{retired} \triangleleft \text{stpools} \\
\text{retired} \triangleleft \text{poolParams} \\
\text{retired} \triangleleft \text{retiring} \\
\text{retired} \triangleleft \text{avgs}
\end{array} \right)$$

(15)

Figure 30: Pool Reap Inference Rule

10.6 Protocol Parameter Update Transition

Finally, reaching the epoch boundary may trigger a change in the protocol parameters. The protocol parameters environment consists of the upcoming epoch number, the new protocol parameters, and delegation and pool states. The state change is a change of the UTxOState, the Acnt states, and the current PParams. The type of this state transition is given in Figure 31.

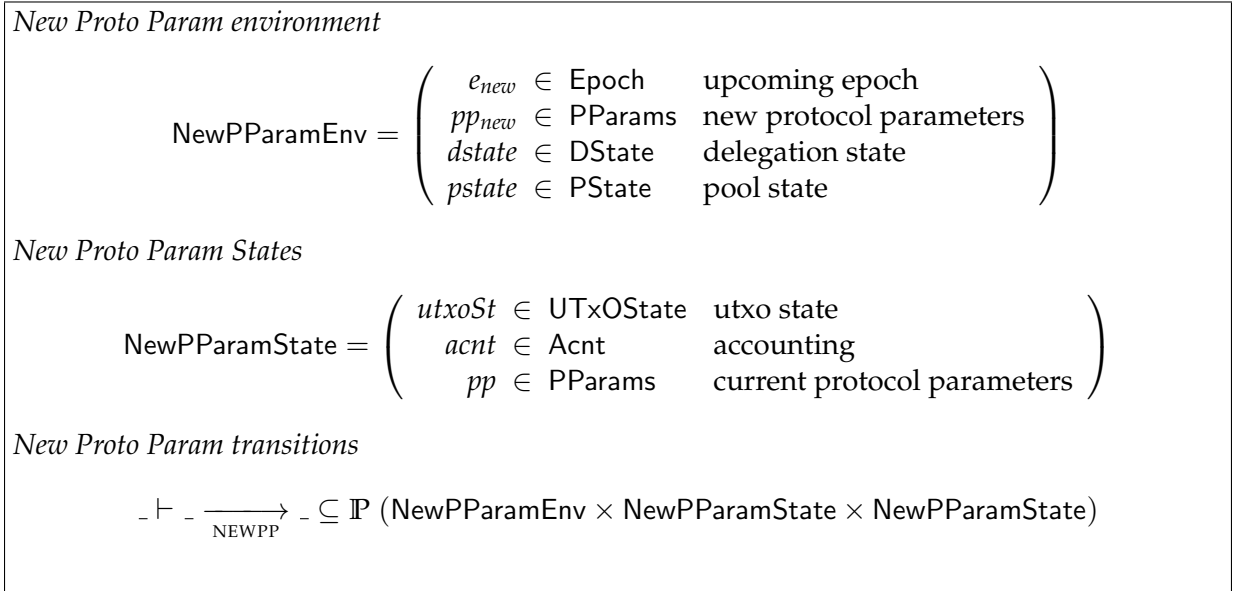


Figure 31: New Proto Param transition-system types

Figure 32 defines the new protocol parameter transition. The transition has two rules, depending on whether or not the new protocol parameters would incur a debt of the system that could not be covered by the reserves. The transition has two rules, each with one precondition. The preconditions are negations of each other, so that exactly one will always be met. This transition results in the following state change:

- If the new protocol parameters mean that **fewer** funds are required in the deposit pot to cover all possible refunds, then Rule 16 meets the precondition. The excess is moved to the reserves and the protocol parameters are updated.
- If the new protocol parameters mean that **more** funds are required in the deposit pot to cover all possible refunds, and the difference is **less** than the reserve pot, then Rule 16 meets the precondition. Funds are moved from the reserve pot to cover the difference and the protocol parameters are updated.
- If the new protocol parameters mean that **more** funds are required in the deposit pot to cover all possible refunds, and the difference is **more** than the reserve pot, then Rule 17 meets the precondition and no state changes.

Note that here, unlike most of the inference rules in this document, the $utxoSt'$ and the $acnt'$ do not come from valid UTxO or accounts transitions in the antecedent. We simply define the consequent transition using these directly (instead of listing all the fields in both states in the consequent transition). It is done this way here for ease of reading.

$$\begin{array}{l}
\text{slot} = \text{firstSlot } e_{\text{new}} \\
\text{obl}_{\text{cur}} = \text{obligation } pp \text{ stkeys stpools slot} \\
\text{obl}_{\text{new}} = \text{obligation } pp_{\text{new}} \text{ stkeys stpools slot} \\
\text{diff} = \text{obl}_{\text{cur}} - \text{obl}_{\text{new}} \\
\\
\text{reserves} + \text{diff} \geq 0 \\
\\
\text{utxoSt}' = \begin{pmatrix} \text{utxo} \\ \text{obl}_{\text{new}} \\ \text{fees} \end{pmatrix} \quad \text{acct}' = \begin{pmatrix} \text{treasury} \\ \text{reserves} + \text{diff} \\ \text{rewardPot} \\ \text{rewards} \end{pmatrix} \\
\text{New-Proto-Param-Accepted} \text{-----} \tag{16} \\
\\
\begin{array}{l}
e_{\text{new}} \\
pp_{\text{new}} \\
\text{dstate} \\
\text{pstate}
\end{array} \vdash \begin{pmatrix} \text{utxoSt} \\ \text{acct} \\ pp \end{pmatrix} \xrightarrow{\text{NEWPP}} \begin{pmatrix} \text{utxoSt}' \\ \text{acct}' \\ pp_{\text{new}} \end{pmatrix} \\
\\
\text{slot} = \text{firstSlot } e_{\text{new}} \\
\text{obl}_{\text{cur}} = \text{obligation } pp \text{ stkeys stpools slot} \\
\text{obl}_{\text{new}} = \text{obligation } pp_{\text{new}} \text{ stkeys stpools slot} \\
\text{diff} = \text{obl}_{\text{cur}} - \text{obl}_{\text{new}} \\
\\
\text{reserves} + \text{diff} < 0 \\
\text{New-Proto-Param-Denied} \text{-----} \tag{17} \\
\\
\begin{array}{l}
e_{\text{new}} \\
pp_{\text{new}} \\
\text{dstate} \\
\text{pstate}
\end{array} \vdash \begin{pmatrix} \text{utxoSt} \\ \text{acct} \\ pp \end{pmatrix} \xrightarrow{\text{NEWPP}} \begin{pmatrix} \text{utxoSt} \\ \text{acct} \\ pp \end{pmatrix}
\end{array}$$

Figure 32: New Proto Param Inference Rule

10.7 Complete Epoch Boundary Transition

Finally, it is possible to define the complete epoch boundary transition type, which is defined in Figure 33. In the environment of this transition, we have the slot number, potentially new protocol parameters, and the blocks made this epoch. The state is made up of the the UTxO state, the accounting state, the delegation state, the pool state, the current protocol parameters, and the snapshots.

Epoch environment

$$\text{EpochEnv} = \left(\begin{array}{ll} e_{new} \in \text{Epoch} & \text{upcoming epoch} \\ pp_{new} \in \text{PParams} & \text{new protocol parameters} \\ blocks \in \text{BlocksMade} & \text{blocks made in the epoch} \end{array} \right)$$

Epoch States

$$\text{EpochState} = \left(\begin{array}{ll} utxoSt \in \text{UTxOState} & \text{utxo state} \\ acnt \in \text{Acnt} & \text{accounting} \\ dstate \in \text{DState} & \text{delegation state} \\ pstate \in \text{PState} & \text{pool state} \\ pp \in \text{PParams} & \text{current protocol parameters} \\ ss \in \text{Snapshots} & \text{snapshots} \end{array} \right)$$

Epoch transitions

$$_ \vdash _ \xrightarrow{\text{EPOCH}} _ \subseteq \mathbb{P} (\text{EpochEnv} \times \text{EpochState} \times \text{EpochState})$$

Figure 33: Epoch transition-system types

The epoch transition rule calls SNAP, POOLREAP, and NEWPP in sequence.

$$\begin{array}{c}
e_{new} \\
pp \\
dstate \\
pstate \\
blocks
\end{array}
\vdash
\begin{array}{c}
ss \\
utxoSt
\end{array}
\begin{array}{c}
\longrightarrow \\
SNAP
\end{array}
\begin{array}{c}
ss' \\
utxoSt'
\end{array}$$

$$\begin{array}{c}
e_{new} \\
pp
\end{array}
\vdash
\begin{array}{c}
acnt \\
dstate \\
pstate
\end{array}
\begin{array}{c}
\longrightarrow \\
POOLREAP
\end{array}
\begin{array}{c}
acnt' \\
dstate' \\
pstate'
\end{array}$$

$$\begin{array}{c}
e_{new} \\
pp_{new} \\
dstate' \\
pstate''
\end{array}
\vdash
\begin{array}{c}
utxoSt' \\
acnt' \\
pp
\end{array}
\begin{array}{c}
\longrightarrow \\
NEWPP
\end{array}
\begin{array}{c}
utxoSt'' \\
acnt'' \\
pp'
\end{array}$$

$$\begin{array}{c}
e_{new} \\
pp_{new} \\
blocks
\end{array}
\vdash
\begin{array}{c}
utxoSt \\
acnt \\
dstate \\
pstate \\
pp \\
ss
\end{array}
\begin{array}{c}
\longrightarrow \\
EPOCH
\end{array}
\begin{array}{c}
utxoSt'' \\
acnt'' \\
dstate' \\
pstate' \\
pp' \\
ss'
\end{array}$$

(18)

Figure 34: Epoch Inference Rule

10.8 Rewards Distribution Calculation

This section defines the reward calculation for the proof of stake leader election. Figure 35 defines the pool reward as described in section 6.5.1 of [Kant et al. \(2018\)](#).

- The function `maxPool` gives the maximum reward a stake pool can receive in an epoch. This is a fraction of the total available rewards for the epoch. The result depends on the pool's relative stake, the pool's pledge, and the following protocol parameters:
 - a_0 , the leader-stake influence
 - n_{opt} , the optimal number of saturated stake pools
- The function `movingAvg` calculates the new moving average for a given stake pool based on its performance and the protocol parameter:
 - α , the moving average weight
- The function `poolReward` gives the total rewards available to be distributed to the members of the given pool. It depends on one additional protocol parameter:
 - γ , the moving average exponent

Maximal Reward Function, called $f(s, \sigma)$ in section 6.5.1 of [Kant et al. \(2018\)](#)

$$\text{maxPool} \in \text{PParams} \rightarrow \text{Coin} \rightarrow [0, 1] \rightarrow [0, 1] \rightarrow \text{Coin}$$

$$\text{maxPool } pp \ R \ \sigma \ p_r = \left\lfloor \frac{R}{1 + a_0} \cdot \left(\sigma' + p' \cdot a_0 \cdot \frac{\sigma' - p' \frac{z_0 - \sigma'}{z_0}}{z_0} \right) \right\rfloor$$

where

$$a_0 = \text{influence } pp$$

$$n_{opt} = \text{nopt } pp$$

$$z_0 = 1/n_{opt}$$

$$\sigma' = \min(\sigma, z_0)$$

$$p' = \min(p_r, z_0)$$

Exponential moving average, called $\langle x \rangle_e$ in 6.5.1 of [Kant et al. \(2018\)](#)

$$\text{movingAvg} \in \text{PParams} \rightarrow \text{HashKey} \rightarrow \mathbb{N} \rightarrow \mathbb{R}^{\geq 0} \rightarrow \text{Avgs} \rightarrow \mathbb{R}$$

$$\text{movingAvg } pp \ hk \ n \ \bar{N} \ avgs = \begin{cases} \frac{n}{\max(\bar{N}, 1)} & hk \notin \text{dom } avgs \\ \alpha \cdot \frac{n}{\max(\bar{N}, 1)} + (1 - \alpha) \cdot \text{prev} & hk \mapsto \text{prev} \in \text{avgs} \end{cases}$$

where

$$\alpha = \text{movingAvgWeight } pp$$

Actual Reward Function, called \hat{f}_j in section 6.5.1 of [Kant et al. \(2018\)](#)

$$\text{poolReward} \in \text{PParams} \rightarrow \text{HashKey} \rightarrow \mathbb{N} \rightarrow \mathbb{R}^{\geq 0} \rightarrow \text{Avgs} \rightarrow \text{Coin} \rightarrow \text{Coin}$$

$$\text{poolReward } pp \ hk \ n \ \bar{N} \ avgs \ maxP = \lfloor \text{avg}^\gamma \cdot \text{maxP} \rfloor$$

where

$$\gamma = \text{movingAvgExp } pp$$

$$\text{avg} = \text{movingAvg } pp \ hk \ n \ \bar{N} \ avgs$$

Figure 35: Functions used in the Reward Calculation

Figure 36 gives the calculation for splitting the pool rewards with its members, as described 6.5.2 of Kant et al. (2018). The portion of rewards allocated to the pool operator and owners is different than that of the members.

- The r_{leader} function calculates the leader reward, based on the pool cost, margin, and proportion of the pool's total stake. Note that this reward will go to the reward account specified in the pool registration certificate.
- The r_{member} function calculates the member reward, proportionally to their stake after the cost and margin are removed.

Pool leader reward, from section 6.5.2 of Kant et al. (2018)

$$r_{\text{leader}} \in \text{Coin} \rightarrow \text{PoolParam} \rightarrow [0, 1] \rightarrow (0, 1] \rightarrow \text{Coin}$$

$$r_{\text{leader}} \hat{f} \text{ pool } s \sigma = \begin{cases} \hat{f} & \hat{f} \leq c \\ c + \left\lfloor (\hat{f} - c) \cdot (m + (1 - m) \cdot \frac{s}{\sigma}) \right\rfloor & \text{otherwise.} \end{cases}$$

where

$$c = \text{poolCost } \text{pool}$$

$$m = \text{poolMargin } \text{pool}$$

Pool member reward, from section 6.5.2 of Kant et al. (2018)

$$r_{\text{member}} \in \text{Coin} \rightarrow \text{PoolParam} \rightarrow [0, 1] \rightarrow (0, 1] \rightarrow \text{Coin}$$

$$r_{\text{member}} \hat{f} \text{ pool } t \sigma = \begin{cases} 0 & \hat{f} \leq c \\ \left\lfloor (\hat{f} - c) \cdot (1 - m) \cdot \frac{t}{\sigma} \right\rfloor & \text{otherwise.} \end{cases}$$

where

$$c = \text{poolCost } \text{pool}$$

$$m = \text{poolMargin } \text{pool}$$

Figure 36: Functions used in the Reward Splitting

Finally, the full reward calculation is presented in Figure 37. The calculation is done pool-by-pool.

- The `rewardOnePool` function calculates the rewards given out to each member of a given pool. The pool leader is identified by the stake key of the pool operator. The function returns both the rewards and the total amount of unrealized potential rewards (ie the difference between the max reward R and what was actually paid out). The unrealized amount will go to the treasury. Note that rewards attached to unregistered reward accounts will end up in the unrealized amount. Involved in the calculation is:
 - $p\text{stake}$, the total amount of stake controlled by the stake pool.
 - $o\text{stake}$, the total amount of stake controlled by the stake pool operator and owners
 - σ , the total proportion of stake controlled by the stake pool.
 - \bar{N} , the expected number of blocks the pool should have produced.
 - $p\text{pledge}$, the pool's pledge in lovelace.

- p_r , the pool's pledge, as a proportion of active stake.
 - $maxP$, maximum rewards the pool can claim if the pledge is met, and zero otherwise.
 - $poolR$, the pool's actual reward, based on its performance.
 - $mRewards$, the member's rewards as a mapping of reward accounts to coin.
 - $lReward$, the leader's reward as coin.
 - $potentialRewards$, the combination of $mRewards$ and $lRewards$.
 - $rewards$, the restriction of $potentialRewards$ to the active reward accounts.
 - $unrealized$, difference between R (max rewards) and the sum of all individual rewards paid out.
- The reward function applies `rewardOnePool` to each registered stake pool, calculating both the full reward mapping and the total unrealized rewards value.

Calculation to reward a single stake pool

rewardOnePool \in PParams \rightarrow Coin \rightarrow \mathbb{N} \rightarrow HashKey \rightarrow PoolParam
 \rightarrow Stake \rightarrow Avgs \rightarrow Coin \rightarrow \mathbb{P} Addr_{rwd} \rightarrow (Addr_{rwd} \mapsto Coin) \times Coin
rewardOnePool pp R n poolHK pool stake avgs tot addr_{rew} = (rewards, unrealized)

where

$$pstake = \sum_{_ \mapsto c \in stake} c$$

$$ostake = \sum_{\substack{hk \mapsto c \in stake \\ hk \in (poolOwners\ pool)}} c$$

$$\sigma = pstake / tot$$

$$\bar{N} = \sigma * SlotsPerEpoch$$

$$pledge = poolPledge\ pool$$

$$p_r = pledge / tot$$

$$maxP = \begin{cases} \maxPool\ pp\ R\ \sigma\ p_r & \text{pledge} \leq \text{ostake} \\ 0 & \text{otherwise.} \end{cases}$$

$$poolR = poolReward\ pp\ hk\ n\ \bar{N}\ avgs\ maxP$$

$$mRewards = \left\{ \text{addr}_{rwd}\ hk \mapsto r_{member}\ poolR\ pool\ \frac{c}{tot}\ \sigma \mid hk \mapsto c \in stake, hk \neq poolHK \right\}$$

$$lReward = r_{leader}\ poolR\ pool\ \frac{ostake}{tot}\ \sigma$$

$$potentialRewards = mReward \cup \{ (poolRAcnt\ pool) \mapsto lReward \}$$

$$rewards = addr_{rew} \triangleleft potentialRewards$$

$$unrealized = R - \left(\sum_{_ \mapsto c \in rewards} c \right)$$

Calculation to reward all stake pools

reward \in PParams \rightarrow BlocksMade \rightarrow Coin \rightarrow \mathbb{P} Addr_{rwd} \rightarrow (HashKey \mapsto PoolParam)
 \rightarrow Avgs \rightarrow Stake \rightarrow (HashKey_{stake} \mapsto HashKey_{pool}) \rightarrow (Addr_{rwd} \mapsto Coin) \times Coin
reward pp blocks R addr_{rew} poolParams avgs stake delegs = (rewards, unrealized)

where

$$tot = \sum_{_ \mapsto c \in stake} c$$

$$pdata = \left\{ hk \mapsto (p, n, poolStake\ hk\ delegs\ stake) \mid \begin{array}{l} hk \mapsto p \in poolParams \\ hk \mapsto n \in blocks \end{array} \right\}$$

$$results = \{ hk \mapsto \text{rewardOnePool}\ pp\ R\ n\ hk\ p\ s\ avgs\ tot\ addr_{rew} \mid hk \mapsto (p, n, s) \in pdata \}$$

$$unrealized = \sum_{_ \mapsto (_, u) \in results} u$$

$$rewards = \bigcup_{_ \mapsto (r, _) \in results} r$$

Figure 37: The Reward Calculation

10.9 Reward Transition

Figure 38 gives the definitions for the reward transition. The figure lists the accounting fields, denoted by *Acnt*, which consists of:

- The value *treasury* tracks the amount of coin currently stored in the treasury. Initially there will be no way to remove these funds.
- The value *reserves* tracks the amount of coin currently stored in the reserves. This pot is used to pay rewards.
- The value *rewardPot* tracks the rewards from the previous epoch that were not paid out.

The figure also defines the reward environment, *RewardEnv*, and the reward state, *RewardState*, which combines the accounting fields described above with the UTxO State, the delegation state, and the pool state. The reward state transition type, like all the transition types in this section, has no signal.

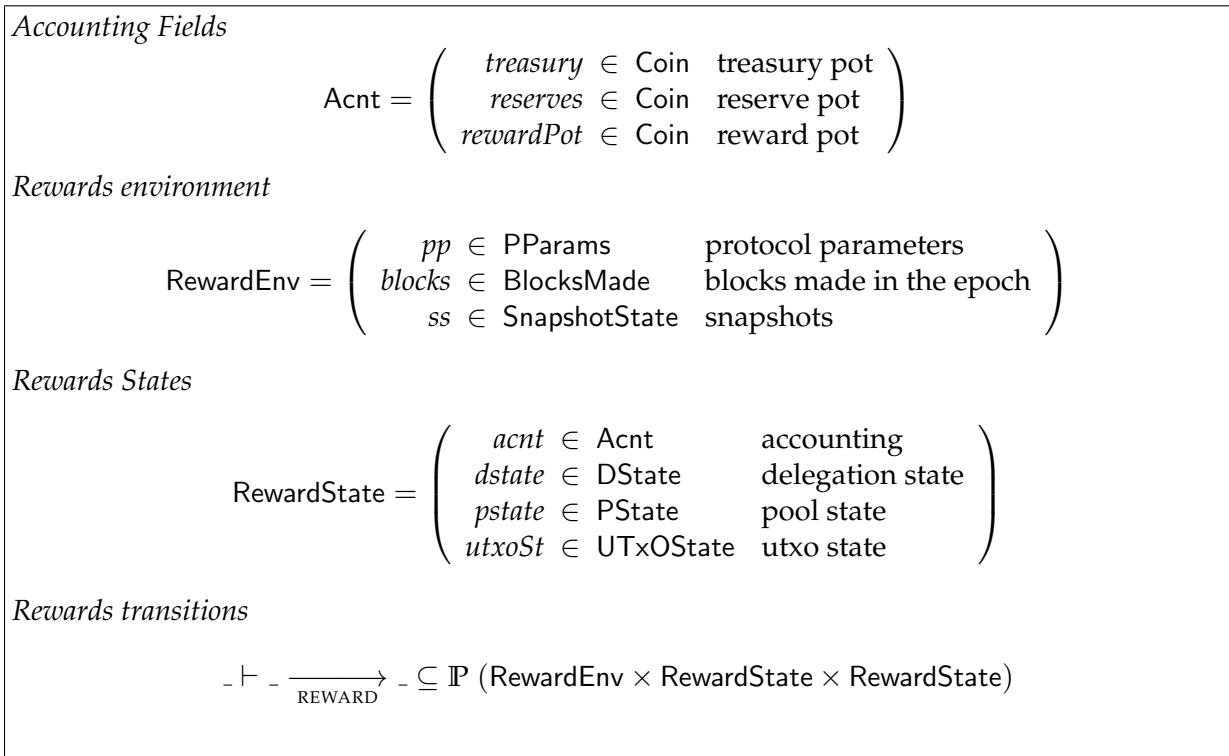


Figure 38: Rewards transition-system types

Figure 39 captures the potential movement of funds in the entire system, taking every transition system in this document into account. Value is moved between accounting pots, but the total amount of value in the system remains constant. In particular, the red subgraph represents the inputs and outputs to the “total pot” used during the reward calculation in Figure 40. The blue arrows represent the movement of funds that pass through the “total pot”.

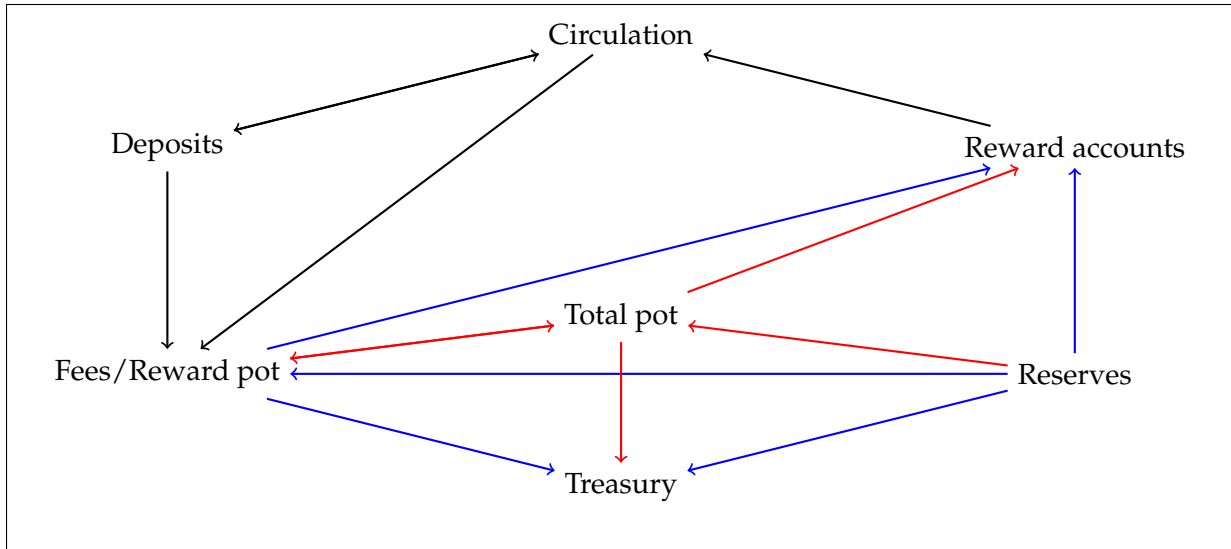


Figure 39: Preservation of Value

Figure 40 defines the reward transition rule. The reward transition has no preconditions.

- First we calculate *totalPot*, the total amount of coin avail for rewards this epoch, , as described in section 6.4 of Kant et al. (2018). It consists of four pots:
 - The fee pot, containing the transaction fees from the epoch.
 - The amount of coin in the deposit pot that is no longer needed, due to decay.
 - The reward pot, which is the left-over rewards from the previous epoch.
 - Some amount of monetary expansion from the reserves, as determined by the ρ protocol parameter.

Note that the fee pot and the decayed amount are taken from the snapshot taken at the epoch boundary. (See Figure28).

- Some proportion of the total pot is moved to the treasury, as determined by the τ protocol parameter. The remaining pot is called the *R*, just as in section 6.5 of Kant et al. (2018).
- The rewards are calculated, using the oldest stake distribution snapshot (the one labeled “go”). As given by maxPool, each pool can receive a maximal amount, determined by its performance. The difference between the maximal amount and the actual amount received is moved to the treasury.
- The reward pot is now set to *R* less the total amount of rewards actually paid out and the unrealized rewards given to the treasury.
- The moving averages are updated. Note the averages were already computed on the previous epoch boundary, as a part of the stake distribution calculation, and could therefore be cached in order to prevent calculating them twice.
- The fee pot is reduced by *feeSS*.

Note that fees are not explicitly removed from any account: the fees come from transactions paying them, and are accounted for whenever transactions are processed, and the deposit decay value comes from returning smaller refunds for deposits than were paid upon depositing.

$$\begin{array}{l}
\begin{pmatrix} - \\ - \\ pstate_{go} \\ poolsSS \\ blocksSS \\ feeSS \end{pmatrix} =_{ss} \\
(stake, delegs) = pstate_{go} \\
expansion = \lfloor (\rho \ pp) \cdot reserves \rfloor \\
totalPot = feeSS + rewardPot + expansion \\
newTreasury = \lfloor (\tau \ pp) \cdot totalPot \rfloor \\
R = totalPot - newTreasury \\
rewards', unrealized = reward \ pp \ blocksSS \ R \ (\text{dom } rewards) \ poolsSS \ avgs \ stake \ delegs \\
newTreasury' = newTreasury + unrealized \\
paidRewards = \left(\sum_{c \in rewards'} c \right) + unrealized \\
avgs' = updateAvg \ pp \ avgs \ blocks \ delegs \ stake \\
\hline
\text{Rewards} \\
\begin{array}{l}
pp \\
blocks \vdash \\
ss
\end{array}
\begin{pmatrix}
treasury \\
reserves \\
rewardPot \\
\\
stkeys \\
rewards \\
delegations \\
ptrs \\
\\
stpools \\
poolParams \\
retiring \\
avgs \\
\\
utxo \\
deposits \\
fees
\end{pmatrix}
\begin{array}{c}
\longrightarrow \\
\text{REWARD}
\end{array}
\begin{pmatrix}
treasury + newTreasury' \\
reserves - expansion \\
R - paidRewards \\
\\
stkeys \\
rewards \cup_+ rewards' \\
delegations \\
ptrs \\
\\
stpools \\
poolParams \\
retiring \\
avgs' \\
\\
utxo \\
deposits \\
fees - feeSS
\end{pmatrix} \\
(19)
\end{array}$$

Figure 40: Rewards inference rules

11 Properties

In this section we discuss the properties which we want the ledger to have. One goal is to include these properties in the executable specification for doing property-based testing or formal verification.

11.1 Validity of a Ledger State

Many properties only make sense when applied to a valid ledger state. In informal terms, a valid ledger state l can only be reached when starting from an initial state l_0 (genesis state) and only executing state transition rules as specified in Section 7 for UTxO or Section 8 for delegation.

$Genesis_{Id} \in$	$TxId$
$Genesis_{Out} \in$	$TxOut$
$Genesis_{UTxO} :=$	$\{Genesis_{Id}, \emptyset\} \mapsto Genesis_{Out}$
$ledgerState \in$	$\left(\begin{array}{c} UTxO \\ DPState \end{array} \right)$
$getUTxO \in$	$ledgerState \rightarrow UTxO$
$getUTxO :=$	$(utxo, _) \rightarrow utxo$

Figure 41: Definitions and Functions for Valid Ledger State

In Figure 41 $Genesis_{Id}$ marks the transaction identifier of the initial coin distribution, where $Genesis_{Out}$ represents the initial UTxO. It should be noted that no corresponding inputs exists, i.e., the transaction inputs are the empty set for the initial transaction. An element of $ledgerState$ is a tuple of UTxO and delegation witness state (DPState).

Definition 11.1 (Valid Ledger State).

$$\begin{aligned} \forall l_0, \dots, l_n \in ledgerState, l_0 = & \left(\begin{array}{c} \{Genesis_{UTxO}\} \\ \left(\begin{array}{c} \emptyset \\ \emptyset \end{array} \right) \end{array} \right) \\ \implies \forall 0 < i \leq n, (\exists tx_i \in Tx, l_{i-1} \xrightarrow[LEDGER]{tx_i} l_i) & \implies validLedgerState l_n \end{aligned}$$

Definition 11.1 defines a valid ledger state reachable from the genesis state via valid UTxO, stake delegation or stake pool transactions. This gives a constructive rule how to reach a valid ledger state.

11.2 Ledger Properties

The following properties state the desired features of updating a valid ledger state.

Property 11.1 (Preserve Balance Modulo Fee).

$$\begin{aligned} \forall l, l' \in ledgerState : validLedgerstate l & \\ \implies \forall tx \in Tx, l \xrightarrow[UTxOW]{tx} l' & \\ \implies destroyed pcutxostKeysrewardstx = created pcstPoolstx & \end{aligned}$$

Property 11.1 states that for each valid ledger l , if a transaction tx is added to the ledger via the state transition rule $utxow$ to the new ledger state l' , the balance of the UTxOs in l equals the balance of the UTxOs in l' in the sense that the amount of created value in l' equals the amount of destroyed value in l . This means that the total amount of value is left unchanged by a transaction.

Property 11.2 (Preserve Balance Restricted to TxIns in Balance of TxOuts).

$$\begin{aligned} & \forall l, l' \in \text{ledgerState} : \text{validLedgerstate } l \\ \implies & \forall tx \in \text{Tx}, l \xrightarrow[\text{UTXOW}]{tx} l' \implies \text{ubalance}(\text{txins } tx \triangleleft \text{getUTxO } l) = \text{ubalance}(\text{outs } tx) + \text{txfee } tx \end{aligned}$$

Property 11.2 states the more detailed relation of the balances change. For ledgers l, l' and a transaction tx as above, the balance of the UTxOs of l restricted to those whose domain is in the set of transaction inputs of tx equals the balance of the transaction outputs of tx minus the transaction fees.

Property 11.3 (Preserve Outputs of Transaction).

$$\begin{aligned} & \forall l, l' \in \text{ledgerState} : \text{validLedgerstate } l \\ \implies & \forall tx \in \text{Tx}, l \xrightarrow[\text{UTXOW}]{tx} l' \implies \forall out \in \text{outs } tx, out \in \text{getUTxO } l' \end{aligned}$$

Property 11.3 states that for every ledger states l, l' and transaction tx as above, all output UTxOs of tx are in the UTxO set of l' , i.e., they are now available as unspent transaction output.

Property 11.4 (Eliminate Inputs of Transaction).

$$\begin{aligned} & \forall l, l' \in \text{ledgerState} : \text{validLedgerstate } l \\ \implies & \forall tx \in \text{Tx}, l \xrightarrow[\text{UTXOW}]{tx} l' \implies \forall in \in \text{txins } tx, in \notin \text{dom}(\text{getUTxO } l') \end{aligned}$$

Property 11.4 states that for every ledger states l, l' and transaction tx as above, all transaction inputs in of tx are not in the domain of the UTxO set of l' , i.e., these are no longer available to spend.

Property 11.5 (Completeness and Collision-Freeness of new Transaction Ids).

$$\begin{aligned} & \forall l, l' \in \text{ledgerState} : \text{validLedgerstate } l \\ \implies & \forall tx \in \text{Tx}, l \xrightarrow[\text{UTXOW}]{tx} l' \implies \forall utxo' \in \text{outs } tx, utxo' \in \text{getUTxO } l' \wedge \\ & (utxo' = ((txId', _) \mapsto _) \implies \forall utxo \in \text{getUTxO } l, utxo = ((txId, _) \mapsto _) \implies txId' \neq txId \end{aligned}$$

Property 11.5 states that for ledger states l, l' and a transaction tx as above, the UTxOs of l' contain all newly created UTxOs and the referred transaction id of each new UTxO is not used in the UTxO set of l .

Property 11.6 (Absence of Double-Spend).

$$\begin{aligned}
\forall l_0, \dots, l_n \in \text{ledgerState}, l_0 &= \left(\begin{array}{c} \{\text{Genesis}_{\text{UTxO}}\} \\ \left(\begin{array}{c} \emptyset \\ \emptyset \end{array} \right) \end{array} \right) \wedge \text{validLedgerState } l_n \\
\implies \forall 0 < i \leq n, tx_i \in \text{Tx}, l_{i-1} &\xrightarrow[\text{LEDGER}]{tx_i} l_i \wedge \text{validLedgerState } l_i \\
&\implies \forall j < i, \text{txins } tx_j \cap \text{txins } tx_i = \emptyset
\end{aligned}$$

Property 11.6 states that for each valid ledger state l_n reachable from the genesis state, each transaction t_i does not share any input with any previous transaction t_j . This means that each output of a transition is spent at most once.

11.3 Ledger State Properties for Delegation Transitions

getStKeys ∈	ledgerState → ℙ HashKey
getStKeys :=	(_, (stKeys, _, _), _) → stkeys
getRewards ∈	ledgerState → Addr _{rwd} ↦ Coin
getRewards :=	(_, (_, rewards, _), _) → rewards
getDelegations ∈	ledgerState → HashKey ↦ HashKey
getDelegations :=	(_, (_, _, delegations), _) → delegations
getStPools ∈	ledgerState → HashKey ↦ DCert _{regpool}
getStPools :=	(_, _, (stpools, _)) → stpools
getRetiring ∈	ledgerState → HashKey ↦ Epoch
getRetiring :=	(_, _, (_, retiring)) → retiring

Figure 42: Definitions and Functions for Stake Delegation in Ledger States

Property 11.7 (Registered Staking Key with Zero Rewards).

$$\begin{aligned}
\forall l, l' \in \text{ledgerState} : \text{validLedgerstate } l \\
\implies \forall c \in \text{DCert}_{\text{regkey}}, l \xrightarrow[\text{DELEGW}]{c} l' \implies \text{author } c = hk \\
\implies hk \in \text{getStKeys } l' \wedge (\text{getRewards } varrewards)[hk] = 0
\end{aligned}$$

Property 11.7 states that for each valid ledger state l , if a delegation transaction of type $\text{DCert}_{\text{regkey}}$ is executed, then in the resulting ledger state l' , the set of staking keys of l' includes the key hk associated with the key registration certificate and the associated reward is set to 0 in l' .

Property 11.8 (Deregistered Staking Key).

$$\begin{aligned} \forall l, l' \in \text{ledgerState} : \text{validLedgerstate } l \\ \implies \forall c \in \text{DCert}_{\text{deregkey}}, l \xrightarrow[\text{DELEGW}]{c} l' \implies \text{author } c = hk \\ \implies hk \notin \text{getStKeys } l' \wedge (\text{dom}(\text{getRewards } l') \cup \text{dom}(\text{getDelegations } l')) \cap \{hk\} = \emptyset \end{aligned}$$

Property 11.8 states that for l, l' as above but with a delegation transition of type $\text{DCert}_{\text{deregkey}}$, the staking key hk associated with the deregistration certificate is not in the set of staking keys of l' and is not in the domain of neither the rewards nor the delegation map of l' .

Property 11.9 (Delegated Stake).

$$\begin{aligned} \forall l, l' \in \text{ledgerState} : \text{validLedgerstate } l \\ \implies \forall c \in \text{DCert}_{\text{delegate}}, l \xrightarrow[\text{DELEGW}]{c} l' \implies \text{author } c = hk \\ \implies hk \in \text{getStKeys } l \wedge (\text{getDelegations } l')[hk] = \text{pool } c \end{aligned}$$

Property 11.9 states that for l, l' as above but with a delegation transition of type $\text{DCert}_{\text{delegate}}$, the staking key hk associated with the deregistration certificate is in the set of staking keys of l and delegates to the staking pool associated with the delegation certificate in l' .

11.4 Ledger State Properties for Staking Pool Transitions

Property 11.10 (Registered Staking Pool).

$$\begin{aligned} \forall l, l' \in \text{ledgerState} : \text{validLedgerstate } l \\ \implies \forall c \in \text{DCert}_{\text{regpool}}, l \xrightarrow[\text{POOL}]{c} l' \implies \text{author } c = hk \\ \implies (\text{getStPools } l')[hk] = c \wedge hk \notin \text{getRetiring } l' \end{aligned}$$

Property 11.10 states that for l, l' as above but with a delegation transition of type $\text{DCert}_{\text{regpool}}$, the key hk is associated with the author of the pool registration certificate in stPools of l' and that hk is not in the set of retiring stake pools in l' .

Property 11.11 (Start Staking Pool Retirement).

$$\begin{aligned} \forall l, l' \in \text{ledgerState}, \text{epoch} \in \text{Epoch} : \text{validLedgerstate } l \\ \implies \forall c \in \text{DCert}_{\text{retirepool}}, l \xrightarrow[\text{POOL}]{c} l' \\ \implies e = \text{retire } c \wedge \text{epoch} < e < \text{epoch} + E_{\text{max}} \wedge \text{author } c = hk \\ \implies (\text{getRetiring } l')[hk] = e \wedge hk \in \text{dom}(\text{getStPools } l) \end{aligned}$$

Property 11.11 states that for l, l' as above but with a delegation transition of type $\text{DCert}_{\text{retirepool}}$, the key hk is associated with the author of the pool registration certificate in stPools of l' and that hk is not in the set of retiring stake pools in l' .

Property 11.12 (Stake Pool Reaping).

$$\begin{aligned}
& \forall l, l' \in \text{ledgerState}, \text{epoch} \in \text{Epoch} : \text{validLedgerstate } l \\
& \implies l \xrightarrow{\text{POOLREAP}} l' \implies \forall \text{retire} = \text{retiring}^{-1} \text{epoch}, \text{retired} \neq \emptyset \\
& \quad \wedge \text{retire} \subseteq \text{dom}(\text{getStPool } l) \wedge \text{retire} \cap \text{dom}(\text{getStPool } l') = \emptyset \\
& \quad \wedge \text{retire} \subseteq \text{dom}(\text{getRetiring } l) \wedge \text{retire} \cap \text{dom}(\text{getRetiring } l') = \emptyset
\end{aligned}$$

Property 11.12 states that for l, l' as above but with a delegation transition of type `poolreap`, there exist registered stake pools in l which are associated to stake pool registration certificates and which are to be retired at the current epoch epoch . In l' all those stake pools are removed from the maps `stpools` and `retiring`.

11.5 Properties of Numerical Calculations

The numerical calculations for refunds and rewards calculation in (see Section 10) are also required to have certain properties. In particular we need to make sure that the functions that use non-integral arithmetic have properties which guarantee consistency of the system. Here, we state those properties and formulate them in a way that makes them possible to use properties-based testing for validation in the executable spec.

Property 11.13 (Minimal Refund). The function `refund` takes a value, a minimal percentage, a decay parameter and a duration. It must guarantee that the refunded amount is within the minimal refund (off-by-one for rounding / floor) and the original value.

$$\begin{aligned}
& \forall d_{\text{val}} \in \mathbb{N}, d_{\text{min}} \in [0, 1], \lambda \in (0, \infty), \delta \in \mathbb{N} \\
& \implies \max(0, d_{\text{val}} \cdot d_{\text{min}} - 1) \leq \left\lfloor d_{\text{val}} \cdot (d_{\text{min}} + (1 - d_{\text{min}}) \cdot e^{-\lambda \cdot \delta}) \right\rfloor \leq d_{\text{val}}
\end{aligned}$$

Property 11.14 (Exponential Moving Average). The function `movingAvg` calculates the exponential moving average, dividing the number of blocks created by the pool by the expected number of slots the pool is elected leader (or 1 if the expected number is below 1). It guarantees that the result is (i) non-negative and (ii) if a previous moving average has already been calculated, the new moving average lies between the minimum and maximum of the old and new calculated value. With $\text{current} := \frac{n}{\max(\bar{N}, 1)}$ this is trivial for (i), for (ii) it is

$$\begin{aligned}
& \forall \alpha \in [0, 1], n \in \mathbb{N}, \bar{N} \in \mathbb{R}^{\geq 0}, \text{prev} \in \mathbb{R}^{\geq 0} \\
& \implies 0 \leq \min(\text{prev}, \text{current}) \leq \alpha \cdot \text{current} + (1 - \alpha) \cdot \text{prev} \leq \max(\text{prev}, \text{current})
\end{aligned}$$

Property 11.15 (Maximal Pool Reward). The maximal pool reward is the expected maximal reward paid to a stake pool. The sum of all these rewards cannot exceed the total available reward, let Pool be the set of active stake pools:

$$\forall R \in \text{Coin} : \sum_{p \in \text{Pools}} \left\lfloor \frac{R}{1 + p_{a_0}} \cdot \left(p_{\sigma'} + p_{p'} \cdot a_0 \cdot \frac{p_{\sigma'} - p_{p'} \cdot \frac{p_{z_0} - p_{\sigma'}}{p_{z_0}}}{p_{z_0}} \right) \right\rfloor \leq R$$

Property 11.16 (Actual Reward). The actual reward for a stake pool in an epoch is calculated by the function `poolReward`. The actual reward per stake pool is non-negative and bounded by the

maximal reward for the stake pool, with avg being the calculated moving average of the stake pool and $maxP$ being the maximal reward for the stake pool, we get:

$$\forall \gamma \in [0, 1] \implies 0 \leq \lfloor avg^\gamma \cdot maxP \rfloor \leq maxP$$



TODO

The property (11.16) requires that $avg \in [0, 1]$, else the actual reward can exceed the maximal reward. This is not true, we need to take into account the rewards for all stake pools.

The two functions r_{leader} and r_{member} are closely related as they do split the reward between the pool leader and the members.

Property 11.17 (Reward Splitting). The reward splitting is done via r_{leader} and r_{member} , i.e., a split between the pool leader and the pool members using the pool cost c and the pool margin m . Therefore the property relates the total reward \hat{f} to the split rewards in the following way:

$$\forall m \in [0, 1], c \in Coin \implies c + \left\lfloor (\hat{f} - c) \cdot (m + (1 - m)) \cdot \frac{s}{\sigma} \right\rfloor + \sum_j \left\lfloor (\hat{f} - c) \cdot (1 - m) \cdot \frac{t_j}{\sigma} \right\rfloor \leq \hat{f}$$

Property 11.18 (Accounting transition). The ACCNT transition rule has the following properties in order to ensure non-negative values:

- $obl \leq deposits$
- $\forall \rho \in [0, 1] : expansion = \lfloor \rho \cdot reserves \rfloor \leq reserves$
- $\forall \tau \in [0, 1] : newTreasury = \lfloor \tau \cdot totalPool \rfloor \leq totalPool$
- $totalPool - newTreasury - paidRewards \geq 0$

where $paidRewards$ is the sum of rewards paid to the stake pool members.

12 Non-Integral Calculations

In the ledger there are several cases where non-integral calculations are required. This does concern the delegation transitions, not value transactions.

12.1 Types of Non-Integral Calculations

The specification employs non-integral calculations for different mathematical operations. Table 1 shows the function and transition rules that use non-integral calculations and which type.

name	page	multiplication	division	exponential function	exponentiation
refund	18	✓		✓	
maxPool	48	✓	✓		
movingAvg	48	✓	✓		
poolReward	48	✓		✓	✓
r _{leader}	49	✓	✓		
r _{member}	49	✓	✓		
rewardOnePool	51	✓	✓		
updateAvg	36	✓	✓		
REWARD	55	✓			

Table 1: Functions with Non-Integral Calculation

The transcendental exponential function is used in reward and refund calculation to model the decay of the deposit values. The pool reward uses exponentiation to calculate a pool's ranking.

The domain for the exponential function are the non-negative reals, more precisely the distribution parameter $\lambda \in (0, \infty)$ multiplied by a discrete non-negative duration δ .

The domain of the base of the exponentiation in poolReward are the non-negative reals resulting from the calculation in movingAvg, the exponent γ is a constant take from the protocol parameters.

12.2 Implementation of Non-Integer Calculations

The large part consists of multiplication and division which can easily be done using fractional arithmetic to the desired precision. The precision necessary is bounded by the ability to represent a single lovelace in all calculations.

12.2.1 Function Simplification

The transcendental function e^x can be approximated using different approaches, depending on the desired accuracy. In general, one uses the exponential laws $e^x = 1/e^{-x}$ and $e^x = \left(e^{\frac{x}{n}}\right)^n$, $n \in \mathbb{N}$ to reduce the approximation to the unit interval and apply fast integral exponentiation afterwards.

Exponentiation is implemented using the law $a^b = e^{\ln(a^b)} = e^{b \ln(a)}$. This therefore requires being able to calculate e^x and $\ln(x)$. The approximation of the natural logarithm can be approximated using different approaches, again, depending on the desired accuracy. Most approximations work for $\ln(x)$, $x \in [1, c)$ with some $c > 0$. One then uses the law $\log_b(x) = \log_b\left(\frac{x}{b^n} b^n\right)$ where $n \in \mathbb{N}$ is chosen in such a way that $\frac{x}{b^n} \in [1, c)$. Using this, one can separate the calculation of the integral and decimal part as follows:

$$\log_b\left(\frac{x}{b^n}b^n\right) = \log_b(b^n) + \log_b\left(\frac{x}{b^n}\right) = n + \log\left(\frac{x}{b^n}\right)$$

12.2.2 Properties of Function Approximation

There are several properties that approximations of the transcendental functions are expected to have. In the following let $\ln'(x)$ be the approximation of $\ln(x)$, $\exp'(x)$ be the approximation of e^x and $x \star y$ the approximation of x^y .

Property 12.1 (monotonicity). Both \exp' and \ln' must be monotone on their respective domains.

In order to guarantee correctness of the approximations, we also require that the mathematical laws are fulfilled. For some small $\epsilon > 0$, define $x \approx y \Leftrightarrow |x - y| < \epsilon$.

Property 12.2 (Mathematical Laws). The following mathematical laws state the requirements for the approximations of the \ln' and \exp' function:

- $\ln'(x \cdot y) \approx \ln'(x) + \ln'(y)$
- $\ln'(x^y) \approx y \cdot \ln'(x)$
- $\ln'(\exp'(x)) \approx \exp'(\ln'(x)) \approx x$
- $x, y \in [0, 1] \implies x \star y \in [0, 1]$
- $x, y, z \in [0, 1], x > 0 \implies (z \star \frac{1}{x}) \star y \approx (z \star y) \star \frac{1}{x}$
- $\exp'(x + y) = \exp'(x) + \exp'(y)$

A Proofs

For the proofs we use the automated theorem prover MetiTarski [Akbarpour and Paulson \(2010\)](#) which is specialized for proofs over real arithmetic, including elementary functions.

Proof. The property (11.13) (p. 60) for the minimal refund can be proven automatically via

```
fof(minimal_refund, conjecture,
! [Dmin, Lambda, Delta, Dval] :
((Dmin : (=0,1=) & Lambda > 0 & Delta > 0 & Dval > 0
=>
Dval*Dmin >= 0 &
(Dval * (Dmin + (1 - Dmin) * exp(-Lambda * Delta))) : (=Dval * Dmin, Dval=))))).
```

```
fof(floor_lower_upper, conjecture,
! [X] :
(X >= 0 => X - 1 <= floor(X) & floor(X) <= X)).
```

`minimal_refund` shows that the resulting value is within the interval $[d_{val} \cdot d_{min}, d_{val}]$ and that $d_{val} \cdot d_{min}$ is non-negative, while `floor_lower_upper` shows that the floor of a value x has an upper bound x and lower bound $x - 1$.

□

Proof. The property (11.14) (p. 60) for the bounded moving average can be proven automatically via

```
fof(simple_moving_average, conjecture,
! [Alpha, Prev, N, Nexpected] :
(Alpha : (=0, 1=) & N >= 0 & Nexpected >= 0 & Prev >= 0
=>
0 <= N/max(Nexpected, 1) &
N/max(Nexpected, 1) <= N &
Alpha*N/max(Nexpected, 1) + (1 - Alpha) * Prev <= max(Prev, N/max(Nexpected, 1)) &
min(Prev, N/max(Nexpected, 1)) <= Alpha*N/max(Nexpected, 1) + (1 - Alpha) * Prev)).
```

□

Proof. For the property (11.17) (p. 61) for reward splitting between we actually show a stronger one, by removing the floor function. Using the fractional values we get an upper bound for the real value, and showing that this upper bound is bounded by \hat{f} we show that the real value is also bounded by \hat{f} . To eliminate the sum, we use the identity $\frac{s + \sum_j t_j}{\sigma} = 1$, see the definition of σ in [Kant et al. \(2018\)](#). Using this, we show for $\hat{f} > c$

$$\begin{aligned}
0 &\leq c + (\hat{f} - c) \cdot (m + (1 - m)) \cdot \frac{s}{\sigma} + \sum_j (\hat{f} - c) \cdot (1 - m) \cdot \frac{t_j}{\sigma} \leq \hat{f} \\
&\Leftrightarrow 0 \leq c + (\hat{f} - c) \cdot m \cdot \frac{s}{\sigma} + (\hat{f} - c) \cdot (1 - m) \cdot \frac{s + \sum_j t_j}{\sigma} \leq \hat{f} \\
&\Leftrightarrow 0 \leq c + (\hat{f} - c) \cdot m \cdot \frac{s}{\sigma} + (\hat{f} - c) \cdot (1 - m) \leq \hat{f}
\end{aligned}$$

This can be proven automatically using


```

fof(reward_splitting, conjecture,
! [C, F, M, S, Sigma] :
(
M : (=0, 1=) & C >= 0 & F > C & Sigma : (0, 1=) & S : (=0, Sigma=)
=>
C + (F - C) * M * S / Sigma + (F - C) * (1 - M) <= F &
0 <= C + (F - C) * M * S / Sigma + (F - C) * (1 - M))).

```

□

References

- Behzad Akbarpour and Lawrence C. Paulson. Metitarski: An automatic theorem prover for real-valued special functions. *J. Autom. Reasoning*, 44(3):175–205, 2010. doi: 10.1007/s10817-009-9149-2. URL <https://doi.org/10.1007/s10817-009-9149-2>.
- IOHK Formal Methods Team. Small step semantics for cardano, 2018. URL <https://github.com/input-output-hk/cardano-chain/blob/master/specs/semantics/latex/small-step-semantics.tex>.
- IOHK Formal Methods Team. ?? - shelley consensus. TODO.
- Philipp Kant, Lars Brünjes, and Duncan Coutts. Design specification for delegation and incentives in cardano, 2018. URL https://github.com/input-output-hk/fm-ledger-rules/tree/master/docs/delegation_design_spec.
- Joachim Zahnentferner. Chimeric ledgers: Translating and unifying utxo-based and account-based cryptocurrencies. *Cryptology ePrint Archive, Report 2018/262*, 2018. URL <https://eprint.iacr.org/2018/262>.