



A Formal Specification of the Cardano Consensus

Deliverable XXXX

Javier Díaz <javier.diaz@iohk.io>

Project: Cardano

Type: Deliverable
Due Date: XXXX

Responsible team: Formal Methods Team
Editor: Javier Díaz, IOHK
Team Leader: James Chapman, IOHK

Version 1.0
XXXX

Dissemination Level		
PU	Public	✓
CO	Confidential, only for company distribution	
DR	Draft, not for general circulation	

Contents

1	Transition Rule Dependencies	2
2	Common Interface	3
2.1	NEWEPOCH Transition	3
3	Blockchain Layer	4
3.1	Block Definitions	4
3.2	TICKN Transition	7
3.3	UPDN Transition	7
3.4	OCERT Transition	7
3.5	Verifiable Random Function	10
3.6	PRTCL Transition	12
3.7	TICKF Transition	14
3.8	CHAINHEAD Transition	15
4	Properties	18
4.1	Header-Only Validation	18
5	Leader Value Calculation	20
5.1	Computing the leader value	20
5.2	Node eligibility	20
	References	21
A	Cryptographic Details	22
A.1	Abstract functions	22

List of Figures

1	STS Rules, Sub-Rules and Dependencies	2
2	New Epoch transition-system types	3
3	Block Definitions	5
4	Helper Functions used for Blocks	6
5	Tick Nonce types	7
6	Tick Nonce rules	7
7	Update Nonce transition-system types	8
8	Update Nonce rules	8
9	Operational Certificate transition-system types	9
10	Operational Certificate rules	9
11	VRF helper functions	10
12	Protocol transition-system types	12
13	Protocol rules	12
14	Tick Forecast transition-system types	14
15	Tick Forecast rules	14
16	Chain Head transition-system types	16
17	Helper functions used in the Chain Head transition	16
18	Chain Head rules	17

Change Log

Rev.	Date	Who	Team	What
1	2024/03/01	Javier Díaz	FM (IOHK)	Initial version (0.1).

A Formal Specification of the Cardano Consensus

Javier Díaz
javier.diaz@iohk.io

March 16, 2026

Abstract

This document provides a formal specification of the Cardano consensus layer.

List of Contributors

...

1 Transition Rule Dependencies

Figure 1 shows all STS rules, the sub-rules they use and possible dependencies. Each node in the graph represents one rule, the top rule being CHAINHEAD. A straight arrow from one node to another one represents a sub-rule relationship.

An arrow with a dotted line from one node to another represents a dependency in the sense that the output of the target rule is an input to the source one, either as part of the source state, the environment or the signal. These dependencies are between sub-rules of a rule.




Figure 1: STS Rules, Sub-Rules and Dependencies

2 Common Interface

2.1 NEWEPOCH Transition

The New Epoch Transition (NEWEPOCH) performs some upkeep when a new epoch begins.

New Epoch Transitions

$$\vdash _ \xrightarrow[\text{NEWEPOCH}]{} _ \subseteq \mathbb{P} (\text{NewEpochState} \times \text{Epoch} \times \text{NewEpochState})$$

Figure 2: New Epoch transition-system types

3 Blockchain Layer

This chapter introduces the view of the blockchain layer as required for the consensus. The main transition rule is CHAINHEAD which calls the subrules TICKF, TICKN and PRTCL.

3.1 Block Definitions

<i>Abstract types</i>	
$h \in \text{HashHeader}$	hash of a block header
$hb \in \text{HashBBody}$	hash of a block body
$bn \in \text{BlockNo}$	block number
$\eta \in \text{Nonce}$	nonce
$vrfRes \in \text{VRFRes}$	VRF result value
<i>Operational Certificate</i>	
$\text{OCert} =$	$\left(\begin{array}{ll} vk_{hot} \in \text{VKey}_{ev} & \text{operational (hot) key} \\ n \in \mathbb{N} & \text{certificate issue number} \\ c_0 \in \text{KESPeriod} & \text{start KES period} \\ \sigma \in \text{Sig} & \text{cold key signature} \end{array} \right)$
<i>Block Header Body</i>	
$\text{BHeader} =$	$\left(\begin{array}{ll} prevHeader \in \text{HashHeader}^? & \text{hash of previous block header} \\ issuerVk \in \text{VKey} & \text{block issuer} \\ vrfVk \in \text{VKey} & \text{VRF verification key} \\ blockNo \in \text{BlockNo} & \text{block number} \\ slot \in \text{Slot} & \text{block slot} \\ vrfRes \in \text{VRFRes} & \text{VRF result value} \\ vrfPrf \in \text{Proof} & \text{VRF proof} \\ bodySize \in \mathbb{N} & \text{size of the block body} \\ bodyHash \in \text{HashBBody} & \text{block body hash} \\ oc \in \text{OCert} & \text{operational certificate} \\ pv \in \text{ProtVer} & \text{protocol version} \end{array} \right)$
<i>Block Types</i>	
$bh \in \text{BHeader}$	$= \text{BHeader} \times \text{Sig}$
<i>Abstract functions</i>	
$\star \in \text{Nonce} \rightarrow \text{Nonce} \rightarrow \text{Nonce}$	binary nonce operation
$\text{headerHash} \in \text{BHeader} \rightarrow \text{HashHeader}$	hash of a block header
$\text{headerSize} \in \text{BHeader} \rightarrow \mathbb{N}$	size of a block header
$\text{slotToSeed} \in \text{Slot} \rightarrow \text{Seed}$	convert a slot to a seed
$\text{nonceToSeed} \in \text{Nonce} \rightarrow \text{Seed}$	convert a nonce to a seed
$\text{prevHashToNonce} \in \text{HashHeader}^? \rightarrow \text{Seed}$	convert an optional header hash to a seed
<i>Accessor Functions</i>	
$\text{blockHeader} \in \text{Block} \rightarrow \text{BHeader}$	$\text{headerBody} \in \text{BHeader} \rightarrow \text{BHeader}$
$\text{headerSig} \in \text{BHeader} \rightarrow \text{Sig}$	$\text{hBVkCold} \in \text{BHeader} \rightarrow \text{VKey}$
$\text{hBVrfVk} \in \text{BHeader} \rightarrow \text{VKey}$	$\text{hBPrevHeader} \in \text{BHeader} \rightarrow \text{HashHeader}^?$
$\text{hBSlot} \in \text{BHeader} \rightarrow \text{Slot}$	$\text{hBBlockNo} \in \text{BHeader} \rightarrow \text{BlockNo}$
$\text{hBVrfRes} \in \text{BHeader} \rightarrow \text{VRFRes}$	$\text{hBVrfPrf} \in \text{BHeader} \rightarrow \text{Proof}$
$\text{hBBodySize} \in \text{BHeader} \rightarrow \mathbb{N}$	$\text{hBBodyHash} \in \text{BHeader} \rightarrow \text{HashBBody}$
$\text{hBOC} \in \text{BHeader} \rightarrow \text{OCert}$	$\text{hBProtVer} \in \text{BHeader} \rightarrow \text{ProtVer}$

Figure 3: Block Definitions

Block Helper Functions $\text{hBLeader} \in \text{BHBody} \rightarrow \mathbb{N}$ $\text{hBLeader } bhb = \text{hash} ("L" \mid (\text{hBVrfRes } bhb))$ $\text{hBNonce} \in \text{BHBody} \rightarrow \text{Nonce}$ $\text{hBNonce } bhb = \text{hash} ("N" \mid (\text{hBVrfRes } bhb))$ **Figure 4:** Helper Functions used for Blocks

Tick Nonce environments

$$\text{TickNonceEnv} = \left(\begin{array}{ll} \eta_c \in \text{Nonce} & \text{candidate nonce} \\ \eta_{ph} \in \text{Nonce} & \text{previous header hash as nonce} \end{array} \right)$$

Tick Nonce states

$$\text{TickNonceState} = \left(\begin{array}{ll} \eta_0 \in \text{Nonce} & \text{epoch nonce} \\ \eta_h \in \text{Nonce} & \text{nonce from hash of previous epoch's last block header} \end{array} \right)$$

Figure 5: Tick Nonce types

3.2 TICKN Transition

The Tick Nonce Transition (TICKN) is responsible for updating the epoch nonce and the previous epoch's hash nonce at the start of an epoch. Its environment is shown in Figure 5 and consists of the candidate nonce η_c and the previous epoch's last block header hash as a nonce η_{ph} . Its state consists of the epoch nonce η_0 and the previous epoch's last block header hash nonce η_h .

The signal to the transition rule TICKN is a marker indicating whether we are in a new epoch. If we are in a new epoch, we update the epoch nonce and the previous hash. Otherwise, we do nothing.

$$\begin{array}{l} \text{Not-New-Epoch} \text{---} \frac{\eta_c \vdash \left(\begin{array}{l} \eta_0 \\ \eta_h \end{array} \right)}{\eta_{ph}} \xrightarrow[\text{TICKN}]{\text{False}} \left(\begin{array}{l} \eta_0 \\ \eta_h \end{array} \right) \quad (1) \\ \\ \text{New-Epoch} \text{---} \frac{\eta_c \vdash \left(\begin{array}{l} \eta_0 \\ \eta_h \end{array} \right)}{\eta_{ph}} \xrightarrow[\text{TICKN}]{\text{True}} \left(\begin{array}{l} \eta_c \star \eta_h \\ \eta_{ph} \end{array} \right) \quad (2) \end{array}$$

Figure 6: Tick Nonce rules

3.3 UPDN Transition

The Update Nonce Transition (UPDN) updates the nonces until the randomness gets fixed. The environment is shown in Figure 7 and consists of the block nonce η . The state is shown in Figure 7 and consists of the candidate nonce η_c and the evolving nonce η_v .

The transition rule UPDN takes the slot s as signal. There are two different cases for UPDN: one where s is not yet `RandomnessStabilisationWindow`¹ slots from the beginning of the next epoch and one where s is less than `RandomnessStabilisationWindow` slots until the start of the next epoch.

Note that in 3, the candidate nonce η_c transitions to $\eta_v \star \eta$, not $\eta_c \star \eta$. The reason for this is that even though the candidate nonce is frozen sometime during the epoch, we want the two nonces to again be equal at the start of a new epoch.

3.4 OCERT Transition

The Operational Certificate Transition (OCERT) environment consists of the set of stake pools *stools*. Its state is the mapping of operation certificate issue numbers. Its signal is a block header.

¹Note that in pre-Conway eras `StabilityWindow` was used instead of `RandomnessStabilisationWindow`.

Update Nonce environments

$$\text{UpdateNonceEnv} = (\eta \in \text{Nonce} \quad \text{new nonce})$$

Update Nonce states

$$\text{UpdateNonceState} = \left(\begin{array}{ll} \eta_v \in \text{Nonce} & \text{evolving nonce} \\ \eta_c \in \text{Nonce} & \text{candidate nonce} \end{array} \right)$$

Update Nonce Transitions

$$- \vdash - \xrightarrow{\text{UPDN}} - \subseteq \mathbb{P} (\text{UpdateNonceEnv} \times \text{UpdateNonceState} \times \text{Slot} \times \text{UpdateNonceState})$$

Figure 7: Update Nonce transition-system types

$$\text{Update-Both} \frac{s < \text{firstSlot} ((\text{epoch } s) + 1) - \text{RandomnessStabilisationWindow}}{\eta \vdash \left(\begin{array}{l} \eta_v \\ \eta_c \end{array} \right) \xrightarrow{\text{UPDN}} \left(\begin{array}{l} \eta_v \star \eta \\ \eta_v \star \eta \end{array} \right)} \quad (3)$$

$$\text{Only-Evolve} \frac{s \geq \text{firstSlot} ((\text{epoch } s) + 1) - \text{RandomnessStabilisationWindow}}{\eta \vdash \left(\begin{array}{l} \eta_v \\ \eta_c \end{array} \right) \xrightarrow{\text{UPDN}} \left(\begin{array}{l} \eta_v \star \eta \\ \eta_c \end{array} \right)} \quad (4)$$

Figure 8: Update Nonce rules

The transition rule OCERT is shown in Figure 10. From the block header body bhb we first extract the following:

- The operational certificate, consisting of the hot key vk_{hot} , the certificate issue number n , the KES period start c_0 and the cold key signature τ .
- The cold key vk_{cold} .
- The slot s for the block.
- The number of KES periods that have elapsed since the start period on the certificate.

Using this we verify the preconditions of the operational certificate state transition which are the following:

- The KES period of the slot in the block header body must be greater than or equal to the start value c_0 listed in the operational certificate, and less than MaxKESEvo -many KES periods after c_0 . The value of MaxKESEvo is the agreed-upon lifetime of an operational certificate, see [SL-D1].
- hk exists as key in the mapping of certificate issues numbers to a KES period m and that period is less than or equal to n . Also, n must be less than or equal to the successor of m .
- The signature τ can be verified with the cold verification key vk_{cold} .
- The KES signature σ can be verified with the hot verification key vk_{hot} .

Operational Certificate Transitions

$$- \vdash - \xrightarrow{\text{OCERT}} - \subseteq \mathbb{P} (\mathbb{P} \text{KeyHash} \times \text{KeyHash}_{\text{pool}} \mapsto \mathbb{N} \times \text{BHeader} \times \text{KeyHash}_{\text{pool}} \mapsto \mathbb{N})$$

Operational Certificate helper function

$$\begin{aligned} \text{currentIssueNo} &\in \mathbb{P} \text{KeyHash} \rightarrow (\text{KeyHash}_{\text{pool}} \mapsto \mathbb{N}) \rightarrow \text{KeyHash}_{\text{pool}} \rightarrow \mathbb{N}^? \\ \text{currentIssueNo } \text{stools } \text{cs } \text{hk} &= \begin{cases} \text{hk} \mapsto n \in \text{cs} & n \\ \text{hk} \in \text{stools} & 0 \\ \text{otherwise} & \diamond \end{cases} \quad (7) \end{aligned}$$

Figure 9: Operational Certificate transition-system types

After this, the transition system updates the operational certificate state by updating the mapping of operational certificates where it overwrites the entry of the key hk with the KES period n .

$$\begin{aligned} (bhb, \sigma) &:= bh & (vk_{\text{hot}}, n, c_0, \tau) &:= \text{hBOCert } bhb & vk_{\text{cold}} &:= \text{hBVkCold } bhb \\ hk &:= \text{hashKey } vk_{\text{cold}} & s &:= \text{hBSlot } bhb & t &:= \text{kesPeriod } s - c_0 \\ & & c_0 &\stackrel{(1)}{\leq} \text{kesPeriod } s && \stackrel{(2)}{<} c_0 + \text{MaxKESEvo} \\ \text{currentIssueNo } \text{stools } \text{cs } \text{hk} &= m & n &\in \{m, m+1\} \\ & & \mathcal{V}_{vk_{\text{cold}}} \llbracket (vk_{\text{hot}}, n, c_0) \rrbracket_{\tau} & & \mathcal{V}_{vk_{\text{hot}}}^{\text{KES}} \llbracket bhb \rrbracket_{\sigma}^t & \\ \text{OCert} \xrightarrow{\quad} & \text{stools } \vdash \text{cs} \xrightarrow{\text{OCERT}} \text{cs} \cup \{hk \mapsto n\} & & & & (5) \end{aligned}$$

Figure 10: Operational Certificate rules

The OCERT rule has the following predicate failures:

1. If the KES period is less than the KES period start in the certificate, there is a *KESBeforeStart* failure.
2. If the KES period is greater than or equal to the KES period end (start + MaxKESEvo) in the certificate, there is a *KESAfterEnd* failure.
3. If the period counter in the original key hash counter mapping is larger than the period number in the certificate, there is a *CounterTooSmall* failure.
4. If the period number in the certificate is larger than the successor of the period counter in the original key hash counter mapping, there is a *CounterOverIncremented* failure.
5. If the signature of the hot key, KES period number and period start is incorrect, there is an *InvalidSignature* failure.
6. If the KES signature using the hot key of the block header body is incorrect, there is an *InvalideKesSignature* failure.
7. If there is no entry in the key hash to counter mapping for the cold key, there is a *NoCounterForKeyHash* failure.




Figure 11: VRF helper functions

3.5 Verifiable Random Function

In this section we define a function `praosVrfChecks` which performs all the VRF related checks on a given block header body. In addition to the block header body, the function requires the epoch nonce, the stake distribution (aggregated by pool), and the active slots coefficient from the protocol parameters. The function checks:

- The validity of the proofs for the leader value and the new nonce.
- The verification key is associated with relative stake σ in the stake distribution.
- The `hBLeader` value of `bhb` indicates a possible leader for this slot. The function `checkLeaderVal` is defined in 5.

The definition of `praosVrfChecks` is shown in Figure 11 and has the following predicate failures:

1. If the VRF key is not in the pool distribution, there is a *VRFKeyUnknown* failure.
2. If the VRF key hash does not match the one listed in the block header, there is a *VRFKeyWrongVRFKey* failure.

3. If the VRF generated value in the block header does not validate against the VRF certificate, there is a *VRFKeyBadProof* failure.
4. If the VRF generated leader value in the block header is too large compared to the relative stake of the pool, there is a *VRFLeaderValueTooBig* failure.




Figure 12: Protocol transition-system types

3.6 PRTCL Transition

The Protocol Transition (PRTCL) calls the transition UPDN to update the evolving and candidate nonces, and checks the operational certificate with OCERT.

The environments for this transition are:

- The stake pool stake distribution pd .
- The epoch nonce η_0 .

The states for this transition consists of:

- The operational certificate issue number mapping.
- The evolving nonce.
- The candidate nonce for the next epoch.




Figure 13: Protocol rules

This transition establishes that a block producer is in fact authorized. Since there are three key pairs involved (cold keys, VRF keys, and hot KES keys) it is worth examining the interaction closely.

- First we check the operational certificate with OCERT. This uses the cold verification key given in the block header. We do not yet trust that this key is a registered pool key. If this

transition is successful, we know that the cold key in the block header has authorized the block.

- Next, in the `vrfChecks` predicate, we check that the hash of this cold key is in the mapping pd , and that it maps to (σ, hk_{vrf}) , where (σ, hk_{vrf}) is the hash of the VRF key in the header. If `praosVrfChecks` returns true, then we know that the cold key in the block header was a registered stake pool at the beginning of the previous epoch, and that it is indeed registered with the VRF key listed in the header.
- Finally, we use the VRF verification key in the header, along with the VRF proofs in the header, to check that the operator is allowed to produce the block.

Tick Forecast Transitions

$$\vdash \frac{}{\text{TICKF}} _ \subseteq \mathbb{P} (\text{NewEpochState} \times \text{Slot} \times \text{NewEpochState})$$

Tick Forecast helper function

$\text{adoptGenesisDelegs} \in \text{EpochState} \rightarrow \text{Slot} \rightarrow \text{EpochState}$

$\text{adoptGenesisDelegs } es \text{ slot} = es'$

where

$(acnt, ss, (us, (ds, ps)), prevPp, pp) := es$

$(rewards, delegations, ptrs, fGenDelegs, genDelegs, i_{rwd}) := ds$

$curr := \{(s, gkh) \mapsto (vkh, vrf) \in fGenDelegs \mid s \leq slot\}$

$fGenDelegs' := fGenDelegs \setminus curr$

$$genDelegs' := \left\{ gkh \mapsto (vkh, vrf) \mid \begin{array}{l} (s, gkh) \mapsto (vkh, vrf) \in curr \\ s = \max\{s' \mid (s', gkh) \in \text{dom } curr\} \end{array} \right\}$$

$ds' := (rewards, delegations, ptrs, fGenDelegs', genDelegs \sqcup genDelegs', i_{rwd})$

$es' := (acnt, ss, (us, (ds', ps)), prevPp, pp)$

Figure 14: Tick Forecast transition-system types

3.7 TICKF Transition

The Tick Forecast Transition (TICKF) performs some chain level upkeep. The state is the epoch specific state necessary for the NEWEPOCH transition.

Part of the upkeep is updating the genesis key delegation mapping according to the future delegation mapping. For each genesis key, we adopt the most recent delegation in $fGenDelegs$ that is past the current slot, and any future genesis key delegations past the current slot is removed. The helper function $\text{adoptGenesisDelegs}$ accomplishes the update.

The TICKF transition rule is shown in Figure 15. The signal is a slot s .

One sub-transition is done: The NEWEPOCH transition performs any state change needed if it is the first block of a new epoch.

$$\begin{array}{c} \vdash nes \xrightarrow[\text{NEWEPOCH}]{\text{epoch } s} nes' \\ (e'_\ell, b'_{prev}, b'_{cur}, es', ru', pd') := nes' \\ es'' := \text{adoptGenesisDelegs } es' \ s \\ forecast := (e'_\ell, b'_{prev}, b'_{cur}, es'', ru', pd') \\ \text{TickForecast} \text{---} \\ \vdash nes \xrightarrow[\text{TICKF}]{s} forecast \end{array} \quad (7)$$

Figure 15: Tick Forecast rules

3.8 CHAINHEAD Transition

The Chain Head Transition rule (CHAINHEAD) is the main rule of the blockchain layer part of the STS. It calls TICKF, TICKN, and PRTCL, as sub-rules.

The transition checks the following things: (via chainChecks and prtSeqChecks from Figure 17):

- The slot in the block header body is larger than the last slot recorded.
- The block number increases by exactly one.
- The previous hash listed in the block header matches the previous block header hash which was recorded.
- The size of the block header is less than or equal to the maximal size that the protocol parameters allow for block headers.
- The size of the block body, as claimed by the block header, is less than or equal to the maximal size that the protocol parameters allow for block bodies.
- The node is not obsolete, meaning that the major component of the protocol version in the protocol parameters is not bigger than the constant MaxMajorPV.

The CHAINHEAD state is shown in Figure 16, it consists of the following:

- The operational certificate issue number map cs .
- The epoch nonce η_0 .
- The evolving nonce η_v .
- The candidate nonce η_c .
- The previous epoch hash nonce η_h .
- The last header hash h .
- The last slot s_ℓ .
- The last block number b_ℓ .

The CHAINHEAD transition rule is shown in Figure 18. It contains a new epoch state nes in the environment and its signal is a block header bh . The transition uses a few helper functions defined in Figure 17.

The CHAINHEAD transition rule has the following predicate failures:

1. If the slot of the block header body is not larger than the last slot, there is a *WrongSlotInterval* failure.
2. If the block number does not increase by exactly one, there is a *WrongBlockNo* failure.
3. If the hash of the previous header of the block header body is not equal to the last header hash, there is a *WrongBlockSequence* failure.
4. If the size of the block header is larger than the maximally allowed size, there is a *Header-SizeTooLarge* failure.
5. If the size of the block body is larger than the maximally allowed size, there is a *BlockSize-TooLarge* failure.
6. If the major component of the protocol version is larger than MaxMajorPV, there is a *ObsoleteNode* failure.

Chain Head states

$$\text{LastAppliedBlock} = \begin{pmatrix} b_\ell \in \text{BlockNo} & \text{last block number} \\ s_\ell \in \text{Slot} & \text{last slot} \\ h \in \text{HashHeader} & \text{latest header hash} \end{pmatrix}$$

$$\text{ChainHeadState} = \begin{pmatrix} cs \in \text{KeyHash}_{\text{pool}} \mapsto \mathbb{N} & \text{operational certificate issue numbers} \\ \eta_0 \in \text{Nonce} & \text{epoch nonce} \\ \eta_v \in \text{Nonce} & \text{evolving nonce} \\ \eta_c \in \text{Nonce} & \text{candidate nonce} \\ \eta_h \in \text{Nonce} & \text{nonce from hash of last epoch's last header} \\ lab \in \text{LastAppliedBlock}^? & \text{latest applied block} \end{pmatrix}$$

Chain Head Transitions

$$_ \vdash _ \xrightarrow{\text{CHAINHEAD}} _ \subseteq \mathbb{P} (\text{NewEpochState} \times \text{ChainHeadState} \times \text{BHeader} \times \text{ChainHeadState})$$

Figure 16: Chain Head transition-system types

Chain Head Transition Helper Functions

$$\text{chainChecks} \in \mathbb{N} \rightarrow (\mathbb{N}, \mathbb{N}, \text{ProtVer}) \rightarrow \text{BHeader} \rightarrow \text{Bool}$$

$$\text{chainChecks } \text{maxpv} (\text{maxBHSize}, \text{maxBBSize}, \text{protocolVersion}) \text{ bh} =$$

$$m \stackrel{(6)}{\leq} \text{maxpv}$$

$$\wedge \text{headerSize } \text{bh} \stackrel{(4)}{\leq} \text{maxBHSize}$$

$$\wedge \text{hBBodySize} (\text{headerBody } \text{bh}) \stackrel{(5)}{\leq} \text{maxBBSize}$$

$$\textbf{where } (m, _) := \text{protocolVersion}$$

$$\text{lastAppliedHash} \in \text{LastAppliedBlock}^? \rightarrow \text{HashHeader}^?$$

$$\text{lastAppliedHash } \text{lab} = \begin{cases} \diamond & \text{lab} = \diamond \\ h & \text{lab} = (_, _, h) \end{cases}$$

$$\text{prtISeqChecks} \in \text{LastAppliedBlock}^? \rightarrow \text{BHeader} \rightarrow \text{Bool}$$

$$\text{prtISeqChecks } \text{lab } \text{bh} = \begin{cases} \text{True} & \text{lab} = \diamond \\ s_\ell \stackrel{(1)}{<} \text{slot} \wedge b_\ell + 1 \stackrel{(2)}{=} \text{bn} \wedge \text{ph} \stackrel{(3)}{=} \text{hBPrevHeader } \text{bhb} & \text{lab} = (b_\ell, s_\ell, _) \end{cases}$$

where

$$\text{bhb} := \text{headerBody } \text{bh}$$

$$\text{bn} := \text{hBBlockNo } \text{bhb}$$

$$\text{slot} := \text{hBSlot } \text{bhb}$$

$$\text{ph} := \text{lastAppliedHash } \text{lab}$$

Figure 17: Helper functions used in the Chain Head transition

$$\begin{array}{c}
bhb := \text{headerBody } bh \qquad s := \text{hBSlot } bhb \\
\\
\text{prt!SeqChecks } lab \ bh \\
\vdash nes \xrightarrow[\text{TICKF}]{s} \text{forecast} \\
(e_1, _ _ _ _ _) := nes \\
(e_2, _ _ _ es, _ _ pd) := \text{forecast} \\
(_ _ _ _ _ pp) := es \\
ne := e_1 \neq e_2 \\
\eta_{ph} := \text{prevHashToNonce } (\text{lastAppliedHash } lab) \\
\\
\text{chainChecks MaxMajorPV } (\text{maxHeaderSize } pp, \text{maxBlockSize } pp, pv \ pp) \ bh \\
\\
\eta_c \quad \eta_{ph} \vdash \begin{pmatrix} \eta_0 \\ \eta_h \end{pmatrix} \xrightarrow[\text{TICKN}]{ne} \begin{pmatrix} \eta'_0 \\ \eta'_h \end{pmatrix} \\
\\
pd \quad \eta'_0 \vdash \begin{pmatrix} cs \\ \eta_v \\ \eta_c \end{pmatrix} \xrightarrow[\text{PRTCL}]{bh} \begin{pmatrix} cs' \\ \eta'_v \\ \eta'_c \end{pmatrix} \\
\\
\text{ChainHead} \text{---} lab' := (\text{hBBlockNo } bhb, s, \text{headerHash } bh) \qquad (8) \\
\\
nes \vdash \begin{pmatrix} cs \\ \eta_0 \\ \eta_v \\ \eta_c \\ \eta_h \\ lab \end{pmatrix} \xrightarrow[\text{CHAINHEAD}]{bh} \begin{pmatrix} cs' \\ \eta'_0 \\ \eta'_v \\ \eta'_c \\ \eta'_h \\ lab' \end{pmatrix}
\end{array}$$

Figure 18: Chain Head rules

4 Properties

This section describes the properties that the consensus layer should have. The goal is to include these properties in the executable specification to enable e.g. property-based testing or formal verification.

4.1 Header-Only Validation

In any given chain state, the consensus layer needs to be able to validate the block headers without having to download the block bodies. Property 4.1 states that if an extension of a chain that spans less than `StabilityWindow` slots is valid, then validating the headers of that extension is also valid. This property is useful for its converse: if the header validation check for a sequence of headers does not pass, then we know that the block validation that corresponds to those headers will not pass either. In these properties, we refer to the CHAIN transition system as defined in [SL-D5].

Property 4.1 (Header only validation). For all states s with slot number t^2 , and chain extensions E with corresponding headers H such that:

$$0 \leq t_E - t \leq \text{StabilityWindow}$$

we have:

$$\vdash s \xrightarrow[\text{CHAIN}]{E}^* s' \implies nes \vdash \tilde{s} \xrightarrow[\text{CHAINHEAD}]{H}^* \tilde{s}'$$

where $s = (nes, \tilde{s})$, t_E is the maximum slot number appearing in the blocks contained in E , and H is obtained from E by applying `blockHeader` to each block in E .

Property 4.2 (Body only validation). For all states s with slot number t , and chain extensions $E = [b_0, \dots, b_n]$ with corresponding headers $H = [h_0, \dots, h_n]$ such that:

$$0 \leq t_E - t \leq \text{StabilityWindow}$$

we have that for all $i \in [1, n]$:

$$nes \vdash \tilde{s} \xrightarrow[\text{CHAINHEAD}]{H}^* s_h \wedge \vdash (nes, \tilde{s}) \xrightarrow[\text{CHAIN}]{[b_0 \dots b_{i-1}]}^* s_{i-1} \implies nes' \vdash \tilde{s}_{i-1} \xrightarrow[\text{CHAINHEAD}]{h_i}^* s'_h$$

where $s = (nes, \tilde{s})$, $s_{i-1} = (nes', \tilde{s}_{i-1})$, t_E is the maximum slot number appearing in the blocks contained in E .

Property 4.2 states that if we validate a sequence of headers, we can validate their bodies independently and be sure that the blocks will pass the chain validation rule. To see this, given an environment e and initial state s , assume that a sequence of headers $H = [h_0, \dots, h_n]$ corresponding to blocks in $E = [b_0, \dots, b_n]$ is valid according to the CHAINHEAD transition system:

$$nes \vdash \tilde{s} \xrightarrow[\text{CHAINHEAD}]{H}^* \tilde{s}'$$

Assume the bodies of E are valid according to the BBODY rules (defined in [SL-D5]), but E is not valid according to the CHAIN rule. Assume that there is a $b_j \in E$ such that it is **the first block** such that does not pass the CHAIN validation. Then:

$$\vdash (nes, \tilde{s}) \xrightarrow[\text{CHAIN}]{[b_0 \dots b_{j-1}]}^* s_j$$

²i.e. the component s_ℓ of the last applied block of s equals t

But by Property 4.2 we know that

$$nes_j \vdash \tilde{s}_j \xrightarrow[\text{CHAINHEAD}]{h_j} \tilde{s}_{j+1}$$

which means that block b_j has valid headers, and this in turn means that the validation of b_j according to the chain rules must have failed because it contained an invalid block body. But this contradicts our assumption that the block bodies were valid.

Values associated with the leader value calculations

$certNat \in \{n n \in \mathbb{N}, n \in [0, 2^{512}]\}$	Certified natural value from VRF
$f \in [0, 1]$	Active slot coefficient
$\sigma \in [0, 1]$	Stake proportion

5 Leader Value Calculation

This section details how we determine whether a node is entitled to lead (under the Praos protocol) given the output of its verifiable random function calculation.

5.1 Computing the leader value

The verifiable random function gives us a 64-byte random output. We interpret this as a natural number $certNat$ in the range $[0, 2^{512})$.

5.2 Node eligibility

As per [DGKR17], a node is eligible to lead when its leader value $p < 1 - (1 - f)^\sigma$. We have

$$p < 1 - (1 - f)^\sigma$$

$$\iff \left(\frac{1}{1 - p} \right) < \exp(-\sigma \cdot \ln(1 - f))$$

The latter inequality can be efficiently computed through use of its Taylor expansion and error estimation to stop computing terms once we are certain that the result will be either above or below the target value.

We carry out all computations using fixed precision arithmetic (specifically, we use 34 decimal bits of precision, since this is enough to represent the fraction of a single lovelace.)

As such, we define the following:

$$p = \frac{certNat}{2^{512}}$$

$$q = 1 - p$$

$$c = \ln(1 - f)$$

and define the function $checkLeaderVal$ as follows:

$$checkLeaderVal \ certNat \ \sigma \ f = \begin{cases} \text{True}, & f = 1 \\ \frac{1}{q} < \exp(-\sigma \cdot c), & \text{otherwise} \end{cases}$$

References

- [BC-D1] IOHK Formal Methods Team. Byron Blockchain Specification, IOHK Deliverable BC-D1, 2019. URL <https://github.com/intersectmbo/cardano-ledger/tree/master/docs/>.
- [SL-D5] IOHK Formal Methods Team. A Formal Specification of the Cardano Ledger, 2019. URL <https://github.com/intersectmbo/cardano-ledger/releases/latest/download/shelley-ledger.pdf>
- [DGKR17] B. M. David, P. Gazi, A. Kiayias, and A. Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake protocol. *IACR Cryptology ePrint Archive*, 2017:573, 2017.
- [SL-D1] IOHK Formal Methods Team. Design Specification for Delegation and Incentives in Cardano, IOHK Deliverable SL-D1, 2018. URL <https://github.com/intersectmbo/cardano-ledger/releases/latest/download/shelley-delegation.pdf>.

A Cryptographic Details

A.1 Abstract functions

- The nonce operation $x \star y$ from Figure 3 is implemented as the BLAKE2b-256 hash of the concatenation of x and y .
- The functions `slotToSeed` and `nonceToSeed` from Figure 3 are implemented as the big-endian encoding of the slot/nonce number in 8 bytes.