

Hydra HeadV1 Specification: Coordinated Head protocol

DRAFT

Sebastian Nagel `sebastian.nagel@iohk.io`

Sasha Bogicevic `sasha.bogicevic@iohk.io`

Franco Testagrossa `franco.testagrossa@iohk.io`

Daniel Firth `daniel.firth@iohk.io`

April 10, 2025

Contents

1	Introduction	1
2	Protocol Overview	2
2.1	Opening the head	2
2.2	The Coordinated Head protocol	3
2.3	Closing the head	4
2.4	Differences	4
3	Preliminaries	5
3.1	Notation	5
3.2	Public key multi-signature scheme	6
3.3	Extended UTxO	6
4	Protocol Setup	9
5	On-chain Protocol	10
5.1	Init transaction	10
5.2	Commit Transaction	12
5.3	Abort Transaction	13
5.4	CollectCom Transaction	14
5.5	Deposit Transaction	15
5.6	Recover Transaction	16
5.7	Increment Transaction	18
5.8	Decrement Transaction	19
5.9	Close Transaction	20
5.10	Contest Transaction	23
5.11	Fan-Out Transaction	25
6	Off-Chain Protocol	26
6.1	Assumptions	27
6.2	Notation	28
6.3	Variables	28
6.4	Protocol flow	29
6.5	Rollbacks and protocol changes	31
7	Security (WIP — Iteration 1)	34
7.1	Proofs	35

1 Introduction

This document specifies the 'Coordinated Hydra Head' protocol to be implemented as the first version of Hydra Head on Cardano - **Hydra HeadV1**. The protocol is derived from variants described in the original paper [8], but was further simplified to make a first implementation on Cardano possible.

Note that the format and scope of this document is (currently) also inspired by the paper and hence does not include a definition of the networking protocols or concrete message formats. It is structured similarly, but focuses on a single variant, and avoids indirections and unnecessary generalizations. The document is kept in sync with the reference implementation available on Github [4]. Red sections indicate that they are currently not covered or missing in the implementation, where blue parts mean a difference in how it is realized.

Add:
network
specifi-
cation
(message
formats)

First, a high-level overview of the protocol and how it differs from legacy variants of the Head protocol is given in Section 2. Relevant definitions and notations are introduced in Section 3, while Section 4 describes protocol setup and assumptions. Then, the actual on-chain transactions of the protocol are defined in Section 5, before the off-chain protocol part specifies behavior of Hydra parties off-chain and ties the knot with on-chain transactions in Section 6. At last, Section 7 gives the security definition, properties and proofs for the Coordinated Head protocol.

2 Protocol Overview

The Hydra Head protocol provides functionality to lock a set of UTXOs on a blockchain, referred to as the *mainchain*, and evolve it inside a so-called off-chain *head*, independently of the mainchain. At any point, the head can be closed with the effect that the locked set of UTXOs on the mainchain is replaced by the latest set of UTXOs inside the head. The protocol guarantees full wealth preservation: no generation of funds can happen off-chain (inside a head) and no responsive honest party involved in a head can ever lose any funds other than by consenting to give them away. In exchange for decreased liveness guarantees (stop any time), it can essentially proceed at network speed under good conditions, thereby reducing latency and increasing throughput. At the same time, the head protocol provides the same capabilities as the mainchain by reusing the same ledger model and transaction formats — making the protocol “isomorphic”.

2.1 Opening the head

To create a head-protocol instance, any party may take the role of an *initiator* and ask other parties, the *head members*, to participate in the head by exchanging public keys and agreeing on other protocol parameters. These public keys are used for both, the authentication of head-related on-chain transactions that are restricted to head members (e.g., a non-member is not allowed to close the head) and for signing off-chain transactions in the head.

The initiator then establishes the head by submitting an *initial* transaction to the mainchain that contains the Hydra protocol parameters and mints special *participation tokens* (*PT*) identifying the head members. The *initial* transaction also initializes a state machine (see Fig. 1) that manages the “transfer” of UTXOs into the head and back. The state machine comprises the four states: *initial*, *open*, *closed*, and *final*. A *state thread token* (*ST*) minted in *initial* marks the head output and ensures contract continuity [6].

Once the initial transaction appears on the mainchain, establishing the initial state *initial*, each head member can attach a *commit* transaction, which locks (on the mainchain) the UTXOs that the party wants to commit to the head, or deliberately acknowledges to commit nothing.

The commit transactions are subsequently collected by the *collectCom* transaction causing a transition from *initial* to *open*. Once the *open* state is confirmed, the head members start running the off-chain head protocol, which evolves the initial UTXO set (the union over all UTXOs committed




Figure 1: Mainchain state diagram for this version of the Hydra protocol.

by all head members) independently of the mainchain. For the case where some head members fail to post a *commit* transaction, the head can be aborted by going directly from initial to final.

2.2 The Coordinated Head protocol

The actual Head protocol starts after the initialization phase with an initial set of UTxOs that is identical to the UTxOs locked on-chain via the *commit* and *collectCom* transactions.

The protocol processes off-chain transactions by distributing them between participants, while each party maintains their view of the local UTxO state. That is, the current set of UTxOs evolved from the initial UTxO set by applying transactions as they are received from the other parties.

To confirm transactions and allow for distribution of the resulting UTxO set without needing the whole transaction history, snapshots are created by the protocol participants. The initial snapshot U_0 corresponds to the initial UTxO set, while snapshots thereafter U_1, U_2, \dots are created with monotonically increasing snapshot numbers.

For this, the next snapshot leader (round-robin) requests his view of a new confirmed state to be signed by all participants as a new snapshot. The leader does not need to send his local state, but only indicate, by hashes, the set of transactions to be included in order to obtain the to-be-snapshotted UTxO set.

The other participants sign the snapshot as soon as they have (also) seen the transactions that are to be processed on top of its preceding snapshot: a party's local state is always ahead of the latest confirmed snapshot.

Signatures are broadcast and aggregated by each party. When all signature parts of the multi-signature are received and verified, a snapshot is considered confirmed. As a consequence, a participant can safely delete (if wished) all transactions that have been processed into it as the snapshot's multi-signature is now evidence that this state once existed during the head evolution.

2.2.1 Updating an open head

Besides processing “normal” transactions, participants can also request to withdraw some UTxO they can spend from the Head and make it available on main chain via a *decrement* 5.8 transaction — the overall process is also called “decommit”.

The inverse of locking more UTxO on the mainchain and making it available within the head is called (incremental) “commit”. First, anyone may create a deposit of one or more UTxO using a *deposit* 5.5 transaction. The head participants will observe this deposit and, once settled, request off-chain agreement to include the deposited UTxO in the form of a snapshot. With that agreement, an *increment* 5.7 transaction can be created and used to update the head state on the mainchain. A deadline is associated with the deposit, which ensures that the UTxO is locked up long enough to be safely consumed into the head without risk of double spending. Should a deposit have not been picked up in time, a *recover* 5.6 transaction allows anyone to unlock the original UTxO.

2.3 Closing the head

The head protocol is designed to allow any head member at any point in time to produce, without interaction, a certificate to close the head. This certificate is created from the latest confirmed, multi-signed snapshot. Using this certificate, the head member may “force close” the head by advancing the mainchain state machine to the *closed* state.

Once in *closed*, the state machine grants parties a contestation period, during which parties may contest the closure by posting the certificate of a newer snapshot on-chain in a contest transaction. Contesting leads back to the state *closed* and each party can contest at most once. After the contestation period has elapsed, the state machine may proceed to the *final* state. The state machine enforces that the outputs of the transaction leading to *final* correspond exactly to the latest UTxO set seen during the contestation period.

2.4 Differences

In the Coordinated Head protocol, off-chain consensus is simplified by not having transactions confirmed concurrently to the snapshots (and to each other) but having the snapshot leader propose, in their snapshot, a set of transactions for explicit confirmation. The parties’ views of confirmed transactions thus progress in sync with each other (once per confirmed snapshot), thus simplifying the close/contest procedure on the mainchain. Also, there is no need for conflict resolution as in Appendix B of [8]. In summary, the differences to the original Head protocol in [8] are:

- No hanging transactions due to ‘coordination’.
- No acknowledgement nor confirmation of transactions.
- No confirmation message for snapshots (two-round local confirmation).
- Allow for incremental commits and decommits while head is open.

3 Preliminaries

This section introduces notation and other preliminaries used in the remainder of the specification.

3.1 Notation

The specification uses set-notation based approach while also inspired by [1] and [6]. Values a are in a set $a \in \mathcal{A}$, also indicated as being of some type $a : \mathcal{A}$, and multidimensional values are tuples drawn from a \times product of multiple sets, e.g. $(a, b) \in (\mathcal{A} \times \mathcal{B})$. An empty set is indicated by \emptyset and sets may be enumerated using $\{a_1 \dots a_n\}$ notation. The $=$ operator means equality and \leftarrow is explicit assignment of a variable or value to one or more variables. Projection is used to access the elements of a tuple, e.g. $(a, b)^{\downarrow 1} = a$. Functions are morphisms mapping from one set to another $x : \mathcal{A} \rightarrow f(x) : \mathcal{B}$, where function application of a function f to an argument x is written as $f(x)$.

Furthermore, given a set \mathcal{A} , let

- $\mathcal{A}^? = \mathcal{A} \cup \diamond$ denotes an option: a value from \mathcal{A} or no value at all indicated by \perp ,
- \mathcal{A}^n be the set of all n-sized sequences over \mathcal{A} ,
- $\mathcal{A}^! = \bigcup_{i=1}^{n \in \mathbb{N}} \mathcal{A}^i$ be the set of non-empty sequences over \mathcal{A} , and
- $\mathcal{A}^* = \bigcup_{i=0}^{n \in \mathbb{N}} \mathcal{A}^i$ be the set of all sequences over \mathcal{A} .

With this, we further define:

- $\mathbb{B} = \{\text{false}, \text{true}\}$ are boolean values
- \mathbb{N} are natural numbers $\{0, 1, 2, \dots\}$
- \mathbb{Z} are integer numbers $\{\dots, -2, -1, 0, 1, 2, \dots\}$
- $\mathbb{H} = \bigcup_{n=0}^{\text{inf}} \{0, 1\}^{8n}$ denotes a arbitrary string of bytes
- $\text{concat} : \mathbb{H}^* \rightarrow \mathbb{H}$ is concatenating bytes, we also use operator \oplus for this
- $\text{hash} : x \rightarrow \mathbb{H}$ denotes a collision-resistant hashing function and $x^\#$ indicates the hash of x
- $\text{bytes} : x \rightarrow \mathbb{H}$ denotes an invertible serialisation function mapping arbitrary data to bytes
- $a||b = \text{concat}(\text{bytes}(a), \text{bytes}(b))$ is an operator which concatenates the $\text{bytes}(b)$ to the $\text{bytes}(a)$
- Lists of values $l \in \mathcal{A}^*$ are written as $l = [x_1, \dots, x_n]$. Empty lists are denoted by $[],$ the i th element x_i is also written $l[i]$ and the length of the list is $|l| = n$. An underscore is also used to indicate a list of values $\underline{x} = l$. Projection on lists are mapped to their elements, i.e. $\underline{x}^{\downarrow 1} = [x_1^{\downarrow 1}, \dots, x_n^{\downarrow 1}]$.
- $\text{sortOn} : i \rightarrow \mathcal{A}^* \rightarrow \mathcal{A}^*$ does sort a list of values on the i th projection.
- **Data** is a universal data type of nested sums and products built up recursively from the base types of \mathbb{Z} and \mathbb{H} .

3.2 Public key multi-signature scheme

A multisignature scheme is a set of algorithms where

- MS-Setup generates public parameters Π , such that
- $(k^{ver}, k^{sig}) \leftarrow \text{MS-KG}(\Pi)$ can be used to generate fresh key pairs,
- $\sigma \leftarrow \text{MS-Sign}(\Pi, k^{sig}, m)$ signs a message m using key k^{sig} ,
- $\tilde{k} \leftarrow \text{MS-AVK}(\Pi, \bar{k})$ aggregates a list of verification keys \bar{k} into a single, aggregate key \tilde{k} ,
- $\tilde{\sigma} \leftarrow \text{MS-ASig}(\Pi, m, \bar{k}, \bar{\sigma})$ aggregates a list of signatures $\bar{\sigma}$ about message m into a single, aggregate signature $\tilde{\sigma}$.
- $\text{MS-Verify}(\Pi, \tilde{k}, m, \tilde{\sigma}) \in \mathbb{B}$ verifies an aggregate signature $\tilde{\sigma}$ of message m under an aggregate verification key \tilde{k} .

The security definition of a multisignature scheme from [9, 10] guarantees that, if \tilde{k} is produced from a tuple of verification keys \bar{k} via MS-AVK, then no aggregate signature $\tilde{\sigma}$ can pass verification $\text{MS-Verify}(\tilde{k}, m, \tilde{\sigma})$ unless all honest parties holding keys in \bar{k} signed m .

Note that in the following, we make the parameter Π implicit and leave out the *ver* suffix for verification key such that $k = k^{ver}$ for better readability.

3.3 Extended UTxO

The Hydra Head protocol is specified to work on the so-called Extended UTxO (EUTxO) ledgers like Cardano.

The basis for EUTxO is Bitcoin’s UTxO ledger model [5, 11]. Intuitively, it arranges transactions in a directed acyclic graph, such as the one in Figure 2, where boxes represent transactions with (red) inputs to the left and (black) outputs to the right. A dangling (unconnected) output is an *unspent transaction output (UTxO)* — there are two UTxOs in the figure.




Figure 2: Example of a plain UTxO graph

The following paragraphs will give definitions of the UTxO model and its extension to support scripting (EUTxO) suitable for this Hydra Head protocol specification. For a more detailed introduction to the EUTxO ledger model, see [6], [1] and [7].

3.3.1 Values

Definition 1 (Values). *Values are sets that keep track of how many units of which tokens of which currency are available. Given a finitely supported function \mapsto , that maps keys to monoids, a value is the set of such mappings over currencies (minting policy identifiers), over a mapping of token names t to quantities q :*

$$\text{val} \in \text{Val} = (c : \mathbb{H} \mapsto (t : \mathbb{H} \mapsto q : \mathbb{Z}))$$

where addition of values is defined as $+$ and \emptyset is the empty value.

For example, the value $\{c_1 \mapsto \{t_1 \mapsto 1, t_2 \mapsto 1\}\}$ contains tokens t_1 and t_2 of currency c_1 and addition merges currencies and token names naturally:

$$\begin{aligned} & \{c_1 \mapsto \{t_1 \mapsto 1, t_2 \mapsto 1\}\} \\ + & \{c_1 \mapsto \{t_2 \mapsto 1, t_3 \mapsto 1\}, c_2 \mapsto \{t_1 \mapsto 2\}\} \\ = & \{c_1 \mapsto \{t_1 \mapsto 1, t_2 \mapsto 2, t_3 \mapsto 1\}, c_2 \mapsto \{t_1 \mapsto 2\}\}. \end{aligned}$$

While the above definition should be sufficient for the purpose of this specification, a full definition for finitely supported functions and values as used here can be found in [7]. To further improve readability, we define the following shorthands:

- $\{t_1, \dots, t_n\} :: c$ for a set positive single quantity assets $\{c \mapsto \{t_1 \mapsto 1, \dots, t_n \mapsto 1\}\}$,
- $\{t_1, \dots, t_n\}^{-1} :: c$ for a set of negative single quantity assets $\{c \mapsto \{t_1 \mapsto -1, \dots, t_n \mapsto -1\}\}$,
- $\{c \mapsto t \mapsto q\}$ for the value entry $\{c \mapsto \{t \mapsto q\}\}$,
- $\{c \mapsto \cdot \mapsto q\}$ for any asset with currency c and quantity q irrespective of token name.

3.3.2 Scripts

Validator scripts are called *phase-2* scripts in the Cardano Ledger specification (see [3] for a formal treatment of these). Scripts are used for multiple purposes, but most often (and sufficient for this specification) as a *spending* or *minting* policy script.

Definition 2 (Minting Policy Script). *A script $\mu \in \mathcal{M}$ governing whether a value can be minted (or burned), is a pure function with type*

$$\mu \in \mathcal{M} = (\rho : \text{Data}) \rightarrow (\gamma : \Gamma) \rightarrow \mathbb{B},$$

where $\rho \in \text{Data}$ is the redeemer provided as part of the transaction being validated and $\gamma \in \Gamma$ is the validation context.

Definition 3 (Spending Validator Script). *A validator script $\nu \in \mathcal{V}$ governing whether an output can be spent, is a pure function with type*

$$\nu \in \mathcal{V} = (\delta : \text{Data}) \rightarrow (\rho : \text{Data}) \rightarrow (\gamma : \Gamma) \rightarrow \mathbb{B},$$

where $\delta \in \text{Data}$ is the datum available at the output to be spent, $\rho \in \text{Data}$ is the redeemer data provided as part of the transaction being validated, and $\gamma \in \Gamma$ is the validation context.

3.3.3 Transactions

We define EUTxO inputs, outputs and transactions as they are available to scripts and just enough to specify the behavior of the Hydra validator scripts. For example outputs addresses and datums are much more complicated in the full ledger model [1, 2].

actual transactions \mathcal{T} are not defined

Definition 4 (Outputs). An output $o \in \mathcal{O}$ stores some value $\text{val} \in \text{Val}$ at some address, defined by the hash of a validator script $\nu^\# \in \mathbb{H} = \text{hash}(\nu \in \mathcal{V})$, and may store (reference) some data $\delta \in \text{Data}$:

$$o \in \mathcal{O} = (\text{val} : \text{Val} \times \nu^\# : \mathbb{H} \times \delta : \text{Data})$$

Definition 5 (Output references). An output reference $\phi \in \Phi$ points to an output of a transaction, using a transaction id (that is, a hash of the transaction body) and the output index within that transaction.

$$\phi \in \Phi = (\mathbb{H} \times \mathbb{N})$$

Definition 6 (Inputs). A transaction input $i \in \mathcal{J}$ is an output reference $\phi \in \Phi$ with a corresponding redeemer $\rho \in \text{Data}$:

$$i \in \mathcal{J} = (\phi : \Phi \times \rho : \text{Data})$$

Definition 7 (Validation Context). A validation context $\gamma \in \Gamma$ is a view on the transaction to be validated:

$$\gamma \in \Gamma = (\mathcal{J}^* \times \mathcal{O}^* \times \text{Val} \times \mathcal{S}^{\leftrightarrow} \times \mathcal{K})$$

where $\mathcal{J} \in \mathcal{J}^*$ is a **set** of inputs, $\mathcal{O} \in \mathcal{O}^*$ is a **list** of outputs, $\text{mint} \in \text{Val}$ is the minted (or burned) value, $(t_{\min}, t_{\max}) \in \mathcal{S}^{\leftrightarrow}$ are the lower and upper validity bounds where $t_{\min} \leq t_{\max}$, and $\kappa \in \mathcal{K}$ is the set of verification keys which signed the transaction.

Informally, scripts are evaluated by the ledger when it applies a transaction to its current state to yield a new ledger state (besides checking the transaction integrity, signatures and ledger rules). Each validator script referenced by an output is passed its arguments drawn from the output it locks and the transaction context it is executed in. The transaction is valid if and only if all scripts validate, i.e. $\mu(\rho, \gamma) = \text{true}$ and $\nu(\delta, \rho, \gamma) = \text{true}$.

3.3.4 State machines and graphical notation

State machines in the EUTxO ledger model are commonly described using the *constraint emitting machine (CEM)* formalism [6], e.g. the original paper describes the Hydra Head protocol using this notation [8]. Although inspired by CEMs, this specification uses a more direct representation of individual transactions to simplify description of non-state-machine transactions and help translation to concrete implementations on Cardano. The structure of the state machine is enforced on-chain through *scripts* which run as part of the ledger's validation of a transaction (see Section 3.3). For each protocol transaction, the specification defines the structure of the transaction and enumerates the transaction constraints enforced by the scripts (tx^\equiv in the CEM formalism).

Add an example graph with a legend

4 Protocol Setup

In order to create a head-protocol instance, an initiator invites a set of participants (the initiator being one of them) to join by announcing to them the protocol parameters.

- For on-chain transaction authentication (Cardano) purposes, each party p_i generates a corresponding key pair (k_i^{ver}, k_i^{sig}) and sends their verification key k_i^{ver} to all other parties. In the case of Cardano, these are Ed25519 keys.
- For off-chain signing (Hydra) purposes, a very basic multisignature scheme (MS, as defined in Section 3.2) based on EdDSA using Ed25519 keys is used:
 - MS-KG is Ed25519 key generation (requires no parameters)
 - MS-Sign creates an EdDSA signature
 - MS-AVK is concatenation of verification keys into an ordered list
 - MS-ASig is concatenation of signatures into an ordered list
 - MS-Verify verifies the "aggregate" signature by verifying each individual EdDSA signature under the corresponding Ed25519 verification key

To help distinguish on- and off-chain key sets, Cardano verification keys are written k_C , while Hydra verification keys are indicated as k_H for the remainder of this document.

- Each party p_i generates a hydra key pair and sends their hydra verification key to all other parties.
- Each party p_i computes the aggregate key from the received verification keys, stores the aggregate key, their signing key as well as the number of participants n .
- Each party establishes pairwise communication channels to all other parties. That is, every network message received from a specific party is checked for (channel) authentication. It is the implementer's duty to find a suitable authentication process for the communication channels.
- All parties agree on a contestation period T .

If any of the above fails (or the party does not agree to join the head in the first place), the party aborts the initiation protocol and ignores any further action. Finally, at least one of the participants posts the *init* transaction onchain as described next in Section 5.

5 On-chain Protocol

The following sections describe the the *on-chain* protocol controlling the life-cycle of a Hydra head, which can be intuitively described as a state machine (see Figure 1). Each transition in this state machine is represented and caused by a corresponding Hydra protocol transaction on-chain: *init* 5.1, *commit* 5.2, *abort* 5.3, *collectCom* 5.4, *increment* 5.7, *decrement* 5.8, *close* 5.9, *contest* 5.10, and *fanout* 5.11.

Besides the main state transitions of the head protocol, there is the related “deposit protocol” with two transactions in support of *increment*: *deposit* 5.5 and *recover* 5.6.

The head protocol defines one minting policy script and three validator scripts:

- μ_{head} governs minting of state and participation tokens in *init* and burning of these tokens in *abort* and *fanout*.
- ν_{initial} controls how UTXOs are committed to the head in *commit* or when the head initialization is aborted via *abort*.
- ν_{commit} controls the collection of committed UTXOs into the head in *collectCom* or that funds are reimbursed in an *abort*.
- ν_{head} represents the main protocol state machine logic and ensures contract continuity throughout *collectCom*, *decrement*, *increment*, *close*, *contest* and *fanout*.

The deposit protocol defines one validator script:

- ν_{deposit} controls that *deposit* transaction output is claimed correctly into a head via *increment* or recovered after the deadline has passed in a *recover* transaction.

5.1 Init transaction

The *init* transaction creates a head instance and establishes the initial state of the protocol and is shown in Figure 3. The head instance is represented by the unique currency identifier *cid* created by minting tokens using the μ_{head} minting policy script which is parameterized by a single output reference parameter $\phi_{\text{seed}} \in \Phi$:

$$\text{cid} = \text{hash}(\mu_{\text{head}}(\phi_{\text{seed}}))$$

Two kinds of tokens are minted:

- A single *State Thread (ST)* token marking the head output. This output contains the state of the protocol on-chain and the token ensures contract continuity. The token name is the well known string `HydraHeadV1`, i.e. $\text{ST} = \{\text{cid} \mapsto \text{HydraHeadV1} \mapsto 1\}$.
- One *Participation Token (PT)* per participant $i \in \{1 \dots |k_{\text{H}}|\}$ to be used for authenticating further transactions and to ensure every participant can commit and cannot be censored. The token name is the participant’s verification key $k_i^{\#} = \text{hash}(k_i^{\text{ver}})$ of the verification key as received during protocol setup, i.e. $\text{PT}_i = \{\text{cid} \mapsto k_i^{\#} \mapsto 1\}$.

Consequently, the *init* transaction

Update figures

Open problem: ensure abort is always possible. e.g. by individual aborts or undoing commits

Open problem: ensure fanout is always possible, e.g. by limiting complexity of U_0




Figure 3: *init* transaction spending a seed UTxO, and producing the head output in state initial and initial outputs for each participant.

- has at least input ϕ_{seed} ,
- mints the state thread token ST, and one PT for each of the $|k_H|$ participants with policy cid,
- has $|k_H|$ initial outputs o_{initial_i} with datum $\delta_{\text{initial}} = \text{cid}$,
- has one head output o_{head} , which captures the initial state of the protocol in the datum

$$\delta_{\text{head}} = (\text{initial}, \text{cid}', \phi'_{\text{seed}}, \underline{k}_H, T)$$

where

- initial is a state identifier,
- cid' is the unique currency id of this instance,
- ϕ'_{seed} is the output reference parameter of μ_{head} ,
- \underline{k}_H are the off-chain multi-signature keys from the setup phase,
- T is the contestation period.

The $\mu_{\text{head}}(\phi_{\text{seed}})$ minting policy is the only script that verifies init transactions and can be redeemed with either a Mint or Burn redeemer:

- When evaluated with the Mint redeemer,
 1. The seed output is spent in this transaction. This guarantees uniqueness of the policy cid because the EUTxO ledger ensures that ϕ_{seed} cannot be spent twice in the same chain. $(\phi_{\text{seed}}, \cdot) \in \mathcal{I}$
 2. All entries of mint are of this policy and of single quantity $\forall \{c \mapsto \cdot \mapsto q\} \in \text{mint} : c = \text{cid} \wedge q = 1$

3. Right number of tokens are minted $|\text{mint}| = |\underline{k}_H| + 1$
 4. State token is sent to the head validator $\text{ST} \in \text{val}_{\text{head}}$
 5. The correct number of initial outputs are present $|(\cdot, \nu_{\text{initial}}, \cdot) \in \mathcal{O}| = |\underline{k}_H|$
 6. All participation tokens are sent to the initial validator as an initial output $\forall i \in [1 \dots |\underline{k}_H|] : \{\text{cid} \mapsto \cdot \mapsto 1\} \in \text{val}_{\text{initial}_i}$
 7. The δ_{head} contains own currency id $\text{cid} = \text{cid}'$ and the right seed reference $\phi_{\text{seed}} = \phi'_{\text{seed}}$
 8. All initial outputs have a cid as their datum: $\forall i \in [1 \dots |\underline{k}_H|] : \text{cid} = \delta_{\text{initial}_i}$
- When evaluated with the Burn redeemer,
 1. All tokens for this policy in mint need to be of negative quantity $\forall \{\text{cid} \mapsto \cdot \mapsto q\} \in \text{mint} : q < 0$.

move to
abort/fanout

Important: The μ_{head} minting policy only ensures uniqueness of cid, that the right amount of tokens have been minted and sent to ν_{head} and ν_{initial} respectively, while these validators in turn ensure continuity of the contract. However, it is **crucial** that all head members check that head output always contains an ST token of policy cid which satisfies $\text{cid} = \text{hash}(\mu_{\text{head}}(\phi_{\text{seed}}))$. The ϕ_{seed} from a head datum can be used to determine this. Also, head members should verify whether the correct verification key hashes are used in the PTs and the initial state is consistent with parameters agreed during setup. See the initialTx behavior in Figure 14 for details about these checks.

5.2 Commit Transaction

A *commit* transaction may be submitted by each participant $\forall i \in \{1 \dots |\underline{k}_H|\}$ to commit some UTXO into the head or acknowledge to not commit anything. The transaction is depicted in Figure 4 and has the following structure:

- One input spending from ν_{initial} with datum δ_{initial} , where value $\text{val}_{\text{initial}_i}$ holds a PT_i , and the redeemer $\rho_{\text{initial}} \in \Phi^*$ is a list of output references to be committed,
- zero or more inputs with reference $\phi_{\text{committed}_j}$ spending output $o_{\text{committed}_j}$ with $\text{val}_{\text{committed}_j}$,
- one output paying to ν_{commit} with value $\text{val}_{\text{commit}_i}$ and datum δ_{commit} .

The ν_{initial} validator with $\delta_{\text{initial}} = \text{cid}$ and $\rho_{\text{initial}} = \underline{\phi}_{\text{committed}}$ ensures that:

1. All committed value is in the output $\text{val}_{\text{commit}_i} \supseteq \text{val}_{\text{initial}_i} \cup (\bigcup_{j=1}^m \text{val}_{\text{committed}_j})$ ¹
2. Currency id and committed outputs are recorded in the output datum $\delta_{\text{commit}} = (\text{cid}, C_i)$, where $C_i = \forall j \in \{1 \dots m\} : [(\phi_{\text{committed}_j}, \text{bytes}(o_{\text{committed}_j}))]$ is a list of all committed UTXO recorded as tuples on-chain.
3. Transaction is signed by the right participant $\exists \{\text{cid} \mapsto k_i^\# \mapsto 1\} \in \text{val}_{\text{initial}} \Rightarrow k_i^\# \in \kappa$
4. No minting or burning $\text{mint} = \emptyset$

The ν_{commit} validator ensures the output is collected by either a *collectCom* in Section 5.4 or *abort* in Section 5.3 transaction of the on-chain state machine, selected by the appropriate redeemer.

¹The \supseteq is important for real world situations where the values might not be exactly equal due to ledger constraints

update
with mul-
tiple com-
mits




Figure 4: *commit* transaction spending an initial output and a single committed output, and producing a commit output.

5.3 Abort Transaction

The *abort* transaction (see Figure 5) allows a party to abort the creation of a head and consists of

- one input spending from ν_{head} holding the ST with δ_{head} ,
- $\forall i \in \{1 \dots |\underline{k}_H|\}$ inputs either
 - spending from ν_{initial} with with $\text{PT}_i \in \text{val}_{\text{initial}_i}$ and $\delta_{\text{initial}_i} = \text{cid}$, or
 - spending from ν_{commit} with with $\text{PT}_i \in \text{val}_{\text{commit}_i}$ and $\delta_{\text{commit}_i} = (\text{cid}, C_i)$,
- outputs $o_1 \dots o_m$ to redistribute already committed UTxOs.

Note that *abort* represents a final transition of the state machine and hence there is no state machine output.

Each spent ν_{initial} validator with $\delta_{\text{initial}_i} = \text{cid}$ and $\rho_{\text{initial}_i} = \text{abort}$ ensures that:

1. The state token of currency cid is getting burned $\{\text{ST} \mapsto -1\} \subseteq \text{mint}$.

Each spent ν_{commit} validator with $\delta_{\text{commit}_i} = (\text{cid}, \cdot)$ and $\rho_{\text{commit}_i} = \text{abort}$ ensures that:

1. The state token of currency cid is getting burned $\{\text{ST} \mapsto -1\} \subseteq \text{mint}$.

The $\mu_{\text{head}}(\phi_{\text{seed}})$ minting policy governs the burning of tokens via redeemer burn that:

1. All tokens in mint need to be of negative quantity $\forall \{\text{cid} \mapsto \cdot \mapsto q\} \in \text{mint} : q < 0$.

The state-machine validator ν_{head} is spent with $\rho_{\text{head}} = (\text{abort}, m)$, where m is the number of outputs to reimburse, and checks:

1. State is advanced from $\delta_{\text{head}} \sim \text{initial}$ to terminal state final :

$$(\text{initial}, \text{cid}, \phi_{\text{seed}}, \underline{k}_H, T) \xrightarrow[m]{\text{abort}} \text{final}.$$

(i.e. to ensure a minimum value on outputs).




Figure 5: *abort* transaction spending the initial state head output and collecting all initial and commit outputs, which get reimbursed by outputs $o_1 \dots o_m$. Note that each PT may be in either, an initial or commit output.

2. All UTxOs committed into the head are reimbursed exactly as they were committed. This is done by comparing hashes of serialised representations of the m reimbursing outputs $o_1 \dots o_m$ ² with the canonically combined committed UTxOs in C_i :

$$\text{hash}\left(\bigoplus_{j=1}^m \text{bytes}(o_j)\right) = \text{combine}([C_i \mid \forall [1 \dots |\underline{k}_H|], C_i \neq \perp])$$

3. Transaction is signed by a participant $\exists\{\text{cid} \mapsto k_i^\# \mapsto -1\} \in \text{mint} \Rightarrow k_i^\# \in \kappa$.
4. All tokens are burnt $|\{\text{cid} \mapsto \cdot \mapsto -1\} \in \text{mint}| = |\underline{k}_H| + 1$.

5.4 CollectCom Transaction

The *collectCom* transaction (Figure 6) collects all the committed UTxOs to the same head. It has

- one input spending from ν_{head} holding the ST with δ_{head} ,
- $\forall i \in \{1 \dots |\underline{k}_H|\}$ inputs spending commit outputs ($\text{val}_{\text{commit}_i}, \nu_{\text{commit}}, \delta_{\text{commit}_i}$) with $\text{PT}_i \in \text{val}_{\text{commit}_i}$ and $\delta_{\text{commit}_i} = (\text{cid}, C_i)$, and
- one output paying to ν_{head} with value $\text{val}'_{\text{head}}$ and datum δ'_{head} .

The state-machine validator ν_{head} is spent with $\rho_{\text{head}} = \text{collect}$ and checks:

1. State is advanced from $\delta_{\text{head}} \sim \text{initial}$ to $\delta'_{\text{head}} \sim \text{open}$, parameters $\text{cid}, \underline{k}_H, T$ stay unchanged and the new state is governed again by ν_{head}

$$(\text{initial}, \text{cid}, \phi_{\text{seed}}, \underline{k}_H, T) \xrightarrow{\text{collect}} (\text{open}, \text{cid}, \underline{k}_H, T, v, \eta)$$

²Only the first m outputs are used for reimbursing, while more outputs may be present in the transaction, e.g for returning change.

where snapshot version is initialized as $v = 0$.

- Commits are collected in η

$$\eta = U^\# = \text{combine}([C_1, \dots, C_n])$$

where $n = |k_H|$ and

$$\text{combine}(\underline{C}) = \text{hash}(\text{concat}(\text{sortOn}(1, \text{concat}(\underline{C}))^\downarrow^2))$$

That is, given a list of committed UTXO \underline{C} , where each element is a list of output references and the serialised representation of what was committed, **combine** first concatenates all commits together, sorts this list by the output references, concatenates all bytes and hashes the result³.

- All committed value captured and no value is extracted

$$\text{val}'_{\text{head}} = \text{val}_{\text{head}} \cup \left(\bigcup_{i=1}^n \text{val}_{\text{commit}_i} \right)$$

- Every participant had the chance to commit, by checking all tokens are present in output⁴

$$|\{\text{cid} \rightarrow . \rightarrow 1\} \in \text{val}'_{\text{head}}| = |k_H| + 1$$

- Transaction is signed by a participant

$$\exists \{\text{cid} \mapsto k_i^\# \mapsto 1\} \in \text{val}_{\text{commit}_i} \Rightarrow k_i^\# \in \kappa$$

- No minting or burning

$$\text{mint} = \emptyset$$

Each spent ν_{commit} validator with $\delta_{\text{commit}_i} = (\text{cid}, \cdot)$ and $\rho_{\text{commit}_i} = \text{collect}$ ensures that:

- The state token of currency cid is present in the output value

$$\text{ST} \in \text{val}'_{\text{head}}$$

5.5 Deposit Transaction

The *deposit* transaction initiates a (incremental) commit by locking funds in ν_{deposit} for later collection by the head protocol. Any transaction paying to ν_{deposit} is a *deposit* transaction as there is no on-chain verification in *deposit* transactions. Consequently, protocol actors **must ensure off-chain** that a valid datum is used when paying to the ν_{deposit} validator.

A valid deposit output is governed by ν_{deposit} with value $\text{val}_{\text{deposit}}$ and datum

$$\delta_{\text{deposit}} = (\text{cid}, t_{\text{recover}}, C)$$

where

³Sorting is required to ensure a canonical representation which can also be reproduced from the UTXO set later in the fanout.

⁴This is sufficient as a Head participant would check off-chain whether a Head is initialized correctly with the right number of tokens.

explain why this is enough?




Figure 6: *collectCom* transaction spending the head output in initial state and collecting from multiple commit outputs into a single *open* head output.

- *cid* is the currency id of the target head protocol instance (see 5.1),
- t_{recover} is a deadline after which the deposit can be recovered, and
- $C \in (\mathcal{J} \times \mathbb{H})^*$ is a list of transaction output references with along with serialized outputs that should become available in the head (similar to commits in 5.2).

Head protocol participants determine **off-chain** whether a deposit output o_{deposit} is eligible for their head by checking

1. *cid* matches their head identifier,
2. t_{recover} is reasonably far in the future, and
3. $\text{val}_{\text{deposit}}$ contains the value of all decoded outputs of C from δ_{deposit} .

explain;
relate to
contesta-
tion pe-
riod?

An example transaction which records all its spent inputs in a deposit output is shown in Figure 7. The $j \in \{1 \dots m\}$ inputs of this example with reference $\phi_{\text{deposited}_j}$ each spend output $o_{\text{deposited}_j}$ with $\text{val}_{\text{deposited}_j}$ would be recorded in the output datum as

$$C = \forall j \in \{1 \dots m\} : [(\phi_{\text{deposited}_j}, \text{bytes}(o_{\text{deposited}_j}))]$$

and the value check would need to satisfy

$$\text{val}_{\text{deposit}} \supseteq \bigcup_{j=1}^m \text{val}_{\text{deposited}_j}$$

5.6 Recover Transaction

If a *deposit* transaction output (see 5.5) was not collected into a head by an *increment* transaction 5.7, the *recover* transaction (Figure 8) allows for restoring the UTxO as recorded in the deposit after the deadline has passed. It consists of




Figure 7: *deposit* transaction spending multiple UTxO into a deposit output.




Figure 8: *recover* transaction restoring UTxO of a deposit output.

- one input spending from ν_{deposit} with datum $\delta_{\text{deposit}} = (\text{cid}, t_{\text{recover}}, C)$, and
- outputs $o_1 \dots o_m$ to recover UTxOs.

The script validator ν_{deposit} is spent with redeemer $\rho_{\text{deposit}} = (\text{Recover}, m)$, where m is the number of outputs to recover, and checks:

1. All UTxOs are recovered exactly as they were deposited. This is done by comparing hashes of serialised representations of the m recovering outputs $o_1 \dots o_m$ with the canonically combined committed UTxOs in C :

$$\text{hash}\left(\bigoplus_{j=1}^m \text{bytes}(o_j)\right) = \text{hash}(\text{concat}(\text{sortOn}(1, C)^{\downarrow 2}))$$

2. Transaction is posted after the deadline

$$t_{\text{min}} > t_{\text{recover}}$$

5.7 Increment Transaction

The *increment* transaction (Figure 9) allows a participant to add a *deposit* output 5.5 to an already open head using a snapshot that approves inclusion. Consequently this transaction consists of:

- one input spending from ν_{head} with value val_{head} holding the ST and datum δ_{head} ,
- one input ϕ_{deposit} spending from ν_{deposit} with value $\text{val}_{\text{deposit}}$ and datum $\delta_{\text{deposit}} = (\text{cid}_{\text{deposit}}, t_{\text{recover}}, C)$,
- one output paying to ν_{head} with value $\text{val}'_{\text{head}}$ and datum δ'_{head} .

The deposit validator ν_{deposit} is spent with $\rho_{\text{deposit}} = (\text{claim}, \text{cid})$ and ensures:

1. Claiming head id matches the deposit datum

$$\text{cid} = \text{cid}_{\text{deposit}}$$

2. Transaction is posted before the deadline

$$t_{\text{max}} \leq t_{\text{recover}}$$

The state-machine validator ν_{head} is spent with $\rho_{\text{head}} = (\text{increment}, \xi, s, \phi_{\text{increment}})$, where ξ is a multi-signature of the increment snapshot which authorizes addition of deposited UTxO, s is the snapshot number and ϕ_{deposit} points to the claimed deposit. The validator checks:

1. State is advanced from $\delta_{\text{head}} \sim \text{open}$ to $\delta'_{\text{head}} \sim \text{open}$, parameters $\text{cid}, \tilde{k}_{\text{H}}, n, T$ stay unchanged and the new state is governed again by ν_{head} :

$$(\text{open}, \text{cid}, \tilde{k}_{\text{H}}, n, T, v, \eta) \xrightarrow[\xi, s, \phi_{\text{increment}}]{\text{increment}} (\text{open}, \text{cid}, \tilde{k}_{\text{H}}, n, T, v', \eta')$$

2. New version v' is incremented correctly

$$v' = v + 1$$

3. Claimed deposit is spent

$$\phi_{\text{increment}} = \phi_{\text{deposit}}$$

4. ξ is a valid multi-signature of the new head state η'

$$\text{MS-Verify}(\tilde{k}_{\text{H}}, (\text{cid} || v || s || \eta' || \eta_{\alpha} || \perp), \xi) = \text{true}$$

where η_{α} is the digest of all deposited UTxO in C sorted by their output references

$$\eta_{\alpha} = \text{hash}(\text{concat}(\text{sortOn}(1, C)^{\downarrow 2}))$$

5. The value in the head output is increased accordingly

$$\text{val}_{\text{head}} \cup \text{val}_{\text{deposit}} = \text{val}'_{\text{head}}$$

6. Transaction is signed by a participant

$$\exists \{\text{cid} \mapsto k_i^{\#} \mapsto 1\} \in \text{val}'_{\text{head}} \Rightarrow k_i^{\#} \in \kappa$$

Only check $\text{val}'_{\text{head}} > \text{val}_{\text{head}}?$

Redundant to snapshot sig?




Figure 9: *increment* transaction spending an open head output, producing a new head output which includes the new UTxO.

5.8 Decrement Transaction

The *decrement* transaction (Figure 10) allows a party to remove a UTxO from an already open head and consists of:

- one input spending from ν_{head} holding the ST with δ_{head} ,
- one output paying to ν_{head} with value $\text{val}'_{\text{head}}$ and datum δ'_{head} ,
- one or more decommit outputs $o_2 \dots o_{m+1}$ with values $\text{val}_2 \dots \text{val}_{m+1}$.

The state-machine validator ν_{head} is spent with $\rho_{\text{head}} = (\text{decrement}, \xi, s, m)$, where ξ is a multi-signature of the decrement snapshot which authorizes removal of some UTxO, s is the used snapshot number and m is the number of outputs to distribute. The validator checks:

1. State is advanced from $\delta_{\text{head}} \sim \text{open}$ to $\delta'_{\text{head}} \sim \text{open}$, parameters $\text{cid}, \underline{k}_{\text{H}}, T$ stay unchanged and the new state is governed again by ν_{head}

$$(\text{open}, \text{cid}, \underline{k}_{\text{H}}, T, v, \eta) \xrightarrow[\xi, s, m]{\text{decrement}} (\text{open}, \text{cid}, \underline{k}_{\text{H}}, T, v', \eta')$$

2. New version v' is incremented correctly

$$v' = v + 1$$

3. ξ is a valid multi-signature of the new snapshot state η'

$$\text{MS-Verify}(\underline{k}_{\text{H}}, (\text{cid} || v || s || \eta' || \eta \alpha || \eta_{\omega}), \xi) = \text{true}$$

where η_{ω} is the digest of all removed UTxO

$$\eta_{\omega} = \text{hash}\left(\bigoplus_{j=2}^{m+1} \text{bytes}(o_j)\right)$$




Figure 10: *decrement* transaction spending an open head output, producing a new head output and multiple decommitted outputs.




Figure 11: *close* transaction spending the open head output and producing a closed head output.

4. The value in the head output is decreased accordingly

$$\text{val}'_{\text{head}} \cup \left(\bigcup_{j=2}^{m+1} \text{val}_j \right) = \text{val}_{\text{head}}$$

Only check $\text{val}'_{\text{head}} < \text{val}_{\text{head}}$?

5. Transaction is signed by a participant

$$\exists \{ \text{cid} \mapsto k_i^{\#} \mapsto 1 \} \in \text{val}'_{\text{head}} \Rightarrow k_i^{\#} \in \kappa$$

Redundant to snapshot sig?

5.9 Close Transaction

In order to close a head, a head member may post the *close* transaction (see Figure 11). This transaction has

- one input spending from ν_{head} holding the ST with δ_{head} ,
- one output paying to ν_{head} with value $\text{val}'_{\text{head}}$ and datum δ'_{head} .

The state-machine validator ν_{head} is spent with $\rho_{\text{head}} = (\text{close}, \text{CloseType})$, where *CloseType* is a hint against which open state to close and checks:

1. State is advanced from $\delta_{\text{head}} \sim \text{open}$ to $\delta'_{\text{head}} \sim \text{closed}$, parameters $\text{cid}, \underline{k}_{\text{H}}, T$ stay unchanged and the new state is governed again by ν_{head}

$$(\text{open}, \text{cid}, \underline{k}_{\text{H}}, T, v, \eta) \xrightarrow[\text{CloseType}]{\text{close}} (\text{closed}, \text{cid}, \underline{k}_{\text{H}}, T, v', s', \eta', \eta_{\alpha} \Delta', \eta_{\omega} \Delta', \mathcal{C}, t_{\text{final}})$$

2. Last known open state version is recorded in closed state

$$v' = v$$

3. Based on the redeemer $\text{CloseType} = \text{Initial} \cup (\text{Any}, \xi) \cup (\text{UnusedInc}, \xi, \eta_{\alpha}) \cup (\text{UnusedDec}, \xi) \cup (\text{UsedInc}, \xi) \cup (\text{UsedDec}, \xi, \eta_{\omega})$, where ξ is a multi-signature of the closing snapshot and η_{α} and η_{ω} are the digests of the UTxO to increment or decrement respectively, six cases are distinguished:

- (a) **Initial:** The initial snapshot is used to close the head and open state was not updated. No signatures are available and it suffices to check

$$v = 0$$

$$s' = 0$$

$$\eta' = \eta$$

- (b) **Any:** Closing snapshot refers to current state version v and both UTxO to increment and decrement must be empty in the closed state.

$$\eta_{\alpha} \Delta' = \eta_{\alpha} = \perp$$

$$\eta_{\omega} \Delta' = \eta_{\omega} = \perp$$

$$\text{MS-Verify}(\underline{k}_{\text{H}}, (\text{cid} || v || s' || \eta' || \eta_{\alpha} || \eta_{\omega}), \xi) = \text{true}$$

- (c) **UnusedInc:** Closing snapshot refers to current state version v and any UTxO to increment must not be recorded in the closed state.

$$\eta_{\alpha} \neq \perp$$

$$\eta_{\alpha} \Delta' = \perp$$

$$\eta_{\omega} \Delta' = \eta_{\omega} = \perp$$

$$\text{MS-Verify}(\underline{k}_{\text{H}}, (\text{cid} || v || s' || \eta' || \eta_{\alpha} || \eta_{\omega}), \xi) = \text{true}$$

where η_{α} is provided by the redeemer ⁵.

- (d) **UnusedDec:** Closing snapshot refers to current state version v and any UTxO to decrement need to be present in the closed state too.

$$\eta_{\alpha} \Delta' = \eta_{\alpha} = \perp$$

$$\eta_{\omega} \Delta' = \eta_{\omega} \neq \perp$$

$$\text{MS-Verify}(\underline{k}_{\text{H}}, (\text{cid} || v || s' || \eta' || \eta_{\alpha} || \eta_{\omega}), \xi) = \text{true}$$

⁵ η_{α} needs to be provided to verify the signature, but is otherwise not relevant for the closed state

this is
hard to
under-
stand

- (e) **UsedInc**: Closing snapshot refers the previous state $v - 1$ and any UTxO to increment must be recorded in the closed state.

$$\eta_\alpha \Delta' = \eta_\alpha \neq \perp$$

$$\eta_\omega \Delta' = \eta_\omega = \perp$$

$$\text{MS-Verify}(k_{\text{H}}, (\text{cid} || v - 1 || s' || \eta' || \eta_\alpha || \eta_\omega), \xi) = \text{true}$$

- (f) **UsedDec**: Closing snapshot refers the previous state $v - 1$ and any UTxO to decrement must not be recorded in the closed state.

$$\eta_\alpha \Delta' = \eta_\alpha = \perp$$

$$\eta_\omega \neq \perp$$

$$\eta_\omega \Delta' = \perp$$

$$\text{MS-Verify}(k_{\text{H}}, (\text{cid} || v - 1 || s' || \eta' || \eta_\alpha || \eta_\omega), \xi) = \text{true}$$

where η_ω is provided by the redeemer⁶.

4. Initializes the set of contesters

$$\mathcal{C} = \emptyset$$

This allows the closing party to also contest and is required for use cases where pre-signed, valid in the future, close transactions are used to delegate head closing.

5. Correct contestation deadline is set

$$t_{\text{final}} = t_{\text{max}} + T$$

6. Transaction validity range is bounded by

$$t_{\text{max}} - t_{\text{min}} \leq T$$

to ensure the contestation deadline t_{final} is at most $2 * T$ in the future.

7. Value in the head is preserved

$$\text{val}'_{\text{head}} = \text{val}_{\text{head}}$$

8. Transaction is signed by a participant

$$\exists \{\text{cid} \mapsto k_i^\# \mapsto 1\} \in \text{val}_{\text{commit}_i} \Rightarrow k_i^\# \in \kappa$$

9. No minting or burning

$$\text{mint} = \emptyset$$




Figure 12: *contest* transaction spending the closed head output and producing a different closed head output.

5.10 Contest Transaction

The *contest* transaction (see Figure 12) is posted by a party to prove the currently closed state is not the latest one. This transaction has

- one input spending from ν_{head} holding the ST with δ_{head} ,
- one output paying to ν_{head} with value $\text{val}'_{\text{head}}$ and datum δ'_{head} .

The state-machine validator ν_{head} is spent with $\rho_{\text{head}} = (\text{contest}, \text{ContestType})$, where *ContestType* is a hint how to verify the snapshot and checks:

1. State is advanced from $\delta_{\text{head}} \sim \text{open}$ to $\delta'_{\text{head}} \sim \text{closed}$, parameters $\text{cid}, \underline{k}_{\text{H}}, T$ stay unchanged and the new state is governed again by ν_{head}

$$(\text{closed}, \text{cid}, \underline{k}_{\text{H}}, T, v, s, \eta, \eta_{\alpha} \Delta, \eta_{\omega} \Delta, \mathcal{C}, t_{\text{final}}) \xrightarrow[\text{ContestType}]{\text{contest}} (\text{closed}, \text{cid}, \underline{k}_{\text{H}}, T, v', s', \eta', \eta_{\alpha} \Delta', \eta_{\omega} \Delta', \mathcal{C}', t'_{\text{final}})$$

2. Last known open state version stays recorded in closed state

$$v' = v$$

3. Contested snapshot number s' is higher than the currently stored snapshot number s

$$s' > s$$

4. Based on the redeemer $\text{ContestType} = (\text{Current}, \xi) \cup (\text{UnusedInc}, \xi, \eta_{\alpha}) \cup (\text{UnusedDec}, \xi) \cup (\text{UsedInc}, \xi) \cup (\text{UsedDec}, \xi, \eta_{\omega})$, where ξ is a multi-signature of the contesting snapshot and η_{α} and η_{ω} are the digests of the UTxO to increment or decrement respectively, five cases are distinguished:

- (a) **Current:** Contesting snapshot refers to current state version v and any UTxO to increment or decrement must be \perp in the closed state.

$$\eta_{\alpha} \Delta' = \eta_{\alpha} = \perp$$

$$\eta_{\omega} \Delta' = \eta_{\omega} = \perp$$

$$\text{MS-Verify}(\underline{k}_{\text{H}}, (\text{cid} || v || s' || \eta' || \eta_{\alpha} || \eta_{\omega}), \xi) = \text{true}$$

⁶ η_{ω} needs to be provided to verify the signature, but is otherwise not relevant for the closed state

- (b) **UnusedInc**: Contesting snapshot refers to current state version v and any UTxO to increment must be \perp in the closed state.

$$\eta_\alpha \neq \perp$$

$$\eta_\alpha \Delta' = \perp$$

$$\eta_\omega \Delta' = \eta_\omega = \perp$$

$$\text{MS-Verify}(\underline{k}_H, (\text{cid}||v||s'||\eta'||\eta_\alpha||\eta_\omega), \xi) = \text{true}$$

where η_α is provided by the redeemer⁷.

- (c) **UnusedDec**: Contesting snapshot refers to current state version v and any UTxO to decrement need to be present in the closed state too.

$$\eta_\alpha \Delta' = \eta_\alpha = \perp$$

$$\eta_\omega \Delta' = \eta_\omega \neq \perp$$

$$\text{MS-Verify}(\underline{k}_H, (\text{cid}||v||s'||\eta'||\eta_\alpha||\eta_\omega), \xi) = \text{true}$$

- (d) **UsedInc**: Contesting snapshot refers the previous state $v-1$ and any UTxO to increment must be recorded in the closed state.

$$\eta_\alpha \Delta' = \eta_\alpha \neq \perp$$

$$\eta_\omega \Delta' = \eta_\omega = \perp$$

$$\text{MS-Verify}(\underline{k}_H, (\text{cid}||v-1||s'||\eta'||\eta_\alpha||\eta_\omega), \xi) = \text{true}$$

- (e) **UsedDec**: Contesting snapshot refers the previous state $v-1$ and any UTxO to decrement must not be recorded in the closed state.

$$\eta_\alpha \Delta' = \eta_\alpha = \perp$$

$$\eta_\omega \neq \perp$$

$$\eta_\omega \Delta' = \perp$$

$$\text{MS-Verify}(\underline{k}_H, (\text{cid}||v-1||s'||\eta'||\eta_\alpha||\eta_\omega), \xi) = \text{true}$$

where η_ω is provided by the redeemer⁸.

5. The single signer $\{k^\#\} = \kappa$ has not already contested and is added to the set of contesters

$$k^\# \notin \mathcal{C}$$

$$\mathcal{C}' = \mathcal{C} \cup k^\#$$

6. Transaction is posted before deadline

$$t_{\max} \leq t_{\text{final}}$$

⁷ η_α needs to be provided to verify the signature, but is otherwise not relevant for the closed state

⁸ η_ω needs to be provided to verify the signature, but is otherwise not relevant for the closed state




Figure 13: *fanout* transaction spending the closed head output and distributing funds with outputs $o_1 \dots o_{m+n+n'}$.

7. Contestation deadline is updated correctly to

$$t'_{\text{final}} = \begin{cases} t_{\text{final}} & \text{if } |\mathcal{C}'| = n, \\ t_{\text{final}} + T & \text{otherwise.} \end{cases}$$

8. Value in the head is preserved

$$\text{val}'_{\text{head}} = \text{val}_{\text{head}}$$

9. Transaction is signed by a participant

$$\exists \{\text{cid} \mapsto k_i^{\#} \mapsto 1\} \in \text{val}_{\text{commit}_i} \Rightarrow k_i^{\#} \in \kappa$$

10. No minting or burning

$$\text{mint} = \emptyset$$

5.11 Fan-Out Transaction

Once the contestation phase is over, a head may be finalized by posting a *fanout* transaction (see Figure 13), which distributes UTxOs from the head according to the latest state. It consists of

- one input spending from ν_{head} holding the ST, and
- outputs $o_1 \dots o_{m+n+n'}$ to distribute UTxOs.

Note that *fanout* represents a final transition of the state machine and hence there is no state machine output.

The state-machine validator ν_{head} is spent with $\rho_{\text{head}} = (\text{fanout}, m, n, n')$, where m , n and n' are outputs to distribute from the closed state, and checks:

1. State is advanced from $\delta_{\text{head}} \sim \text{closed}$ to terminal state **final**:

$$(\text{closed}, \text{cid}, \underline{k}_{\text{H}}, T, v, s, \eta, \eta_{\alpha} \Delta, \eta_{\omega} \Delta, \mathcal{C}, t_{\text{final}}) \xrightarrow[\text{m, n, n'}]{\text{fanout}} \text{final}$$

2. The first m outputs are distributing funds according to η . That is, the outputs exactly correspond to the UTxO canonically combined $U^\#$ (see Section 5.4):

$$\eta = U^\# = \text{hash}\left(\bigoplus_{j=1}^m \text{bytes}(o_j)\right)$$

3. The following n outputs are distributing funds according to $\eta_{\alpha\Delta}$. That is, the outputs exactly correspond to the UTxO canonically combined $U_{\alpha\Delta}^\#$ (see Section 5.4):

$$\eta_{\alpha\Delta} = U_{\alpha\Delta}^\# = \text{hash}\left(\bigoplus_{j=m}^{m+n} \text{bytes}(o_j)\right)$$

4. The next n' outputs are distributing funds according to $\eta_{\omega\Delta}$. That is, the outputs exactly correspond to the UTxO canonically combined $U_{\omega\Delta}^\#$ (see Section 5.4):

$$\eta_{\omega\Delta} = U_{\omega\Delta}^\# = \text{hash}\left(\bigoplus_{j=m'}^{m+n'} \text{bytes}(o_j)\right)$$

5. Transaction is posted after contestation deadline $t_{\min} > t_{\text{final}}$.
6. All tokens are burnt $|\{\text{cid} \mapsto \cdot \mapsto -1\} \in \text{mint}| = n + 1$.

The $\mu_{\text{head}}(\phi_{\text{seed}})$ minting policy governs the burning of tokens via redeemer **burn** that:

1. All tokens in mint need to be of negative quantity $\forall \{\text{cid} \mapsto \cdot \mapsto q\} \in \text{mint} : q < 0$.

6 Off-Chain Protocol

This section describes the actual Coordinated Hydra Head protocol, an even more simplified version of the original publication [8]. See the protocol overview in Section 2 for an introduction and notable changes to the original protocol. While the on-chain part already describes the full life-cycle of a Hydra head on-chain, this section completes the picture by defining how the protocol behaves off-chain and notably the relationship between on- and off-chain semantics. Participants of the protocol are also called Hydra head members, parties or simply protocol actors. The protocol is specified as a reactive system that processes three kinds of inputs:

1. On-chain protocol transactions as introduced in Section 5, which are posted to the mainchain and can be observed by all actors:
 - **initialTx**: initiates a head
 - **commitTx**: commits UTxO to an initializing head
 - **collectComTx**: opens a head
 - **depositTx**: prepares UTxO to be incremented
 - **incrementTx**: adds UTxO to an open head
 - **decrementTx**: removes UTxO from an open head

- `closeTx`: closes a head
 - `contestTx`: contests a closed head
2. Off-chain network messages sent between protocol actors (parties):
 - `reqTx`: to request a transaction to be included in the next snapshot
 - `reqDec`: to request exclusion of UTXO via a decommit transaction
 - `reqSn`: to request a snapshot to be created & signed by every head member
 - `ackSn`: to acknowledge a snapshot by replying with their signatures
 3. Commands issued by the participants themselves or on behalf of end-users and clients
 - `init`: to start initialization of a head
 - `close`: to request closure of an open head

The behavior is fully specified in Figure 14, while the following paragraphs introduce notation, explain variables and walk-through the protocol flow.

6.1 Assumptions

On top of the statements of the protocol setup in Section 4, the off-chain protocol logic relies on these assumptions:

- Every network message received from a specific party is checked for authentication. An implementation of the specification needs to find a suitable means of authentication, either on the communication channel or for individual messages. Unauthenticated messages must be dropped.
- The head protocol gets correctly (and with completeness) notified about observed transactions on-chain belonging to the respective head instance. **Observation may be delayed varying by type of transaction. See 6.5 for a discussion on eventual consistency of transactions.**
- All inputs are processed to completion, i.e. run-to-completion semantics and no preemption.
- Inputs are deduplicated. That is, any two identical inputs must not lead to multiple invocations of the handling semantics.
- Given the specification, inputs may pile up forever and implementations need to consider these situations (i.e. potential for DoS). A valid reaction to this would be to just drop these inputs. Note that, from a security standpoint, these situations are identical to a non-collaborative peer and closing the head is also a possible reaction.
- The lifecycle of a Hydra head on-chain does not cross (hard fork) protocol update boundaries. Note that these inputs are announced in advance hence it should be possible for implementations to react in such a way as to expedite closing of the head before such a protocol update. This further assumes that the contestation period parameter is picked accordingly.

move/merge
with pro-
tocol
setup?

6.2 Notation

missing;
apply tx

- **on event** specifies how the protocol reacts on a given input *event*. Further information may be available from the constituents of *event* and origin of the input.
- **require** p means that boolean expression $p \in \mathbb{B}$ must be satisfied for the further execution of a routine, while discontinued on $\neg p$. A conservative protocol actor could interpret this as a reason to close the head.
- **wait** p is a non-blocking wait for boolean predicate $p \in \mathbb{B}$ to be satisfied. On $\neg p$, the execution of the routine is stopped, queued, and reactivated at latest when p is satisfied.
- **multicast** msg means that a message msg is (channel-) authenticated and sent to all participants of this head, including the sender.
- **postTx** tx has a party create transaction tx , potentially from some data, and submit it on-chain. See Section 5 for individual transaction details.
- **output event** signals an observation of *event*, which is used in the security definition and proofs of Section 7. This keyword can be ignored when implementing the protocol.

6.3 Variables

Besides parameters agreed in the protocol setup (see Section 4), a party's local state consists of the following variables:

- \hat{v} : Last seen open state version.
- \hat{s} : Sequence number of latest seen snapshot.
- $\hat{\Sigma} \in (\mathbb{N} \times \mathbb{H})^*$: Accumulator of signatures of the latest seen snapshot, indexed by parties.
- $\hat{\mathcal{L}}$: UTxO set representing the local ledger state resulting from applying $\hat{\mathcal{T}}$ to $\bar{S}.U$ to validate requested transactions.
- $\hat{\mathcal{T}} \in \mathcal{T}^*$: List of transactions applied locally and pending inclusion in a snapshot (if this party is the next leader).
- $U_\alpha \in (\mathcal{J} \times \mathcal{O})^*$: UTxO set pending to be added to the head.
- $tx_\omega \in \mathcal{T}$: Pending decrement transaction, whose outputs are to be withdrawn from the head.
- $\bar{\mathcal{S}}$: Snapshot object of the latest confirmed snapshot which contains:

$\bar{\mathcal{S}}.v$	snapshot version
$\bar{\mathcal{S}}.s$	snapshot number
$\bar{\mathcal{S}}.\mathcal{T}$	list of transactions relating this snapshot to the previous
$\bar{\mathcal{S}}.U$	snapshotted UTxO set
$\bar{\mathcal{S}}.U_\alpha$	pending UTxO to increment
$\bar{\mathcal{S}}.U_\omega$	pending UTxO to decrement
$\bar{\mathcal{S}}.\sigma$	multisignature

where constructor $\text{snObj}(v, n, \mathcal{T}, U, U_\alpha, U_\omega)$ initializes a new snapshot object with $\bar{\mathcal{S}}.\sigma = \emptyset$.

6.4 Protocol flow

6.4.1 Initializing the head

init. Before a head can be initialized, all parties need to exchange and agree on protocol parameters during the protocol setup phase (see Section 4), so we can assume the public Cardano keys k_C^{setup} , Hydra keys \tilde{k}_H^{setup} , as well as the contestation period T^{setup} are available. One of the clients then can start head initialization using the `init` command, which will result in an *init* transaction being posted.

initialTx. All parties will receive this *init* transaction and validate announced parameters against the pre-agreed *setup* parameters, as well as the structure of the transaction and the minting policy used. This is a vital step to ensure the initialized Head is valid, which cannot be checked completely on-chain (see also Section 5.1).

commitTx. As each party p_j posts a *commit* transaction, the protocol records observed committed UTxOs of each party C_j . With all committed UTxOs known, the η -state is created (as defined in Section 5.4) and the *collectCom* transaction is posted. Note that while each participant may post this transaction, only one of them will be included in the blockchain as the mainchain ledger prevents double spending. Should any party want to abort, they would post an *abort* transaction and the protocol would end at this point.

collectComTx. Upon observing the *collectCom* transaction, the parties compute $U_0 \leftarrow \bigcup_{j=1}^n C_j$ using previously observed C_j and initialize $\hat{\mathcal{L}} = U_0$. The seen transaction set is initialized empty $\hat{\mathcal{T}} = \emptyset$, seen head state version $\hat{v} = 0$, as well as snapshot number $\hat{s} = 0$. No UTxO to increment $U_\alpha = \emptyset$ and no decrement transaction $\text{tx}_\omega = \perp$ is pending, and the last confirmed snapshot is initialized accordingly $\bar{\mathcal{S}} \leftarrow \text{snObj}(0, 0, [], U_0, \emptyset, \emptyset)$.

6.4.2 Processing transactions off-chain

Transactions are announced and captured in so-called snapshots. Parties generate snapshots in a strictly sequential round-robin manner. The party responsible for issuing the i^{th} snapshot is the *leader* of the i^{th} snapshot. Leader selection is round-robin per the k_H from the protocol setup. While the frequency of snapshots in the general Head protocol [8] was configurable, the Coordinated Head protocol does specify a snapshot to be created after each transaction.

reqTx. Upon receiving request (`reqTx`, `tx`), the transaction is applied to the *local* ledger state $\hat{\mathcal{L}} \circ \text{tx}$. If not applicable yet, the protocol does **wait** to retry later or eventually marks this transaction as invalid (see assumption about events piling up). After applying and if there is no current snapshot in flight ($\hat{s} = \bar{\mathcal{S}}.s$) and the receiving party p_i is the next snapshot leader, a message to request snapshot signatures `reqSn` is sent.

reqDec. Upon receiving request (`reqDec`, `txω`), the transaction is checked against the *local* ledger state and if it is not applicable yet or another commit or decommit is pending still, the protocol does **wait** to retry later or eventually marks the decommit as invalid. After applying `tx`, its outputs

are removed from *local* ledger state $\hat{\mathcal{L}}$ so that they are not available any more and the decommit transaction is kept in the local state (tx_ω). If there is no current snapshot in flight ($\hat{s} = \bar{\mathcal{S}}.s$) and the receiving party \mathbf{p}_i is the next snapshot leader, a message to request snapshot signatures reqSn containing the decrement transaction tx_ω is sent.

depositTx. Upon observing a *deposit* transaction as settled⁹ and no other commit or decommit is pending still, each party keeps track of the observed deposited UTxO as the pending increment UTxO set $U_\alpha = U$. If other commits or decommits are pending, the protocol **waits** and retries updating state later. If the observing party is the next snapshot leader, it may request a new snapshot by sending a reqSn including the UTxO to increment U_α .

smelly
and frag-
ile

reqSn. Upon receiving request $(\text{reqSn}, v, s, \underline{\text{tx}}_{\text{req}}, U_\alpha, \text{tx}_\omega)$ ¹⁰ from party \mathbf{p}_j , the receiving \mathbf{p}_i **requires** that only a commit or decommit is currently pending, and that v refers to the current open state version, s is the next snapshot number and that party \mathbf{p}_j is responsible for leading its creation. Party \mathbf{p}_i may has to **wait** until the previous snapshot is confirmed ($\bar{\mathcal{S}}.s = \hat{s}$). If the decommit transaction tx_ω is not \perp , the transaction is **required** to be applicable to the last confirmed UTxO set $\bar{\mathcal{S}}.U$ and decommitted transaction outputs must be removed, yielding the still active UTxO set U_{active} . Then, all requested transactions $\underline{\text{tx}}_{\text{req}}$ are **required** to be applicable to U_{active} , otherwise the snapshot is rejected as invalid. Only then, \mathbf{p}_i increments their seen-snapshot counter \hat{s} , resets the signature accumulator $\hat{\Sigma}$, and computes the UTxO set of the new local snapshot as $U \leftarrow U_{\text{active}} \circ \underline{\text{tx}}_{\text{req}}$. Then, \mathbf{p}_i creates a signature σ_i using their signing key $k_{\text{H}}^{\text{sig}}$ on a message comprised by the cid, the new snapshot number \hat{s} , the new η resulting from canonically combining U (see Section 5.9 for details), and either η_α or η_ω derived from commit UTxO U_α or decommit transaction tx_ω respectively. The signature is sent to all head members via message $(\text{ackSn}, \hat{s}, \sigma_i)$. Finally, the local ledger state $\hat{\mathcal{L}}$ and pending transaction set $\hat{\mathcal{T}}$ get pruned by re-applying all locally pending transactions $\hat{\mathcal{T}}$ to the just requested snapshot's UTxO set iteratively and ultimately yielding a “pruned” version of $\hat{\mathcal{T}}$ and $\hat{\mathcal{L}}$.

define
leader

ackSn. Upon receiving acknowledgment $(\text{ackSn}, s, \sigma_j)$, all participants **require** that it is from an expected snapshot (either the last seen \hat{s} or $+1$), potentially **wait** for the corresponding reqSn such that $\hat{s} = s$ and **require** that the signature is not yet included in $\hat{\Sigma}$. They store the received signature in the signature accumulator $\hat{\Sigma}$, and if the signature from each party has been collected, \mathbf{p}_i aggregates the multisignature $\tilde{\sigma}$ and **require** it to be valid (constructing the signed message as in reqSn). If everything is fine, the snapshot can be considered confirmed by creating the snapshot object $\bar{\mathcal{S}} \leftarrow \text{snObj}(\hat{v}, \hat{s}, \hat{\mathcal{T}}, \hat{U}, U_\alpha, \text{outputs}(\text{tx}_\omega))$ and storing the multi-signature $\tilde{\sigma}$ in it for later reference. In case there is a pending decommit, any participant can now submit a *decrement* transaction by providing the just confirmed snapshot with its digests of the active UTxO set η and the to be removed UTxO set η_ω . If, however, there was a pending commit, any participant can now submit an *incrementTx* by providing the confirmed snapshot with its digests of the active UTxO set η and the UTxO set to be added η_α . Lastly, if \mathbf{p}_i is the next snapshot leader and there are already transactions to snapshot in $\hat{\mathcal{T}}$, a corresponding reqSn is distributed.

⁹Protocol actors might use different techniques and delays to determine transaction finality. See also 6.5.

¹⁰Snapshot requests with only transaction identifiers and output references are possible if all parties keep an index of previously seen transactions and their identifiers.

decrementTx. Upon observing the *decrement* transaction, which removed outputs U from the head, the corresponding pending decrement transaction is cleared and the observed version v is used for future snapshots by setting $\hat{v} \leftarrow v$. Note that the version of the open head state is incremented on each *decrement* transaction as described in Section 5.8.

incrementTx. Upon observing the *increment* transaction, which added outputs U to the head, the local ledger state $\hat{\mathcal{L}}$ is extended with the newly added UTxO while the pending increment state U_α is cleared. Also the observed version v is used for future snapshots by setting $\hat{v} = v$. Note that the version of the open head state is incremented on each *increment* transaction as described in Section 5.7

6.4.3 Closing the head

close. In order to close a head, a client issues the `close` input which uses the latest confirmed snapshot \bar{s} to create the new η -state from the last confirmed UTxO set, the digest of either increment or decrement UTxO set (η_α or η_ω), and the certificate ξ using the corresponding multi-signature. With these, the *close* transaction can be constructed and posted. See Section 5.9 for details about this transaction.

closeTx/contestTx. When a party observes the head getting closed or contested, the η -state extracted from the *close* or *contest* transaction represents the latest head status that has been aggregated on-chain so far (by a sequence of *close* and *contest* transactions). If the last confirmed (off-chain) snapshot is newer than the observed (on-chain) snapshot number s_c , an updated η -state, along with the digest of either increment or decrement UTxO set (η_α or η_ω), and certificate ξ is constructed and posted in a *contest* transaction (see Section 5.10).

6.5 Rollbacks and protocol changes

The overall life-cycle of the Head protocol is driven by on-chain inputs (see introduction of Section 6) which stem from observing transactions on the mainchain. Most blockchains, however, do only provide *eventual* consistency. The consensus algorithm ensures a consistent view of the history of blocks and transactions between all parties, but this so-called *finality* is only achieved after some time and the local view of the blockchain history may change until that point.

On Cardano with its Ouroboros consensus algorithm, this means that any local view of the mainchain may not be the longest chain and a node may switch to a longer chain, onto another fork. This other version of the history may not include what was previously observed and hence, any tracking state needs to be updated to this “new reality”. Practically, this means that an observer of the blockchain sees a *rollback* followed by rollforwards.

For the Head protocol, this means that chain events like `closeTx` may be observed a second time. Hence, it is crucial, that the local state of the Hydra protocol is kept in sync and also rolled back accordingly to be able to observe and react to these events the right way, e.g. correctly contesting this `closeTx` if need be.

The rollback handling can be specified fully orthogonal on top of the nominal protocol behavior, if the chain provides strictly monotonically increasing points p on each chain event via a new or wrapped `rollforward` event and `rollback` event with the point to which a rollback happened:

Explain varying finality of deposit vs. increment

Explain why rollbacks are no problem to increment/decrement

Write about contestation deadline vs. rollbacks

rollforward. On every chain event that is paired or wrapped in a rollforward event (`rollback, p`) with point p , protocol participants store their head state indexed by this point in a history Ω of states $\Delta \leftarrow (\hat{v}, \hat{s}, \hat{u}, \hat{\Sigma}, \hat{\mathcal{L}}, \hat{\mathcal{F}}, \bar{s})$ and $\Omega' = (p, \Delta) \cup \Omega$.

rollback. On a rollback (`rollback, prb`) to point p_{rb} , the corresponding head state Δ need to be retrieved from Ω , with the maximal point $p \leq p_{rb}$, and all entries in Ω with $p > p_{rb}$ get removed.

This will essentially reset the local head state to the right point and allow the protocol to progress through the life-cycle normally. Most stages of the life-cycle are unproblematic if they are rolled back, as long as the protocol logic behaves as in the nominal case.

A rollback “past open” is a special situation though. When a Head is open and snapshots have been signed, but then a *collectCom* and one or more *commit* transactions were rolled back, a bad actor could choose to commit a different UTxO and open the Head with a different initial UTxO set, while the already signed snapshots would still be (cryptographically) valid. To mitigate this, all signatures on snapshots need to incorporate the initial UTxO set by including η_0 .

version-counting is less powerful

Coordinated Hydra Head

<pre> on (init) from client $n \leftarrow \underline{k}_H^{setup}$ $\tilde{k}_H \leftarrow \text{MS-AVK}(\underline{k}_H^{setup})$ $\underline{k}_C \leftarrow \underline{k}_C^{setup}$ $T \leftarrow T^{setup}$ postTx (init, $n, \tilde{k}_H, \underline{k}_C, T$) on (initialTx, cid, $\phi_{seed}, n, \tilde{k}_H, \underline{k}_C^\#, T$) from chain require $\tilde{k}_H = \text{MS-AVK}(\underline{k}_H^{setup})$ require $\underline{k}_C^\# = [\text{hash}(k) \mid \forall k \in \underline{k}_C^{setup}]$ require $T = T^{setup}$ require cid = hash($\mu_{head}(\phi_{seed})$) </pre>	<pre> on (commitTx, j, U) from chain $C_j \leftarrow U$ if $\forall k \in [1..n] : C_k \neq \perp$ $\eta \leftarrow \text{combine}([C_1 \dots C_n])$ postTx (collectCom, η) on (collectComTx, η_0) from chain $U_0 \leftarrow \bigcup_{j=1}^n U_j$ $\hat{\mathcal{L}} \leftarrow U_0$ $\bar{s} \leftarrow \text{snObj}(0, 0, [], U_0, \emptyset, \emptyset)$ $\hat{v}, \hat{s} \leftarrow 0$ $\hat{\mathcal{T}} \leftarrow \emptyset$ $\text{tx}_\omega \leftarrow \perp$ $U_\alpha \leftarrow \emptyset$ </pre>
<pre> on (reqTx, tx) from p_j wait $\hat{\mathcal{L}} \circ \text{tx} \neq \perp$ $\hat{\mathcal{L}} \leftarrow \hat{\mathcal{L}} \circ \text{tx}$ $\hat{\mathcal{T}} \leftarrow \hat{\mathcal{T}} \cup \{\text{tx}\}$ if $\hat{s} = \bar{s}.s \wedge \text{leader}(\bar{s}.s + 1) = i$ multicast (reqSn, $\hat{v}, \bar{s}.s + 1, \hat{\mathcal{T}}, U_\alpha, \text{tx}_\omega$) on (reqDec, tx) from p_j wait $U_\alpha = \emptyset \wedge \text{tx}_\omega = \perp \wedge \hat{\mathcal{L}} \circ \text{tx} \neq \perp$ $\hat{\mathcal{L}} \leftarrow \hat{\mathcal{L}} \circ \text{tx} \setminus \text{outputs}(\text{tx})$ $\text{tx}_\omega \leftarrow \text{tx}$ if $\hat{s} = \bar{s}.s \wedge \text{leader}(\bar{s}.s + 1) = i$ multicast (reqSn, $\hat{v}, \bar{s}.s + 1, \hat{\mathcal{T}}, U_\alpha, \text{tx}_\omega$) on (depositTx, U) from chain wait $\text{tx}_\omega = \perp \wedge U_\alpha = \emptyset$ $U_\alpha = U$ if $\hat{s} = \bar{s}.s \wedge \text{leader}(\bar{s}.s + 1) = i$ multicast (reqSn, $\hat{v}, \bar{s}.s + 1, \hat{\mathcal{T}}, U_\alpha, \text{tx}_\omega$) on (reqSn, $v, s, \underline{\text{tx}}_{\text{req}}, U_\alpha, \text{tx}_\omega$) from p_j require $\text{tx}_\omega = \perp \vee U_\alpha = \emptyset$ require $v = \hat{v} \wedge s = \hat{s} + 1 \wedge \text{leader}(s) = j$ wait $\hat{s} = \bar{s}.s$ require $\bar{s}.U \circ \text{tx}_\omega \neq \perp$ $U_{\text{active}} \leftarrow \bar{s}.U \circ \text{tx}_\omega \setminus \text{outputs}(\text{tx}_\omega)$ require $U_{\text{active}} \circ \underline{\text{tx}}_{\text{req}} \neq \perp$ $U \leftarrow U_{\text{active}} \circ \underline{\text{tx}}_{\text{req}}$ $\hat{s} \leftarrow s$ $\eta \leftarrow \text{combine}(U)$ $\eta_\alpha \leftarrow \text{combine}(U_\alpha)$ $\eta_\omega \leftarrow \text{combine}(\text{outputs}(\text{tx}_\omega))$ $\sigma_i \leftarrow \text{MS-Sign}(k_H^{sig}, (\text{cid} \parallel v \parallel \hat{s} \parallel \eta \parallel \eta_\alpha \parallel \eta_\omega))$ $\hat{\Sigma} \leftarrow \emptyset$ multicast (ackSn, \hat{s}, σ_i) $\forall \text{tx} \in \underline{\text{tx}}_{\text{req}} : \text{output}(\text{seen}, \text{tx})$ $\hat{\mathcal{L}} \leftarrow U$ $X \leftarrow \hat{\mathcal{T}}$ $\hat{\mathcal{T}} \leftarrow \emptyset$ for $\text{tx} \in X : \hat{\mathcal{L}} \circ \text{tx} \neq \perp$ $\hat{\mathcal{T}} \leftarrow \hat{\mathcal{T}} \cup \{\text{tx}\}$ $\hat{\mathcal{L}} \leftarrow \hat{\mathcal{L}} \circ \text{tx}$ </pre>	<pre> on (ackSn, s, σ_j) from p_j require $s \in \{\hat{s}, \hat{s} + 1\}$ wait $\hat{s} = s$ require $(j, \cdot) \notin \hat{\Sigma}$ $\hat{\Sigma}[j] \leftarrow \sigma_j$ if $\forall k \in [1..n] : (k, \cdot) \in \hat{\Sigma}$ $\tilde{\sigma} \leftarrow \text{MS-ASig}(\underline{k}_H^{setup}, \hat{\Sigma})$ $\eta \leftarrow \text{combine}(\hat{U})$ $\eta_\alpha \leftarrow \text{combine}(U_\alpha)$ $U_\omega \leftarrow \text{outputs}(\text{tx}_\omega)$ $\eta_\omega \leftarrow \text{combine}(U_\omega)$ require MS-Verify($\tilde{k}_H, (\text{cid} \parallel \hat{v} \parallel \hat{s} \parallel \eta \parallel \eta_\alpha \parallel \eta_\omega), \tilde{\sigma}$) $\bar{s} \leftarrow \text{snObj}(\hat{v}, \hat{s}, \hat{\mathcal{T}}, \hat{U}, U_\alpha, U_\omega)$ $\bar{s}.\sigma \leftarrow \tilde{\sigma}$ $\forall \text{tx} \in \mathcal{T}_{\text{req}} : \text{output}(\text{conf}, \text{tx})$ if $\bar{s}.U_\omega \neq \perp$ postTx (decrementTx, $\hat{v}, \hat{s}, \eta, \eta_\alpha, \eta_\omega$) if $\bar{s}.U_\alpha \neq \perp$ postTx (incrementTx, $\hat{v}, \hat{s}, \eta, \eta_\alpha, \eta_\omega$) if $\text{leader}(s + 1) = i \wedge \hat{\mathcal{T}} \neq \emptyset$ multicast (reqSn, $\hat{v}, \bar{s}.s + 1, \hat{\mathcal{T}}, U_\alpha, \text{tx}_\omega$) on (decrementTx, U, v) from chain $\text{tx}_\omega \leftarrow \perp$ $\hat{v} \leftarrow v$ on (incrementTx, U, v) from chain $\hat{\mathcal{L}} \leftarrow \hat{\mathcal{L}} \cup U$ $U_\alpha \leftarrow \emptyset$ $\hat{v} \leftarrow v$ </pre>

<pre> on (close) <i>from client</i> $\eta \leftarrow \text{combine}(\bar{S}.U)$ $\eta_\alpha \leftarrow \text{combine}(\bar{S}.U_\alpha)$ $\eta_\omega \leftarrow \text{combine}(\bar{S}.U_\omega)$ $\xi \leftarrow \bar{S}.\sigma$ postTx (<i>close</i>, \hat{v}, $\bar{S}.v$, $\bar{S}.s$, η, η_α, η_ω, ξ) </pre>	<pre> on (closeTx, η) \vee (contestTx, s_c, η) <i>from chain</i> if $\bar{S}.s > s_c$ $\eta \leftarrow \text{combine}(\bar{S}.U)$ $\eta_\alpha \leftarrow \text{combine}(\bar{S}.U_\alpha)$ $\eta_\omega \leftarrow \text{combine}(\bar{S}.U_\omega)$ $\xi \leftarrow \bar{S}.\sigma$ postTx (<i>contest</i>, \hat{v}, $\bar{S}.v$, $\bar{S}.s$, η, η_α, η_ω, ξ) </pre>
--	--

Figure 14: Head-protocol machine for the *coordinated head* from the perspective of party p_i .

7 Security (WIP — Iteration 1)

Adversaries:

Active Adversary. An *active adversary* \mathcal{A} has full control over the protocol, i.e., he is fully unrestricted in the above security game.

Network Adversary. A *network adversary* \mathcal{A}_0 does not corrupt any head parties, eventually delivers all sent network messages (i.e., does not drop any messages), and does not cause the **close** event. Apart from this restriction, the adversary can act arbitrarily in the above experiment.

Random variables:

- \hat{S}_i : the set of transactions tx for which party p_i , *while uncorrupted*, output (**seen**, tx);
- \bar{C}_i : the set of transactions tx for which party p_i , *while uncorrupted*, output (**conf**, tx);
- $\bar{\Sigma}_i$: latest snapshot (s, U) that party p_i performed *while uncorrupted*: output (**snap**, (s, U));
- H_{cont} : the set of (at the time) uncorrupted parties who produced ξ upon close/contest request and ξ was applied to correct η ; and
- \mathcal{H} : the set of parties that remain uncorrupted.

Security conditions / events:

- **CONSISTENCY (HEAD)**: In presence of an active adversary, the following condition holds at any point in time: For all i, j , $U_0 \circ (\bar{C}_i \cup \bar{C}_j) \neq \perp$, i.e., no two uncorrupted parties see conflicting transactions confirmed.
- **OBLIVIOUS LIVENESS (HEAD)**: Consider any protocol execution in presence of a network adversary wherein the head does not get closed for a sufficiently long period of time, and consider an honest party p_i who enters transaction tx by executing (**newTx**, tx) *each time after having finished a snapshot*.

In figure:
combine
on UTxO
slightly
different
than on
commits

The secu-
rity analy-
sis is still
sketchy,
with the
goal to
make it
more for-
mal in
upcoming
iterations

Add se-
curity ex-
periment

above this
section
there is
no secu-
rity game

Then the following eventually holds: $\text{tx} \in \bigcap_{i \in [n]} \bar{C}_i \vee \forall i : U_0 \circ (\bar{C}_i \cup \{\text{tx}\}) = \perp$, i.e., every party will observe the transaction confirmed or every party will observe the transaction in conflict with their confirmed transactions.¹¹

- **SOUNDNESS (CHAIN):** In presence of an active adversary, the following condition is satisfied: $\exists \tilde{S} \subseteq \bigcap_{i \in \mathcal{H}} \hat{S}_i : U_{\text{final}} = U_0 \circ \tilde{S} \neq \perp$, i.e., the final UTxO set results from applying a set of transactions to U_0 that have been seen by all honest parties (whereas each such transaction applies conforming to the ledger rules).
- **COMPLETENESS (CHAIN):** In presence of an active adversary, the following condition holds: For \tilde{S} as above, $\bigcup_{p_i \in H_{\text{cont}}} \bar{C}_i \subseteq \tilde{S}$, i.e., all transactions seen as confirmed by an honest party at the end of the protocol are considered.

Note that the original version of the coordinated head satisfies a stronger version of liveness which is important for the 'user experience' in the protocol:

- **LIVENESS (HEAD):** Consider any protocol execution in presence of a network adversary wherein the head does not get closed for a sufficiently long period of time, and consider an honest party p_i who enters transaction tx by executing (`newTx`, tx).

Then the following eventually holds: $\text{tx} \in \bigcap_{i \in [n]} \bar{C}_i \vee \forall i : U_0 \circ (\bar{C}_i \cup \{\text{tx}\}) = \perp$, i.e., every party will observe the transaction confirmed or every party will observe the transaction in conflict with their confirmed transactions.¹²

7.1 Proofs

Consistency.

Lemma 1 (Consistency). *The coordinated head protocol satisfies the CONSISTENCY property.*

Proof. Observe that $\bar{C}_i \cup \bar{C}_j \subseteq \hat{S}_i$ since no transaction can be confirmed without every honest party signing off on it. Since parties do not sign conflicting transactions (see `reqSn`, 'wait'), we have $U_0 \circ \bar{C}_i \neq \perp$, $U_0 \circ \bar{C}_j \neq \perp$, and $U_0 \circ \hat{S}_i \neq \perp$. Thus, since $\bar{C}_i \cup \bar{C}_j \subseteq \hat{S}_i$ it follows that $U_0 \circ (\bar{C}_i \cup \bar{C}_j) \neq \perp$ \square

Oblivious Liveness. For all lemmas towards oblivious liveness, we assume the presence of a network adversary, and that the head does not get closed for a sufficiently long period of time. We call this the *liveness condition*.

Lemma 2. *Under the liveness condition, any snapshot issued as (`reqSn`, s, T) will eventually be confirmed in the sense that every party holds a valid multisignature on it.*

Proof. Consider a party p_i receiving message (`reqSn`, s, T). We demonstrate that p_i executes the code past the 'wait' instruction of the `reqSn` routine.

¹¹In particular, *liveness* expresses that the protocol makes progress under reasonable network conditions if no head parties get corrupted.

¹²In particular, *liveness* expresses that the protocol makes progress under reasonable network conditions if no head parties get corrupted.

- Passing the ‘require’ guard: Note that the snapshot leader sends the request only if $\hat{s} = \bar{s}$, and for $s = \hat{s} + 1$. Thus, $\hat{s}_i = \hat{s}$ since p_i has already signed the snapshot for \hat{s} . The ‘require’ guard is thus satisfied for p_i .
- Passing the ‘wait’ guard: Since the snapshot leader sees $\hat{s} = \bar{s}$, also p_i will eventually see $\hat{s}_i = \bar{s}_i$. Furthermore, since all leaders are honest, it holds that $\hat{\mathcal{U}} \circ \mathcal{T}_{res} \neq \perp$ by construction.

This implies that every party will eventually sign and acknowledge the newly created snapshot. Finally, the ‘require’ and ‘wait’ guards of the `ackSn` code will be passed by every party since an `ackSn` for snapshot number s can only be received for $s \in \{\hat{s}, \hat{s} + 1\}$ as an acknowledgement can only be received for the current snapshot being worked on by p_i or a snapshot that is one step ahead—implying that everybody will hold a valid multisignature on the snapshot in consideration. \square

Lemma 3 (Eternal snapshot confirmation). *Under the liveness condition, as long as new transactions are issued, for any $k > 0$, every party eventually confirms a snapshot with sequence number $s = k$.*

Proof. By Lemma 2, any requested snapshot eventually gets confirmed, implying that the next leader observes $\hat{s} = \bar{s}$ and thus, in turn, issues a new snapshot. Thus, for any k , a snapshot is eventually confirmed. \square

Lemma 4 (Oblivious Liveness). *The coordinated head protocol satisfies the OBLIVIOUS LIVENESS property.*

Proof. Consider the first point in time where a transaction `tx` enters the system by some party p_i issuing `(newTx, tx)`, and consider the next point in time t when p_i issues a snapshot.

By Lemma 3, this snapshot will eventually be issued and confirmed by all parties.

Let $\hat{\mathcal{T}}$ be the transactions to be considered by p_i ’s snapshot: $\hat{\mathcal{L}} = \bar{\mathcal{U}} \circ \hat{\mathcal{T}}$ where $\bar{\mathcal{U}}$ is the snapshot prior to p_i ’s. Since p_i issues `(reqTx, tx)` after each snapshot, we have that, either,

- $\text{tx} \in \hat{\mathcal{T}}$, in which case $\text{tx} \in \bigcap_{i \in [n]} \bar{C}_i$ after everybody has completed this snapshot, or,
- $\text{tx} \notin \hat{\mathcal{T}}$, in which case $\hat{\mathcal{L}} \circ \text{tx} = \perp$ (`tx` is still in the wait queue of `(reqTx, tx)`). After everybody has completed this snapshot, it thus holds that $\forall i : U_0 \circ \bar{C}_i = \hat{\mathcal{L}}$, and thus, that $\forall i : U_0 \circ (\bar{C}_i \cup \{\text{tx}\}) = \perp$.

In both cases, the lemma follows. \square

Soundness and completeness.

Lemma 5 (Soundness). *The basic head protocol satisfies the SOUNDNESS property.*

Proof. Let T be the set of transactions such that $U_{\text{final}} = U_0 \circ T$. Since U_{final} is multi-signed, it holds that $T \subseteq \hat{S}_i$ (T is *seen*) by every honest party in the head. Furthermore, since honest signatures are only issued for valid transaction, $U_{\text{final}} \neq \perp$ (i.e., U_{final} is a valid state), and soundness follows. \square

Lemma 6 (Completeness). *The basic head protocol satisfies the COMPLETENESS property.*

Proof. Consider all parties $p_i \in H_{\text{cont}}$. Since the close/contest process finally accepts the latest multi-signed snapshot, it holds that $U_{\text{final}.s} \geq \max_{p_i \in H_{\text{cont}}} (\bar{s}_i)$, and thus that $\bigcup_{p_i \in H_{\text{cont}}} \bar{C}_i \subseteq \bigcap_{p_i \in \mathcal{H}} \hat{S}_i$, and completeness follows. \square

References

- [1] Extended UTXO-2 model. <https://github.com/hydra-supplementary-material/eutxo-spec/blob/master/extended-utxo-specification.pdf>.
- [2] A formal specification of the cardano ledger. <https://github.com/input-output-hk/cardano-ledger/releases/latest/download/shelley-ledger.pdf>.
- [3] A formal specification of the cardano ledger integrating plutus core. <https://github.com/input-output-hk/cardano-ledger/releases/latest/download/alonzo-ledger.pdf>.
- [4] Hydra repository. <https://github.com/input-output-hk/hydra>.
- [5] Nicola Atzei, Massimo Bartoletti, Stefano Lande, and Roberto Zunino. A formal model of Bitcoin transactions. In *Financial Cryptography and Data Security - 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26 - March 2, 2018, Revised Selected Papers*, pages 541–560, 2018.
- [6] Manuel M. T. Chakravarty, James Chapman, Kenneth MacKenzie, Orestis Melkonian, Michael Peyton Jones, and Philip Wadler. The extended UTxO model. In *4th Workshop on Trusted Smart Contracts*, 2020. http://fc20.ifca.ai/wtsc/WTSC2020/WTSC20_paper_25.pdf.
- [7] Manuel M. T. Chakravarty, James Chapman, Kenneth M. Mackenzie, Orestis Melkonian, Jann, Müller, Michael Peyton Jones, Polina Vinogradova, Philip Wadler, Joachim, and Zahnentferner. Utxoma: Utxo with multi-asset support. 2020.
- [8] Manuel MT Chakravarty, Sandro Coretti, Matthias Fitzi, Peter Gazi, Philipp Kant, Aggelos Kiayias, and Alexander Russell. Hydra: Fast isomorphic state channels. *Cryptology ePrint Archive*, 2020.
- [9] Kazuharu Itakura and Katsuhiko Nakamura. A public-key cryptosystem suitable for digital multisignatures. *NEC Research & Development*, (71):1–8, 1983.
- [10] Silvio Micali, Kazuo Ohta, and Leonid Reyzin. Accountable-subgroup multisignatures: Extended abstract. pages 245–254, 2001.
- [11] Joachim Zahnentferner. An abstract model of UTxO-based cryptocurrencies with scripts. *IACR Cryptology ePrint Archive*, 2018:469, 2018.