

Polkadot Runtime

Protocol Specification

Contents

1	Availability and Validity Verification	5
1.1	Introduction	5
1.2	Preliminaries	6
1.3	Overall process	6
1.4	Primary Validation	7
1.4.1	Primary validity announcement	8
1.4.2	Inclusion of candidate receipt on the relay chain	8
1.4.3	Primary Validation Disagreement	8
1.5	Availability	9
1.6	Distribution of Pieces	9
1.7	Announcing Availability	9
1.7.1	Processing on-chain availability data	10
1.8	Publishing Attestations	12
1.9	Secondary Approval checking	12
1.9.1	Approval Checker Assignment	12
1.9.2	VRP computation	12
1.9.3	One-Shot Approval Checker Assignemnt	13
1.9.4	Extra Approval Checker Assignment	13
1.9.5	Additional Checking in Case of Equivocation	13
1.10	The Approval Check	14
1.10.1	Verification	15
1.10.2	Process validity and invalidity messages	15
1.10.3	Invalidity Escalation	16

Chapter 1

Availability and Validity Verification

1.1 Introduction

Validators are responsible for guaranteeing the validity and availability of PoV blocks. There are two phases of validation that takes place in the AnV protocol.

The primary validation check is carried out by parachain validators who are assigned to the parachain which has produced the PoV block as described in Section 1.4. Once parachain validators have validated a parachain's PoV block successfully, they have to announce that according to the procedure described in Section 1.4.1 where they generate a candidate receipt that includes the parachain header with the new state root and the XCMP message root. This candidate receipt and attestations, which carries signatures from other parachain validators is put on the relay chain.

As soon as the proposal of a PoV block is on-chain, the parachain validators break the PoV block into erasure-coded pieces as described in Section ?? and distribute them among all validators. See Section ?? for details on how this distribution takes place.

Once validators have received erasure-coded pieces for several PoV blocks for the current relay chain block (that might have been proposed a couple of blocks earlier on the relay chain), they announce that they have received the erasure coded pieces on the relay chain by voting on the received pieces, see Section 1.7 for more details.

As soon as $> 2/3$ of validators have made this announcement for any parachain block we *act on* the parachain block. Acting on parachain blocks means we update the relay chain state based on the candidate receipt and considered the parachain block to have happened on this relay chain fork.

After a certain time, if we did not collect enough signatures approving the availability of the parachain data associated with a certain candidate receipt we decide this parachain block is unavailable and allow alternative blocks to be built on its parent parachain block, see ??.

The secondary check described in Section 1.9, is done by one or more randomly assigned validators to make sure colluding parachain validators may not get away with validating a PoV block that is invalid and not keeping it available to avoid the possibility of being punished for the attack.

During any of the phases, if any validator announces that a parachain block is invalid then all

validators obtain the parachain block and check its validity, see Section ?? for more details.

All validity and invalidity attestations go onto the relay chain, see Section 1.8 for details. If a parachain block has been checked at least by certain number of validators, the rest of the validators continue with voting on that relay chain block in the GRANDPA protocol. Note that the block might be challenged later.

1.2 Preliminaries

Definition 1 *In the remainder of this chapter we assume that ρ is a Polkadot Parachain and B is a block which has been produced by ρ and is supposed to be approved to be ρ 's next block. By R_{rho} we refer to runtime code of parachain ρ as a WASM Blob.*

Definition 2 *The witness proof of block B , denoted by π_B , is the set of all the external data which has gathered while the ρ runtime executes block B . The data suffices to re-execute R_{rho} against B and achieve the final state indicated in the $H(B)$.*

This witness proof consists of light client proofs of state data that are generally Merkle proofs for the parachain state trie. We need this because validators do not have access to the parachain state, but only have the state root of it.

Definition 3 *Accordingly we define the **proof of validity block** or **PoV** block in short, PoV_B , to be the tuple:*

$$(B, \pi_B)$$

Definition 4 *The extra validation data v_B is an extra input to the validation function, i.e. additional data from the relay chain state that is needed.*

This extra validation data includes things like the previous parachain block header, likely including the previous state root. Parachain validators get this extra validation data from the current relay chain state. Note that a PoV block can be paired with different extra validation data depending on when and which relay chain fork it is included in. Future validators would need this extra validation data because since the candidate receipt was included on the relay chain the needed relay chain state may have changed.

Definition 5 *Accordingly we define the **erasure coding blob** or **blob** in short, \bar{B} to be the tuple:*

$$(B, \pi_B, v_B)$$

Note that in the code the blob is referred to as "AvailableData".

1.3 Overall process

The Figure 1.3 demonstrates the overall process of assuring availability and validity in Polkadot **TODO: complete the Diagram.**

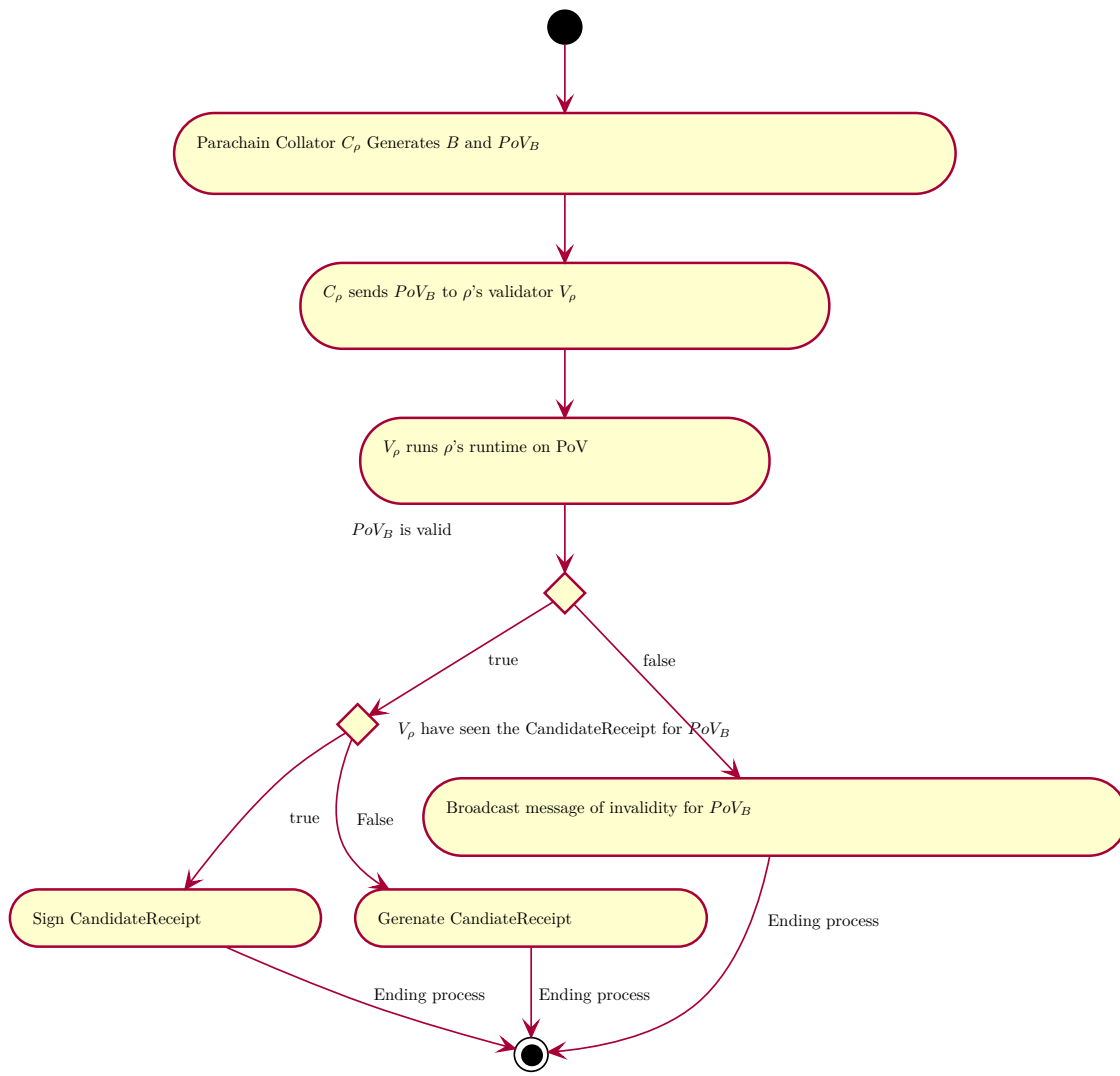


Figure 1.1: Overall process to achieve availability and validity in Polkadot

1.4 Primary Validation

Primary validity checking refers to the process of parachain validators as defined in Definition ?? validating a parachain's PoV block as explained in Algorithm 1.

Algorithm 1 PRIMARYVALIDATION

Input: B, π_B , relay chain parent block $B_{relayparent}$

- 1: Retrieve v_B from the relay chain state at $B_{relayparent}$
 - 2: Run Algorithm 2 using B, π_B, v_B
-

Algorithm 2 VALIDATEBLOCK

Input: B, π_B, v_B

- 1: retrieve the runtime code R_ρ that is specified by v_B from the relay chain state.
 - 2: check that the initial state root in π_B is the one claimed in v_B
 - 3: Execute R_ρ on B using π_B to simulate the state.
 - 4: If the execution fails, return fail.
 - 5: Else return success, the new header data h_B and the outgoing messages M .
-

1.4.1 Primary validity announcement

Validator v needs to perform Algorithm 3 to announce the result of primary validation to the Polkadot network.

In case that validation has been successful, the announcement will either be in the form of sending the candidate receipt for block B as defined in Definition 6 to the relay chain or confirm a candidate receipt sent in from another parachain validators for this block according to Algorithm 5. However, if the validation fails, v reacts by executing Algorithm 6.

Definition 6 *Candidate Receipt is a proposal for B , TBS.*

Algorithm 3 PRIMARYVALIDATIONANNOUNCEMENT

Input:

- 1: TBS
-

1.4.2 Inclusion of candidate receipt on the relay chain

Definition 7 *Parachain Block Proposal, noted by P_{rho}^B is a candidate receipt for a parachain block B for a parachain ρ along with signatures for at least $2/3$ of \mathcal{V}_ρ .*

A block producer which observe a Parachain Block Proposal as defined in definition 7 may/should include the proposal in the block they are producing according to Algorithm 7 during block production procedure.

1.4.3 Primary Validation Disagreement

Parachain validators need to keep track of candidate receipts (see Definition 6) and validation failure messages of their peers. In case, there is a disagreement among the parachain validators about \bar{B} , all parachain validators must invoke Algorithm 8

Algorithm 4 SENDPOVCANDIDATERECEIPT

Input:

1: TBS

Algorithm 5 CONFIRMCANDIDATERECEIPT

Input:

1: TBS

1.5 Availability

When a $v \in \mathcal{V}_p$ observes that a block containing parachain block candidate receipt is included in a relay chain block RB_p then it must invoke Algorithm 9.

Definition 8 *The erasure encoder/decoder $encode_{k,n}/decoder_{k,n}$ is defined to be the Reed-Solomon encoder defined in [?].*

Definition 9 *The set of erasure encode pieces of \bar{B} , denoted by:*

$$Er_B := (e_1, m_1), \dots, (e_n, m_n)$$

is defined to be the output of the Algorithm 9.

1.6 Distribution of Pieces

Following the computation of Er_B , v must construct the \bar{B} Availability message defined in Definition 10. And distribute them to target validators designated by the Availability Networking Specification [?].

Definition 10 *PoV erasure piece message $M_{PoV_{\bar{B}}}(i)$ is TBS*

1.7 Announcing Availability

When validator v receives its designated piece for \bar{B} it needs to broadcast Availability vote message as defined in Definition 11

Definition 11 *Availability vote message M_{PoV}^{Avail, v_i} TBS*

Algorithm 6 ANNOUNCEPRIMARYVALIDATIONFAILURE

Input:

1: TBS

Algorithm 7 INCLUDEPARACHAINPROPOSAL(P_{rho}^B)

Input:1: TBS

Algorithm 8 PRIMARYVALIDATIONDISAGREEMENT

Input:1: TBS

Some parachains have blocks that we need to vote on the availability of, that is decided by $> 2/3$ of validators voting for availability. For 100 parachain and 1000 validators this will involve putting 100k items of data and processing them on-chain for every relay chain block, hence we want to use bit operations that will be very efficient. We describe next what operations the relay chain runtime uses to process these availability votes.

For each parachain, the relay chain stores the following data:

1) availability status, 2) candidate receipt, 3) candidate relay chain block number where availability status is one of {no candidate, to be determined, unavailable, available} .

For each block, each validator v signs a message

Sign(bitfield b_v , block hash h_b)

where the i th bit of b_v is 1 if and only if

1. the availability status of the candidate receipt is "to be determined" on the relay chain at block hash h_b **and**
2. v has the erasure coded piece of the corresponding parachain block to this candidate receipt.

These signatures go into a relay chain block.

1.7.1 Processing on-chain availability data

This section explains how the availability attestations stored on the relay chain, as described in Section ??, are processed as follows:

Based on the result of Algorithm ?? the validator node should mark a parachain block as either available or eventually unavailable according to definitions 12 and ??

Definition 12 *Parachain blocks blocks for which the corresponding blob is noted on the relay chain to be available, meaning that the candidate receipt has been voted to be available by 2/3 validators.*

After a certain time-out in blocks since we first put the candidate receipt on the relay chain if there is not enough votes of availability the relay chain logic decides that a parachain block is unavailable, see 10.

Algorithm 9 ERASURE-ENCODE(\bar{B}, n)

Input: \bar{B} : blob defined in Definition 51: TBS

Algorithm 10 Relay chain's signature processing

- 1: The relay chain stores the last vote from each validator on chain. For each new signature, the relay chain checks if it is for a block in this chain later than the last vote stored from this validator. If it is the relay chain updates the stored vote and updates the bitfield b_v and block number of the vote.
 - 2: For each block within the last t blocks where t is some timeout period, the relay chain computes a bitmask bm_n (n is block number). This bitmask is a bitfield that represents whether the candidate considered in that block is still relevant. That is the i th bit of bm_n is 1 if and only if for the i th parachain, (a) the availability status is to be determined and (b) candidate block number $\leq n$
 - 3: The relay chain initialises a vector of counts with one entry for each parachain to zero. After executing the following algorithm it ends up with a vector of counts of the number of validators who think the latest candidates is available.
 1. The relay chain computes b_v and bm_n where n is the block number of the validator's last vote
 2. For each bit in b_v and bm_n
 - add the i th bit to the i th count.
 - 4: For each count that is $> 2/3$ of the number of validators, the relay chain sets the candidates status to "available". Otherwise, if the candidate is at least t blocks old, then it sets its status to "unavailable".
 - 5: The relay chain acts on available candidates and discards unavailable ones, and then clears the record, setting the availability status to "no candidate". Then the relay chain accepts new candidate receipts for parachains that have "no candidate: status and once any such new candidate receipts is included on the relay chain it sets their availability status as "to be determined".
-

Definition 13 *An unavailable parachain block is TBS*

/syedSo to be clear we are not announcing unavailability we just keep it for grand pa vote

1.8 Publishing Attestations

We have two type of attestations, primary and secondary. Primary attestations are signed by the parachain validators and secondary attestations are signed by secondary checkers and include the VRF that assigned them as a secondary checker into the attestation. Both types of attestations are included in the relay chain block as a transaction. For each parachain block candidate the relay chain keeps track of which validators have attested to its validity or invalidity.

1.9 Secondary Approval checking

Once a parachain block is acted on we carry the secondary validity/availability checks as follows. A scheme assigns every validator to one or more PoV blocks to check its validity, see Section 1.9.3 for details. An assigned validator acquires the PoV block (see Section ??) and checks its validity by comparing it to the candidate receipt. If validators notices that an equivocation has happened an additional validity/availability assignments will be made that is described in Section 1.9.5.

1.9.1 Approval Checker Assignment

Validators assign themselves to parachain block proposals as defined in Definition 7. The assignment needs to be random. Validators use their own VRF to sign the VRF output from the current relay chain block as described in Section 1.9.2. Each validator uses the output of the VRF to decide the block(s) they are revalidating as a secondary checker. See Section ?? for the detail.

In addition to this assignment some extra validators are assigned to every PoV block which is described in Section ??.

1.9.2 VRF computation

Every validator needs to run Algorithm 11 for every Parachain ρ to determines assignments. **TODO: Fix this. It is incorrect so far.**

Algorithm 11 VRF-FOR-APPROVAL(B, z, s_k)

Input: B : the block to be approved

z : randomness for approval assignment

s_k : session secret key of validator planning to participate in approval

1: $(\pi, d) \leftarrow \text{VRF}(H_h(B), s_k(z))$

2: **return** (π, d)

Where VRF function is defined in [?].

1.9.3 One-Shot Approval Checker Assignment

Every validator v takes the output of this VRF computed by $11 \bmod$ the number of parachain blocks that we were decided to be available in this relay chain block according to Definition 12 and executed. This will give them the index of the PoV block they are assigned to and need to check. The procedure is formalised in 12.

Algorithm 12 ONESHOTASSIGNMENT

Input:

1: TBS

1.9.4 Extra Approval Checker Assignment

Now for each parachain block, let us assume we want $\#VCheck$ validators to check every PoV block during the secondary checking. Note that $\#VCheck$ is not a fixed number but depends on reports from collators or fishermen. Lets us $\#VDefault$ be the minimum number of validator we want to check the block, which should be the number of parachain validators plus some constant like 2. We set

$$\#VCheck = \#VDefault + c_f * \text{total fishermen stake}$$

where c_f is some factor we use to weight fishermen reports. Reports from fishermen about this

Now each validator computes for each PoV block a VRF with the input being the relay chain block VRF concatenated with the parachain index.

For every PoV bock, every validator compares $\#VCheck - \#VDefault$ to the output of this VRF and if the VRF output is small enough than the validator checks this PoV blocks immediately otherwise depending on their difference waits for some time and only perform a check if it has not seen $\#VCheck$ checks from validators who either 1) parachain validators of this PoV block 2) or assigned during the assignment procedure or 3) had a smaller VRF output than us during this time.

More fisherman reports can increase $\#VCheck$ and require new checks. We should carry on doing secondary checks for the entire fishing period if more are required. A validator need to keep track of which blocks have $\#VCheck$ smaller than the number of higher priority checks performed. A new report can make us check straight away, no matter the number of current checks, or mean that we need to put this block back into this set. If we later decide to prune some of this data, such as who has checked the block, then we'll need a new approach here.

Algorithm 13 ONESHOTASSIGNMENT

Input:

1: TBS

1.9.5 Additional Checking in Case of Equivocation

In the case of a relay chain equivocation, i.e. a validator produces two blocks with the same VRF, we do not want the secondary checkers for the second block to be predictable. To this end we use the block hash as well as the VRF as input for secondary checkers VRF. So each secondary checker

is going to produce twice as many VRFs for each relay chain block that was equivocated. If either of these VRFs is small enough then the validator is assigned to perform a secondary check on the PoV block. The process is formalized in Algorithm 14

Algorithm 14 EQUIVOCATEDASSIGNMENT

Input:

1: TBS

1.10 The Approval Check

Once a validator has a VRF which tells them to check a block, they announce this VRF and attempt to obtain the block. It is unclear yet whether this is best done by requesting the PoV block from parachain validators or by announcing that they want erasure coded pieces.

Retrieval

There are two fundamental ways to retrieve a parachain block for checking validity. One is to request the whole block from any validator who has attested to its validity or invalidity. Assigned approval checker v sends RequestWholeBlock message specified in Definition ?? to parachain validator in order to receive the specific parachain block. Any parachain validator receiving must reply with PoVBlockRespose message defined in Definition 14

Definition 14 *PoV Block Respose Message TBS*

The second method is to retrieve enough erasure coded pieces to reconstruct the block from them. In the latter cases an announcement of the form specified in Definition has to be gossiped to all validators indicating that one needs the erasure coded pieces.

Definition 15 *Erasure coded pieces request message TBS*

On their part, when a validator receive a erasure coded pieces request message it response with the message specified in Definition 16.

Definition 16 *Erasure coded pieces response message TBS*

Assigned approval checker v must retrieve enough erasure pieces of the block they are verifying to be able to reconstruct the block and the erasure pieces tree.

Reconstruction

After receiving $2f + 1$ of erasure pieces every assigned approval checker v needs to recreate the entirety of the erasure code, hence every v will run Algorithm ?? to make sure that the code is complete and the subsequently recover the original \bar{B} .

Algorithm 15 RECONSTRUCT-POV-ERASURE(S_{Er_B})

Input: $S_{Er_B} := (e_{j_1}, m_{j_1}), \dots, (e_{j_k}, m_{j_k})$ such that $k > 2f$

```

1:  $\bar{B} \rightarrow \text{ERASURE-DECODER}(e_{j_1}, \dots, e_{j_k})$ 
2: if ERASURE-DECODER failed then
3:   ANNOUNCE-FAILURE
4:   return
5: end if
6:  $Er_B \rightarrow \text{ERASURE-ENCODER}(\bar{B})$ 
7: if VERIFY-MERKLE-PROOF( $S_{Er_B}, Er_B$ ) failed then
8:   ANNOUNCE-FAILURE
9:   return
10: end if
11: return  $\bar{B}$ 

```

1.10.1 Verification

Once the parachain block has been obtained or reconstructed the secondary checker needs to execute the PoV block. We declare a the candidate receipt as invalid if one of the following three conditions hold: 1) While reconstructing if the erasure code does not have the claimed Merkle root, 2) the validation function says that the PoV block is invalid, or 3) the result of executing the block is inconsistent with the candidate receipt on the relay chain.

The procedure is formalized in Algorithm

Algorithm 16 REVALIDATINGRECONSTRUCTEDPOV

Input:1: TBS

If everything checks out correctly, we declare the block is valid. This means gossiping an attestation, including a reference that identifies candidate receipt and our VRF as specified in Definition 17.

Definition 17 *Secondary approval attestation message TBS*

1.10.2 Process validity and invalidity messages

When a Block produced receive a Secondary approval attestation message, it execute Algorithm 17 to verify the VRF and may need to judge when enough time has passed.

Algorithm 17 VERIFYAPPROVALATTESTATION

Input:1: TBS

These attestations are included in the relay chain as a transaction specified in

Definition 18 *Approval Attestation Transaction TBS*

Collators reports of unavailability and invalidity specified in Definition **TODO: Define these messages** also go onto the relay chain as well in the format specified in Definition

Definition 19 *Collator Invalidation Transaction TBS*

Definition 20 *Collator unavailability Transaction*

1.10.3 Invalidation Escalation

When for any candidate receipt, there are attestations for both its validity and invalidity, then all validators acquire and validate the blob, irrespective of the assignments from section by executing Algorithm ?? and 16.

We do not vote in GRANDPA for a chain were the candidate receipt is executed until its vote is resolved. If we have n validators, we wait for $> 2n/3$ of them to attest to the blob and then the outcome of this vote is one of the following:

If $> n/3$ validators attest to the validity of the blob and $\leq n/3$ attest to its invalidity, then we can vote on the chain in GRANDPA again and slash validators who attested to its invalidity.

If $> n/3$ validators attest to the invalidity of the blob and $\leq n/3$ attest to its validity, then we consider the blob as invalid. If the relay chain block where the corresponding candidate receipt was executed was not finalised, then we never vote on it or build on it. We slash the validators who attested to its validity.

If $> n/3$ validators attest to the validity of the blob and $> n/3$ attest to its invalidity then we consider the blob to be invalid as above but we do not slash validators who attest either way. We want to leave a reasonable length of time in the first two cases to slash anyone to see if this happens.