Technical Report:
Attestation Aggregation and Packing Problem

Satalia & Sigma Prime

November 24, 2022

Abstract

The Attestation Aggregation and Packing Problem (AAPP) arises in the context of Ethereum's Proof of Stake (PoS) consensus protocol. This document aims at formally defining the problem.

Lighthouse, Sigma Prime's Ethereum consensus client, uses a greedy heuristic to generate solutions to the AAPP. The goal of the present investigation is to assess how far the solutions generated by Lighthouse's greedy heuristic are from the optimal solutions. Ultimately, the goal is to either confirm that the current algorithm is sufficiently good for practical purposes, or to design an algorithm with better quality guarantees.

Contents

1 Background

This section aims to provide a high-level discussion of some concepts behind blockchain technology, specifically those relevant to the definition of the AAPP.

1.1 Decentralised trustless architecture

At its core, and irrespective of its practical applications, a blockchain is a database designed to be immutable and tamper-proof. Unlike other technologies which rely on *trusted central authorities* to guarantee the integrity of the stored data, a blockchain achieves the same goal using a *decentralised trustless architecture*.

This decentralised trustless architecture is enabled by a combination of *cryptography*, which allows every participant in a blockchain to independently verify the integrity of the stored data, and *consensus protocols*, which allow participants to reliably agree on a single version of the blockchain. The necessity to agree on the history of the blockchain is a side effect of the distributed nature of the blockchain, which means that the participants' view of the blockchain can become out of sync.

1.2 Blocks and chains

The blockchain is the historical record of the transactions happened so far. Transactions are grouped into blocks, and valid blocks become part of the history when they are appended to the blockchain. Because of the lack of a centralised authority, any participant could virtually produce a block and broadcast it to the rest of the participants.

1.3 Consensus protocols

Consensus protocols are rules that all participants in a blockchain must respect. Such protocols are designed to provide incentives for participants to do useful work for the blockchain (e.g., producing new blocks, recording new transactions), and to protect the blockchain from attacks that aim at tampering with the blockchain's history. Participants who play by the rules are rewarded by the protocol, and participants who don't are penalised. In this section we give an overview of two common types of consensus protocols: Proof-of-Work and Proof-of-Stake. We cover the former for historical reasons, and the latter because it's the context for the AAPP.

1.3.1 Proof-of-Work

A Proof-of-Work (PoW) consensus protocol discourages tampering by making it unreasonably hard to modify the history of the blockchain. In order to modify the blockchain, a participant needs to demonstrate that they have carried out some *work* that is expensive and cannot be avoided. While doing the work is hard, verifying that that work has been done is easy. This asymmetry is at the core of the security offered by PoW protocols.

Many blockchains, including Bitcoin and Ethereum (pre-Merge), are based on PoW consensus protocols.

1.3.2 Proof-of-Stake

A Proof-of-Stake (PoS) consensus protocol discourages bad behaviour by requiring that participants put forward, or *stake*, a large sum of native cryptocurrency in order to earn the right to produce a new block. If a participant performs legitimate work, they get rewarded with some of the native currency, but if they fail to do so, they lose the resources that they have staked. The current version of Ethereum's (post-Merge) consensus protocol is of this kind.

1.3.3 Incentives

Due to the fact that it relies on a distributed trustless architecture, a blockchain requires participants to carry out useful work, e.g., producing new valid blocks, and to avoid disruptive behaviour, e.g., tampering with the blockchain history. To this end, consensus protocols make extensive use of (positive or negative) reward mechanisms to incentivise participants to align with the interests of the chain. Rewards are quantified in cryptocurrency, e.g., Bitcoin, Ether, etc., native to the blockchain.

1.4 Agreeing on one history

The distributed nature of the blockchain leads to some more complications. Because participants can produce and share new blocks all the time, the chain doesn't always look like a linear *sequence of blocks*. For instance, two participants may produce two "competing" blocks with the same predecessor, resulting in a *tree of blocks*. It should be stressed that this doesn't necessarily mean that one of the two participants is trying to do something nefarious, these situations can be often caused by network latency.

The trouble with trees of blocks is that each branch of the three corresponds to a different "truth" about the state of the blockchain, which needs to be reconciled in order for the chain to operate correctly and efficiently. From the perspective of the participants, this means agreeing on which branch of the tree represents the real truth. Blocks that are not part of such branch are called *uncle blocks*.

1.4.1 Longest chain

In some blockchains, a convention among the participants is that, in presence of a tree of blocks, the *longest branch* of the tree represents the one true history. In order to incentivise participants to commit to the longest branch, consensus protocols reward block producers asymmetrically, i.e., non-uncle blocks are rewarded significantly more than uncle blocks.

1.4.2 Voting on one history

Aside from the introduction of Proof-of-Stake, Ethereum's new consensus protocols introduces some novelties also around the mechanism to agree on the true history of the blockchain. Instead of adopting the longest branch as the one true blockchain, the protocol relies on *validators*, i.e., participants who agreed to stake their currency and provide services to the blockchain. Among other things, validators propose new blocks, and cast votes to elect what they believe is the *head*, i.e., the latest block of the correct branch, of the blockchain.

Such "votes" are called *attestations*, and are the subject of the AAPP defined in this document.

2 Attestations in Ethereum

As discussed in the previous section, a validator votes for the head of the chain by producing an attestation. A fundamental property of attestations is that, under certain conditions they can be aggregated.

An attestation is defined by the following data

- a set of **attesters** (possibly 1) that "back" it,

- some **attestation data** (among other things, the slot in which the attestation was produced, the hash of the block believed to be the head of the chain, and the hashes of two - source and destination - blocks, which are related to the mechanism by which older parts of the blockchain are finalised),

- a Boneh–Lynn–Shacham (BLS) **signature** produced using the private keys of the attesters.

The signature allows participants to verify the integrity of an attestation. Moreover, a crucial property of the BLS signature scheme is that it allows to *aggregate* multiple signatures into a single signature. This is very useful, as it enables participants to aggregate multiple attestations into a single one, which means that more votes can be included in a block. In order to be able to aggregate two attestations, the following conditions must hold

1. their attestation data must be identical,

2. their sets of attesters must be disjoint.

An aggregated attestation can contain a **maximum of 2048** attesters.

2.1 Epochs, slots, committees

Note: this document was written pre-Merge, but its focus was post-Merge Ethereum. In the following, the reader should assume that we are referring to post-Merge Ethereum.

Time in Ethereum is split in successive *epochs* of 6.4 minutes. Each epoch is further split into 32 *slots* of 12 seconds. Each block has a corresponding slot, however not all slots necessarily have a corresponding block. In fact, there should be only one block per slot, and validators are penalised for proposing more than one.

At each epoch one validator (the *proposer*) is selected[1] to propose a new block at a given slot. When its assigned slot starts, the proposer has one chance to produce a valid block.

At each epoch, the validators are shuffled and grouped into *committees*. Each committee has $n \in [128, 2048]$ validators. A slot has $m \in [1, 64]$ committees. The majority of the validators in a committee become *attesters* and will have to produce an attestation. Among these a small amount, on expectation 16, will become *aggregators* and their role will be to aggregate attestations.

2.2 Timeline

This section gives a brief overview of the timelines involved. These are not crucial to the problem formulation, but help understanding where the problem comes from, and why there is only a very limited amount of time available to solve it.

At each epoch, the work of each validator (be it a proposer, an attester, or an aggregator) needs to be carried out in the time of the slot assigned to them. The time before and after the assigned slot is still used to do useful work, mostly around gathering information, or finishing work that wasn't completed on time.

Before the start of their assigned slot, a **proposer** will start gathering attestations produced up to 32 slots earlier, as well as transactions, to include in the block. When the slot starts, the proposer will produce and share the block as quickly as possible, to maximise the likelihood that it will be considered by the attesters.

The **attesters** will wait for the block to be received, or up to 4 seconds from the start of the slot, and then make a decision on the head of the block. Depending on their *fork choice rule*, they may choose the newly created block, the block produced at the previous slot, or some other block. Within 8 seconds from the start of the slot, they will send their attestation to a sub-network where aggregators are listening for attestations.

Aggregators will gather all the attestations, and try to aggregate as many as possible into one. Before the end of the slot, the aggregators will broadcast

---

[1]Since we're operating in a decentralised and trustless context, nobody is actually "selecting" a validator to propose a block. Instead, the validators will independently and deterministically derive their role based on the agreed state of the blockchain.

the aggregated attestations with the network. These aggregated attestations (and possibly some unaggregated ones) will be then gathered by the proposer of the next slot, in a continuous loop.

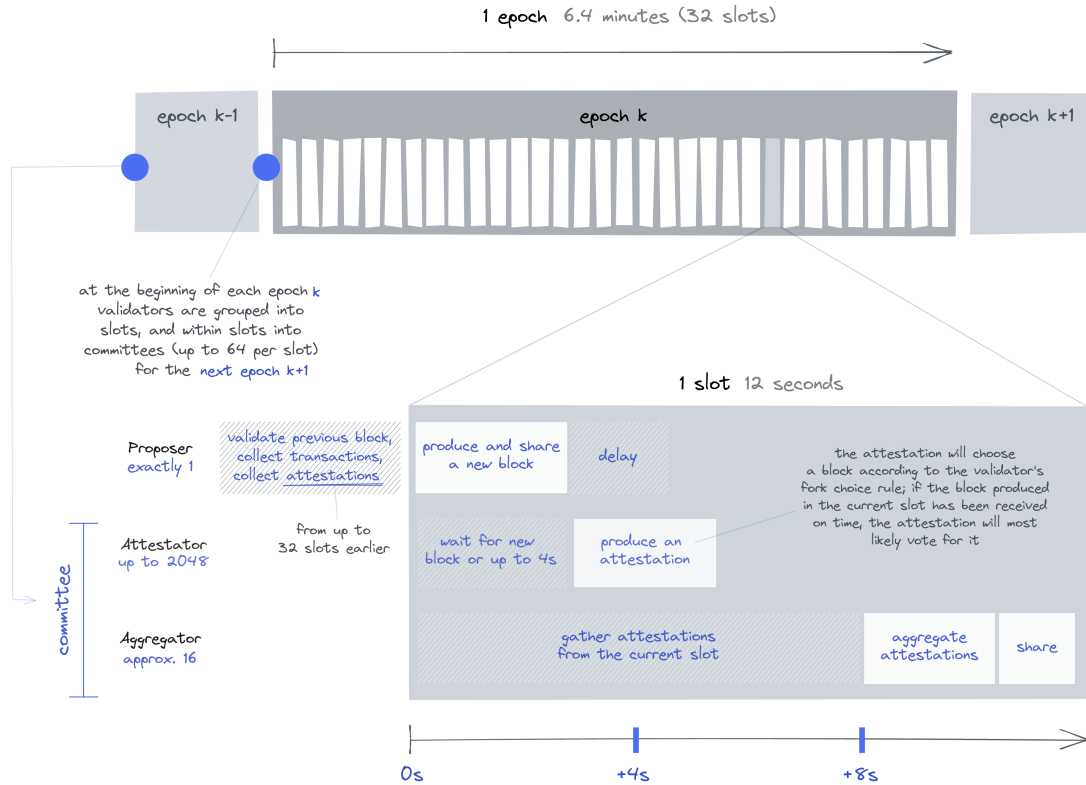The diagram in Figure 1 summarises the above timeline.

Figure 1: Timeline for the work of proposers, attesters, and aggregators.

3 The Attestation Aggregation and Packing Problem

As covered in the previous section, one key responsibility of the proposer is to collect attestations from the network, bundle them into a block (along with transactions), and then share the newly minted block with the network. Because the space in a block is limited, the proposer faces the problem of deciding *which* attestations to include in their block. Moreover, because different attestations contribute differently to the efficient operation of the blockchain, and because some attestations have overlaps, the problem of choosing a set of attestations to

6

include is an *optimisation* problem. We call this optimisation problem the Attestation Aggregation and Packing Problem (AAPP), and this section provides a more formal definition for it in terms of entities, constraints, and objectives.

The creation of a block B^s, where $s \in \mathbb{N}^{\geq 32}$ represents the index of the slot[2] at which the block is proposed, is (from the perspective of an aggregator) unequivocally associated with one and only one instance of the AAPP. For this reason, and to simplify the notation, we will therefore treat the index s as implicit, and drop it from the problem definition.

3.1 Entities

Let V be the set of all **validators**, and A be the set of all **attestations** available to the proposer for inclusion in a given block B. Given an attestation $a \in A$, we define the following properties

- $d_a \in D$ the attestation data of the a,

- $V_a \subseteq V$ a set of attesters agreeing on d_a.

To keep the problem formulation general, we intentionally avoid specifying what attestation data d_a (and its domain D) looks like. Instead, we focus on what can be done with it, e.g., compute the epoch in which it was created[3].

Let the set $E \subseteq \mathbb{N}^{\geq 0}$ denote the epochs[4] of the blockchain, the **epoch** at which each attestation data $d_a \in D$, $a \in A$ was created[5] is a function

$$e : D \to E. \tag{1}$$

Note that A is assumed to satisfy a number of properties that reflect the context of the particular problem instance being considered. For instance, attestations that are too old to be included in B are excluded from A. Moreover, it is assumed that

$$e(d_a) = e(d_b) \Rightarrow V_a \cap V_b = \emptyset, \quad \forall a, b \in A, \tag{2}$$

in other words a validator can produce at most one attestation data per epoch.

Each individual vote produced by an attester $v \in V$ at a given epoch $e \in E$ carries a **reward**. Such a reward is a function

$$r : E \times V \to \mathbb{N}. \tag{3}$$

This is the reward that the block proposer will collect by including the corresponding vote in the produced block. Like for $e(\cdot)$, the way $r(\cdot, \cdot)$ is computed is irrelevant to this problem formulation.

---

[2]For convenience, we avoid considering the first 32 slots of the blockchain, as some special handling happens at the start of a blockchain.

[3]In other words, we can think of each element $d \in D$ as implementing some kind of `AttestationData` *trait*.

[4]For the purpose of solving the AAPP for a given block B, it is sufficient for E to include the epoch during which the block is being created and the one immediately before, however this doesn't change the present problem formulation.

[5]While the practicalities of computing the $e(d_a)$ function for a given d_a is not strictly relevant to the problem formulation, we can count on d_a carrying information about the slot $s \in \mathbb{N}^{\geq 32}$ in which that d_a was created. Under this assumption, $e(d_a) = \lfloor s/32 \rfloor$.

3.2 Solution

A solution S to the AAPP for the block B is a *set of subsets* of A. Formally

$$S = \{P \mid P \subseteq A\}. \tag{4}$$

Note that, according to the Ethereum consensus protocol, a block contains attestations, not *sets* of attestations. This particular formulation allows us to consider, as candidates for inclusion in B, not only attestations in A, but also attestations that can be *produced* from attestations in A by means of aggregation. In this sense, one can think of each set $P \in S$ as a set of attestations that are meant to be aggregated before being included in the proposed B. Of course, aggregation is subject to conditions, which we explicitly model as constraints.

3.3 Constraints

While the previous section defines the general structure of a solution S, a *feasible* solution needs to satisfy some additional constraints.

Aggregation constraints. For a solution S to yield a valid block, the attestations within each element of S must be compatible for aggregation. Recall that two attestations $a, b \in A$ can only be aggregated if their aggregation data is identical, and their sets of attesters are disjoint. More formally, the following two properties must hold

$$d_a = d_b \qquad\qquad \forall a, b \in P, P \in S \tag{5}$$
$$V_a \cap V_b = \emptyset \qquad\qquad \forall a, b \in P, P \in S. \tag{6}$$

Because each set of attestations in a feasible S represents a valid aggregated attestation, it is convenient to lift the properties of attestations introduced above to sets of attestations in the natural way. Specifically, let S be a feasible solution, and $P \in S$ be a set of attestations, we will denote by

$$V_P = \bigcup_{a \in P} V_a \tag{7}$$

the set of attesters of the aggregated attestation represented by P, and by

$$d_P \tag{8}$$

the attestation data shared by each attestation $a \in P$.

Capacity constraint. A feasible solution S must respect the following capacity constraint

$$|S| \leq N \tag{9}$$

where $N = 128$ according to the Ethereum consensus protocol at the time of writing. This corresponds to the maximum number of attestations that can be included in a newly proposed block.

3.4 Objective

The quality of a solution S is a function of the sets of attestations that are chosen to be part of it. Let

$$P_e = \{P \mid e(d_P) = e, P \in S\}, \quad \forall e \in E \tag{10}$$

be a partition (by epoch) of the solution, and let denote by

$$V_{P_e} = \bigcup_{P \in P_e} V_P, \quad \forall e \in E. \tag{11}$$

the set of attesters covered by each set $P \in S$. The quality of a solution can be then defined as

$$R(S) = \sum_{e \in E} \sum_{v \in V_{P_e}} r(e, v). \tag{12}$$

In other words, the quality of a solution is the sum of the rewards that can be collected for including attestations, and considering each attester at most once per epoch. Given this definition, the objective function of the AAPP is then

$$\textbf{maximise } R(S). \tag{13}$$

4 Existing algorithm

Currently, the Lighthouse client developed by Sigma Prime solves the above problem using a two-stage approach. The first stage deals with the aggregation of attestations. The second stage deals with the packing of attestations. Both stages rely on greedy heuristics.

4.1 Aggregation stage

The aggregation stage maintains a pool of attestations, and each time a new attestation is received, a greedy heuristic attempts to aggregate it with everything else in the pool. In case of success, the attestations already in the pool will be aggregated with the new attestation. In case of failure, the new attestation will be added to the pool as a separate attestation.

4.2 Packing stage

The packing stage starts with an empty block and, at each iteration, greedily includes the attestation that contributes most to the reward of the block. This algorithm is a well-known approximation algorithm for the Maximum Covering Problem with a bound on the number of sets to include (Maximum k-Coverage Problem).

The approximation guarantees of this algorithm on the original problem carry over to its application in the current context. Specifically, if the set of attestations to include is final, i.e., there isn't any further aggregation beyond the aggregation stage, the algorithm produces a packing whose value is guaranteed to be at least ≈ 0.632 of the optimal value [1].

5 Conclusions

We have introduced the context around the Attestation Aggregation and Packing Problem (AAPP) and provided a formal definition of it as an optimisation problem. In the next report, we present a series of approaches to solve this problem exactly.

References

[1] Dorit S Hochbaum and Anu Pathria. Analysis of the greedy approach in problems of maximum k-coverage. *Naval Research Logistics (NRL)*, 45(6):615–627, 1998.