

Timelord algorithm

A timelord keeps track of the current peak of beacon chain, which includes an infused block at a certain height, and signage points from the peak onward. A timelord might receive new blocks to infuse, new peaks (blocks which are already infused), or new signage points.

How does a timelord decide which challenges to create proofs of time on, given a limited number of available processors? While ASICs are likely to develop in the future, at the moment the fastest classgroup VDF implementations are on general purpose hardware. Furthermore, even after the development of ASICs, it's important to emphasize that any user with a CPU can be a timelord, to provide fallbacks in the case that the ASIC timelords go down, or becomes malicious, etc.

In general, timelords work on the heaviest chain. They create proofs of time at the signage points, and broadcast these to the network as they are reached. Timelords also infuse blocks as often as they can. When the timelord receives an infused block which has a greater weight than the current peak, they switch to it immediately.

Timelords also run the three VDF chains in parallel. Therefore, at least three fast CPU cores are necessary to advance the blockchain at an efficient rate. Extra CPU cores will be necessary to create proofs, but they do not have to be as fast.

If the timelord receives a challenge with less weight than their current peak, they ignore it. If the timelord receives a challenge point later in the current chain, they do the safe thing: ignore it. The reason is that by switching to a point further in the future, the timelord might be skipping the infusion of blocks, and thus orphaning valid blocks.

If the timelord receives a block for infusion which is late (we have already reached the challenge point at which the block should have been infused), the timelord ignores the block, since infusing to it would allow attackers to instigate a withholding attack, in an attempt to re-org or DoS the chain.

Therefore, the main operation of the timelord involves keeping a cache of future blocks to infuse, broadcasting challenge points when they are reached, and infusing blocks when they reach their challenge points.

If the timelord receives a challenge with the same weight as the current peak, they choose the unfinished block which they saw first (that is, the block that has not been infused yet), as opposed to choosing the infused block (peak) which they saw first. This also disincentivizes the withholding of blocks.