Proposer Builder Separation: The Infinite Buffet

Mark Mackey
Sigma Prime | Lighthouse

@ethDreamer

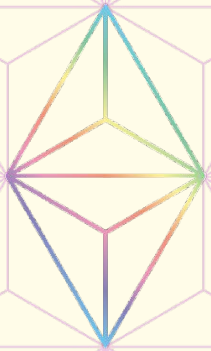

sigma prime

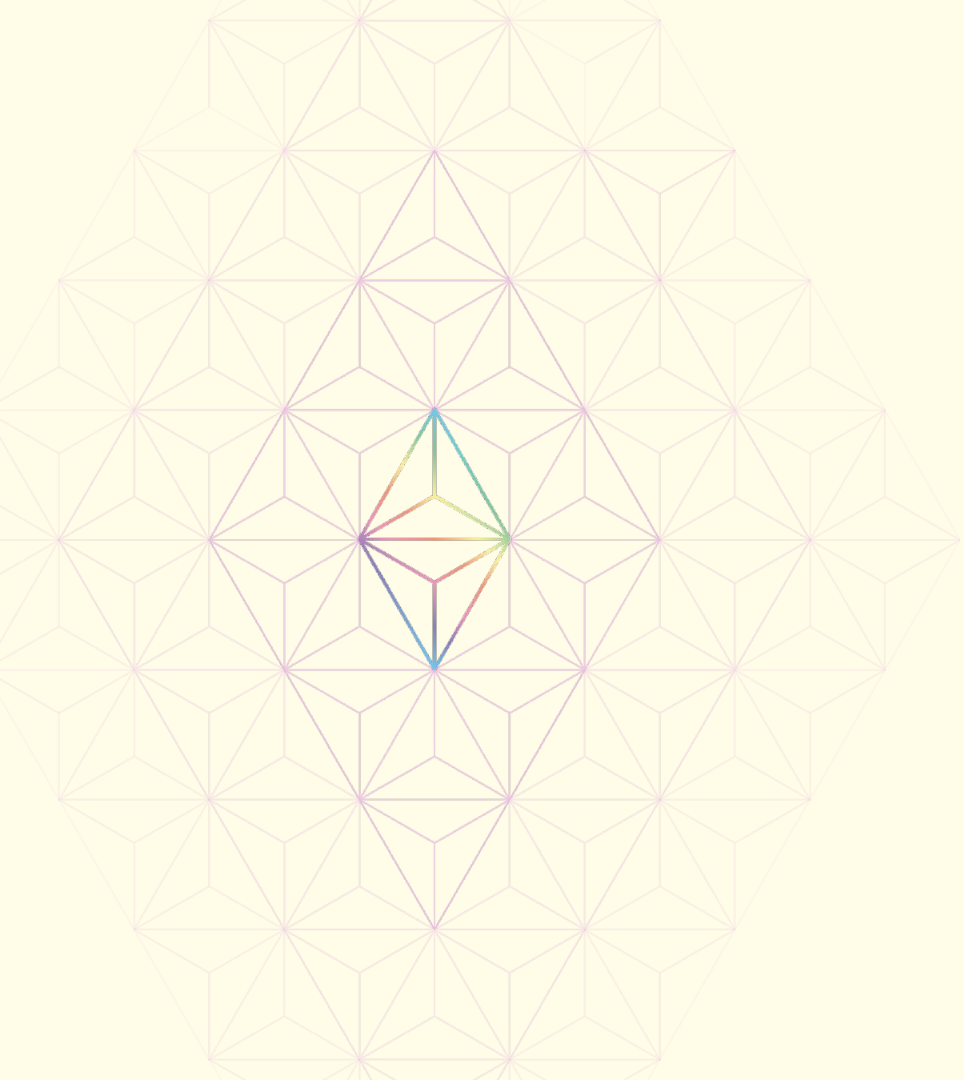

Background

2009 - 2016

January 3 2009

July 30 2015

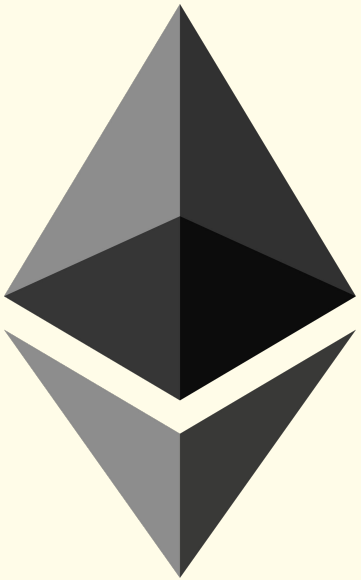

The first DEXs

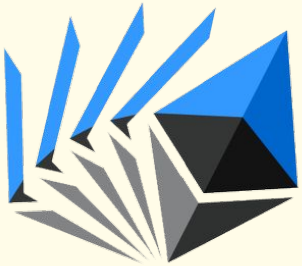

August 2016

EtherDelta

February 2017

2017

November 2 2018

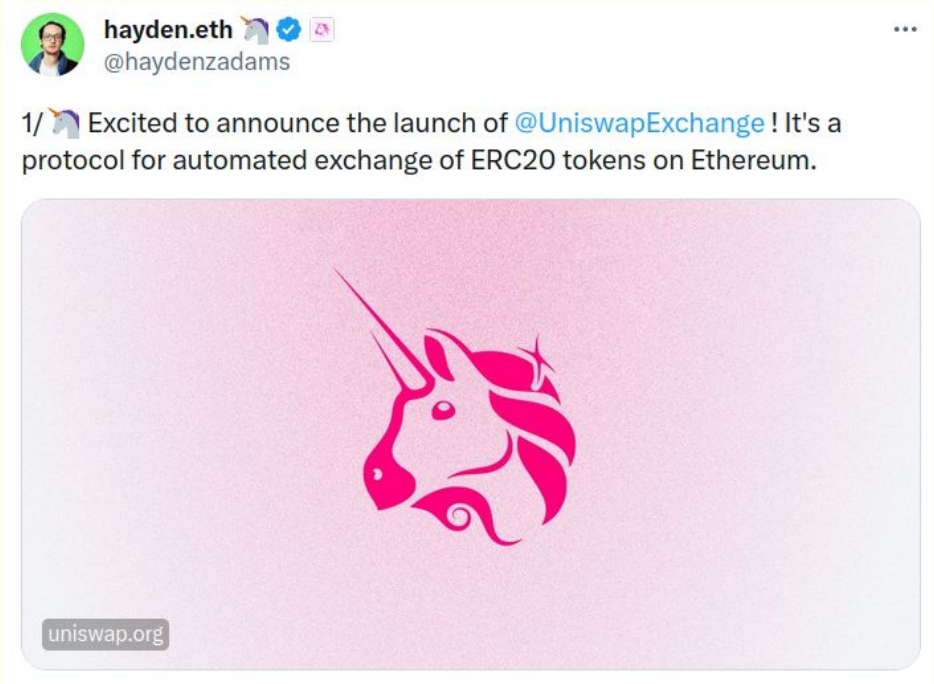

April 10 2019

Computer Science > Cryptography and Security

[Submitted on 10 Apr 2019]

Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges

Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, Ari Juels

Miner-extractable value (MEV): We introduce the notion of MEV, value that is extractable by miners directly from smart contracts as cryptocurrency profits. One particular source of MEV is *ordering optimization (OO) fees*, which result from a miner's control of the ordering of transactions in a particular epoch. PGAs and pure revenue opportunities provide one source of OO fees. We show that MEV creates systemic consensus-layer vulnerabilities.

Our results raise many important questions, some about the arbitrage community itself. For example, it would benefit arbitrageurs to collude with miners, but we observe no such collusion: Preliminary experiments show that bot transactions are equally distributed across mining pools. Are there incentives to avoid collusion, such as concern about the exogenous impact of miner malfeasance coming to light?

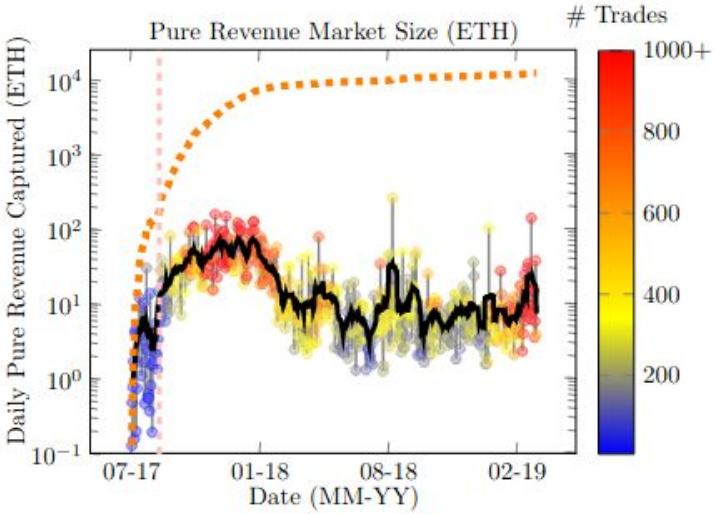

DeFi Summer 2020


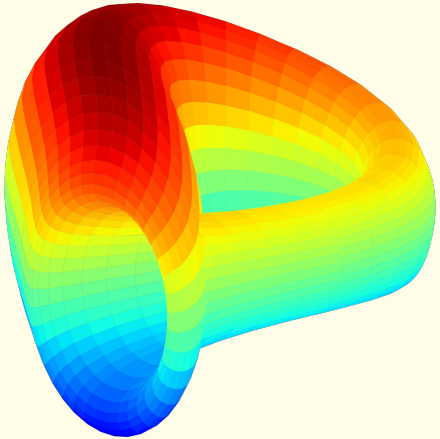






Q3 2020

MEV Inte

Paradigm

About | Portfolio | Opportunities | Contact

ource

rk Forest

antopoulos

On the Instability of Bitcoin Without the Block Reward

Miles Carlsten
carlsten@cs.princeton.edu

Harry Kalodner
kal princeton.edu

yanan
eton.edu

S. Matthew Weinberg
sm @princeton.edu

ABSTRACT

Bitcoin provides two incentives for min
and transaction fees. The former accoun
jority of miner revenues at the beginning
it is expected to transition to the latter as
dwindle. There has been an implicit be
miners are paid by block rewards or tran
not affect the security of the block chain.

We show that this is not the case. Our key
with only transaction fees, the variance of the bl
very high due to the exponentially distributed blo
time, and it becomes attractive to fork a "wealthy" block
to "steal" the rewards therein. We show that this results
in an equilibrium with undesirable properties for Bitcoin's
security and performance, and even non-equilibria in some
circumstances. We also revisit selfish mining and show that

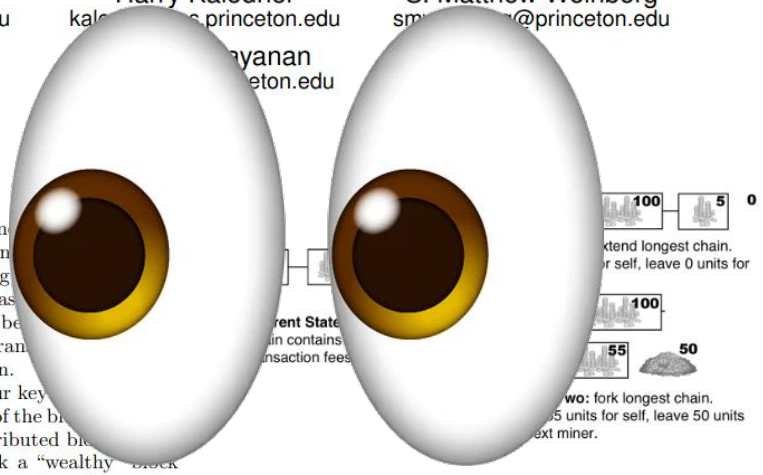

Figure 1: One possible state of the block chain and two possible actions a miner could take.

ERC-20 Market Ca

Ethereum ($ETH) vs. ER

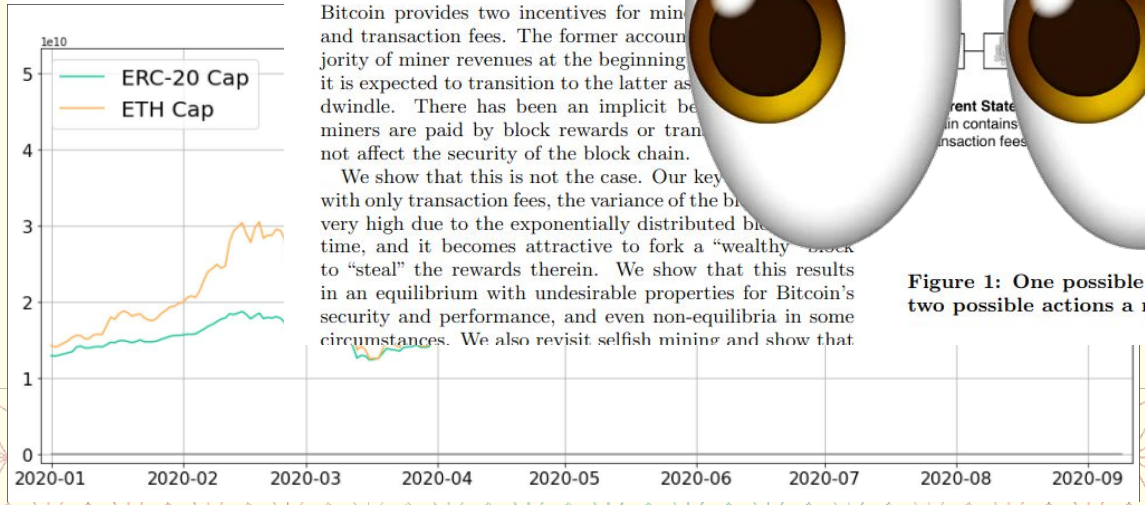

November 2020

“Competition for MEV opportunities leads to Ethereum consensus security instability due to the creation of incentives for … **permissioned communication infrastructure between traders and miners**. Such an infrastructure erodes the neutrality, transparency, decentralization, and permissionlessness of Ethereum today”

Flashbots—Frontrunning the MEV Crisis

Alex Obadia

22.11.20 · 9 min read

2 replies

Flashbots is a research and development organization formed to mitigate the negative externalities and existential risks posed by miner-extractable value (MEV) to smart-contract blockchains. We propose a permissionless, transparent, and fair ecosystem for MEV extraction to reinforce the Ethereum ideals.

Before MEV / Flashbots

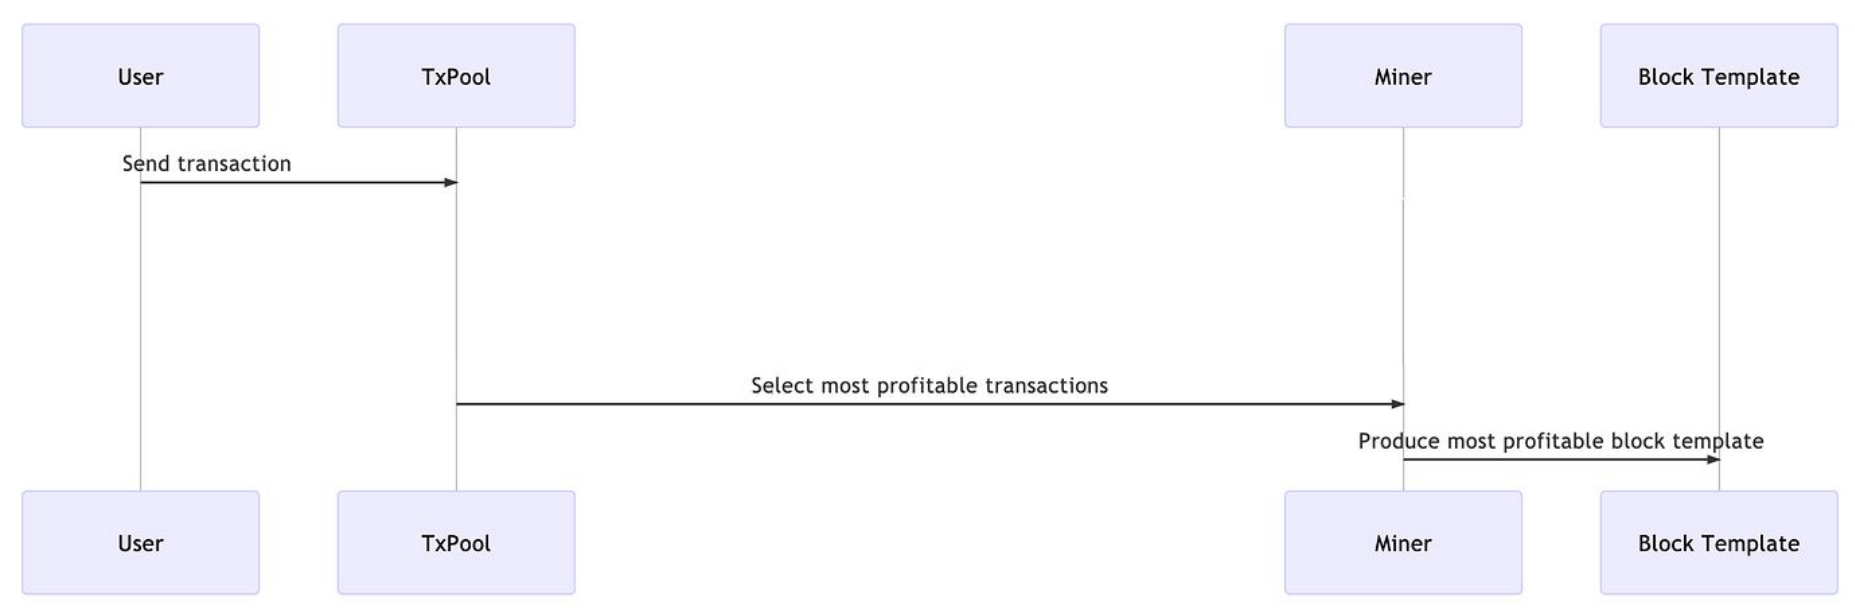

Flashbots Architecture

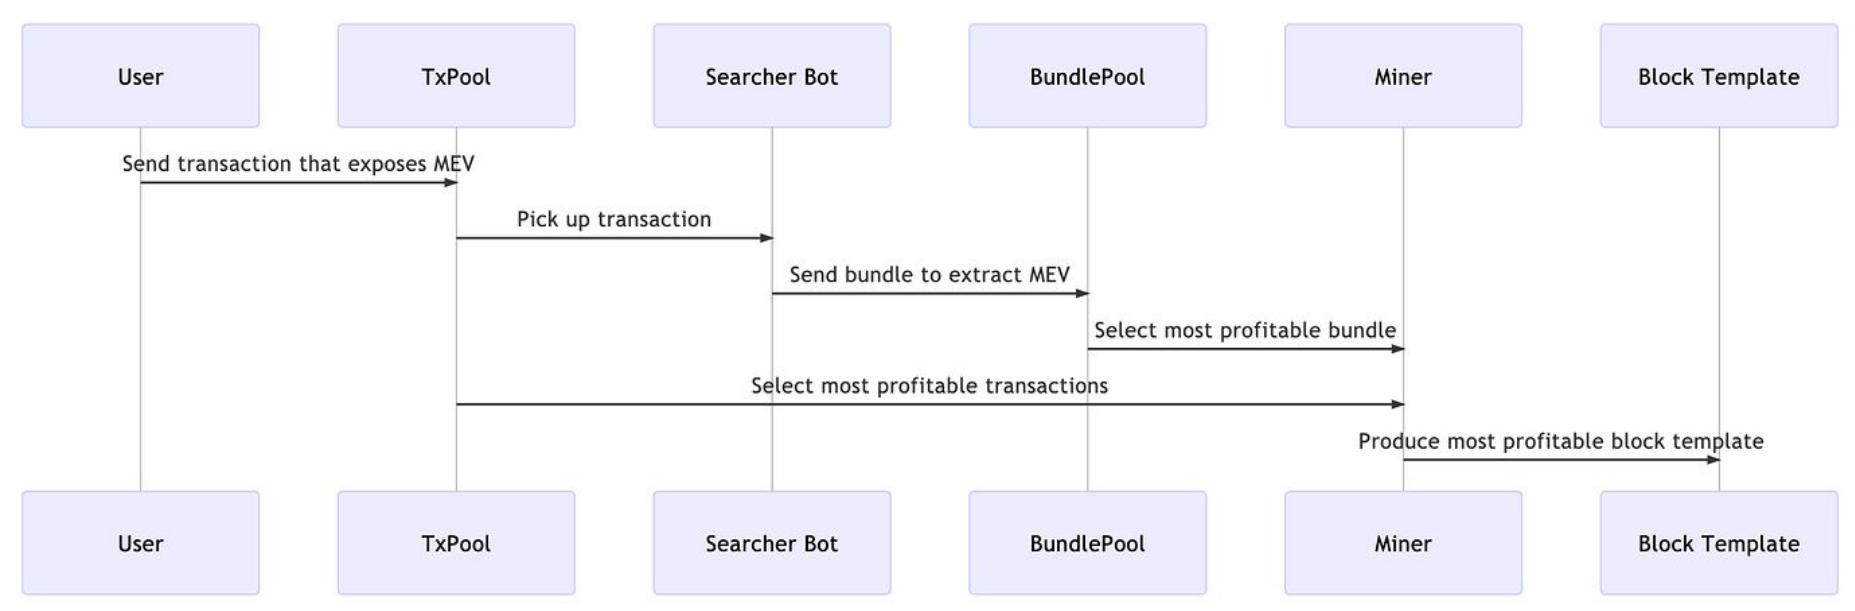

Q1 2021

Flashbots Transparency Report — January 2021

Stephane Gosselin

The first Flashbots bundle was mined in block 11550019 on Dec-29–2020 04:33:22 PM UTC, kicking off what is poised to be the year of MEV.

Flashbots Transparency Report — March 2021

Stephane Gosselin

March 2021 has been the month of adoption and alignment for Flashbots. While one of the largest ethereum mining pools announced it directly extracts MEV by running its own sandwich bots but got rekt'd by a naughty salmonella upon launch, most miners have chosen a different path — to participate in a fair, and efficient sealed-bid MEV auction via flashbots alpha. At the time of writing, **12 mining pools** accounting for **over 58%** of ethereum network hashrate are mining on flashbots.

June 2021 Proposer Builder Separation

Proposer/block builder separation-friendly fee market designs

Economics

vbuterin Jun '21

Special thanks to Justin Drake and the Flashbots team for feedback and discussion.

A major risk threatening the ongoing decentralization of consensus networks is the economics around miner extractable value (MEV), sophisticated tricks to extract profit from the ability to choose the contents of the next block. A simple example of MEV is arbitraging all on-chain decentralized exchanges against price movements that have happened since the previous block. While normal PoS rewards are reasonably egalitarian, as single validators earn the same rate of return as powerful pools, there are significant economies of scale in finding sophisticated MEV extraction opportunities. A pool that is 10x bigger will have 10x more opportunities to extract MEV but it will also be able to spend much more effort on making proprietary optimizations to extract more out of each opportunity. In addition to this problem, MEV also complicates decentralized pooling, as in a decentralized pool there would still need to be one entity packaging and proposing the block, and they can easily secretly extract MEV without sharding that revenue with the pool itself.

The best-known solution is **proposer/block-builder separation**. Instead of the block **proposer** trying to produce a revenue-maximizing block by themselves, they rely on a market where outside actors that we call **block-builders** produce **bundles** consisting of complete block contents and a fee for the proposer, and the proposer chooses the bundle with the highest fee. The proposer's choice is reduced to picking the highest-fee bundle, an algorithm so simple that in a decentralized pool it can even be done inside an MPC to prevent cheating.

PBS: Desired Properties

Desired properties for a proposer/builder separated block proposal design

We will focus on five major desired properties:

- **Untrusted proposer friendliness**: there's minimal or no risk that a proposer will screw over a block builder, so block builders have no incentive to prefer proposers that have some off-chain reputation or personal connection to the builder (as that would favor large pools).
- **Untrusted builder friendliness**: there's minimal or no risk that a block builder will screw over a proposer, so proposers have no incentive to favor builders that have some off-chain reputation or personal connection to the proposer (as that would make it harder for new builders to enter the market). If deposits are needed to accomplish this, they should be maximally low.
- **Weak proposer friendliness**: the mechanism should not require proposers to have either (i) high bandwidth or other computational resources or (ii) high technical sophistication
- **Bundle un-stealability**: proposers should not be able to take bundles proposed by block builders and extract transactions from them to make their own bundles, preventing the block builder from earning a profit (and possibly harming them even further)
- **Consensus-layer simplicity and safety**: the mechanism should continue to be safe and ideally be covered by the same analysis as the existing block proposal mechanism from a consensus-layer perspective

MEV-Boost: Merge ready Flashbots Architecture

The Merge ■ mev

thegostep 2 Nov '21

MEV-Boost: Merge ready Flashbots Architecture

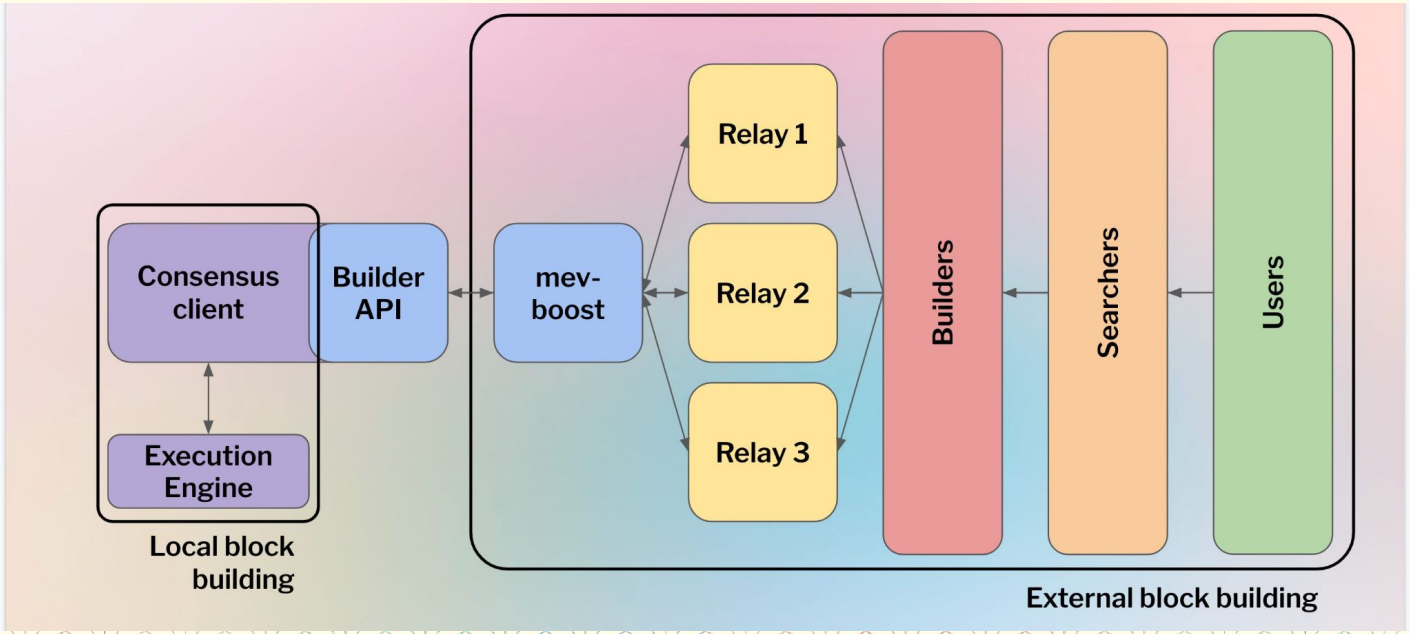


Sep 15 06:43:05.010 INFO Ready for the merge
Sep 15 06:43:08.002 WARN Did not advance head state
SnapshotCache(0x60e751f7d2cf0ae24b195bda37e9add56a7d8c4b75469c01Bc0f911518c3bae0)
state_advance
Sep 15 06:43:13.480 INFO
Sep 15 06:43:13.483 INFO Proof of Stake Activated
service: beacon
Sep 15 06:43:13.483 INFO
Sep 15 06:43:13.485 INFO
Sep 15 06:43:13.485 INFO

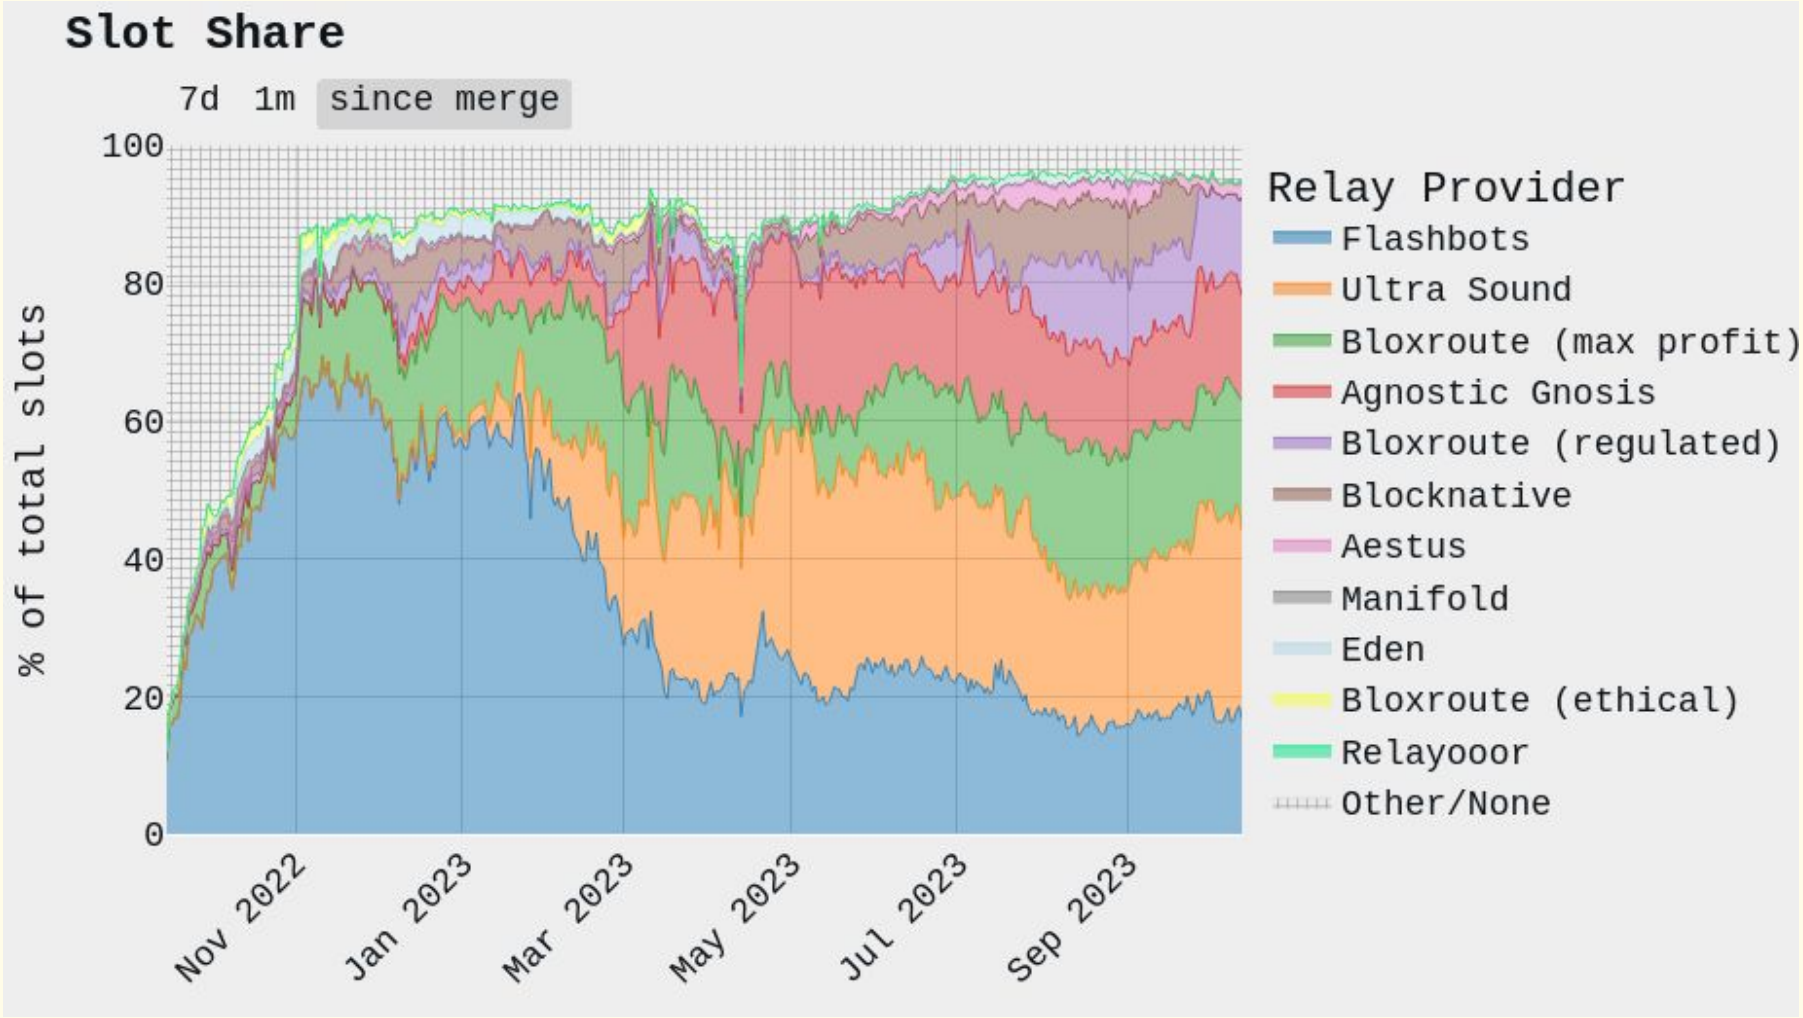
Slot Share
7d
1m
since merge
% of total slots
100
80
60
40
20
0
Nov 2022
Jan 2023
Mar 2023
May 2023
Jul 2023
Sep 2023
Relay Provider
Flashbots
Ultra Sound
Bloxroute (max profit)
Agnostic Gnosis
Bloxroute (regulated)
Blocknative
Aestus
Manifold
Eden
Bloxroute (ethical)
Relayooor
Other/None

Block-building today (in Proof-of-Stake)

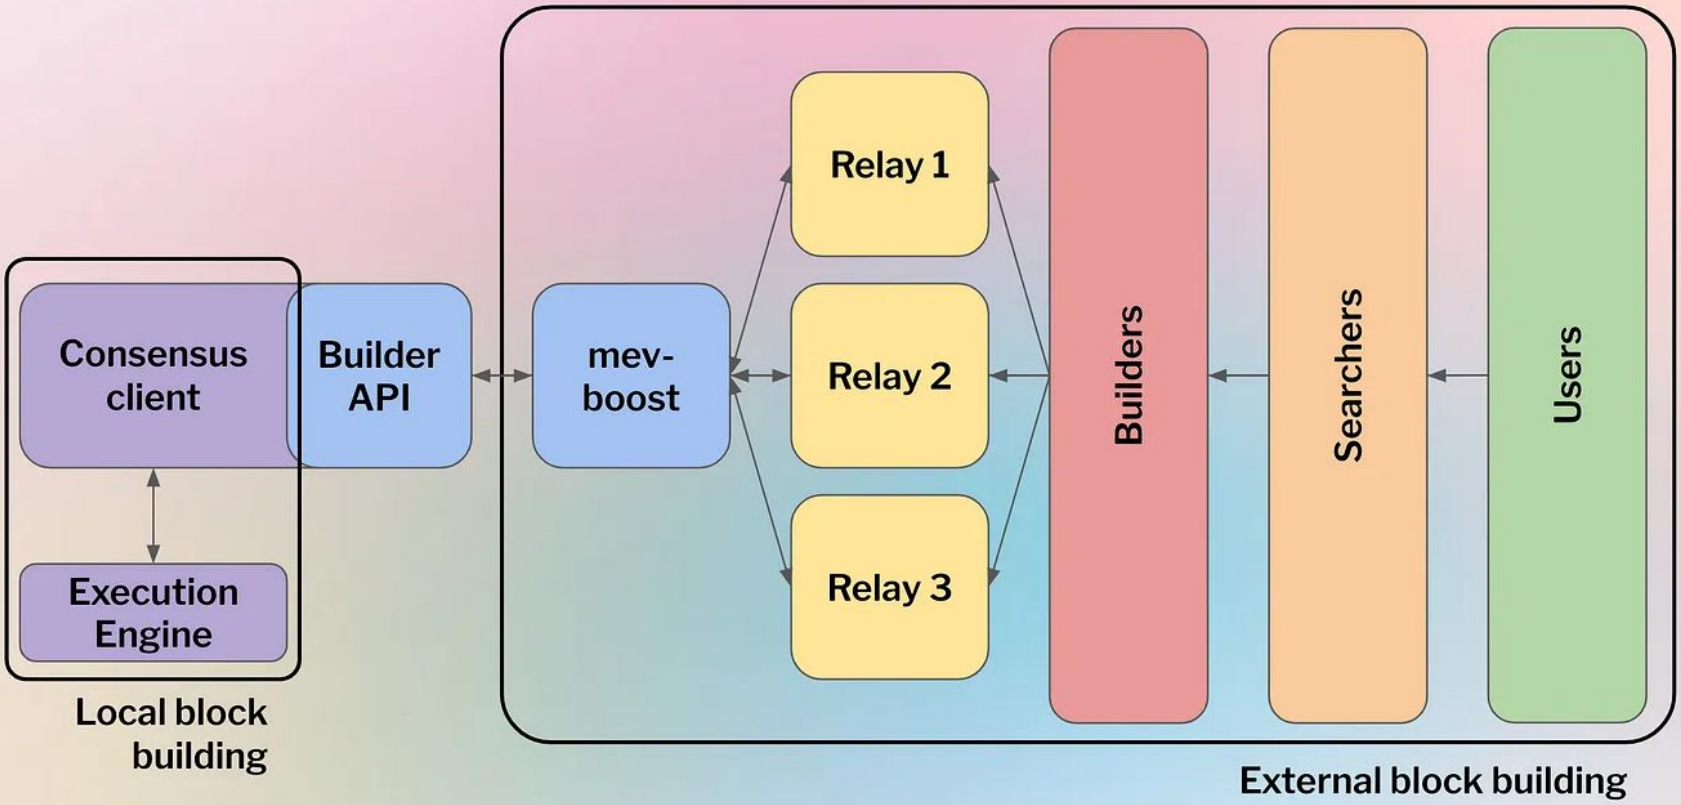

Block-building today (in Proof-of-Stake)

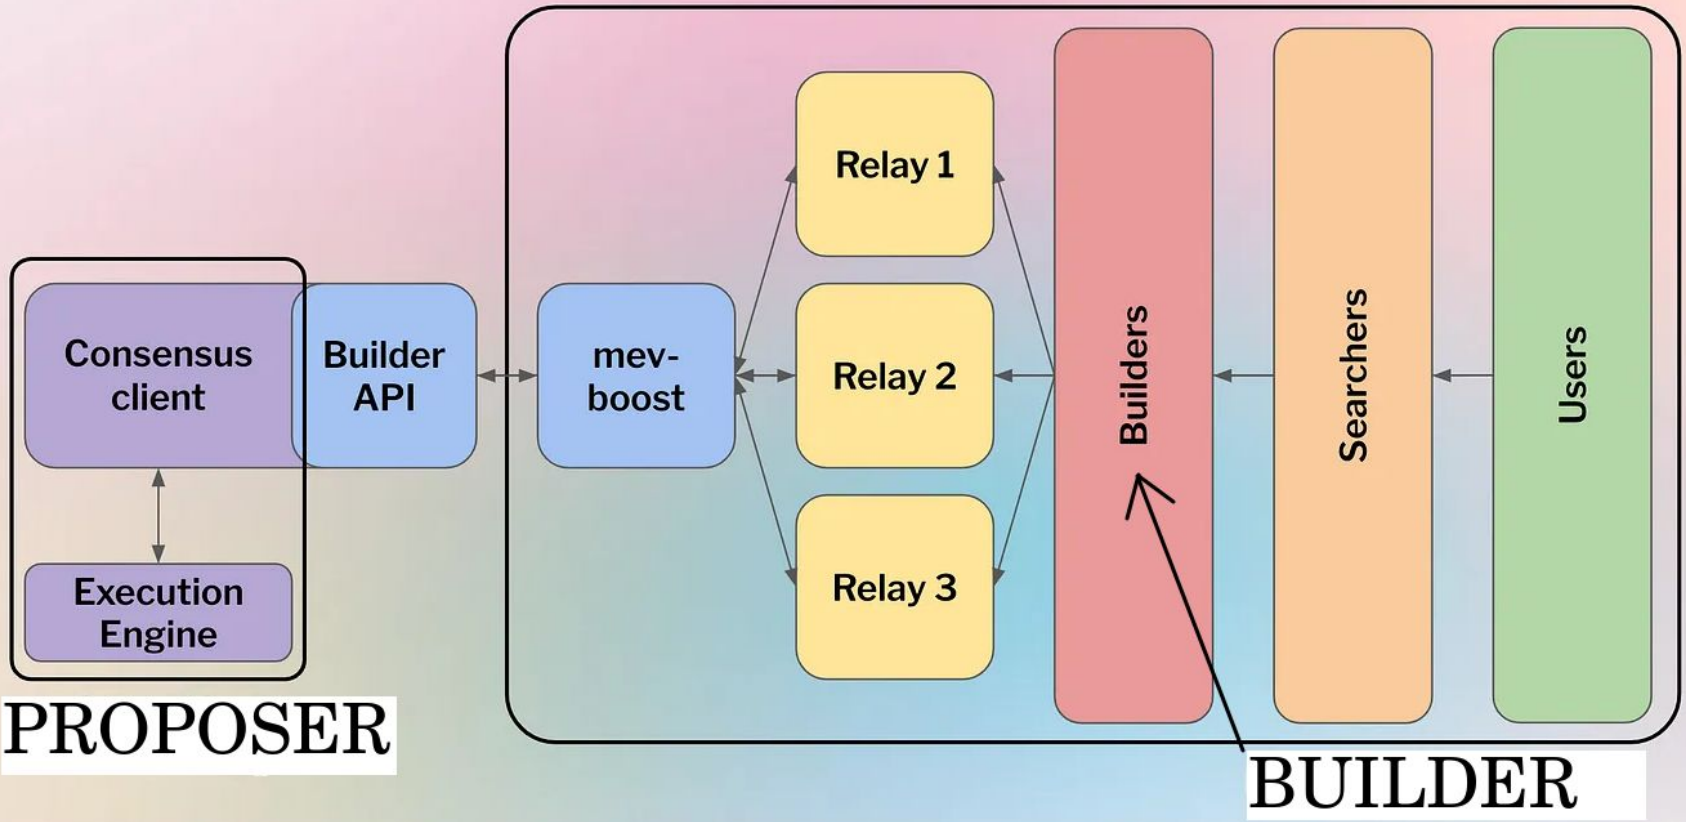

PBS: Desired Properties

Desired properties for a proposer/builder separated block proposal design

We will focus on five major desired properties:

- **Untrusted proposer friendliness**: there's minimal or no risk that a proposer will screw over a block builder, so block builders have no incentive to prefer proposers that have some off-chain reputation or personal connection to the builder (as that would favor large pools).
- **Untrusted builder friendliness**: there's minimal or no risk that a block builder will screw over a proposer, so proposers have no incentive to favor builders that have some off-chain reputation or personal connection to the proposer (as that would make it harder for new builders to enter the market). If deposits are needed to accomplish this, they should be maximally low.
- **Weak proposer friendliness**: the mechanism should not require proposers to have either (i) high bandwidth or other computational resources or (ii) high technical sophistication
- **Bundle un-stealability**: proposers should not be able to take bundles proposed by block builders and extract transactions from them to make their own bundles, preventing the block builder from earning a profit (and possibly harming them even further)
- **Consensus-layer simplicity and safety**: the mechanism should continue to be safe and ideally be covered by the same analysis as the existing block proposal mechanism from a consensus-layer perspective

Block-building today (in Proof-of-Stake)

DONE

Block-building today (in Proof-of-Stake)

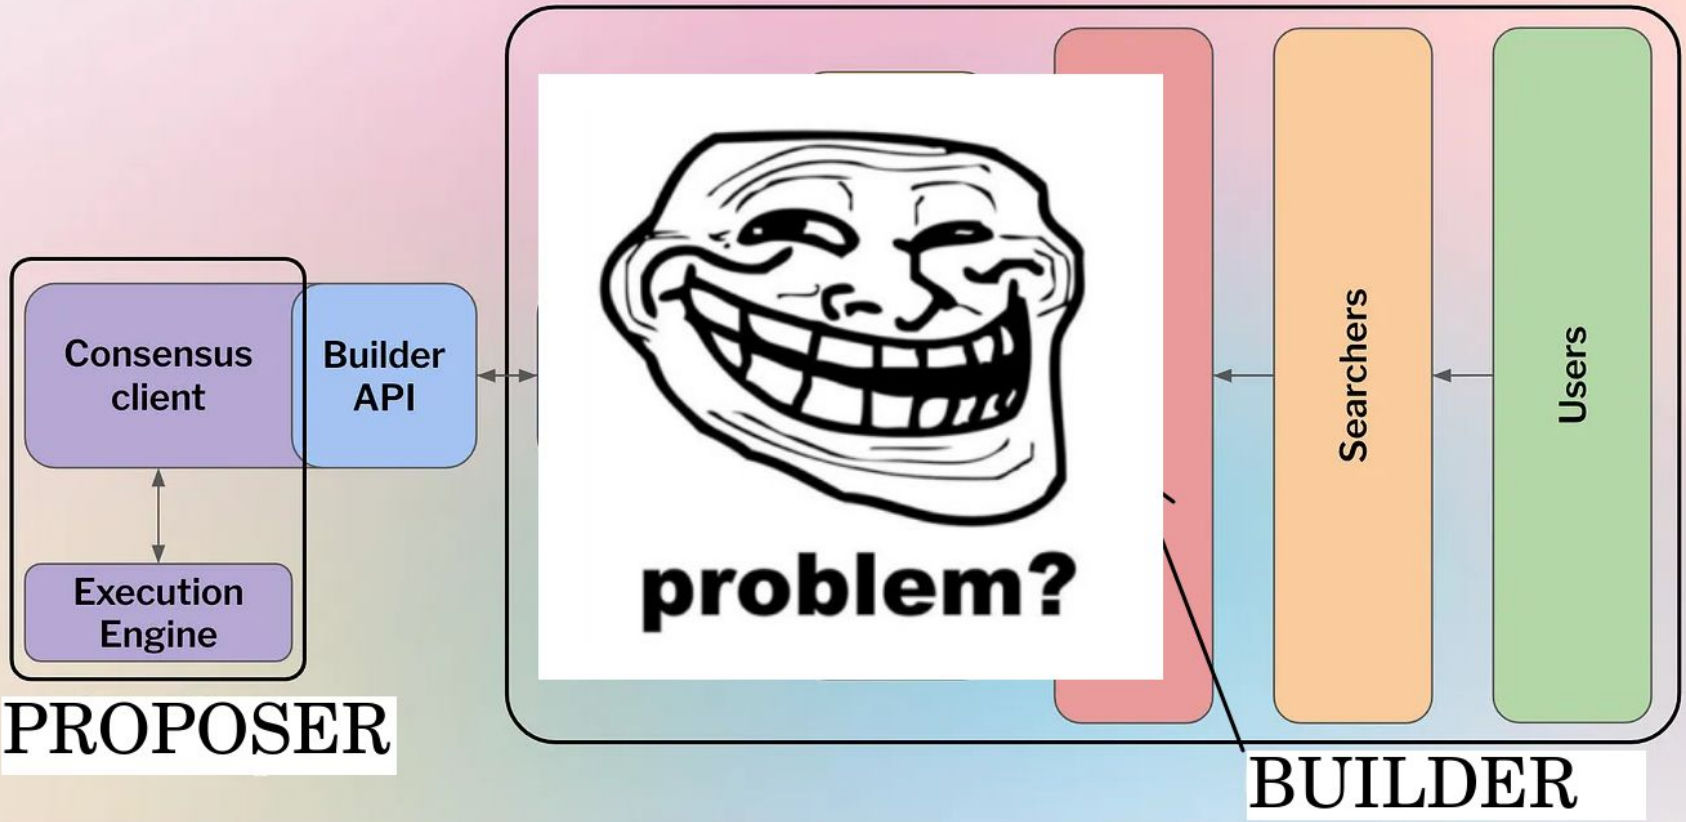

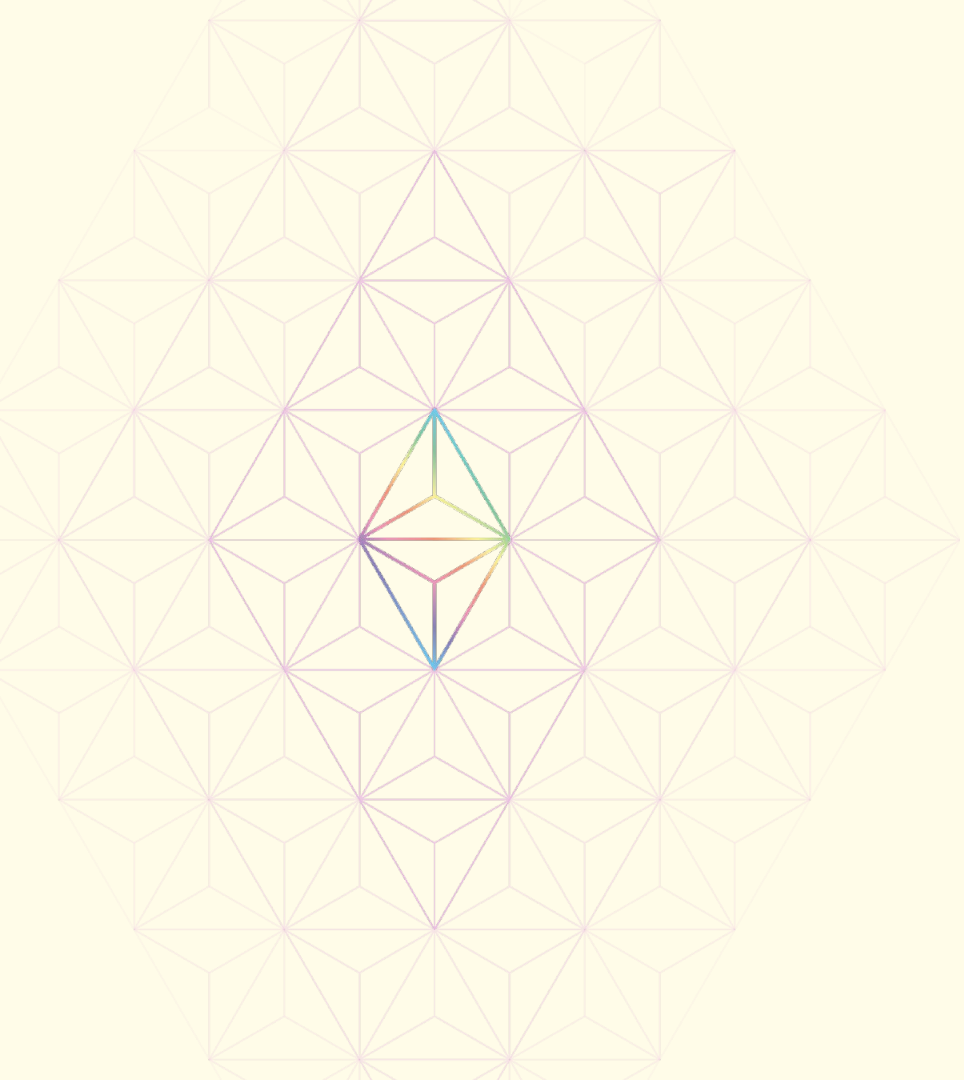

Problems with the Status Quo

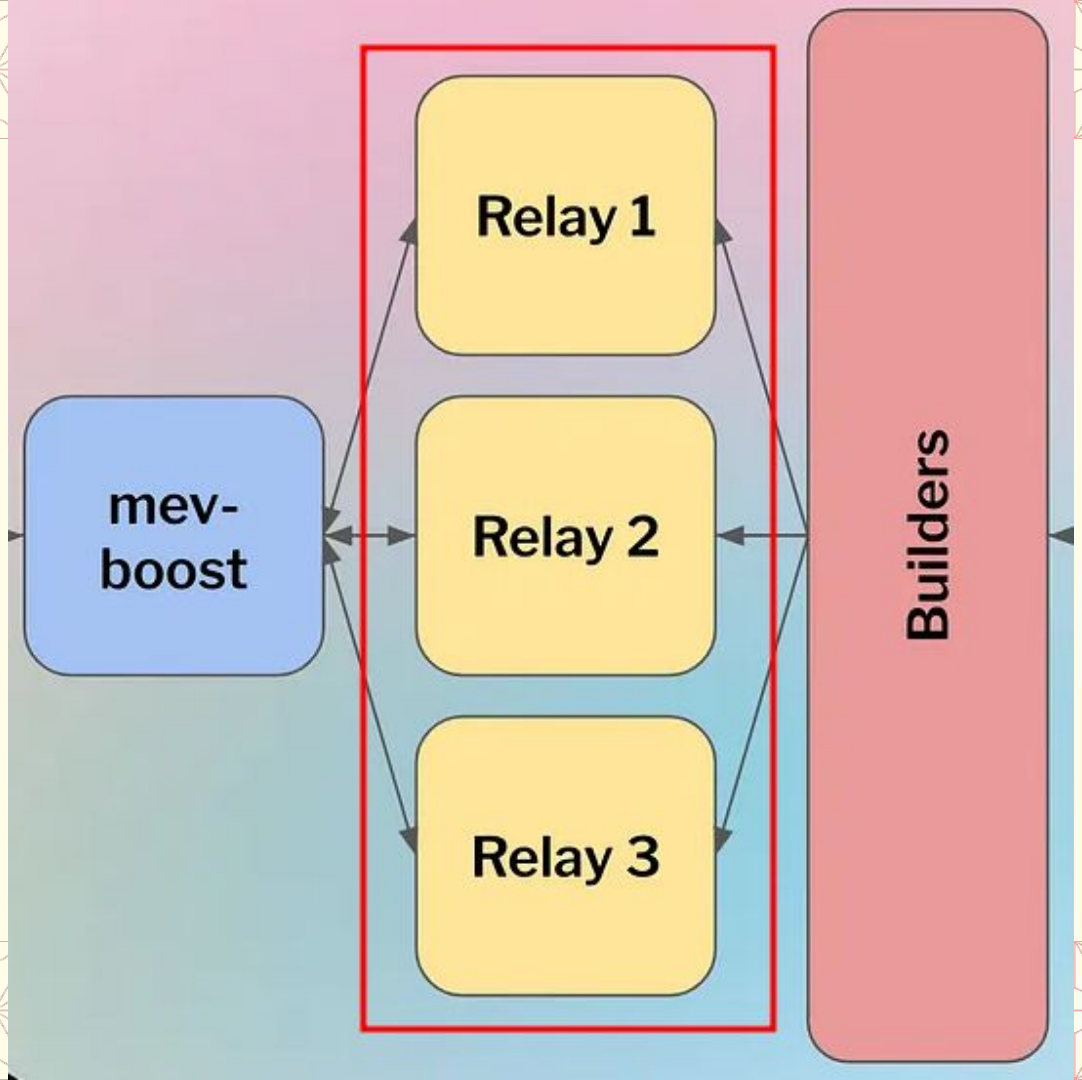

Relays

Relays Oppose Ethereum's Values

Trustlessness

- Builders trust relays:
  - MEV-stealing/unbundling
- Proposers trust relays:
  - Block validity enforcement
  - Payment verification/Collateral Escrow
  - Payload Reveal

Desired properties for a proposer/builder separated block proposal design

We will focus on five major desired properties:

- **Untrusted proposer friendliness**: there's minimal or no risk that a proposer will screw over a block builder, so block builders have no incentive to prefer proposers that have some off-chain reputation or personal connection to the builder (as that would favor large pools).
- **Untrusted builder friendliness**: there's minimal or no risk that a block builder will screw over a proposer, so proposers have no incentive to favor builders that have some off-chain reputation or personal connection to the proposer (as that would make it harder for new builders to enter the market). If deposits are needed to accomplish this, they should be maximally low.
- **Weak proposer friendliness**: the mechanism should not require proposers to have either (i) high bandwidth or other computational resources or (ii) high technical sophistication
- **Bundle un-stealability**: proposers should not be able to take bundles proposed by block builders and extract transactions from them to make their own bundles, preventing the block builder from earning a profit (and possibly harming them even further)
- **Consensus-layer simplicity and safety**: the mechanism should continue to be safe and ideally be covered by the same analysis as the existing block proposal mechanism from a consensus-layer perspective

Relays Oppose Ethereum's Values

Decentralization

There are only a small number of relays

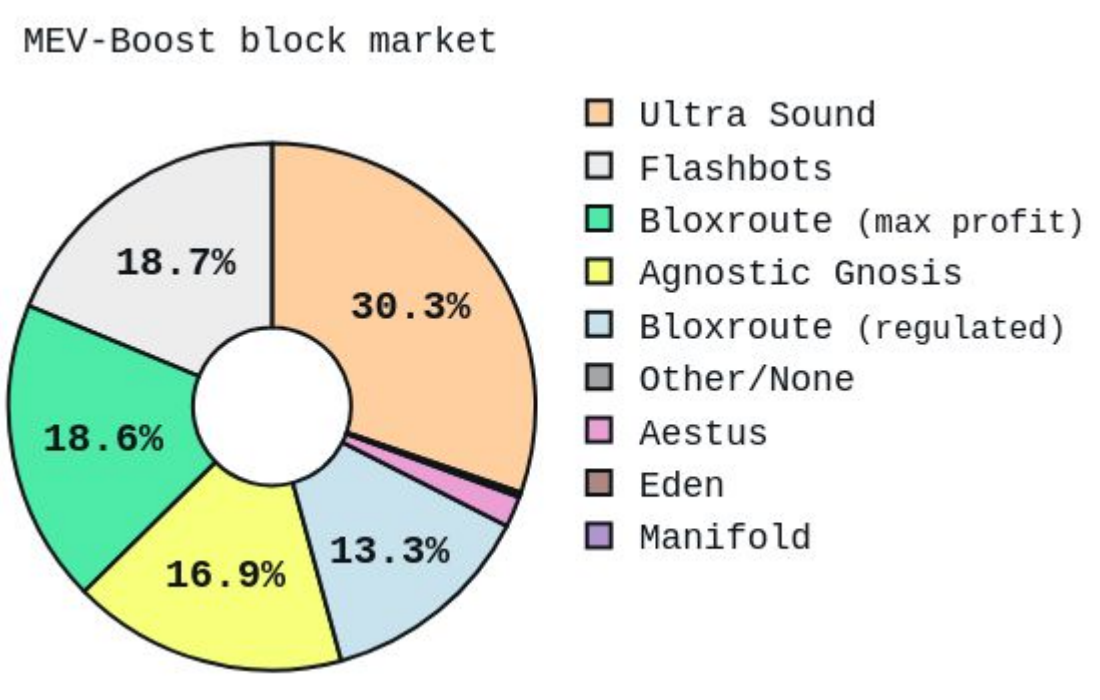

Relays Oppose Ethereum's Values

Censorship Resistance

Relays can censor blocks

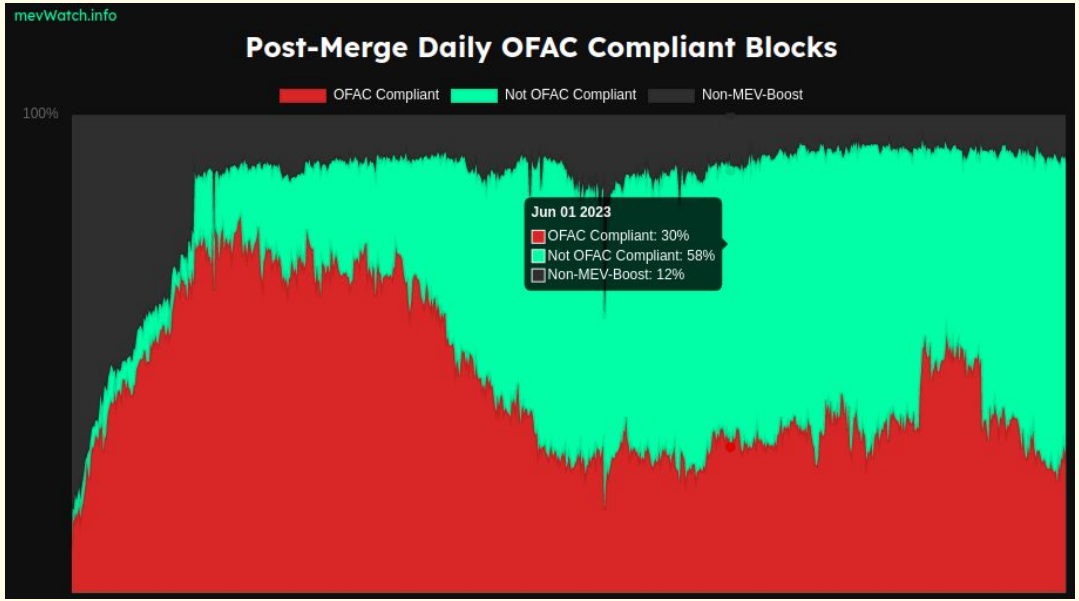

Relays Are Expensive Public Goods

- Relay operational costs range from $20k - $100k annually depending on desired performance
- This does not include engineering & DevOps costs associated with running a highly-available production service
- No clear funding model for future operation besides grants & guilds

Relays are Targets for Attacks

- The Low-Carb Crusader
  - April 3rd, 2023
  - $30M stolen from flashbots builder
- This attack can be extended to general equivocation attacks on builders
- Relay's response to these unbundling attacks cause consensus instability: 5x increase in reorged blocks

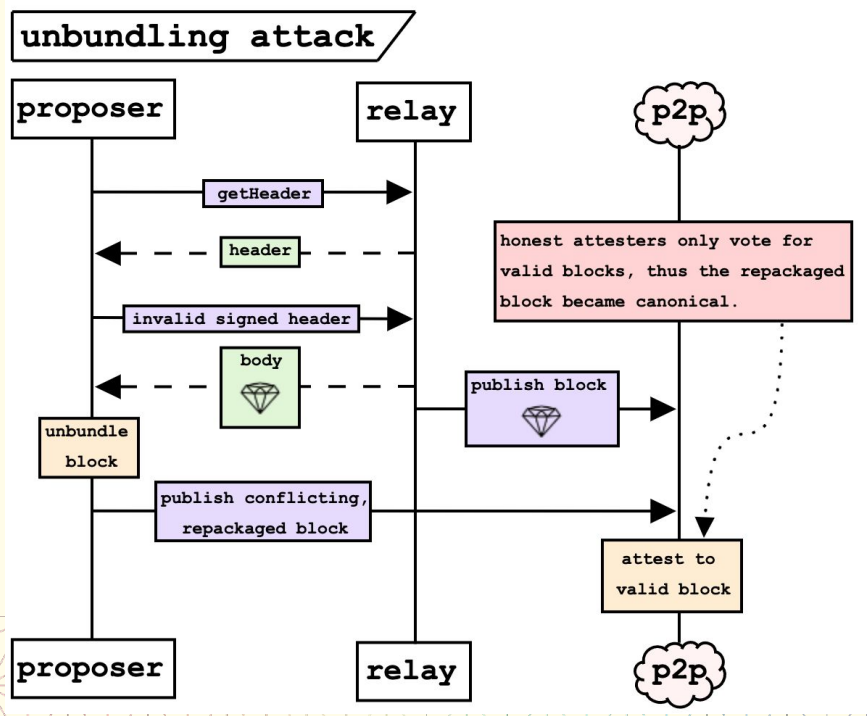

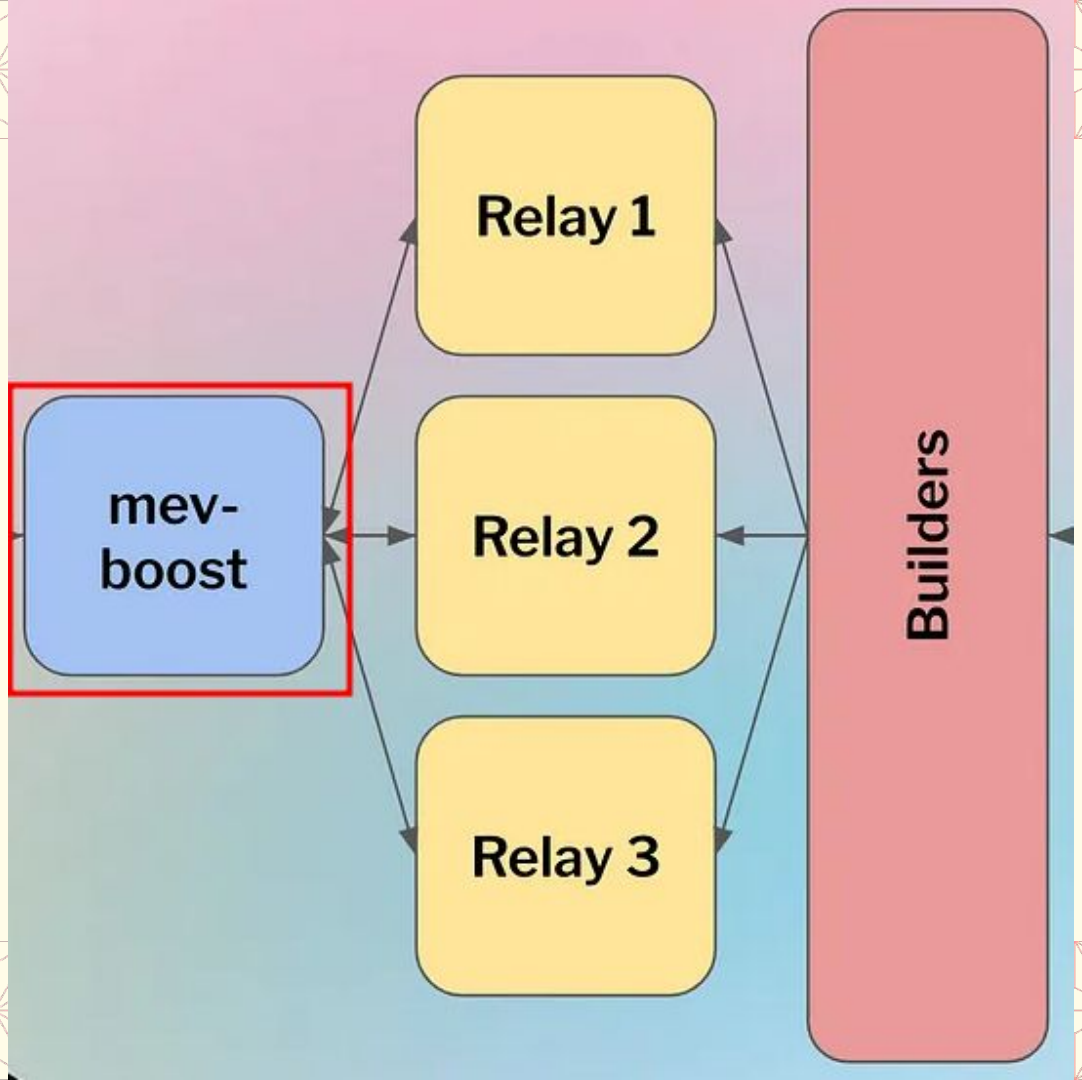

mev-boost

Out-of-protocol Software Is Brittle

- Mev-boost does not inherit the benefits of client diversity & the full consensus specification process
- There are significant core-dev coordination costs involved in maintaining compatibility between beacon clients & relays for every fork
  - builder spec
  - relay spec
- During Shapella upgrade, a bug in the Prysm client code that interacts with mev-boost resulted in a 10x spike in missed slots

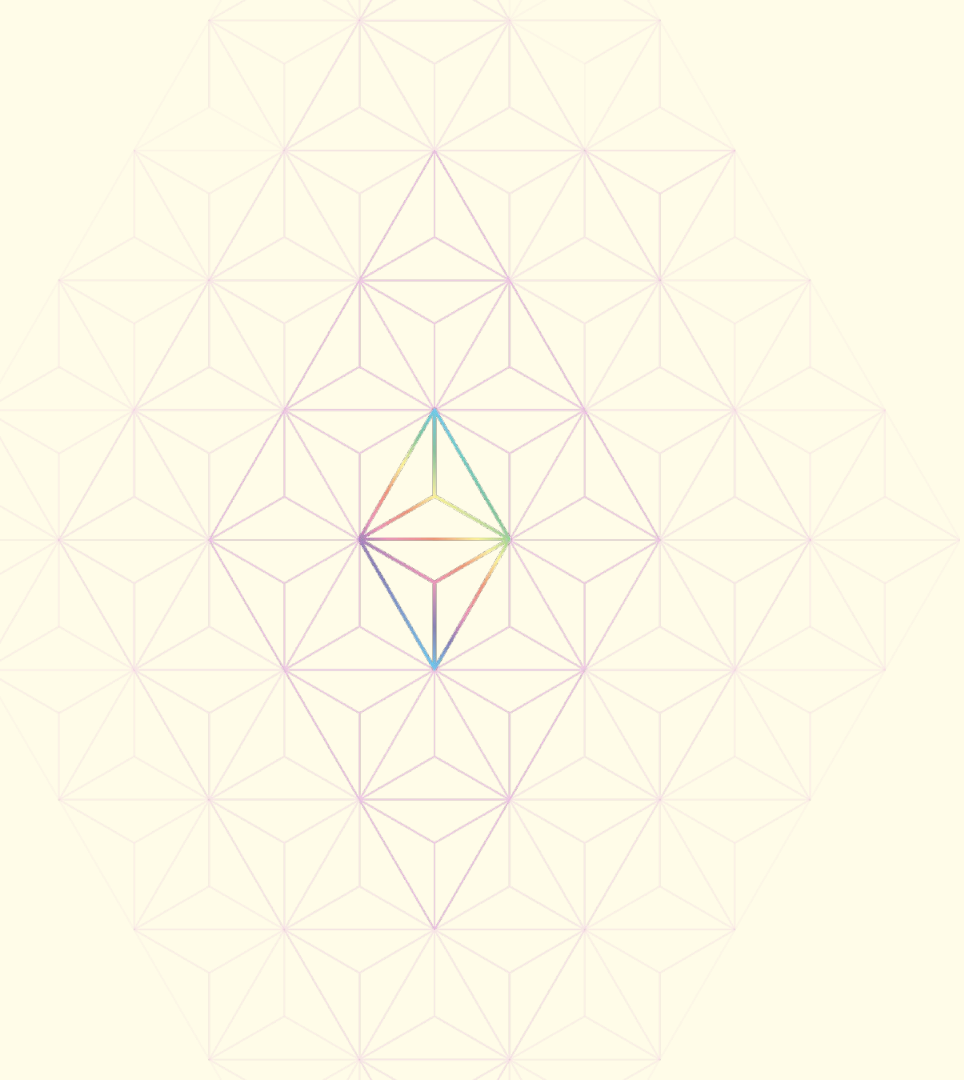

Enshrined Proposer Builder Separation

Eliminate Relays!

- No conflict with ETH values
- No expensive services to maintain

Eliminate out-of-protocol software

Inherit benefits of client-diversity & spec tests

Equivocation Attacks

These can be eliminated with ePBS

Additional Perks

- Compatibility with future roadmap
- MEV-smoothing
- MEV-burn
- PEPC

ePBS Design Goals & Common Components

- A commit-reveal design
  - Builders submit **binding** bids
  - Proposer publishes commitment to builder bid
  - Builder reveals block
  - Proposer imposed cr-lists

ePBS design space

For extensive ePBS literature links see *"Bookmarks relevent for Proposer-Builder Separation researchers"* 56. We define the following properties as desirable:

1. **honest builder publication safety** – If an honest builder wins the auction, the builder (i) *must* have an opportunity to create a block, and (ii) *must* be confident that any payload contents they release become canonical (i.e., protection from unbundling & equivocation attacks 18 from the proposer).
2. **honest builder payment safety** – If an honest builder payment is processed, the builder *must* be able to publish a block that becomes canonical.
3. **honest proposer safety** – If an honest proposer commits to a block on-time, they *must* receive a payment at least as large as specified by the bid they selected.
4. **permissionless** – Any builder can participate in the auction and any validator can outsource block production.
5. **censorship resistance** – There *must* be a mechanism by which honest proposers can force through transactions they suspect are being censored without significantly sacrificing on their own rewards ("If we rely on altruism, don't make altruism expensive" 7 –Vitalik).
6. **roadmap compatibility** – The design *must* be compatible with future roadmap upgrades (SSF, mev-burn, distributed block-building, SSLE, DAS, etc).

- **Weak proposer friendliness**: the mechanism should not require proposers to have either (i) high bandwidth or other computational resources or (ii) high technical sophistication
- **Consensus-layer simplicity and safety**: the mechanism should continue to be safe and ideally be covered by the same analysis as the existing block proposal mechanism from a consensus-layer perspective

Design : Two-slot ePBS

Today

Two-slot

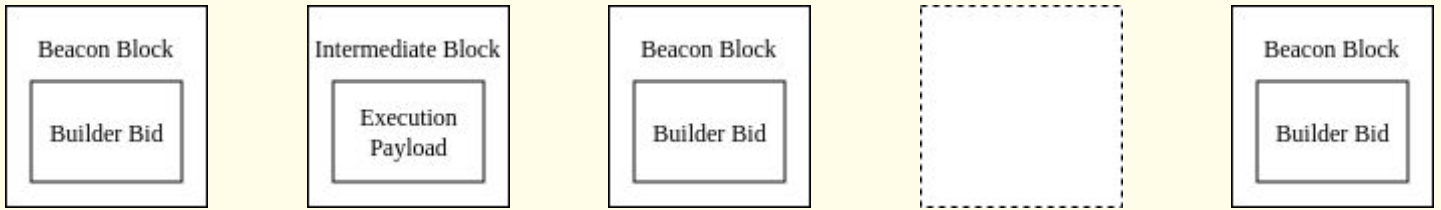

Design : Two-slot ePBS

Today

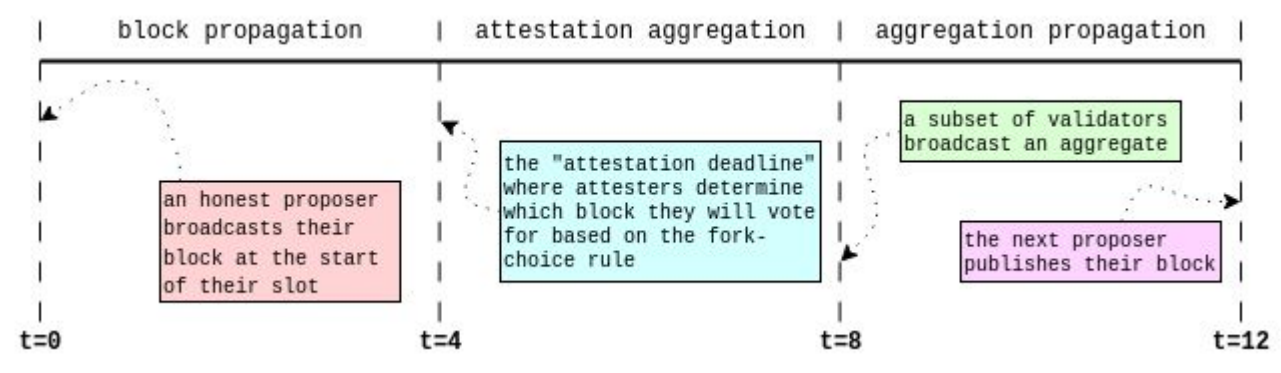

Two-slot

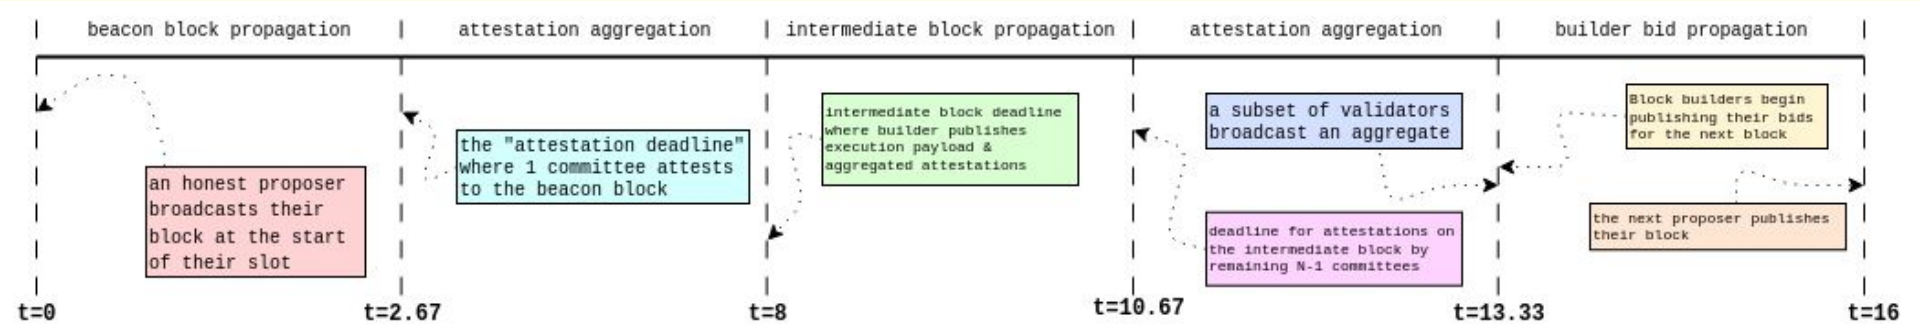

Background : Proposer Boost

CHANGED 3 YEARS AGO Like Bookmark Subscribe

4

Proposal for mitigation against balancing attacks to LMD GHOST

High-level mitigation idea

To more comprehensively mitigate both kinds of attack, the general strategy is to empower honest proposers to impose their view of the fork-choice, but without giving them too much power and making committees irrelevant.

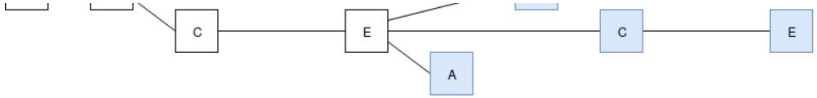

3

Committee-per-slot (as in eth2):

2

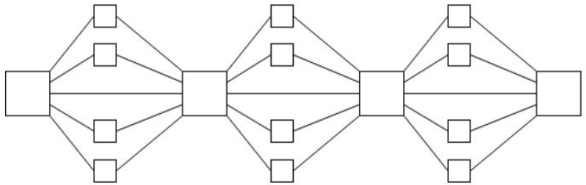

Background : Proposer Boost

Essentially there is a tradeoff between protecting against smaller adversaries with favorable attacking conditions (i.e. controlling multiple blocks in a row) vs. protecting against larger

Conclusion

$W_{PB} = 0.4$ optimally protects against reorgs by 20%-adversaries with $l = 2$, while minimally trading off for $l = 3$ by protecting against $\beta \leq 0.133$ adversaries, as opposed to the optimal $l = 3$ protection against $\beta \leq 0.143$ adversaries as achieved by $W_{PB} = 0.43$.

By choosing $W_{PB} = 0.4$ over higher boost values such as $W_{PB} = 0.43$ we are slightly better protected against larger adversaries in $l = 2$ settings, at the expense of being slightly less protected in $l = 3$ settings. A reasonable tradeoff.

We think any value $W_{PB} \in [0.4, 0.43]$ is reasonable, with a minor preference for $W_{PB} = 0.4$.

Design : Two-slot ePBS

Two Slot

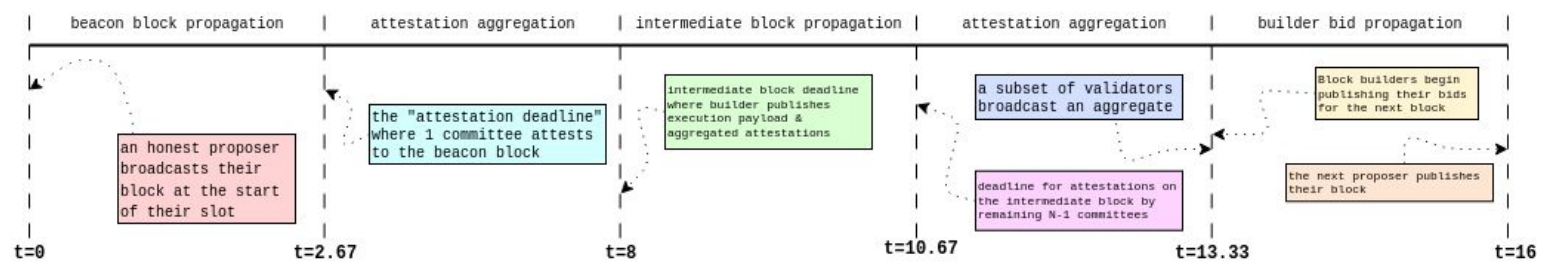

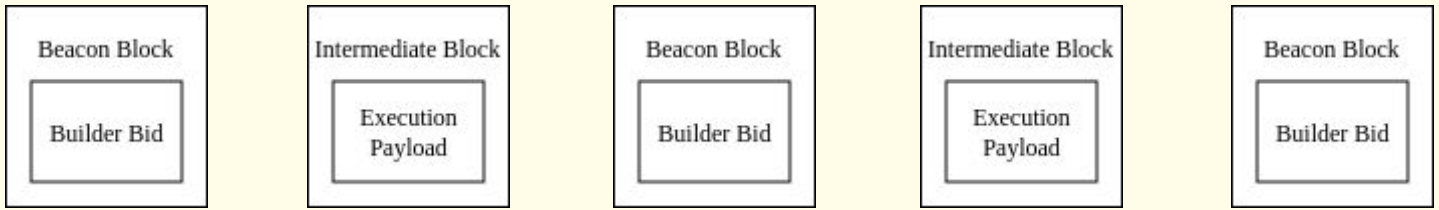

Attack : Proposer Griefs Builder

- **Untrusted proposer friendliness**: there's minimal or no risk that a proposer will screw over a block builder, so block builders have no incentive to prefer proposers that have some off-chain reputation or personal connection to the builder (as that would favor large pools).

1. **honest builder publication safety** – If an honest builder wins the auction, the builder (i) *must* have an opportunity to create a block, and (ii) *must* be confident that any payload contents they release become canonical (i.e., protection from unbundling & equivocation attacks 18 from the proposer).

Attack : Proposer Griefs Builder

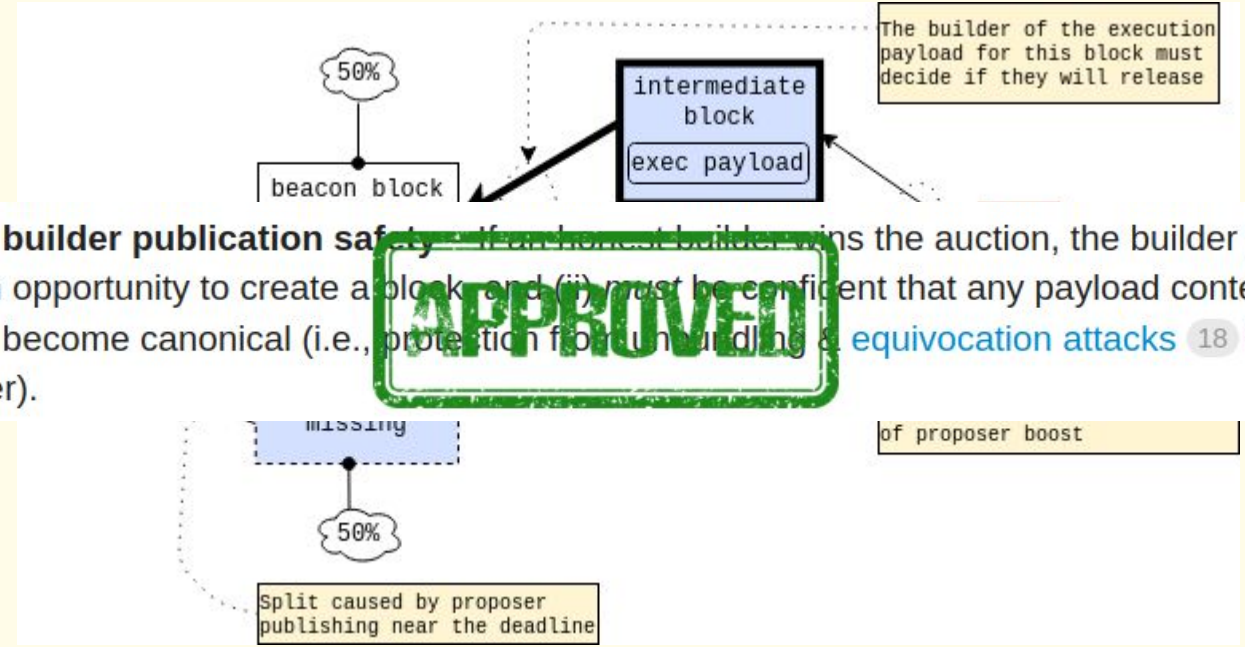

honest builder publication safety — If an honest builder wins the auction, the builder (i) *must* have an opportunity to create a block, and (ii) *must* be confident that any payload contents they release become canonical (i.e., protection from unbundling & equivocation attacks 18 from the proposer).

APPROVED

Problem : Ex-Ante Reorg Probabilities

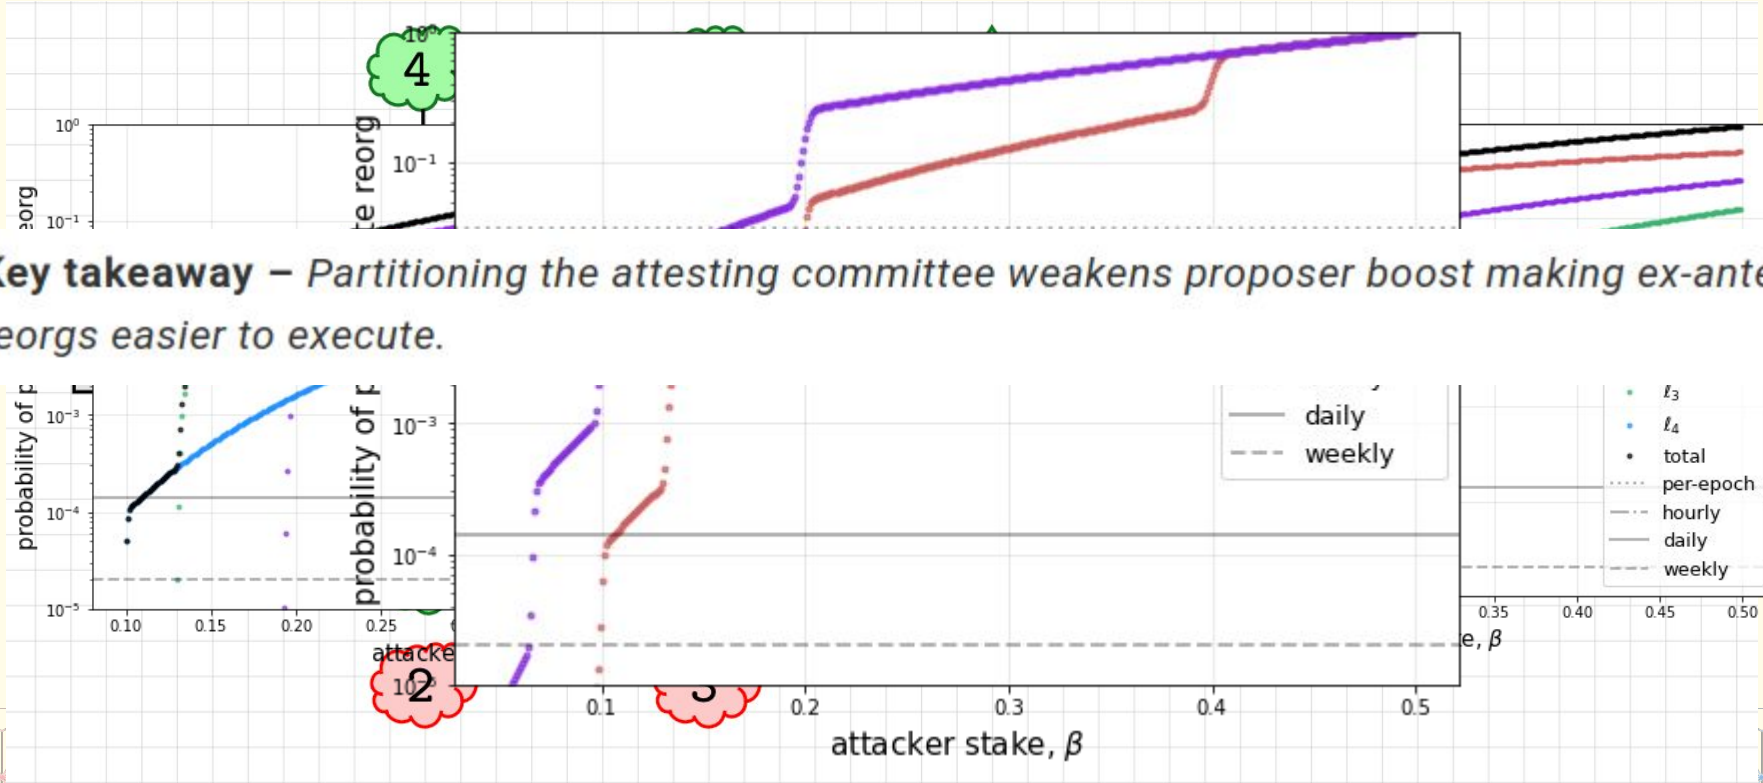

Key takeaway – *Partitioning the attesting committee weakens proposer boost making ex-ante reorgs easier to execute.*

Design : Payload-timeliness Committee (PTC)

Today

PTC

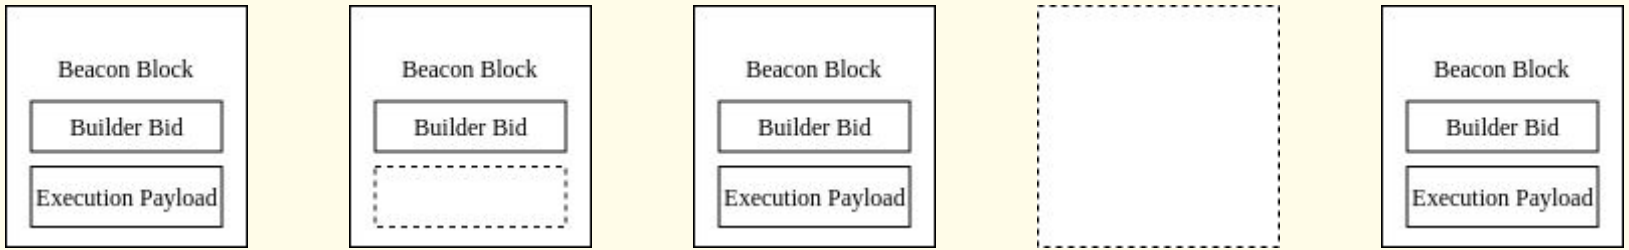

Design : Payload-timeliness Committee (PTC)

Today

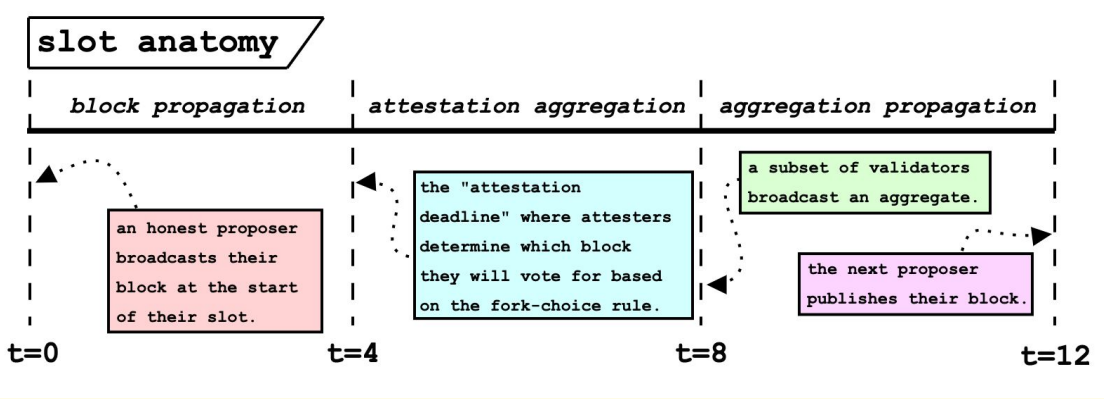

PTC

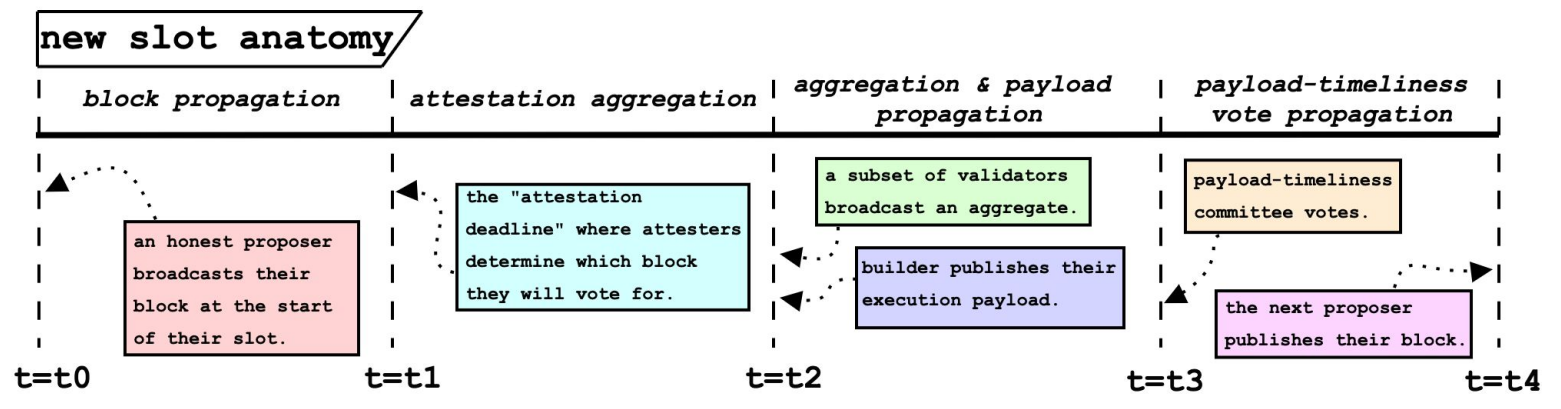

Design : Payload-timeliness Committee (PTC)

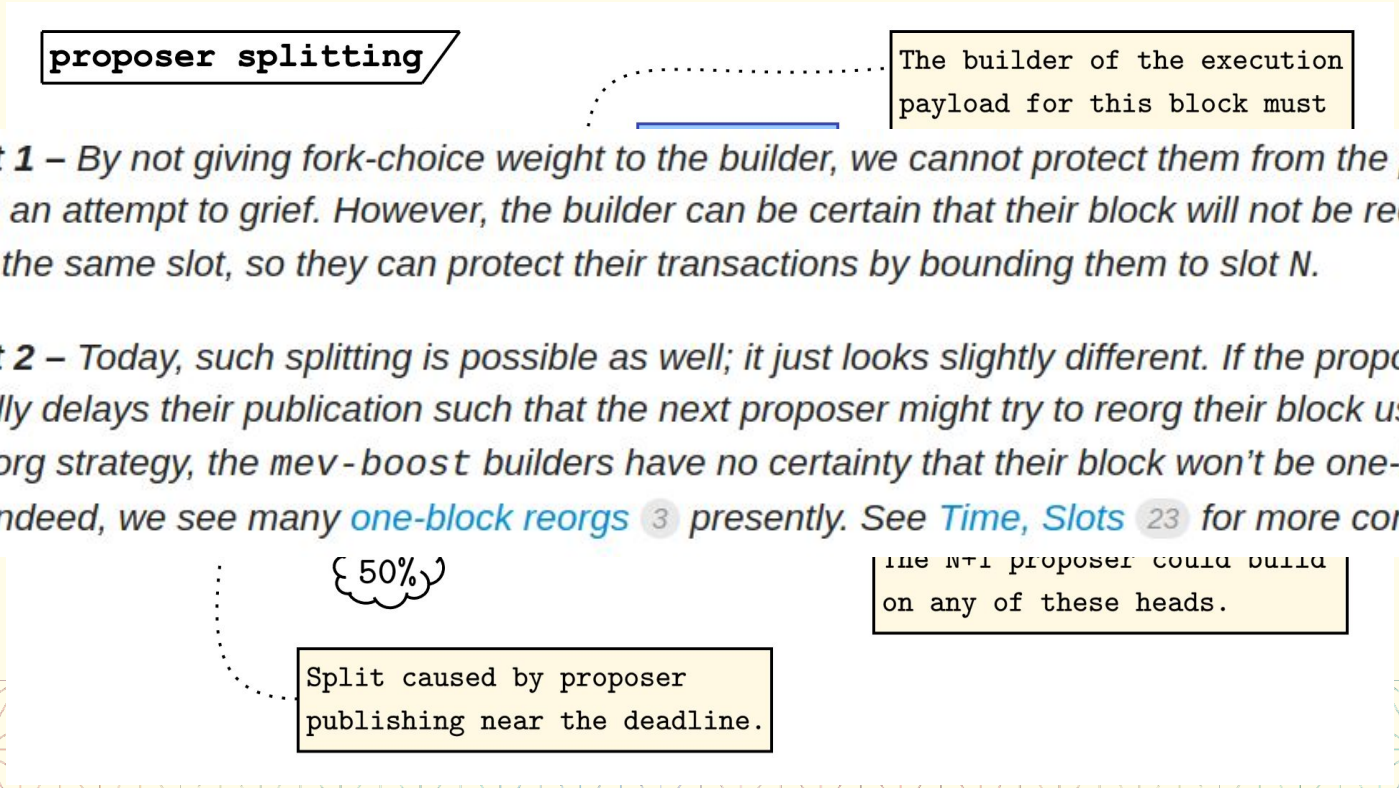

Censorship Resistance in ePBS

CHANGED 2 YEARS AGO

Thu, Oct 21, 2021 11:09 AM

State of research: increasing censorship resistance of transactions under proposer/build S)

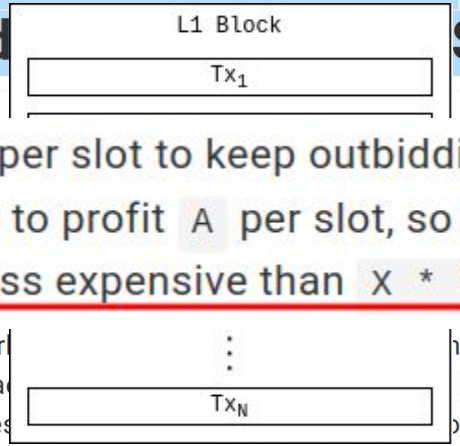

their bid to `M + P`, and lose `P - A` per slot to keep outbidding the non-censoring builders (and additionally sacrifice the opportunity to profit `A` per slot, so total loss is `P` per slot). This is still expensive, but note that it is much less expensive than `X * 66.7` per slot.

In the current transaction mar ner, post-merge: a validator) directly chooses which transa looking at which transactions in the mempool pay the highes poser in a position to use sophisticated strategies to choose which transactions to include, or even include their own, to take advantage of opportunities such as DEX arbitrage and liquidations (hereinafter just called "**MEV**" for simplicity) to maximize their profits. The complexity of these strategies creates a high fixed cost in running an effective miner or validator, and advantages centralized pools that take on this task on behalf of their participants.

Censorship Resistance in ePBS

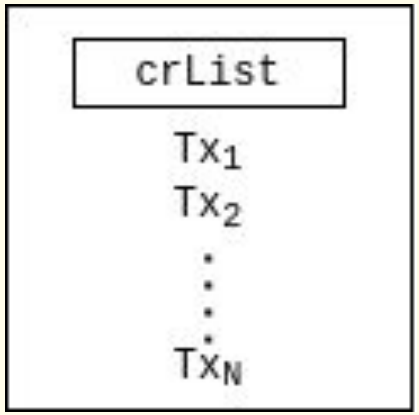

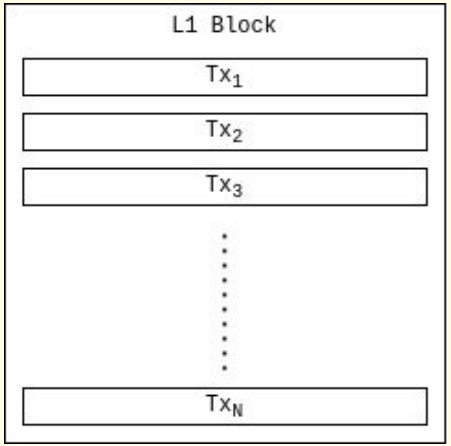

Builder publishes body. Verification of the body requires checking that for each `(sender, gaslimit)` in `crListSummary`, either `block.gasused + gaslimit > block.gaslimit` or `sender` is the sender of some transaction in the block

MEV Burn

Burning MEV through block proposer auctions

Economics ■ mev ■ proposer-builder-separation

domothy Oct '22

Special thanks to @JustinDrake and @pseudotheos 39 for all the discussion and theory-crafting surrounding this topic. Personal opinions are expressed in the first person singular.

Introduction

This document aims to explore the idea of letting multiple proposers compete every slot, with the winner being the one committing to burning the most ETH. A possible implementation is presented, and the incentives are played out to argue this would result in most (if not all) of the MEV being burned.

The core idea of this proposal is to auction off the “right to build a block” on the consensus layer through an in-protocol burn auction. Once a winner is selected, the execution block they propose will be one that provably burns at least as much ETH as their bid. The idea is that the highest bid will naturally approach the maximal extractable value (MEV) in the block, therefore most of the MEV should be directly burned.

MEV Burn

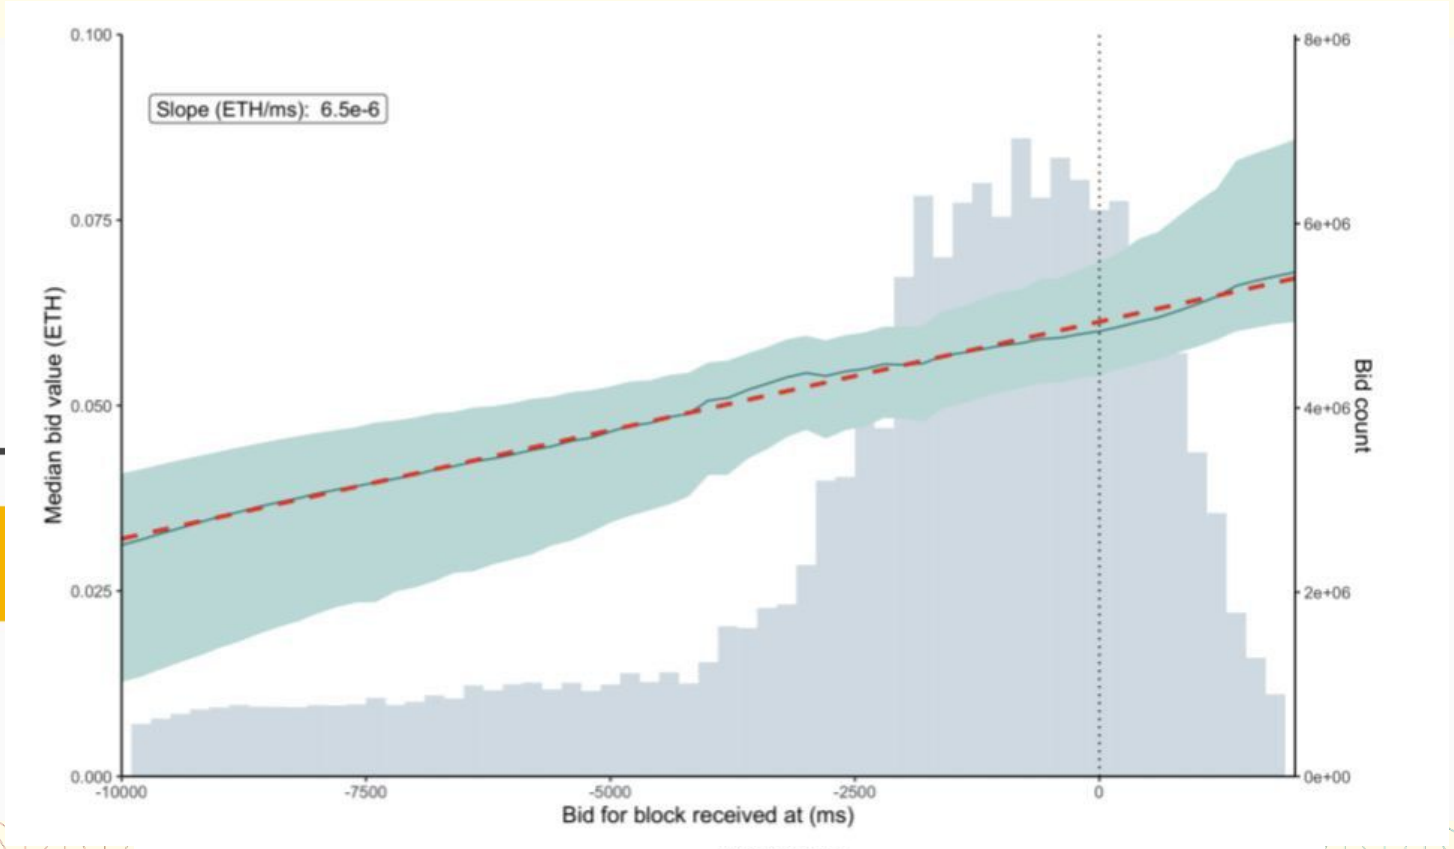
Slope (ETH/ms): 6.5e-6
Median bid value (ETH)
0.100
0.075
0.050
0.025
0.000
Bid count
8e+06
6e+06
4e+06
2e+06
0e+00
-10000
-7500
-5000
-2500
0
Bid for block received at (ms)

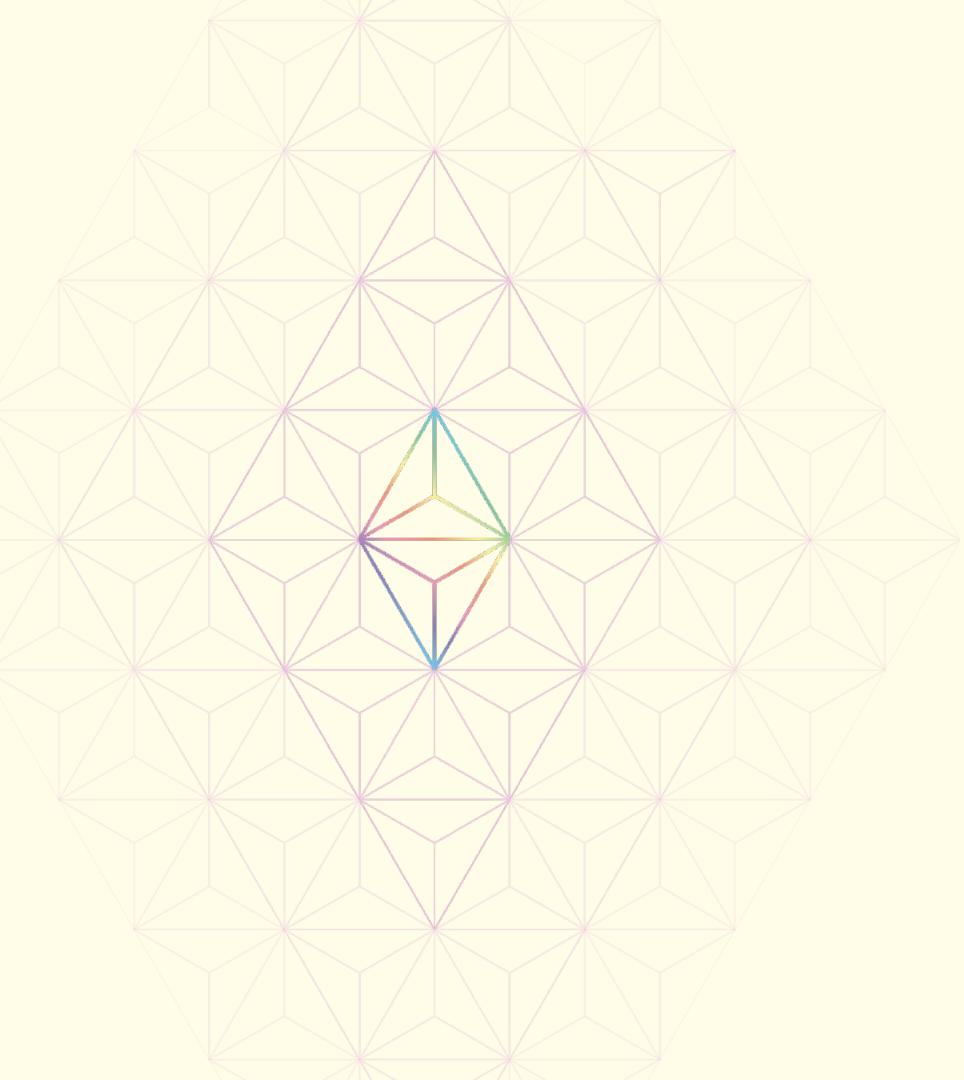

Switching Gears

The Space of Agreements

- The current design of mev-boost (and many ePBS proposals) offer only one agreement:

  the right to build a specific block

- What about other possible agreements between proposers and third parties?
  - the right to a slot (not a specific block)
  - the right to some partial block (prefix or suffix) but not the whole thing
  - the right to the whole block respecting some inclusion list
  - blockspace futures
  - validity proofs for the canonical zk-EVM
  - combinatorial preferences

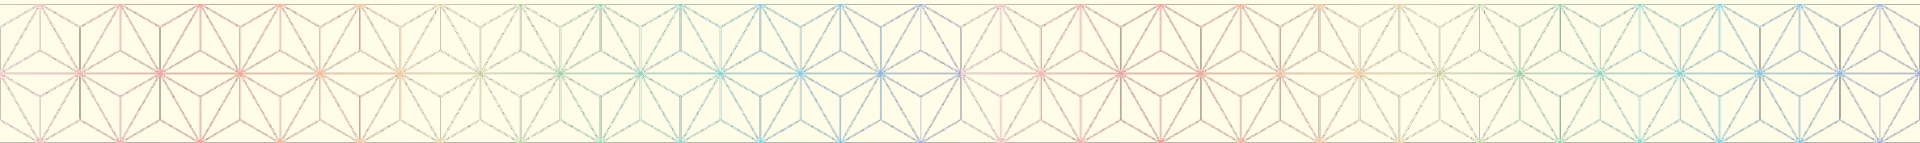

PEPC : Protocol-Enforced Proposer Commitments

- Allow the proposer to commit to arbitrary agreements via enshrined Eigenlayer (stakers can delegate a contract that is allowed to slash them)

- Two Flavors: Optimistic & Pessimistic
  - Optimistic : attesters stand to be slashed if they attest to a block where the proposer violated the agreement
  - Pessimistic : agreements are expressed as EVM code & the state-transition function enforces agreements

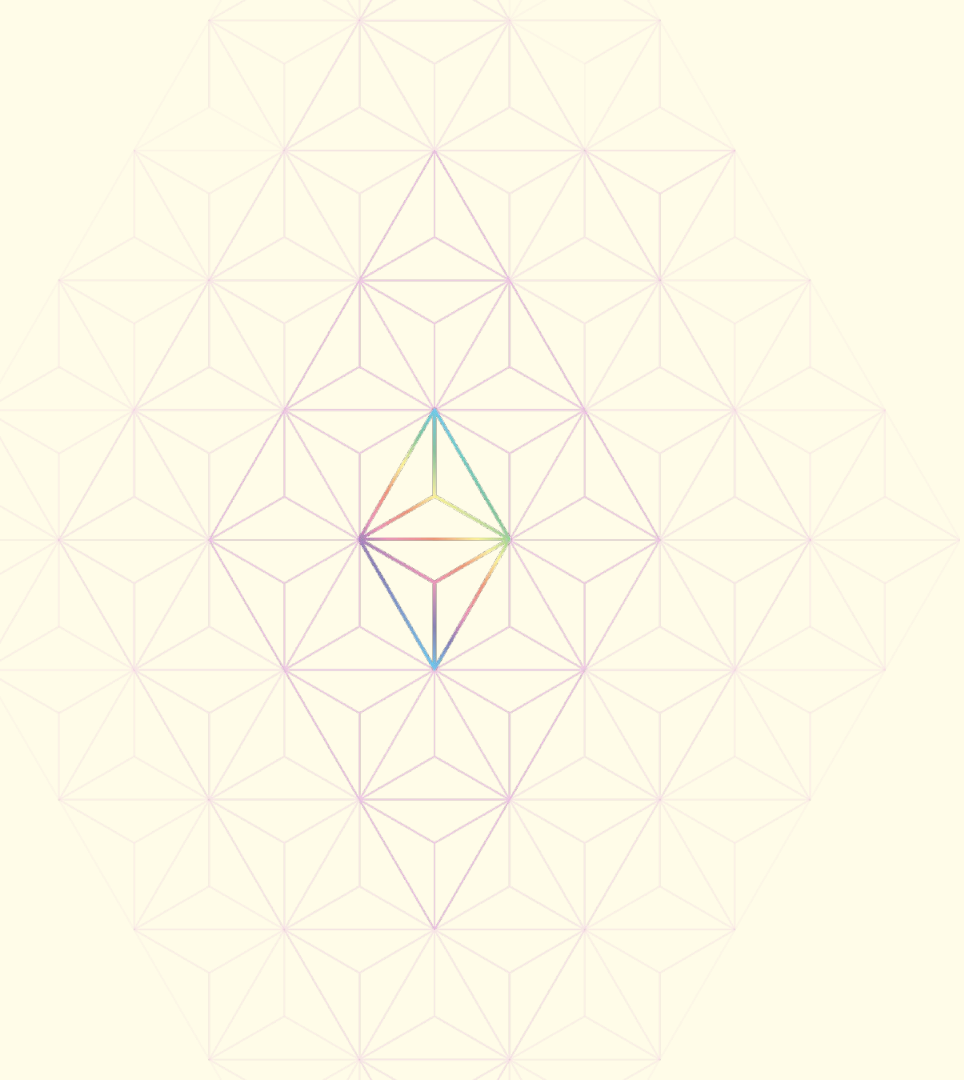

Revisiting Relays

Eliminate R

- No conflict
- No expensi

Equivocatic

These can be

Relays in a post-ePBS world

Proof-of-Stake ■ mev ■ proposer-builder-separation

mikeneuder 1 Aug 4

Relays in a post-ePBS world

.

by mike 7, jon 16, hasu 1, tomasz 11, chris 5, & toni 2
based on discussions with justin 3, caspar 3, & stokes 2
august 4, 2023

ol

rsity

oadmap

Relays / Off-chain Agreements After ePBS

- Services provided by relays today (mev-boost):
  - Services for builders:
    - MEV-stealing/unbundling protection 
    - Bid Cancellation 
  - Services for proposers:
    - Block validity enforcement 
    - Payment verification 
    - Collateral Escrow 

Payment Verification : Top of Block Payments

- The builder bid includes a special “top of block” transaction to pay the proposer. This transaction is **only valid if the proposer selects their bid**.

- Once selected, this tx is applied to the state before the builder payload is applied, regardless of whether or not the builder releases their payload

- Builder block is constructed on top of the post-state **after applying the proposer payment transaction**

- Builder must have enough collateral in their account to cover the bid

Relay Advantages over ePBS

- More flexible payments
  - For large MEV opportunities, the builder may not have the ETH required to cover their bid until after the block is executed
  - This isn’t a concern if the MEV is burned as we can enforce the burn after the block is executed

Re

- M
- L

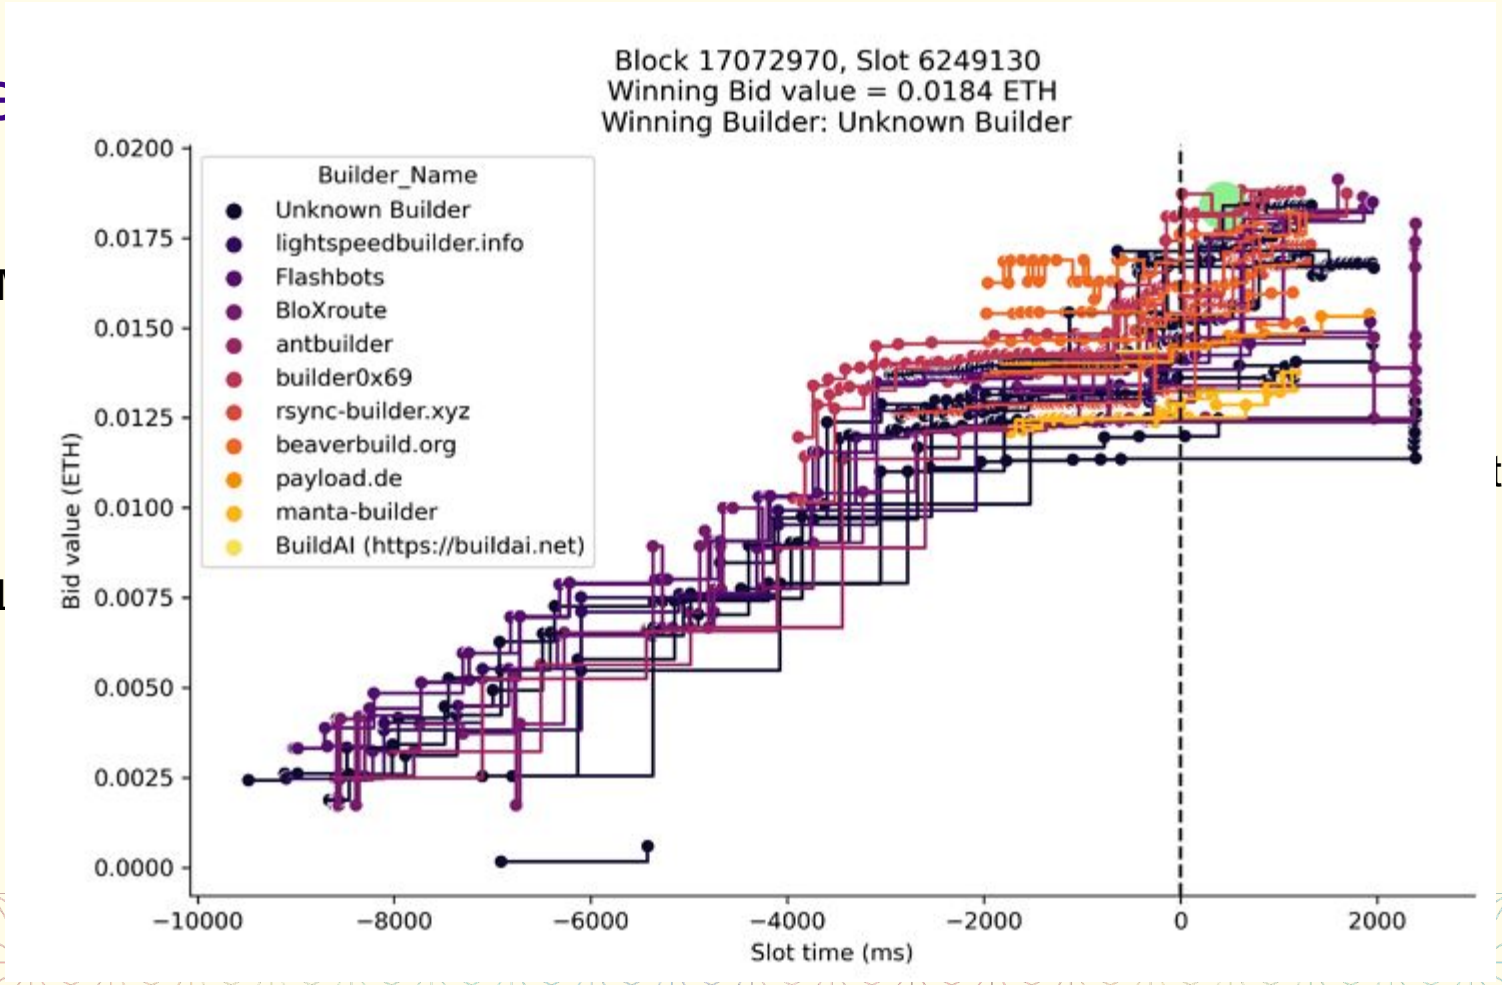
Block 17072970, Slot 6249130
Winning Bid value = 0.0184 ETH
Winning Builder: Unknown Builder
Builder_Name
Unknown Builder
lightspeedbuilder.info
Flashbots
BloXroute
antbuilder
builder0x69
rsync-builder.xyz
beaverbuild.org
payload.de
manta-builder
BuildAI (https://buildai.net)
Bid value (ETH)
0.0200
0.0175
0.0150
0.0125
0.0100
0.0075
0.0050
0.0025
0.0000
−10000
−8000
−6000
−4000
−2000
0
2000
Slot time (ms)

Relay Advantages over ePBS

- More flexible payments
  - For large MEV opportunities, the builder may not have the ETH required to cover their bid until after the block is executed
  - This isn’t a concern if the MEV is burned as we can enforce the burn after the block is executed
- Lower latency connection between builder and proposer
- Bid Cancellations
  - Can’t do it =(

Relays May Continue to Exist after ePBS

- Vertically-integrated builder/relay
- Relay as a Service (RaaS)
- Public goods relays

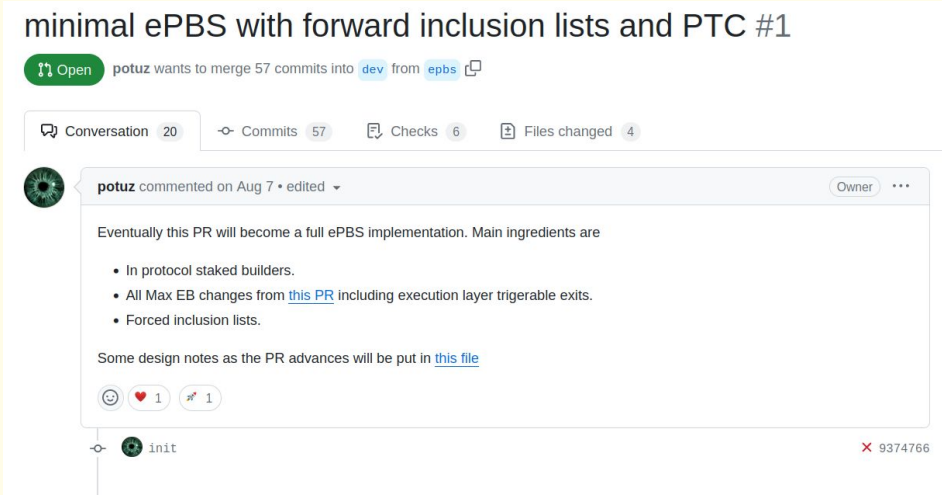

The Bull Case for Enshrinement

- Relays are expensive: relay specific benefits might not offset the cost of running them

The Bull Case for Enshrinement

- Relays are expensive: relay specific benefits might not offset the cost of running them
- ePBS significantly reduces the cost of altruism

The Bull Case for Enshrinement

- Relays are expensive: relay specific benefits might not offset the cost of running them
- ePBS significantly reduces the cost of altruism
- ePBS delineates in-protocol PBS and out-of-protocol mev-boost

The Bull Case for Enshrinement

- Relays are expensive: relay specific benefits might not offset the cost of running them
- ePBS significantly reduces the cost of altruism
- ePBS delineates in-protocol PBS and out-of-protocol mev-boost
- ePBS removes the neutral relay funding issues

The Bull Case for Enshrinement

- Relays are expensive: relay specific benefits might not offset the cost of running them
- ePBS significantly reduces the cost of altruism
- ePBS delineates in-protocol PBS and out-of-protocol mev-boost
- ePBS removes the neutral relay funding issues
- ePBS is future-compatible with mev-burn, inclusion-lists, and L1 zk-EVM proof generation

The Bull Case for Enshrinement

- Relays are expensive: relay specific benefits might not offset the cost of running them
- ePBS significantly reduces the cost of altruism
- ePBS delineates in-protocol PBS and out-of-protocol mev-boost
- ePBS removes the neutral relay funding issues
- ePBS is future-compatible with mev-burn, inclusion-lists, and L1 zk-EVM proof generation
- ePBS backstops the builder market in the case of relay outages

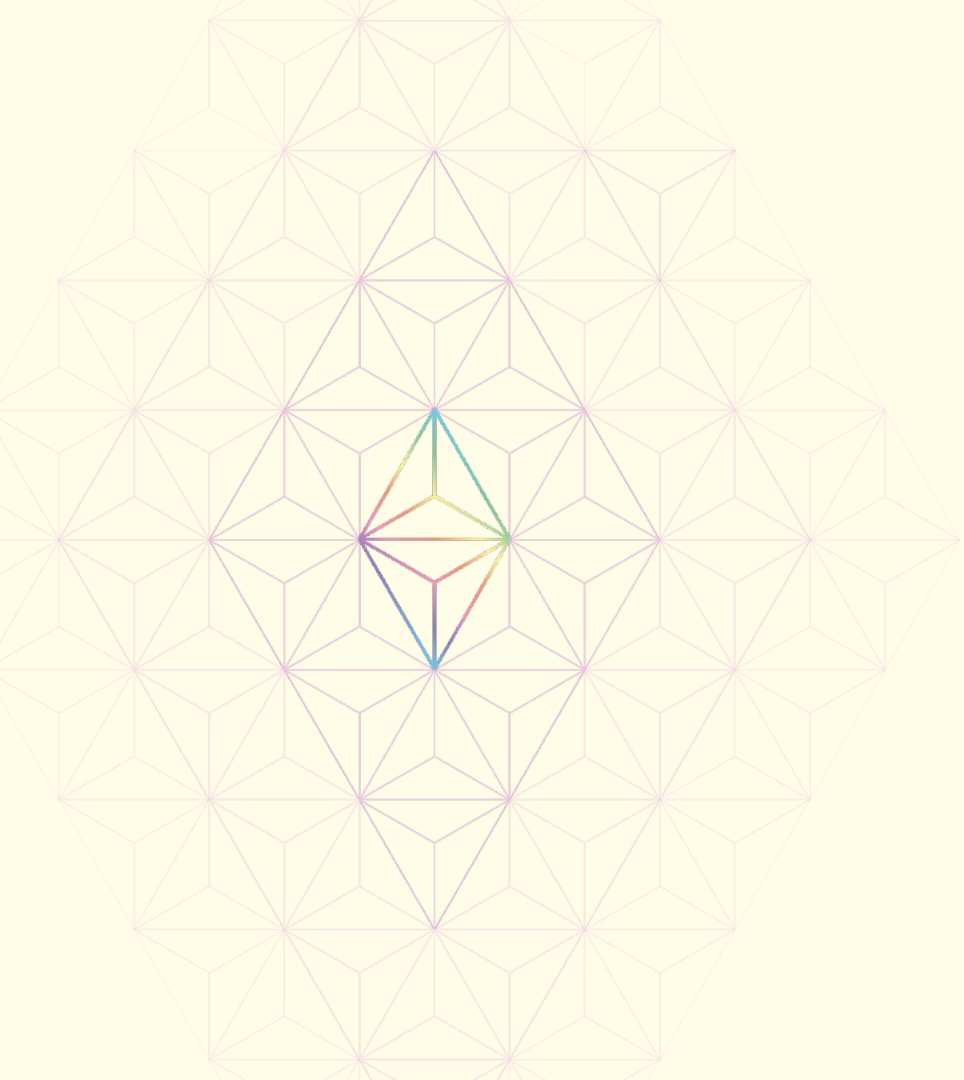

New Relay Designs

The Optimistic Roadmap

- Progressively remove relay responsibilities
- Converge on a system similar to existing ePBS proposals before enshrinement
- Optimistic relaying reduces hardware & networking requirements for competitive relays
- Accessibility: by keeping independent relays competitive & designing the optimistic architecture in the open, we increase visibility into the existing block-building market. This helps answer research questions

The Optimistic Roadmap : Relays Today

- Builders Transmit Full Blocks to Relays

- Relay must simulate each block before bid is added to the auction
  - Verify Proposer Payment
  - Verify Block Validity

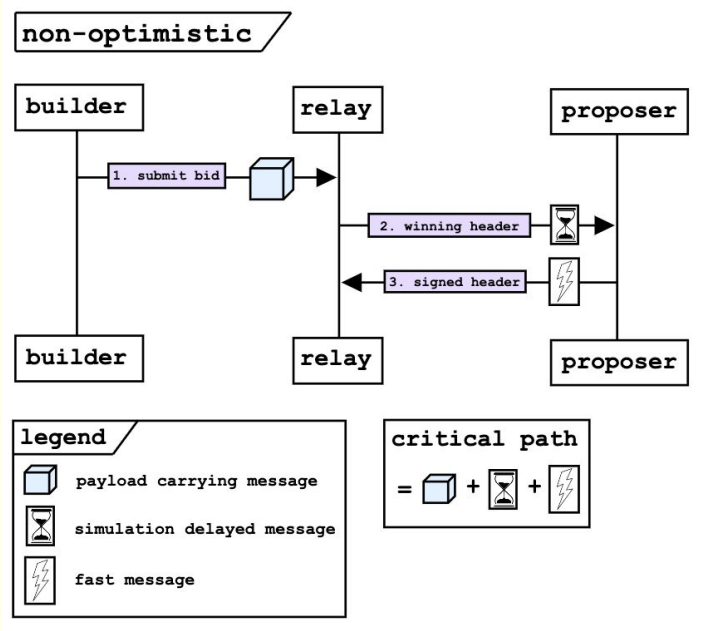

The Optimistic Roadmap : Optimistic Relay V1

Asynchronous block validation

- Relays escrow collateral from the builders
- Relays don’t validate the block before adding bid to the auction
- Pay proposer via collateral if the block is invalid

- ultra sound relay already implements this

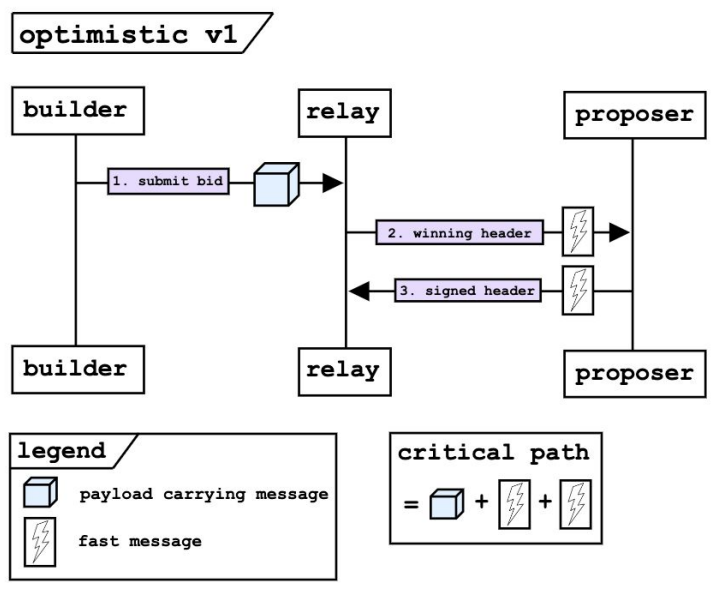

The Optimistic Roadmap : Optimistic Relay V2

- Header-only parsing
  - Again Relays escrow collateral from builders
  - Builders update bids by transmitting headers to relay (bodies transmitted separately)

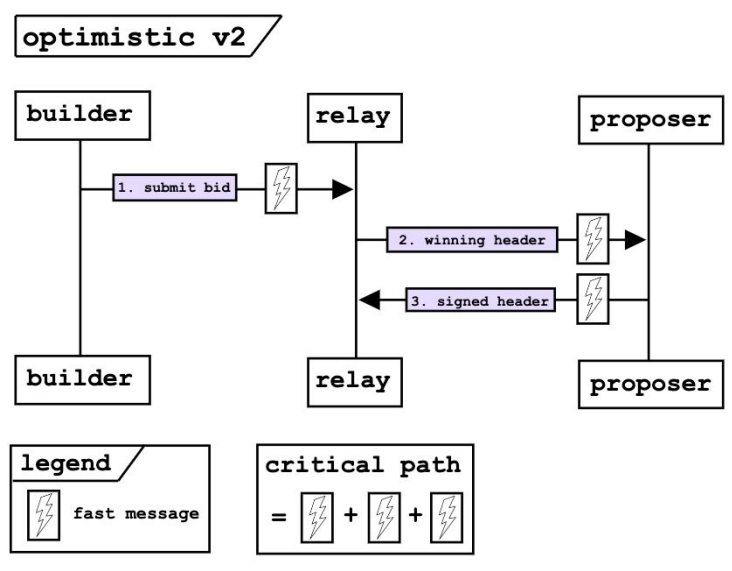

The Optimistic Roadmap : Optimistic Relay V3

- Relay as an oracle
  - Again Relays escrow collateral from builders
  - Communication between builders and proposers happens via gossip
  - Relay is only used as an oracle to pay the proposer if the builder didn’t release their block on time

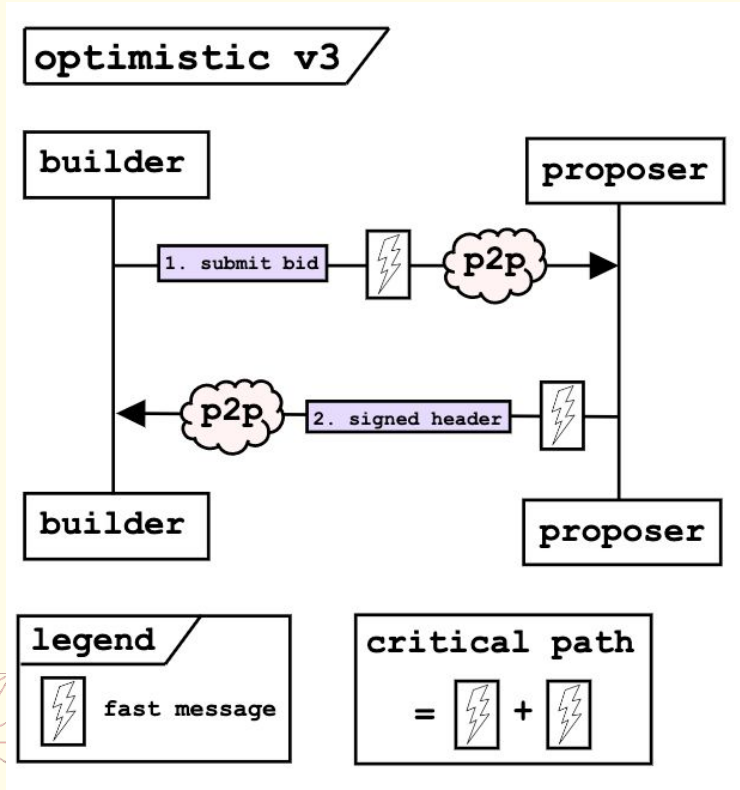

PEPC-DVT

- Use DVT to split up a validator key. 50% of the key stays with the validator and the remaining shares are split among some Distributed Validator Nodes
- Validator is allowed to enter into agreements specified by EVM code with third parties
- Distributed Validator Nodes refuse to sign any messages that violate agreements
- This prevents validator from violating agreements when the incentive is higher than the penalty

mev-boost+/mev-boost++

- Proposer restakes via EigenLayer (sets withdrawal credentials to EigenLayer contract)
- MEV-Boost+ : enable partial block auctions
  - Proposer signs commitment to partial block instead of full block
  - Builder reveals partial block to proposer, and proposer adds their own transactions at the end of the block
  - Proposer’s signed commitment can be used to slash them if they don’t include the partial block at the top of the block (or if they re-order the partial block)
- MEV-Boost++ : replace the relay with a programmable data availability layer

Execution Tickets

Execution Tickets

Proof-of-Stake Economics mev

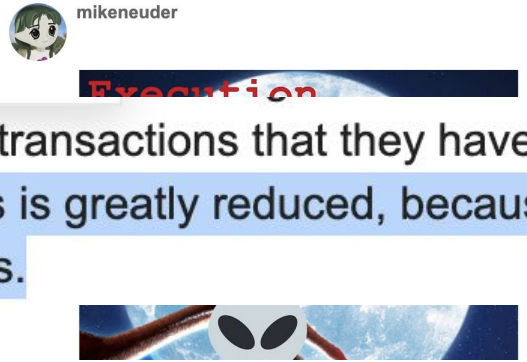

mikeneuder

High-level view

We begin by presenting the design in its most distilled form. The execution ticket mechanism can be succinctly described as:

of the transactions that they have seen. Importantly, the incentive to play proposer timing games is greatly reduced, because the value of the beacon block should not increase as time passes.

(the final list of executed transactions for a block), but instead has an inclusion list that specifies

2. ***Removes MEV from validator rewards.***
   - By explicitly decoupling MEV from the validator rewards, the original incentives of the consensus mechanism are restored. Their rewards are only derived from their performance on the consensus duties of keeping the chain alive. Note that warping of the validator incentives

Note that the economic security of beacon blocks remains the same, thus a finality reversion still has the settlement assurances of 1/3 of all staked ETH.

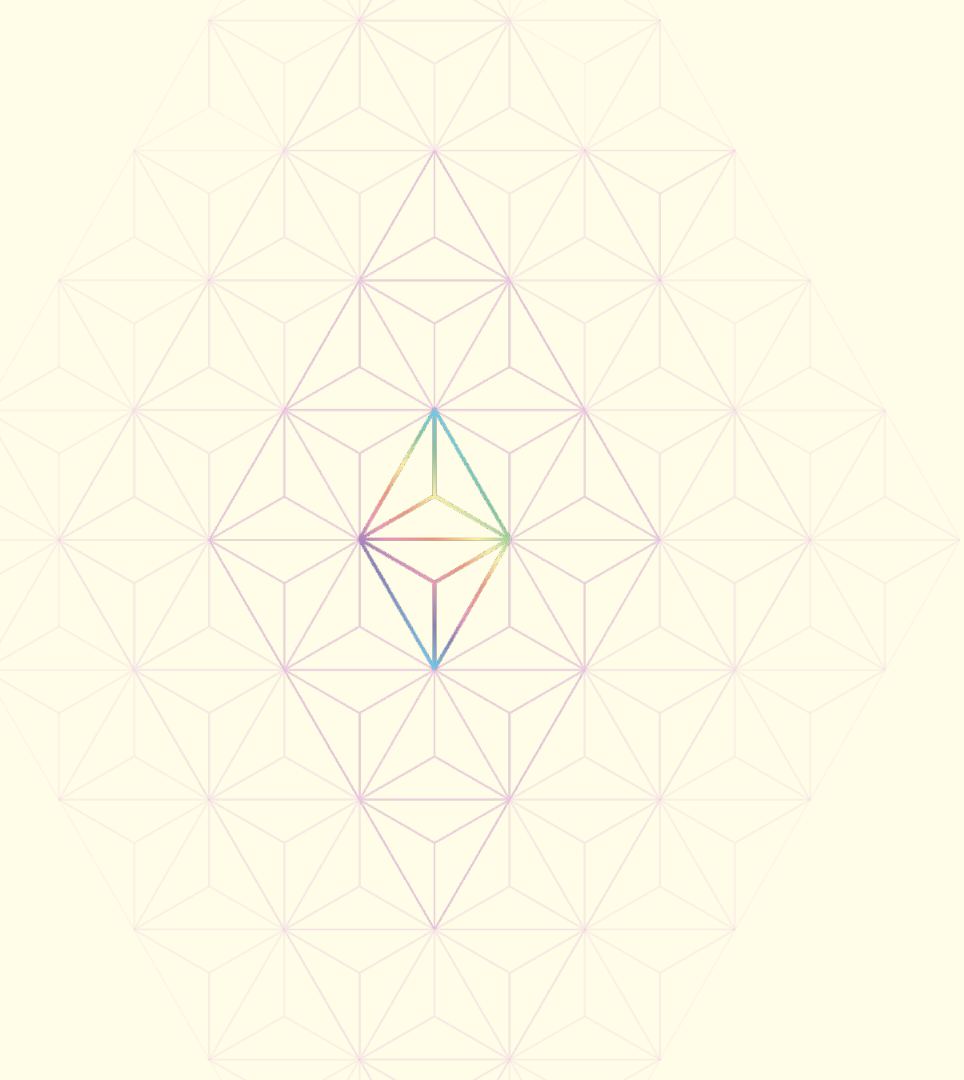

Open Questions

Bypassability

- Even if we enshrine PBS, there is no way to enforce that all proposers and builders make use of the protocol instead of out-of-protocol solutions

- What percentage of proposers do we think would think would bypass the enshrined mechanism?

- What percentage of proposers do we need to justify enshrinement?

What do we aim to achieve by enshrining?

- Just a “relay of last resort”?
- Do we want to encourage validators to disconnect from mev-boost even if it may provide higher payments?
- Are we enshrining as a way to achieve future roadmap goals?
- Does eliminating “neutral relays” accelerate vertically integrated builder relays?

How much can we rely on altruism from the social layer?

- If we define “protocol-aligned” behavior as making use of the in-protocol mechanism, how likely are large ETH holders to sacrifice a small amount of MEV to not bypass ePBS?

- Will they see this as protecting the long term health and decentralization of Ethereum?

- Is the social layer something we want to lean on in this context?

Could MEV be addressed with different tools?

- Several MEV mitigation strategies are being experimented with
  - SUAVE
  - CoW swap
  - MEVBlocker

- Others are being researched (encrypted mempool)

How important is L1 ePBS in a future with L2s & OFAs?

- If most L1 activity migrates to L2s / OFAs and other app-layer MEV mitigation tools gain popularity, there may be less MEV exposed on L1
- It may be less necessary to enshrine PBS if this happens
- How likely is this?
- Should we enshrine anyway?

What priority should ePBS have over other protocol upgrades?

- There are many other important Ethereum upgrades to focus on
  - Single Slot Finality
  - Verkle Trees
  - Single Secret Leader Election

Thank you

Questions?

Lighthouse
github.com/sigp/lighthouse

Mark Mackey
Sigma Prime | Lighthouse
@ethDreamer