CHORUS

Quarterly Insights

PREPARED BY

Felix Lutsch, Xavier Meegan, Michael Moser, Umberto Natale, Luis Nunez, Gabriella Sofia, Thalita Franklin

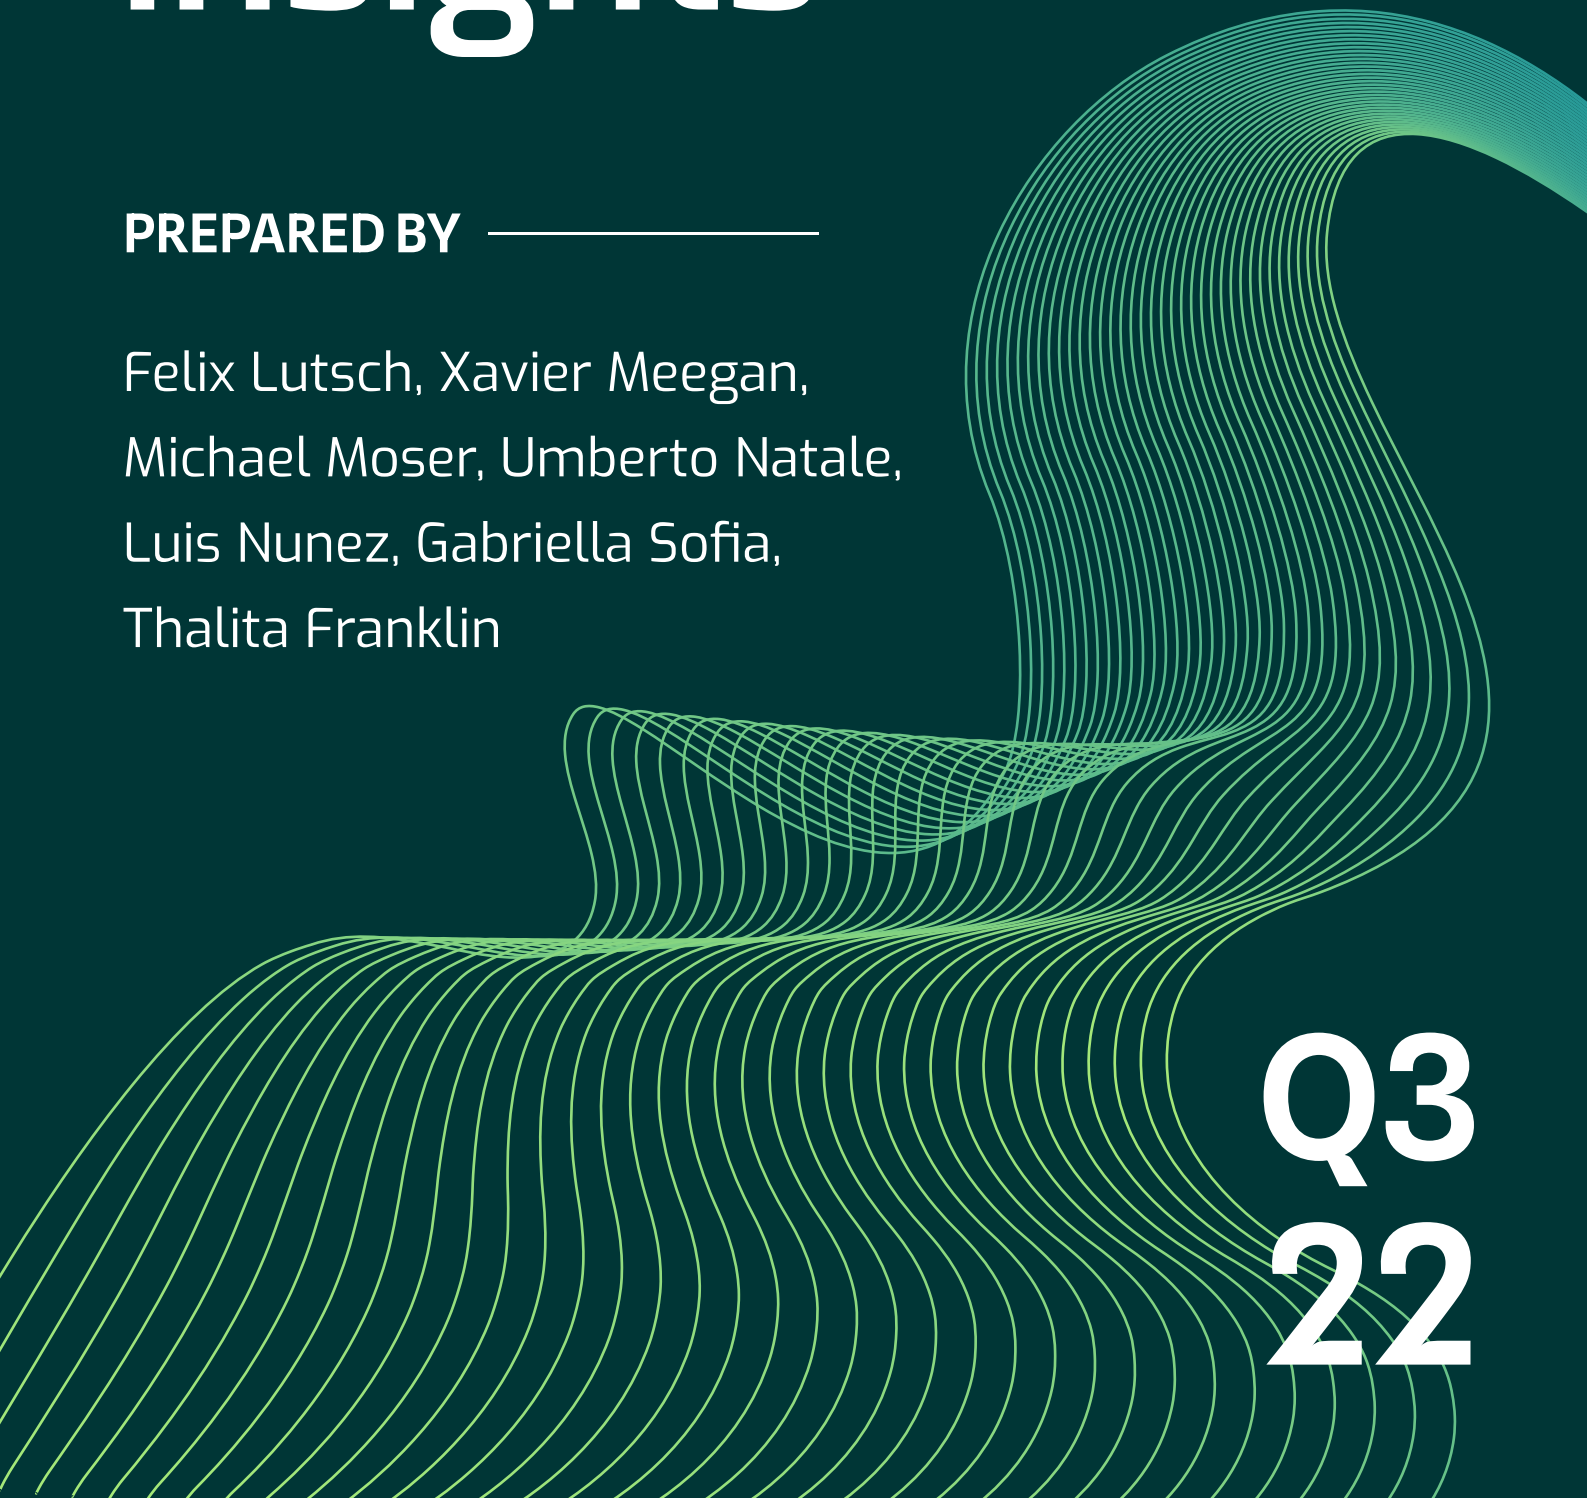

Q3 22

Table of Contents

State of Distributed Validator Technology in Ethereum

The future of Ethereum asks for a reliable, stable and decentralized base layer. After successfully completing the Merge upgrade in the past quarter, the ecosystem has turned to address a number of pain points that have been festering in the network, including a growing risk of censorship (either through regulation or self-imposed measures) and centralization. These latest concerns reflect the most recent trends in the Ethereum staking ecosystem, where validator pools and centralized exchanges are creating a seemingly uncomfortable majority. Layer 2s have become a staple in discussions of Ethereum decentralization, however the Execution Layer is not the only aspect of the architecture that will need to adapt for the future of Ethereum. We must also scale the capabilities of our Consensus Layer.

Now what do we mean by “scaling the Consensus Layer”. This layer of the logical architecture of Ethereum deals with the enforcement of network rules, describing what nodes within the network should do to reach consensus about the broadcasted transactions of the Execution Layer (in our post-Merge world). Furthermore it deals with the generation and verification of blocks.

The rise of Liquid Staking is to blame for the abstraction between depositing ETH and controlling a node. This results in the trust assumption that a single operator will conduct their validator duties properly, and act in accordance with a high number of depositors. What we see is that a single staking node directly controls how any one validator behaves. This is, of course, not a very scalable system. On one side, it stagnates the growing number of nominally active validators to a shrinking number of true operators. These operators have to be vetted (like in the case of the Lido protocol) to ensure the liveness of the chain, or they must put a high amount of collateral upfront (for example, in the case of Rocket Pool). There are several risks that are borne by stakers or staking services, especially in the case of a big scale slashing event. An example of such an event could be the result of an unavailable client that takes down a high number of nodes. In that case, the losses to the staking operator would be significant, same as the individual delegators under the socialized risk of Lido.

To address these and other risks, one of the most important technologies in development for the Ethereum roadmap is

Distributed Validator Technology or DVT. In essence, it is a novel way to distribute the operation of an Ethereum validator by splitting a validator key into a multisig construct. The distributed keys created during this process can then be assigned to different non-trusting nodes that operate under consensus to perform validator duties and generate rewards.

Some clear benefits include: not having a single point of failure if any of the distributed nodes were to go down, stakers are able to store validator keys offline instead of having to have them online at all times (only the multisig must be online), and in general, a much smaller slashing risk vector. Another not-so-clear advantage is that it solves the trust problem arising from Liquid Staking that we previously described. A single bad actor would not matter much, allowing something like Lido to expand its operator network to a much larger set. Thus, a more decentralized, fault tolerant and secure network.

Currently there are two projects working on DVT: Obol Network and SSV Network. Both of them had its inception on the same EF grant a few years back, but resulted in somewhat different views on how to solve the problem of DVT from their founders.

Obol has created a middleware client called Charon, which is in the process of getting tested. When validator clients send output messages with consensus votes, the Charon multi-sig combines them and relays the vote to the four Beacon clients. They then

communicate it to the Beacon chain for validation. Because of a more trusted setup, this is currently a good option for institutional clients or CEXs.

SSV Network is focused more on creating an infrastructure layer for decentralized staking. The idea is that each user would select up to 4 node operators from the protocol to multi-sig vote for that specific validator, in a permissionless fashion that relies on reputation. As both projects rush to Mainnet, it will be interesting to see the impact they will cause after launch, and how they will be adopted by Lido, Rocket Pool, and others. Even though the goal continues to be in the realm of thousands, or tens of thousands, of node operators, and this is a bit outside the scope of DVT, it is definitely a first step in the right direction.

Launch of Uqbar Testnet and Ziggurat Developer Suite at Assembly Miami

At Urbit Assembly, the Uqbar Testnet and Ziggurat Developer Suite (a unified contract testing and editing suite) was officially launched. We are an investor in Uqbar and deeply interested in the intersection of web3 and Urbit (the operating system that Uqbar is built on).In the last quarter, we have released three articles detailing why web3 needs Urbit, what the synergies of Urbit and blockchain are, and how Uqbar stands as a unified execution environment for on-chain and off-chain data.

Uqbar is building a decentralized network on an operating system known as Urbit, an identity-driven, peer-to-peer, deterministic system that uses a functional programming language called Hoon. Uqbar is a decentralized network that leverages the best properties of Urbit's operating system and blockchains to create a self-sovereign, private, functional, unified and composable environment for on-chain and off-chain data. Uqbar itself acts as an execution layer as part of a modular blockchain stack, which is likely to settle transactions on Ethereum to begin with. With the advent of the

Uqbar testnet, we are getting a first glimpse into a variety of use-cases that simply do not exist on the internet or in blockchain today. This is because of the sheer execution power that Uqbar has as a result of being built on top of Urbit. In particular, because of the unified and zero-knowledge environment that Uqbar capitalizes on, creative use-cases such as gaming and visual art experiences will improve monumentally in crypto as Uqbar's native networking and unified functional programming experience orchestrates unprecedented composability. Today, NFTs often only represent data that is hosted elsewhere (e.g. an NFT has a hyperlink, which directs to a JPEG file hosted on IPFS). Tomorrow, on Uqbar, NFTs will represent data that is hosted natively, on Urbit or Uqbar, fully scalable, online and available. With the launch of Uqbar's Testnet and Ziggurat Developer suite, we are getting our first look into what is possible with Uqbar's toolkit. There are already 12 teams building applications on top of Uqbar, ranging from DEXs to RPGs. However, this is just the beginning for Uqbar. We assign a high probability to applications going-live on Uqbar's test-net that push the boundaries of what is possible in both web2 and web3 today. We're closely observing teams building on top of Uqbar and looking forward to experimenting as a sequencer in the network in the near future.

The Rust Blockchain Wars: Solana Down, Aptos Up

Solana was down for 6h on Oct 1st, when the network failed to reach consensus. Solana validator operators successfully completed a cluster restart of Mainnet Beta at 7 AM UTC. An official post-mortem has not yet been published, but instead of excessive traffic, this time the bug was caused by a combination of two factors:

1. A validator running two different instances at the same time, and;
2. A restriction in the consensus algorithm that prevents nodes from switching to a fork that includes a duplicated block.

Validators must be uniquely identified in the network by an identity key. One of Solana validators produced duplicated blocks during its leader slot and subsequent leaders disagreed upon considering or not the duplicated blocks, producing different network states to be voted on - a new fork. The fork with the majority of votes was the one including the blocks produced by the duplicated validator. Thus the nodes that have voted to a different fork should switch to this winner fork. The network was manually restarted to consider the blocks produced by the duplicated validator.

Solana has been investing efforts in preventing stability issues, although unpredictable combinations of factors due to extensive technical debt are hard to address. A much anticipated action that can be taken in this direction is the brand new implementation of new Solana clients. At the same time, the market has been closely watching the launch of new L1s that bring innovative solutions to the flaws of the well known networks. With a recent mainnet launch, the Aptos Blockchain gained attention from investors and from the community, essentially for bringing new approaches to improve:

- Security: Move language, designed specifically for blockchain use cases, differentiates accounts, tokens and coins, instead of treating digital assets and users as plain numbers. Developers are expected to focus on business logic and user experience, instead of manually having to close the doors that a mathematical operation can open when executing a transfer, for example.
- Scalability: optimistically parallelization of transaction execution through Block-STM, deals with dependencies across multiple transactions as they emerge - without the need of having dependencies specified in upfront. The blockchain is also designed in a modular way to allow updates with no downtime to users.
- Decentralization: PoS mechanism allows any validator with a minimum amount of stake to participate in consensus. Voting power is limited according to a maximum parameter (yield

throttling). Minimum and Maximum values, as well as other proposals, are discussed and voted following an on-chain governance process.

Often compared to the Solana Blockchain because of the high number of transactions per second (TPS) and a subsecond time to finality (TTF), Aptos also incorporates some known concepts from the Ethereum Blockchain, as transactions are buffered in a mempool and the cost of execution depends on the complexity and the gas price.

Aptos launched their mainnet on Monday 17 October. However, no decentralised applications have been deployed on Aptos yet, therefore it is hard to evaluate the network performance at this stage versus Solana.

MEV on Osmosis: Insights from building a bot

This quarter, we published @ChorusOneMev, a Twitter bot providing automated updates on arbitrage MEV extraction. Transparency is a pillar of our MEV strategy: to shed light on what, for a majority of users, amounts to an opaque process – how much profit is extracted, by whom, and how. To this end, the first chain covered by the bot is Osmosis, the prime exchange venue for the Cosmos-powered internet of blockchains. At inception, the bot was the only public source on MEV extraction on Osmosis. While there is significant interest in the field, it remains terra incognita for most. The goal of this article is to give an introduction to arbitrage MEV and our approach to it, and to share some learnings on searching and the idiosyncrasies of Osmosis along the way. It will conclude with a practical example - a complete architectural overview of our twitter bot.

Purpose of the bot

One goal of the project is to transparentize the MEV extraction process by highlighting profits, successful searchers, and transaction complexity. By leveraging this, we can actively engage to minimise negative MEV externalities for users through governance, awareness, technical, and research initiatives. The other goal is to create a data repository and infrastructural scaffolding for in-house searching (MEV extraction). It is important to understand what particularly successful searchers do well, and where most value extraction takes place.

There are two related measures of success – on the one hand, to maximise rewards, on the other hand, to do so in a manner that is sustainable and low impact on the network. In the absence of mempool prioritisation, many bots resort to spam to place transactions competitively. This and similarly destructive patterns are to be avoided by a responsible searcher. Exceptionally well performing bots typically minimise negative externalities, as these reflect as waste, for example via misspent gas fees.

One major upside of arbitrage MEV is efficient price discovery. Conceptually, searchers are rewarded by the market for providing accurate pricing. We foresee a wide opportunity space in the future, aligning prices between different networks, and networks and off-chain venues, for all assets with a price tag – from cryptocurrencies

to NFTs. MEV is a new field, and being early, informed, and principled is central to us. Below we describe our methodology that we currently employ for calculating MEV available on Osmosis.

An introduction to on-chain arbitrage

Arbitrage refers to the simultaneous buying and selling of an asset at different price points, generating a profit from the price difference. The classical example is an investor buying a share for $100 on exchange A, and immediately selling it for $101 on exchange B. There can be many reasons for such a price difference, like a legal moat, or a lack of liquidity.
In liquid and established markets, arbitrage opportunities are rapidly taken by sophisticated institutions.

On the blockchain, buy- and sell-transactions are often not executed against an order book as in the above example, but use automated market makers (AMM). An AMM is a lean way of enabling price discovery. While an order book needs to be constantly updated to provide competitive pricing, an AMM derives the asset price from a simple formula.

A common example of an AMM is a smart contract (called a pool) which holds two assets – x and y – and satisfies the relationship: $x*y = k$. Here, k signifies a constant balance of the assets in the pool. The price is derived through this relationship – the value of asset x in terms of y is simply: y / x.

For example, at the time of writing the OSMO/ATOM pool on Osmosis holds 3m ATOM and 36m OSMO. It follows that the price of 1 ATOM in terms of OSMO is: 36m / 3m = 12 OSMO.

A user may exchange OSMO for ATOM by adding OSMO to the pool. The user may then remove a proportionate amount of ATOM, such that the constant k does not change.

As there is no order book, the AMM only updates its pricing when someone trades against it. As such, it is kept aligned with other exchange venues through arbitrage.

The deterministic properties of a blockchain allow for a particularly attractive type of arbitrage called “atomic arbitrage”.
This refers to a swap transaction across several AMM pools, aligning their pricing, and only succeeding if the arbitrageur is guaranteed to emerge with a profit.

The properties of an atomic arbitrage are:

- The input asset equals the output asset.
- The output amount equals the input amount.

A typical example captured by our bot is a transaction where an arbitrageur swapped 91.26 OSMO to CHEQ, CHEQ to ATOM, and ATOM to OSMO, ending up with a balance of 91.86 OSMO,
for a $ 0.64 profit.

Observations on MEV on Osmosis

Most arbitrage MEV profits are realised during special situations, like the LUNA & UST crash.
During such events, searchers on OSMOSIS can earn outsized profits – the single most profitable transaction yielded a profit of ~$80k. A second attractive property is that little inventory is needed to successfully perform arbitrage on Osmosis – the currently most successful address has realised a profit exceeding $ 700k, while consistently holding an inventory smaller than $15k.

A noteworthy idiosyncrasy of Osmosis is the centrality of OSMO as a medium of exchange, due to protocol incentives. The quickest path between any two pools typically involves an intermediary exchange through OSMO, as it interfaces with most pools. Therefore, most MEV is extracted in OSMO, and searchers typically hold inventory in OSMO. The upside of this is more efficient price discovery than on a DEX with the same fees but no similarly dominant asset – the ubiquity of OSMO allows for less hops. Any exchange could make use of this pattern by strategically - and potentially selectively - placing incentives to optimise and direct liquidity, price discovery, and trade volume.

Building an arbitrage tracking bot - architectural overview

The first step is to fetch the raw block data. As our update is provided once daily, the most efficient way of doing this is to fetch the corresponding data in one go at the time of processing.
This allows for very simple coherence – the bot can check that the necessary data is in by following the sequence from the first to the last block in the set. It is also computationally more efficient than an on-going fetch. Therefore, the bot starts by estimating the block range that corresponds to the last 24 hours. This is done by estimating the likely range given the average block time on Osmosis, and then honing in through samples. Usually, it takes ~2-3s to get it right within a 5 minute error window.

The second step is to actually fetch the relevant blocks from one of Chorus One's RPC nodes.
The most convenient way of doing this is via osmosisd, which can parse the raw block data to JSON. This is the most time-intensive step of the process, and typically takes two to four hours, depending on network traffic.

Thirdly, the bot scans over all transactions in the blocks, and composes a dataframe of all swaps. A single transaction can contain several atomic swaps, this needs to be accounted for.

At this stage, the bot flags transactions that fit the arbitrage criteria laid out in the last section of this article for further processing. It proceeds to save the dataframe to a master database of all swap transactions, which can be efficiently queried later, for diverse data analytics tasks.

At this – fourth - point, a variety of housekeeping matters are addressed. The transaction data does not include readable asset names, but lengthy identifiers – IBC denoms. These need to be converted to a more recognizable identifier, this can be done via a RPC call. Similarly, asset amounts are given at different exponents. To process statistics correctly, the appropriate denom needs to be fetched from an RPC on a per-asset basis. Most of this can be done from the cache, but there needs to be a fallback mechanism to a live query to cover new assets as they appear.

As a fifth step, the bot converts the extracted MEV – the difference between the output amount and the input amount – to USD. This is done via Coingecko's API, at the time of processing.
This step can account for significant differences between different MEV trackers, irrespective of the network tracked. As cryptocurrencies are highly volatile, the timing of the price fetch can significantly skew output statistics. This is an important question to ask when analysing MEV statistics not presented in native (i.e. token) units.

Lastly, the bot computes various aggregate statistics and a leaderboard, which are then posted to the Twitter profile - @ChrousOneMEV. This includes information on the pools most frequently traded (pools featuring LUNC and UST tend to feature frequently in recent weeks), profitability, and transaction complexity (e.g. how many hops are traded on average).

Our Review of Cosmoverse Medellin

Cosmoverse'21 took place on November 5th and 6th, 2021 in Lisbon, Portugal. The conference gathered together developers, community members, investors, and entrepreneurs from all over the world to explore the Cosmos ecosystem's future. There were 50 speakers and approximately 700 people at the conference. Cosmoverse'22, completely organised by the community and with over 1300 attendees, proved that the Cosmos community is the most growing and vibrant community in the crypto industry. This time, the event raised awareness of the Cosmos ecosystem within the South American crypto community, as well as addressed current developments in interchain security, IBC, and Atom 2.0.

Cosmos' multi-hub design does not attempt to unify a global common state across all chains, but instead allows users to connect to each hub and establish their own settings in their particular 'app-chain.' Perhaps this is one of the reasons why we saw such a diversified community attendance in Medellin, from MEV and market maker experts to NFT and GameFi degens. The event circled around the announcement of the Atom 2.0 whitepaper and the new projects who aim to thrive in a resilient interchain community. The whitepaper has been announced as a proposal and it is still up for discussion. Here are the basics:

A novel architecture was introduced building on the foundations of the current Cosmos ecosystem, the Cosmos Stack.

Hub-specific Functionality	Interchain Scheduler		Interchain Allocator
Secure Economic Scaling	Interchain Security		Interchain Staking
Cosmos Stack	Tendermin	IBC Relayers	Cosmos SDK CosmWasm + Cosm JS

In order to secure a resilient economic staking, Interchain Security and Liquid staking have been studied by the community for implementation for many years. Interchain Security allows emerging projects to utilise and share the Cosmos Hub security, which reduces the marginal cost of security when compared to launching an independent validator set for an app-chain. A cheaper path to market will accelerate innovation, but it will also increase the chances to perform MEV. With more chain interaction, a greater chance of performing arbitrage emerges. The Interchain Scheduler tackles this issue by providing a MEV market on-chain, providing transparency. This has the potential to allow chains to regulate block space. The Scheduler has the potential of acting like a 'tax' frameworks to sponsor different parties in the ecosystem. These will mostly reward stakers who will no longer receive inflation rewards (during the transitional 36 months). Stakers will also obtain ICS rewards from new chains onboarding. We will need to see the pros frameworks to sponsor different parties in the ecosystem. These will

mostly reward stakers who will no longer receive inflation rewards (during the transitional 36 months). Stakers will also obtain ICS rewards from new chains onboarding. We will need to see the pros and cons of this change, if the rewards are lower than in actuality, governance can decide to distribute more rewards from the Treasury Pool to stakers.

Up until now, the Atom token has to balance between locked staked assets to provide security to the network, and available liquidity to use throughout the IBC ecosystem and provide utility. A safe implementation of liquid staking has the potential of unlocking liquidity without considerably affecting security, as liquid tokens will be provided for use in exchange of the staking of assets, eliminating the opportunity cost.

Lastly, the Interchain Allocator has been introduced to fund new projects leveraging from the returns of staked assets and monetary inflation. These profits will then be redistributed to the involved community. The whitepaper proposal is still up for discussion and an enriched conversation is already emerging in the forum.

Mesh Security

Later that evening, Sunny Aggarwal, founder of Osmosis, presented Mesh security. Mesh Security could be seen as a future development of Interchain Security (ICS). Interchain security is an implementation where a subset of the Cosmos Hub's validator set secures an emerging chain, lowering its security entry cost by paying a fee to the hub. Mesh Security applies this concept bidirectionally:

Mesh Security

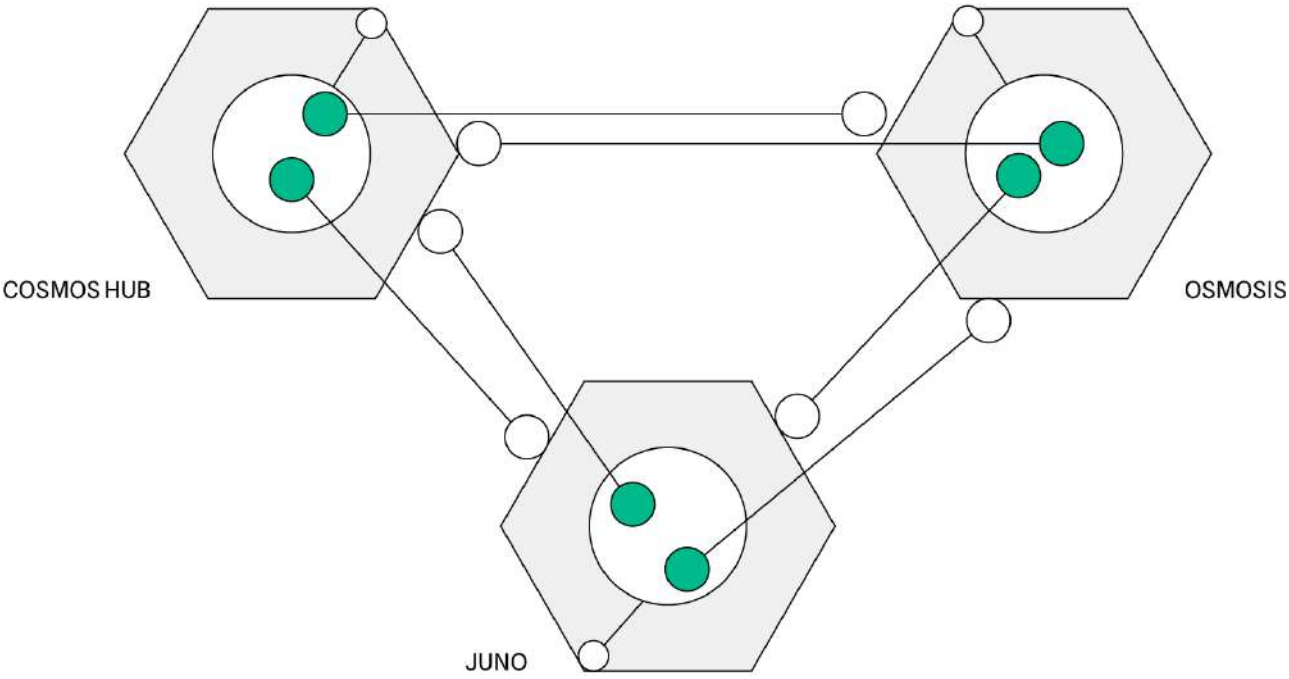

In this new scenario, each chain is both a provider and a consumer of security. Validators will cross-stake and submit transactions on both chains, and their identities will be correlated. Their stake and participation is slashed on both chains if they generate a hazardous

participation is slashed on both chains if they generate a hazardous block on either. This evolution seems beneficial for the ecosystem as a whole, but raises the question of what the position of the Cosmos Hub would be if this were to be implemented.

Networks that caught our attention:

Crescent's reincarnation as a DeFi Hub would greatly extend blockchain utility and connection by allowing users to move investments from other blockchains to Crescent via IBC for more investing options. It focuses on three primary features: Crescent DEX, Crescent Boost, and Crescent Derivatives. These enable a multi-chain asset marketplace with capital-efficient liquidity incentivization.

Now, they're developing a new chain called "Crescent," which will migrate these features so they may innovate even more. Liquid farming (with IBC), fund management to lessen the difficulty of becoming a Liquidity Provider (LP), and inventive market makers are among the new features. Crescent will provide a basic layer for the interoperability of DeFi protocols in a mix of these new capabilities nicknamed "L2Fi."

Sei is a trading-optimized L1. Innovative technical features include a native order matching engine that improves the scalability of orderbook exchanges, a new "Twin-Turbo Consensus" that reduces latency, and the establishment of Sei as the only Cosmos chain that uses market-based parallelization. Sei has achieved 22k orders per

second and 0.45 second block timings using optimistic block processing and the Twin Turbo Consensus, making it the fastest Cosmos SDK chain. This opens the door for decentralized exchanges and trading order-books to thrive on-chain. Protocols such as Vortex have created a perpetual's protocol by leveraging this quick finality. Sei also demonstrated an intriguing future, with over 50 projects from Cosmos and non-Cosmos environments committing to launch their DeFi protocols on Sei. An exciting example is Nitro, the first Solana VM who will port to Cosmos through Sei.

Liquid staking is one of the core values of Cosmos. We have seen a thrilling competition which will make out of Cosmos the go-to for liquid staking utilities. We have a close relationship and are excited about the Quicksilver developments. Quicksilver enables higher yield to stakers and increases liquidity and security supporting Cosmos DeFi. Importantly, Quicksilver enables all these whilst respecting sovereignty. This is achieved by allowing users to select their validator, signaling intent to allow delegators to dictate where their voting power should go and providing governance by proxy. Additionally, stakeholders will be able to transfer their assets directly to Quicksilver without the need of unbonding where rewards are auto-compounding.

Agoric is a Proof-of-Stake chain that uses Hardened JavaScript smart contracts to develop and deploy DeFi quickly. Agoric uses a JavaScript First approach to provide a familiar developer experience while also providing a composable library and a secure architecture.

To fully appreciate Agoric's potential, consider that there are 6000 Solidity developers and 18000 web3 developers, compared to 17M JavaScript developers. This provides a familiar entry point for web2 developers into web3. LH2 Staking, Calypso (an interchain Dashboard), AgreeWe for DAO tooling, and Crabble, an NFT lending platform, are some Mainnet Dapps. Agoric is now in Mainnet 1 phase, with a functional BLD token and the imminent release of the Inter Stable Token (IST), a new stable asset for the Cosmos ecosystem regulated by the Inter Protocol, developed in collaboration with Zaki Manian.

MEV As A Business Model For Networks

A core enabler for the growth of crypto networks has historically been the issuance of new tokens to compensate active participants. From the Bitcoin block reward, to the distribution of inflated tokens to stakers in Proof-of-Stake, to yield farming enabling DeFi protocols to bootstrap liquidity.

All these protocols usually have in common that token incentives are targeted to be temporary; distributed to bootstrap the protocol. There usually is a fee, or an expectation of a fee, that a protocol is or will charge for the services it provides to users. These fees are expected to at some point compensate participants and reduce the need for dilution to pay for the services validators and their delegators, or - in DeFi protocols - other parties such as liquidity providers, provide.

In the past, the most common place these earnings have been expected to come from in smart contract networks was from gas fees paid to execute computation on the network. But improved scalability and the impact of high fees on the user experience have put this model somewhat into question. If networks can scale enough to serve all of their users' needs, do they have to keep

blockspace artificially scarce to create a sustainable business model?

Enter MEV; a relatively recent phenomenon in the crypto space. The realization that there exists opportunities for block proposers to profit from inserting or reordering transactions in their blocks. Examples include e.g. arbitrage transactions balancing prices between DEXes, or - more exploitative - front-running DEX users. MEV as a research field is experiencing a rapid evolution. At a high-level it looks as if leading protocol teams and the ecosystem at large are on a trajectory towards minimizing negative externalities and implementing solutions to create markets and mechanisms that ultimately aim to redistribute profits from these activities to network participants. This suggests that some of the yield required to ultimately pay for network security can be expected to come from MEV.

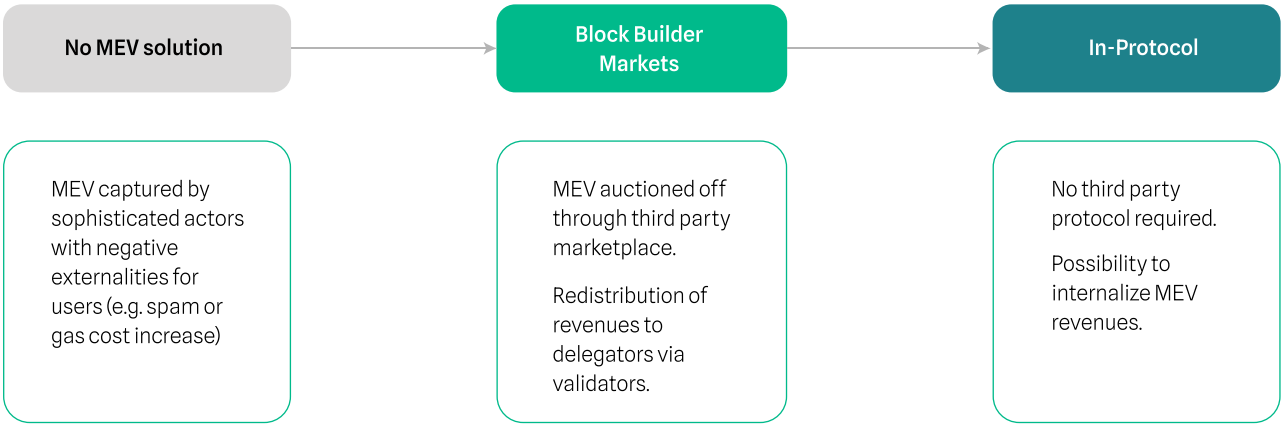

Possible trajectory of MEV solutions

Post-merge Ethereum is an interesting case study for an early look of how MEV reward redistribution can account for a significant portion of the staking rewards in the most mature crypto ecosystem.

At the time of writing, a bit over 50% of the maximally possible APR achievable when staking ETH comes from new issuance of ETH, with 35% coming from fees and 11%, or roughly 1% APR, in additional APR for validators that operate the block builder solution MEV-boost and receive whatever searchers are willing to pay to get their transaction bundles included in their blocks. This shows us that MEV already plays a significant role in funding the security of the Ethereum network and may over time enable further reductions in issuance.

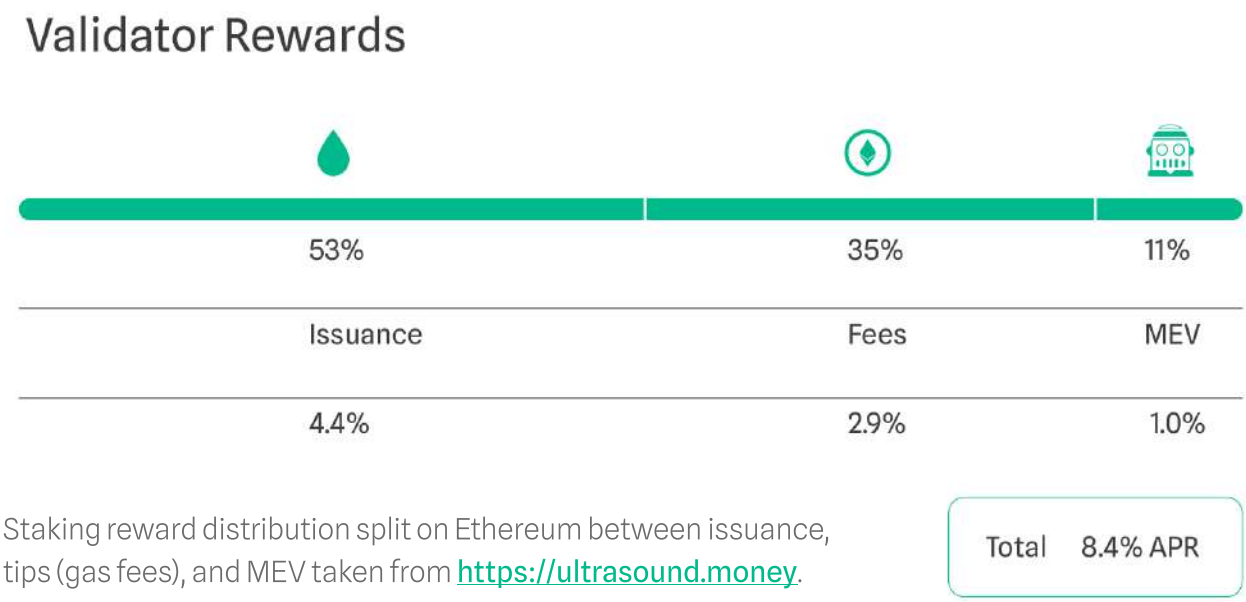

Staking reward distribution split on Ethereum between issuance, tips (gas fees), and MEV taken from https://ultrasound.money.

We believe MEV may be a turning point in the trajectory of business models for PoS networks and opens a novel design space for protocol designers and network participants. Join us at the Staking Summit in Lisbon, where Felix will present further insights from our research into this topic.

Connect with Chorus One

Reach out to our Ventures team
ventures@chorus.one

Speak to our Sales team
sales@chorus.one

Contact our Research team
research@chorus.one

Visit our website
chorus.one