HOUSE OF CHIMERA

House of Chimera Research: Blockchain developments on Syscoin

Our research paper aims to provide insights into the latest blockchain developments in the cryptocurrency industry while suggesting a way forward in a rapidly maturing landscape.

September, 2022

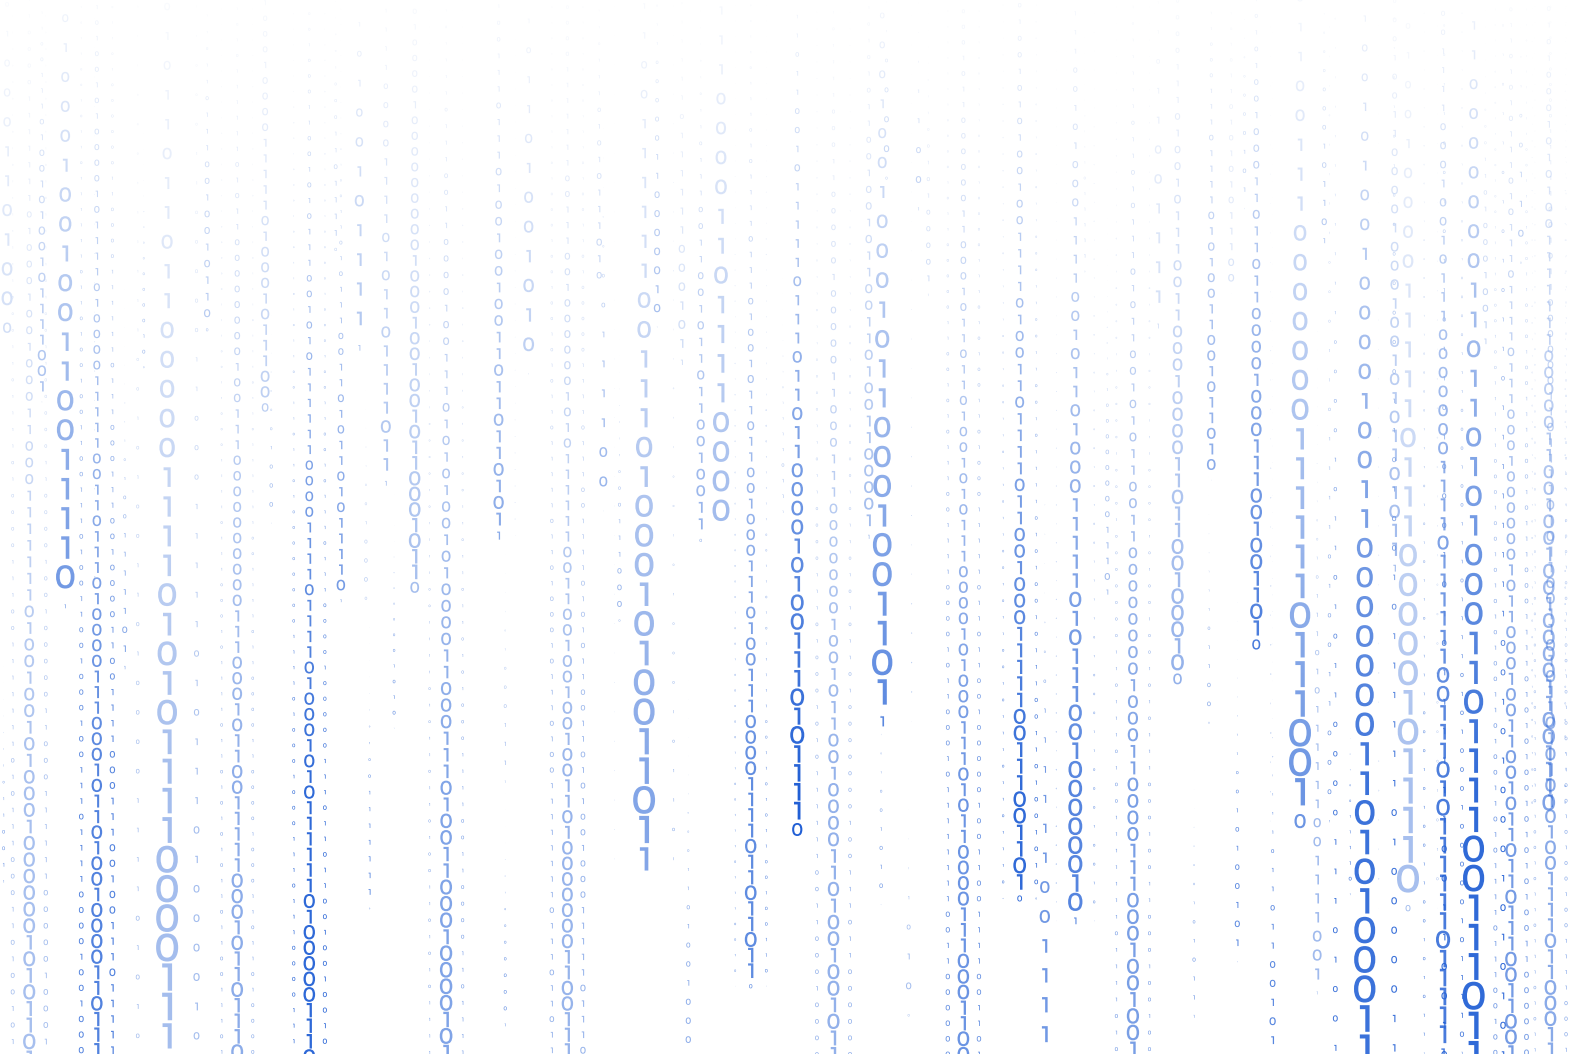

HOUSE OF CHIMERA

Contents

Introduction

The blockchain industry is rapidly evolving; new techniques are emerging, and cryptocurrencies are more integrated than ever. The dynamic environment leads to innovative advancements in security, decentralisation, and scalability. Traditional companies increasingly use blockchain technology for comprehensive use cases. This research aims to provide an outlook on the latest developments in the blockchain industry.

Diederick Jacobs

Founder

House of Chimera

HOUSE OF CHIMERA

Photo by Carlos Muza on Unsplash

Distributed Ledger Technology pushes the boundaries of information technology

Blockchain technology has shown the ability to innovate and disrupt traditional industries; however, blockchain technology remains complex for the masses.

Distributed Ledger Technology (DLT) (i.e. Blockchains) is attracting significant attention owing to its ability to disrupt traditional industries. The overall media coverage of blockchain technology has been considerable, leading to increased interest by companies. Currently, a dozen prominent financial companies (e.g. Mastercard, JPMorgan, Blackrock) have been interacting with the blockchain technology or cryptocurrencies. The reasons for utilising blockchain technology do differ for each company. However, market inefficiencies, substantial overhead costs, and information asymmetry are a few of the imperfections of the financial industry. To put this into perspective, in 2008, JPMorgan acquired investment bank Bear Stearns; however, the number of shares offered to the acquirer was larger than the shares outstanding in the books of Bear Stearns. JPMorgan could not clarify the accounting errors; therefore, they had to bear the damage from the excess digital shares. The incident was mainly caused by inadequate accounting and overall tracking of assets. However, tracking ownership over an extended period is challenging if ownership changes rapidly. This is one of the issues that a public blockchain solves. According to Oxford Language Dictionary, the term 'Blockchain' is defined as follows:"A system in which a record of transactions made in Bitcoin or another cryptocurrency are maintained across several computers that are linked in a peer-to-peer network".

Bitcoin was the first proposed cryptocurrency, which was invented in 2008 and implemented in 2009. A public blockchain is a blockchain that stores data publicly with digital signatures in a decentralised network. Essentially, a public blockchain is a decentralised open database that stores information electronically in a digital format. The main difference between traditional databases and a blockchain is how the data is structured and stored. A database usually structures data into tables containing all the data. A unique characteristic of a blockchain is that every block is linked and forms a chain of information (hence blockchain). When a block is filled, the data in the block is immutable and, therefore, cannot be altered. To be precise, modifying the blockchain and all the blocks through a reorganisation is possible. However, this would imply that the malicious actor needs at least 50.1% of the processing power or tokens, depending on the consensus method (e.g. Proof-of-work, proof-of-stake). In the event of a blockchain reorganisation, this would mean that the malicious actor is changing the structure of the blockchain. This could mean that previously valid blocks are deleted, leading to reversed transactions. In that case, a Double Spending (DS) issue could occur whereby users spent their assets in the past; however, considering transactions are reversed, they have gotten their assets back.

A public blockchain is fully transparent, which provides an entirely auditable and valid ledger of transactions. herefore, the transparency of blockchain offers users the convenience of looking through their history of all transactions but also imposes accountability on any network user. Consequently, blockchain use cases are versatile; blockchain could enable consumers to track anything across a supply chain, proving product genuineness, reviewing workers' rights, and even allowing them to examine food ingredients.

[1] Nofer, M., Gomber, P., Hinz, O., & Schiereck, D. (2017). Blockchain. Business & Information Systems Engineering, 59(3), 183-187.

[2] Dai, H. N., Maharjan, S., Zheng, Z., Hung, P. C., Xu, Q., & Sun, W. (2021). IEEE Access Special Section Editorial: Blockchain-Enabled Trustworthy Systems. IEEE Access.

[3] Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Decentralized Business Review, 1-9.

[4] Kritikos, M. (2018). What if blockchain offered a way to reconcile privacy with transparency?

[5] Damoska Sekuloska, J., & Erceg, A. (2022). Blockchain Technology toward Creating a Smart Local Food Supply Chain. Computers, 11(6), 95.

Blocks have specific storage characteristics (e.g. total block capacity), which allows data to be stored. Typically, every block stores a cryptographic hash code, previous block hash and its data. An average Bitcoin block has a block size of 1 MB, which roughly translates to 1600 to 2000 transactions. Additionally, on average, the Bitcoin block creation time is ten minutes, meaning a block is created approximately every ten minutes, which translates to transactions per second (TPS) of approximately 4 to 5. To put this in perspective, the Dogecoin network utilises block sizes of 1 MB but has a block creation time of 60 second, which translates to approximately a TPS of 33. The differences in blockchain characteristics are issued by the consensus of all network participants. Thus, the network participants have an agreement on the rules they all follow. One of the reasons behind the Bitcoin block creation time is to ensure decentralisation; by having a relatively high block creation time, everyone can download the latest block and therefore, be synchronised with all other nodes without needing an exceptional internet connection.

Currently, a dozen significant financial companies (e.g. Mastercard, JPMorgan, Blackrock) have been interacting with the blockchain technology of cryptocurrencies.

[6] Alam, T. (2019). Blockchain and its Role in the Internet of Things (IoT), 5(1), 151-157.

[7] Github. Retrieved August 28, 2022, from: https://github.com/bitcoin/bitcoin/commit/a30b56ebe76ffff9f9cc8a66671861794l3c6349.

[8] Blockchain.com. Retrieved August 28, 2022, from: https://www.blockchain.com/charts/n-transactions-per-block.

Modular blockchains, the new era of the distributed ledger technology

The core tasks for every blockchain are similar: coming to a consensus, providing security, guaranteeing data availability, and executing transactions. Consensus refers to the process whereby validators agree on which data can be verified as genuine and accurate and, therefore, should be included in the blockchain. Data availability is the theoretical guarantee that a validator has published all available data for a block and made it available for all other network participants The last core task of a blockchain is to execute and settle transactions. The execution refers to nodes participating in consensus and executing transactions using their copy of the blockchain to attest before validating blocks. The settlement provides finality to transactions, a guarantee that a transaction has been added to the blockchain and is, therefore immutable. Commonly, all these tasks are performed on the same layer, which would make a monolithic blockchain (e.g. ETH 1.0). The alternative is modular blockchains (e.g. Syscoin), whereby at least one of these core tasks is performed on a different layer. Thus, every layer specialises in a specific set of tasks and solely performs these assigned tasks. Consequently, this leads to higher throughput, flexibility, and enhanced developer accessibility.

Literature has shown that blockchains can be used in different applications and industries. However, blockchains do suffer from limitations restricting the overall practical use cases.

The blockchain trilemma is a well-known and relatively long-standing cryptographic problem; it was first described by Vitalik Buterin, the Co-Founder of Ethereum. Vitalik implies that trade-offs are inevitable between the three primary blockchain properties: Scalability, Decentralisation, and Security. Scalability is the overall ability of a blockchain to handle an increasing number of transactions and, therefore, meet the consumer's overall transaction demand. Decentralisation can be defined as the transfer of supervision and decision-making from a centralised actor (i.e. corporate, institute) to a dispersed network, whereby there is no need for trust dependencies in small or large centralised actors. The security of the blockchain is the overall capability of safely storing data without the possibility of any alteration in the future.

Blockchain networks' key issue is achieving the perfect balance between decentralisation, security, and scalability without making any significant trade-offs. Several scalability solutions have been proposed in the literature, such as traditional Sharding, Lightning Network, increasing block sizes, and a relatively new scalability method which is called rollups.

Ethereum was initially developed as a monolithic blockchain; at the time of writing, it is transitioning into a modular framework through an on-chain scalability method: Danksharding.

---

[9] Alchemy. Retrieved August 28, 2022, from: https://www.alchemy.com/overviews/modular-vs-monolithic-blockchains.

[10] Etherium. Retrieved August 28, 2022, from: https://www.alchemy.com/overviews/modular-vs-monolithic-blockchains.

[11] Alchemy. Retrieved August 28, 2022, from: https://www.alchemy.com/overviews/modular-vs-monolithic-blockchains.

[12] Aggarwal, S., Chaudhary, R., Aujla, G. S., Kumar, N., Choo, K. K. R., & Zomaya, A. Y. (2019). Blockchain for smart communities: Applications, challenges and opportunities. Journal of Network and Computer Applications, 144, 13-48.

[13] Buterin, V. (2021). Why sharding is great: demystifying the technical properties.

[14] Luu, L., Narayanan, V., Zheng, C., Baweja, K., Gilbert, S., & Saxena, P. (2016). A secure sharding protocol for open blockchains. Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, 17-30.

[15] Poon, J., & Dryja, T. (2016). The bitcoin lightning network: Scalable off-chain instant payments, 1-59.

[16] BitcoinABC. Retrieved August 29,2022, from: https://www.bitcoinabc.org/2018-09-07-bitcoin-abc-and-the-block-size-limit/.

[17] Sedlmeir, J., Buhl, H. U., Fridgen, G., & Keller, R. (2021). Recent developments in blockchain technology and their impact on energy consumption, 1-11.

Danksharding is not similiar to traditional sharding, whereby the blockchain is divided into multiple parallel chains. Sharding utilises numerous parallel chains, thus, transactions can be handled simultaneously. Consequently, drastically increasing the throughput and, therefore, the scalability of the chain. However, Danksharding significant differs from traditional sharding. Danksharding does not split the chain but utilises a proposer-builder separation (PBS) mechanism.

The block proposers provide the block builders with a list of data (i.e. crList), which indicates the transactions which should be included in the block. In a monolithic chain, miners commonly combine and achieve these tasks. The block builders can reorder data (i.e. transactions) to maximise Miner.

Extractable Value (MEV). However, the block builders need a proposer to create a block. This is done by offering the proposer a share of the builders' revenue; therefore, a competitive fee market is created. Both parties are incentivised to work as efficiently as possible, reducing users' transactional fees. Once the proposer has chosen a builder, that particular builder creates the whole block. Therefore, only a single actor produces the block instead of many, which typically happens in traditional sharding. The main advantage of having solely one actor creating the block is that data validation can be done in aggregate. Thus, preventing any shard block confirmation delays and shard blob confirmations do not have to be tracked; they are confirmed on the main chain. This is achieved by coupling sharding and execution blocks; thus, there is no need for a separate shard for transactions (figure 1).

In a traditional sharded chain, each produced block has to be confirmed within a shard. This is achieved through a shard committee; commonly, if the block contains signatures of 66% of the total voting power of the committee, the block is confirmed (figure 1).

Figure 1 Traditional Sharding vs Danksharding (Add source: Ethereum.org)

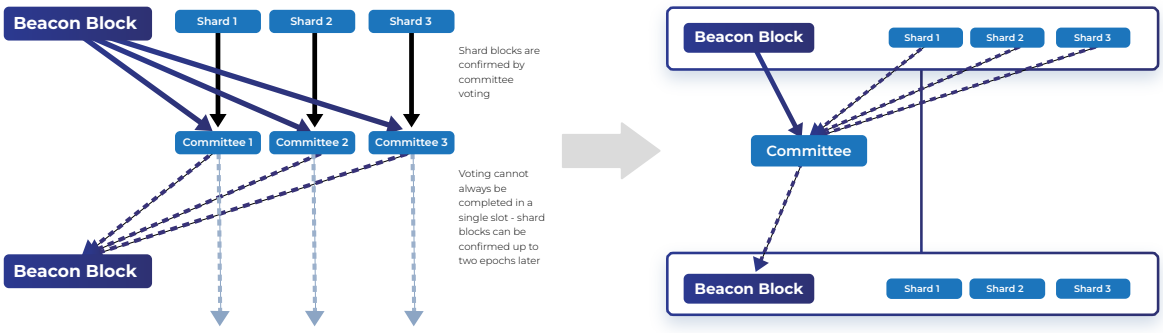

Source: Ethereum.org

[18] Sel, D., Zhang, K., & Jacobsen, H. A. (2018, December). Towards solving the data availability problem for sharded ethereum. In Proceedings of the 2Nd Workshop on Scalable and Resilient Infrastructures for Distributed Ledgers (pp. 25-30).

Getting the perfect balanced layered blockchain: The recipe for success

Blockchain technology provides significant advantages in terms of transparency, security, resource efficiency, and eliminating intermediaries. However, blockchain technology does not solely consist of a data layer whereby transactions are digitally signed, and information (e.g. transaction data) is stored in blocks. But instead, it utilises different layers for scalability, security, interoperability, and applications. A blockchain has a layered architecture that can be categorised into four categories (figure 2). Every layer has its specific purpose, and these layers can be imagined as layered onion, whereby the onion's core refers to as Layer 0 (L0).

L0 is the base layer, which includes hardware and software that builds the backbone of the blockchain ecosystem. To put this in perspective, the base layer ensures interoperability, allowing blockchains to interact with each other.

The first layer (i.e. Layer 1) is the blockchain responsible for carrying out a wide array of tasks and maintaining the blockchain network's fundamental operations. This layer is essentially the engine of the blockchain technology stack; the consensus method, block times, dispute resolutions, and programming languages are all part of this layer.

The second layer (i.e. Layer 2) enhances the scalability of the main chain (i.e. layer 1), essentially providing extra processing power. Layer 2 is a third-party integration developed on top of layer 1 and manages all the transactional validations. In contrast, layer 1 remains responsible for adding and creating blocks to the blockchain. Although it is possible to increase the throughput of layer 1, some practical blockchain limitations have been set through consensus. Therefore, the validators must accept significant changes in a public blockchain to embrace any alterations via a hard fork. A hard fork is required if there is a substantial modification to the blockchain, which makes previously deemed invalid transactions valid, or vice versa. To put this in perspective, increasing block size would require a hard fork. To successfully implement a radical change, the network participants have to form consensus. Otherwise, the blockchain would split into two chains, potentially impacting the decentralisation of both chains.

The third layer, is the most recognisable layer for most cryptocurrency users. This layer allows participants to interact with user interfaces, such as Decentralised Applications (dApps) (e.g. Uniswap, Metamask).

Figure 2 Blockchain layered architecture

Summary

A public blockchain is a blockchain that stores data publicly with digital signatures in a decentralised network. Essentially, a public blockchain is a decentralised open database that stores information electronically in a digital format. Bitcoin was the first proposed cryptocurrency, invented in 2008 and implemented in 2009. Bitcoin is considered the first widely-used blockchain application. The core tasks for every blockchain are similar: coming to a consensus, providing security, guaranteeing data availability, and executing transactions. Commonly, all these tasks are performed on the same layer, which would make a monolithic blockchain (e.g. ETH 1.0). The alternative is modular blockchains (e.g. Syscoin), whereby at least one of these core tasks is performed on a different layer. A blockchain has a layered architecture, whereby there are 4 layers. Each layer specialises in a specific set of tasks and solely performs these assigned tasks. Consequently, this leads to higher throughput, flexibility, and enhanced developer accessibility.

The blockchain trilemma is a well-known and relatively long-standing cryptographic problem; it implies that trade-offs are inevitable between the three primary blockchain properties: Scalability, Decentralisation, and Security. The blockchain trilemma is not solved yet; however, several scalability solutions have been proposed in the literature, such as traditional Sharding, Lightning Network, increasing block sizes, and a relatively new scalability method which is called Rollups. Ethereum was initially developed as a monolithic blockchain; at the time of writing, it is transitioning into a modular framework through an on-chain scalability method: Danksharding. Danksharding is a sharding design that integrates the concept of a merged market fee through a proposer-builder separation (PBS) mechanism.

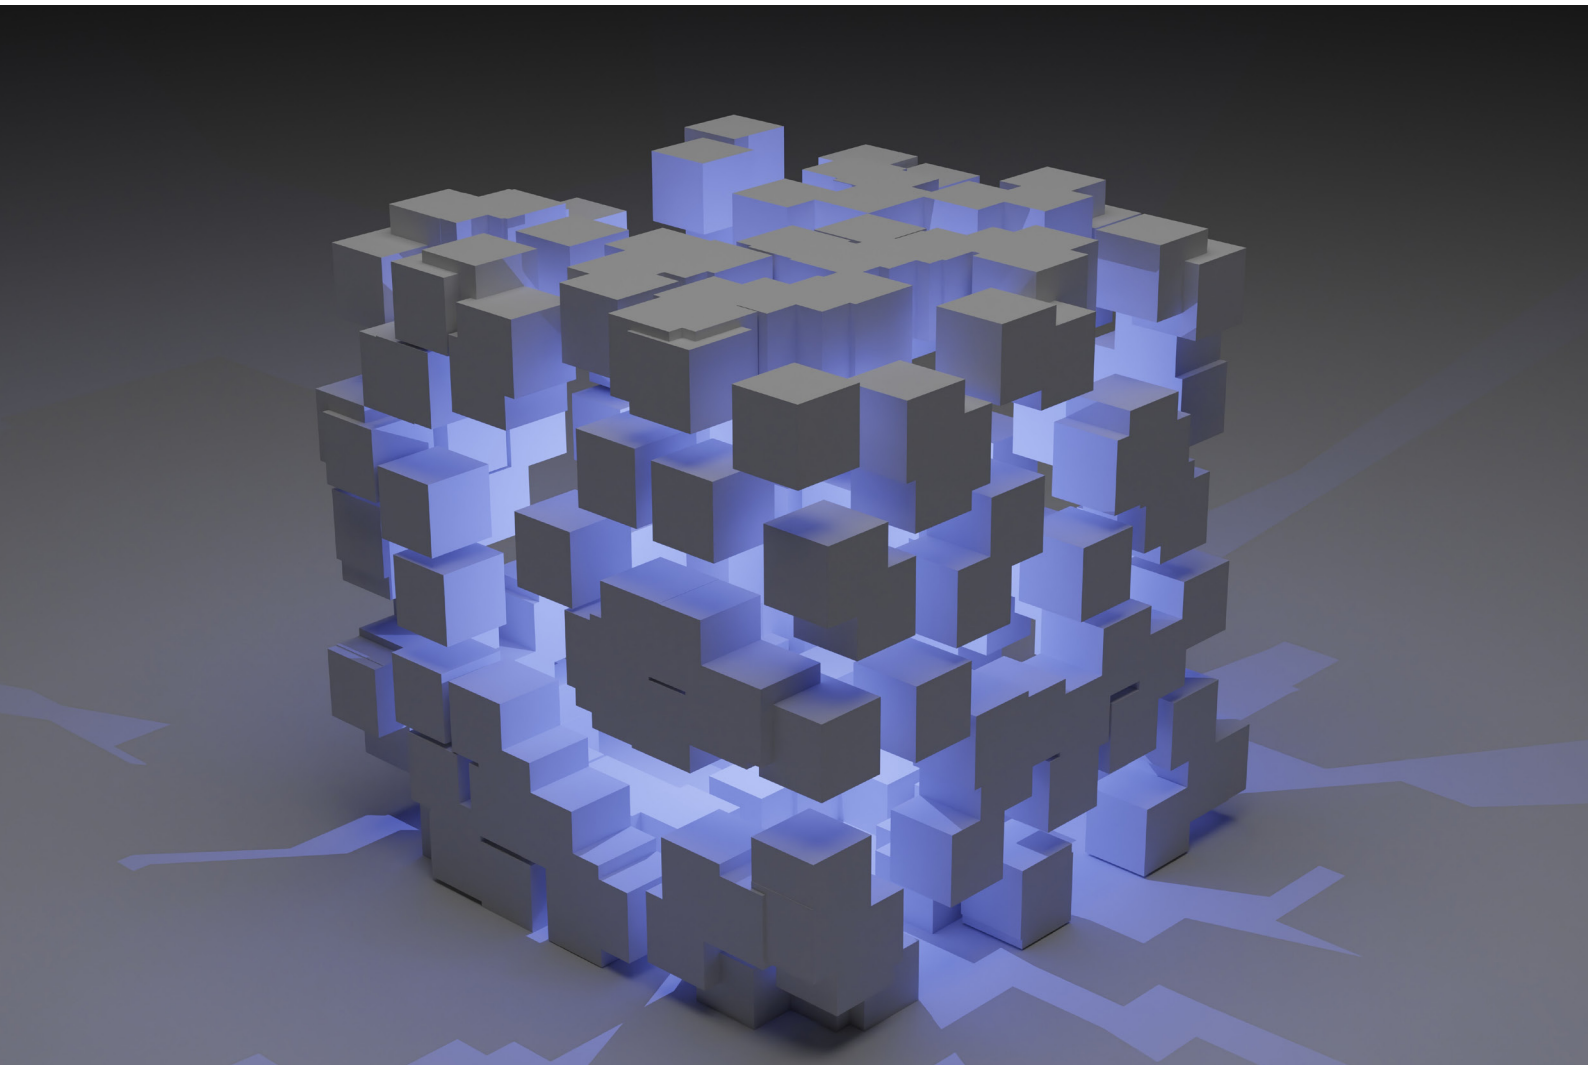

Photo by GuerrillaBuzz Crypto PR on Unsplash

Scaling blockchains through off-chain scalability methods.

Enhancing the scalability of blockchains without significant security or centralisation has evolved as a significant challenge.

The Danksharding scalability method is an on-chain scalability method, making Ethereum a more scalable chain. However, more alternatives exist to transform a monolithic chain into a modular framework. Off-chain scalability methods process transactions outside the main chain, therefore; maintaining the state of the main chain while applying the last state that has been processed in the other chain (i.e. layer 2).

For instance, Polygon Network is a Layer 2 built on the Ethereum Network and utilises Plasma to increase scalability drastically. Plasma is a framework of child chains, which have an independent consensus mechanism and produce their own blocks A Child-chain has a parent-child structure whereby the main chain (e.g. Ethereum) is the parent, all transactional functions of the blockchain are performed in the child chains, and the result is posted to the parent chain.

Plasma contracts enable the bridging of assets from the parent chain to the child chains, similarly to sidechain functions. However, Plasma relies on and benefits from the parent chain's security. In contrast, a sidechain is solely responsible for its security. Plasma allows for high- throughput and drastically increases the scalability. However, the overall use case is minimal. Plasma does not support general computation (i.e. smart contracts); therefore, it can be used solely to transfer assets.

Furthermore, Plasma does utilise off-chain storage and does not post transactional data on the parent chain. Therefore, withdrawals can take several days (i.e. fraud-proof challenge period). The fraud proofs ensure that in a case of malicious activity (i.e. invalid state transitions), users can report (i.e. challenge) dishonest nodes to safeguard their funds and exit the transaction. These proofs involve posting the entire valid state transitions of the child chain onto the parent chain. Liquidity providers can mitigate the challenge period; however, there is an associated capital expense to hedge the operational risk.

Gnosis, previously xDai, is an example of a sidechain on Ethereum. Sidechains are independent blockchains, having their own consensus method, ecosystem design and parameters (e.g. different block creation times) and security mechanisms. Thus, sidechains are isolated from the main chain, meaning that in case of a cryptographic break, the damage is limited to the sidechain itself Sidechains utilise a two-way peg to move assets from the main chain to the side chain. The two-way peg (i.e. blockchain bridge) facilitates communication between blockchains, allowing for the transfer of information, thus, the transfer of assets. The flexibility of a sidechain provides for a broad set of use cases due to being an independent blockchain, as often, the blockchain rules of the main chain are challenging to adjust unless the majority of the participants form a consensus. However, the flexibility of a sidechain comes with a significant security risk. Sidechains can set their own rules; as an independent blockchain, this could lead to vulnerabilities. If the developer configuration of the sidechain is flawed, this could lead to exploits and vulnerabilities. Additionally, sidechains often

[19] Kim, S., Kwon, Y., & Cho, S. (2018). A survey of scalability solutions on blockchain. International Conference on Information and Communication Technology Convergence (ICTC), 1204-1207.

[20] Wiki. Retrieved August 30, 2022, from: https://docs.polygon.technology/docs/home/blockchain-basics/sidechain/#:~:text=Plasma%20is%20a%20framework%20of,some%20other%20%E2%80%9Cmain%E2%80%9D%20blockchain.

[21] Ethereum. Retrieved September 2, 2022, from: https://ethereum.org/en/developers/docs/scaling/plasma/.

[22] Ibid.

[23] Back, A., Corallo, M., Dashjr, L., Friedenbach, M., Maxwell, G., Miller, A.,& Wuille, P. (2014). Enabling blockchain innovations with pegged sidechains, 72, 201-224.

[24] Virconlegal. Retrieved September 2, 2022, from: https://virconlegal.com/comprehensive-analysis-of-recent-side-chain-hacks/.

suffer from a centralised two-way peg, as one usually has to trust a single entity to manage the asset transfer from the main to the side chain and vice versa.

Rollups, the most promising off-chain scalability method.

Rollups are one of the several off-chain scalability systems; however, it is perceived by many as the most promising scalability method for public blockchains. The main advantage of rollups is that they can achieve a significant amount of TPS without making significant trade-offs regarding scalability, decentralisation, or security. Rollups handle transactions off-chain; however, they post transactional data on the main chain. To decrease on-chain storage, rollups bundle transactions and compress them before sending them to the main chain. Therefore, the amount of data posted on the main chain is significantly reduced, increasing the overall network resource efficiency and scalability.

Rollups rely on the main chain's security, considering the smart contracts are deployed on the main chain (figure 3). Different operators in a rollup structure interact with the smart contracts.

Aggregators publish transaction data and other information through the smart contract. Verifiers utilise it to dispute transactions if needed. Users can interact with the rollups through the rollup smart contract, which lives on the main chain. Therefore, users have to deposit funds on the main chain, and the deposit provides the user with a proportional amount of the same token on the rollup. Whenever users perform transactions, the aggregator selects a set of transactions, executes them, and posts the compressed transaction data on the main chain. Therefore, the published data becomes public and immutable. Users can exit the rollup by interacting with the smart contract and thus, receive their proportional amount of assets back on the main chain. There are currently two primary types of rollups: Optimistic and Zero Knowledge Roll-ups.

Figure 3 Schematic rollup overview

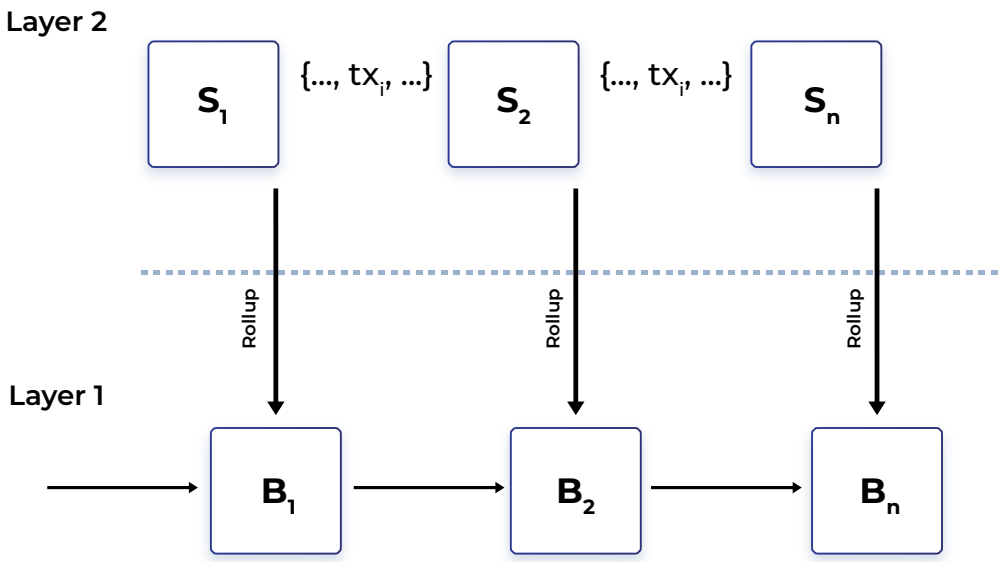

[25] Singh, A., Click, K., Parizi, R. M., Zhang, Q., Dehghantanha, A., & Choo, K. K. R. (2020). Sidechain technologies in blockchain networks: An examination and state-of-the-art review. Journal of Network and Computer Applications, 149.

[26] Thibault, L. T., Sarry, T., & Hafid, A. S. (2022). Blockchain Scaling using Rollups: A Comprehensive Survey. IEEE Access.

[27] Ibid.

Zero Knowledge Rollups

Zero-knowledge (ZK) rollups are similar to Plasma; however, despite the similarities, the design of ZK rollups is relatively complex and has a few significant differences. Plasma utilises fraud proofs to ensure security on the child chain, where all the data is stored off-chain. This leads to a complex data availability problem, which will be highlighted later in this paper. ZK Rollups utilise ZK-proofs; which are used as cryptographic proof for a bundle of transactions that are processed off-chain. The ZK proof is used to prove the validity of transactions, and every batch has its validity proof submitted to the main chain (i.e. Layer 1). The ZK-proofs can be handled off- and on-chain. A significant concern is that transactions could get stuck in the case of off-chain ZK-proofs; however, multiple built-in censorship resistance exiting mechanisms prevent user transactions from getting stuck. Censorship resistance means users can exit the layer 2 (i.e. ZK-roll up) without the requirement of coordinating the layer 2 consensus.

The ZK-proof consists of a compressed transactional data summary; therefore, the overall stored amount of data is minimal compared to typical layer 1 transactional data flows. Despite the compression, the amount of information is sufficient to confidently validate the data without revealing the data. The amount of stored on-chain data can be lowered by indexing instead of transaction addresses. Therefore, decreasing the amount of on-chain data even further. One of the significant effects of having a low amount of on-chain data is the exponential increase in scalability, lowering transaction costs by increasing the overall network capacities.

In September 2018, Vitalik provided a rough calculation of the expected transaction throughput of ZK-rollups for Ethereum. At that time, the on-chain 'calldata' cost was 68 gas per byte. Call data is, simply put, all the data passed to the smart contract at execution. To put this in perspective, if one performs an on-chain transaction on a public blockchain, there is data that needs to be included (e.g. destination, amount, signatures). The advantage of call data is that it behaves as memory; therefore, no unnecessary copies are needed to ensure the data is unaltered. The expected gas for a regular ZK-rollup transfer was 892 gas, according to Vitalik. Considering the gas limit for an Ethereum block was 8,000,000 gas, the approximate transactional throughput was around 550 TPS However, the calculation is relatively outdated and is not accounting for potential gas reductions. Since September 2018, Ethereum has significantly decreased the gas cost of calldata by approximately a factor of four.

Consequently, the expected gas usage for a regular ZK-rollup transfer decreased to 208. Furthermore, the block size limit of Ethereum was increased, through a dynamic block size mechanism whereby the block size can increase from 15,000,000 to 30,000,000 bytes based on market demand. Assuming that the overall gas cost of ZK-proof verification and the ZK-rollup overhead remained stable, the expected TPS increased to roughly 5,500 to 10,000 TPS, based on an average block time of 13 seconds.

---

[28] Schaffner, T., & Schaer, F. (2021). Scaling public blockchains—A comprehensive analysis of optimistic and zero-knowledge rollups (Doctoral dissertation, Master's Thesis, Center for Innovative Finance, University of Basel).

[29] Research. Retrieved September 3, 2022, from: https://ethresear.ch/t/on-chain-scaling-to-potentially-500-tx-sec-through-mass-tx-validation/3477.

[30] Ethereum Improvement Proposals. Retrieved September 3, 2022, from: https://eips.ethereum.org/EIPS/eip-4488

ZK-Proofs

ZK proof is a cryptographic method that allows a party (the prover) to prove to another party (the verifier) that a given statement is true without providing additional information. A ZK proof has three core characteristics (figure 4), Completeness, Soundness, and Zero-knowledge. To make the idea less abstract, there is a relatively simple concept given by Chalkais.

Figure 4 ZK core characteristics

Completeness	This property is satisfied if both verifier and prover are honest
Soundness	This implies that a dishonest prover can't convince an honest verifier
Zero-knowledge	This property makes zero-knowledge proofs unique, convincing an involved party without providing any knowledge.

What if your friend is color blind?

Imagine that your friend is colour-blind and cannot distinguish the colour red from green. Your friend has a green and a red ball that is otherwise identical. Your job is to convince your friend that the balls differ in colour while revealing nothing else. To persuade him, you would ask your friend to show you the balls and put them behind his back afterwards. Then, he may switch the balls behind his back and show you a single ball. The question arises: Did he switch the position of the ball behind his back? You, as prover, should be able to tell him if he did, assuming you are not colour-blind. Hence, you could convince him that the balls differ in colour with a high probability of success (i.e. 99%).

However, let's assume that the prover (i.e. you) has malicious intent, and consequently, you are lying to your friend. The two balls are the same colour; according to the Law of Large numbers of Bernoulli. the expected probability of you guessing right is approximately 50% after a high amount of games (e.g. 1000 switches). Since a high likelihood of success by guessing is improbable, your friend can assume that you are stating the truth: the balls differ in colour.

The example provided above describes an interactive ZK proof, considering there is interaction between the prover (i.e. you) and verifier (i.e. your friend) through the entire process until the verifier is convinced.

[31] Chalkias, K. (2021). Demonstrate how Zero-Knowledge Proofs work without using maths.

[32] Dekking, F. M., Kraaikamp, C., Lopuhaä, H. P., & Meester, L. E. (2005). A Modern Introduction to Probability and Statistics: Understanding why and how (488). London: Springer.

Non-interactive ZK Proofs

Interactive proving has limited usefulness considering involved parties are required to be available and interact repeatedly. Additionally, if the verifier were convinced of a prover's honesty, the proof would be unavailable for independent verification. Therefore, the overall use case is limited.

Non-interactive ZK proofs require the prover to perform and complete a set of challenges provided by a simulated verifier based on the prover's commitments. The main advantage of non-interactive ZK proofs is that the system is automated, and the prover's claims get verified through a set of challenges instead of another person. Therefore, non-interactive ZK proofs are generally more efficient for cryptocurrency applications as part of a transactional layer. The users can complete transactions without direct interaction between involved parties. A popular non-interactive proof is the Zero-Knowledge Succinct Non-Interactive Argument of Knowledge (ZK-SNARK).

One of the significant drawbacks of the ZK-technology is that it is relatively computational heavy to generate ZK-proofs, making near-instant withdrawals impossible. Additionally, this also means that expensive machines are needed to create ZK-proofs. Furthermore, the developments in quantum computing could lead to problems with ZK rollups, primarily if the proof does not utilise collision-resistant hashes. Collision Resistant Hashing (CRH) makes it computationally infeasible to find two inputs that map to the same output. ZK-starks use collision-resistant hashes; therefore, quantum computing may not represent a threat to this type of proof. However, ZK-SNARKs are not collisions-resistant, which is an issue considering the popularity of ZK-SNARKS.

Optimistic Rollups

Optimistic (OP) rollups are relatively similar in design to ZK-rollups. The significant difference between OP and ZK roll ups is the usage of different cryptographic proofs. ZK-rollups utilise ZK-Proofs, while OP-rollups utilise fraud proofs which are similiar to Plasmas' fraud proofs. Therefore, optimistic rollups can be considered a mixture of ZK-rollups and Plasma. Unlike Plasma, OP rollups do post transactional data and state roots on the main chain, thus utilising the main chain as a data layer.

The fraud proofs are significantly less resource-demanding than ZK-proofs, solving one of the significant drawbacks of ZK-rollups. The OP rollup design utilises multiple aggregators (i.e. block producers), which have to deposit a bond when posting a new rollup block (i.e. new state root).

[33] Blum, M., Feldman, P., & Micali, S. (2019). Non-interactive zero-knowledge and its applications. Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali, 329-349.

[34] Rothblum, R. D., & Vasudevan, P. N. (2022). Collision-Resistance from Multi-Collision-Resistance. Cryptology ePrint Archive.

[35] Panther. Retrieved September 6, 2022, from: https://blog.pantherprotocol.io/zk-snarks-vs-zk-starks-differences-in-zero-knowledge-technologies/.

Whenever a user identifies an invalid state root (i.e. a created block including invalid transactions), the user can challenge the aggregator by posting the valid state root with required Merkle proofs to prove it. If the user is correct, the bond and the created block of the aggregator get slashed. A Merkle proof confirms a set of transactions by a branch hash within a Merkle Hash root. Hence, a Merle proof reflects the combined hashes of a block of data. To incentivise users to challenge the aggregators, a part of the slashed bond will be given as a reward to the user, and the remaining part will be burned by the Ethereum Virtual Machine (EVM).

Additionally, this creates censorship resistance. Considering malicious aggregators cannot censor or ignore a single transaction, assuming the mechanism does work efficiently (i.e. users challenge aggregators efficiently). No aggregator will publish a new block on top of an invalid block, which could lead to slashing, and therefore, a loss of capital. Thus, the assumption can be made that every aggregator will verify the validity of the previous block to ensure the block is valid.

The fraud-proof mechanism has a significant drawback: the relatively long challenge period; within this period, the user funds are stuck in the smart contract until the period is over. Therefore, the user's funds are stuck for approximately a week until the transaction becomes final. The throughput of OP rollups is lower than ZK rollups because transactions are not bundled, meaning that all transactions need to include the sender's signature. Therefore, the overall byte size of the transaction significantly increases from 13 bytes to approximately 75 bytes. Consequently, the gas costs increase, leading to a lower TPS due to finite block size. The current expected TPS of OP-rollups is approximately 700 TPS on Ethereum. However, according to Vitalik, OP-roll can utilise aggregate BLS signatures for replacing the current ESCDA signatures This would significantly decrease byte usage, meaning more transactions can be put in a block. BLS signatures work similar to ZK-proofs and would replace all individual signatures with one BLS signature that validates all involved transactions. The expected fixed gas fee for a BLS signature is approximately 113,000 bytes. However, the overall byte size of a transaction will decrease back to 13 bytes; consequently, the OP rollup TPS increases to approximately 3.000 on Ethereum.

Figure 5 Layer 2 solutions overview

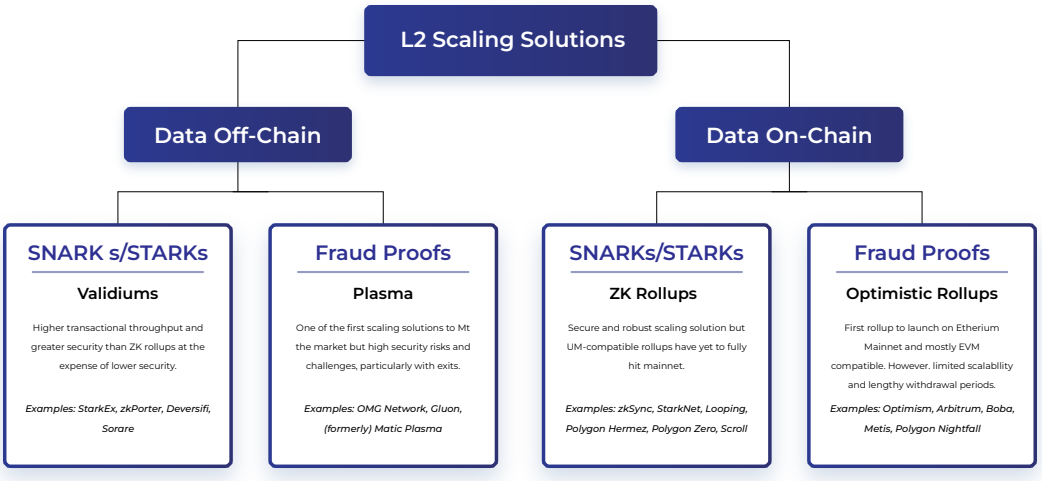

Source: https://www.ambergroup.io

[36] EthHub. Retrieved September 6, 2022, from: https://docs.ethhub.io/ethereum-roadmap/layer-2-scaling/optimistic_rollups/.

[37] Medium. Retrieved September 6, 2022, from: https://medium.com/privacy-scaling-explorations/bls-wallet-bundling-up-data-fb5424d3bdd3.

Summary

There are a dozen of methods to transform a monolithic chain into a modular framework. On-chain scalability methods such as Danksharding are able to transform a monolithic chain. An alternative is Off-chain scalability methods that process transactions outside the main chain, therefore maintaining the state of the main chain while applying the last state that has been processed on the other chain (i.e. layer 2). There is a broad set of off-chain scalability methods, each with different characteristics.

Plasma is a framework of child chains that have an independent consensus mechanism and produce their own blocks. A Child-chain has a parent-child structure whereby the main chain (e.g. Ethereum) is the parent, all transactional functions of the blockchain are performed in the child chains, and the result is posted to the parent chain. Plasma contracts enable bridging assets from the parent chain to the child chains, similar to sidechain functions. However, Plasma relies on and benefits from the parent chain's security. In contrast, a sidechain is solely responsible for its security. Plasma does utilise off-chain storage and does not post transactional data on the parent chain. Therefore, withdrawals can take several days (i.e. fraud-proof challenge period.

Sidechains are independent blockchains and have their own consensus method, ecosystem design and parameters (e.g. different block creation times), and security mechanisms. Thus, sidechains are isolated from the main chain, meaning that in case of a cryptographic break, the damage is limited to the sidechain itself. Sidechains utilise a two-way peg to move assets from the main chain to the side chain. The two-way peg (i.e. blockchain bridge) facilitates communication between blockchains, allowing for the transfer of information and, thus, the transfer of assets.

Rollups handle transactions off-chain; however, they post transactional data on the main chain. To decrease on-chain storage, rollups bundle transactions and compress them before sending them to the main chain. Therefore, the amount of data posted on the main chain is significantly reduced, increasing the overall network resource efficiency and scalability. Rollups rely on the main chain's security, considering the smart contracts are deployed on the main chain. Users can interact with the rollups through the rollup smart contract, which lives on the main chain. There are currently two primary types of rollups: Optimistic and Zero Knowledge Roll-ups.

Zero-knowledge (ZK) rollups are similar to Plasma; however, despite the similarities, the design of ZK rollups is relatively complex and has a few significant differences. Plasma utilises fraud proofs to ensure security on the child chain; ZK Rollups utilise ZK-proofs, which are used as cryptographic proof (i.e. ZK-Proof) for a bundle of transactions that are processed off-chain. The ZK-proof consists of a compressed transactional data summary; therefore, the overall stored amount of data is minimal compared to typical layer 1 transactional data flows. Despite the compression, the amount of information is sufficient to confidently validate the data without revealing the data.

Optimistic (OP) rollups are a mixture of ZK-rollups and Plasma. Unlike Plasma, OP rollups do post transactional data and state roots on the main chain, thus utilising the main chain as a data layer. OP Roll-ups utilize Fraud proofs, instead of ZK-proofs. The OP rollup design uses multiple aggregators, which have to deposit a bond when posting a new rollup block. Whenever a user identifies an invalid state root, the user can challenge the aggregator by posting the valid state root with required Merkle proofs to prove it. If the user is correct, the bond and the created block of the aggregator get slashed.

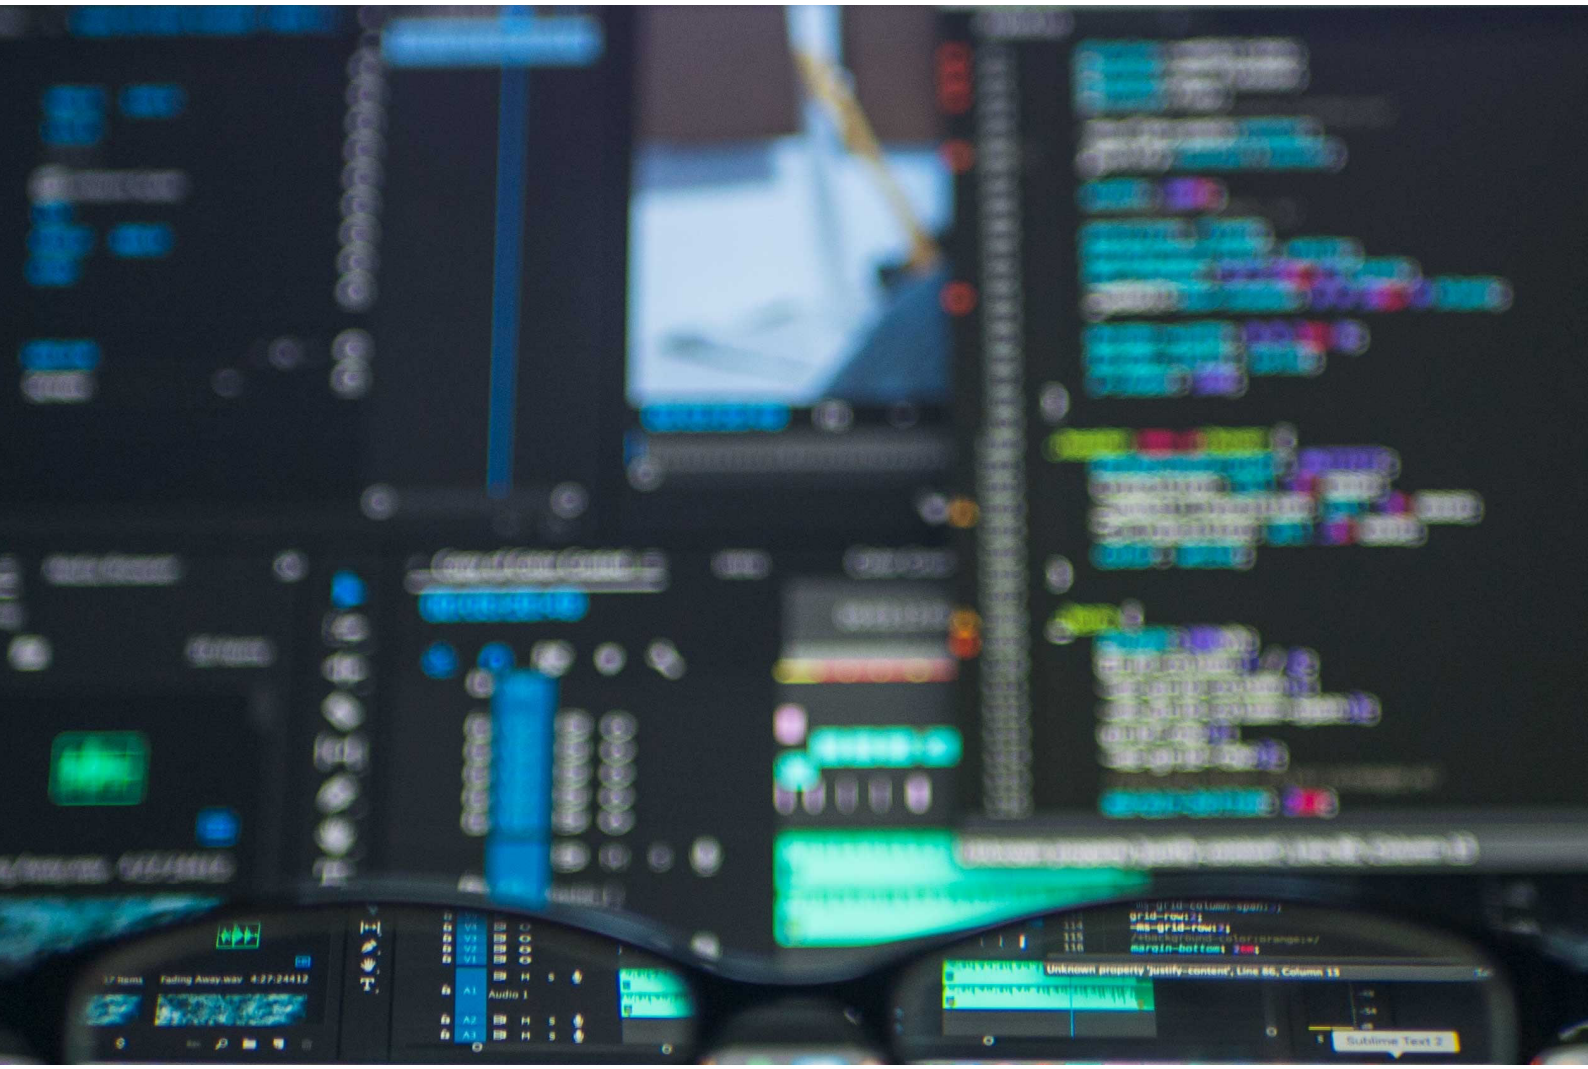

Photo by Kevin Ku on Unsplash

Data availability, the backbone of the blockchain technology

Data Availability can be performed on and off-chain; however, what is considered more efficient?

Data Availability is the theoretical guarantee that a block proposer publishes all transaction data for a block and makes that data available to other network participants. A block consists of two major parts:

- The block header, this portion of the block, contains information about the block itself (e.g. timestamp, cryptographic nonce).
- The block body represents all the transactional data.

When a block proposer proposes a new block, it must publish the whole block, including the block body. Consensus has to be formed, therefore, other network participants will download the block's data to verify and validate the data. If this was not the case, any block proposer could insert malicious transactions in blocks without any repercussions. The data availability problem is relevant for scalability methods due to the obligation on the main chain to allow anyone to validate the computations performed off the main chain.

On-chain data availability is the most common approach, whereby block proposers are forced to publish all block data and have other nodes validating it. On-chain data availability is very apparent in monolithic blockchain platforms, whereby all core tasks are performed on one layer (i.e. consensus, executing and settlement, data availability). However, a few scalability methods also utilise on-chain data availability, meaning that layer 2 does post transactional data on the layer 1.

In contrast, a particular set of scalability methods uses off-chain data availability; commonly, the data availability layer is offloaded by another layer. Therefore, block producers do not publish transactional data to the main chain but instead provide a cryptographic commitment to prove the availability of the data. However, solely a cryptographic proof is not sufficient to guarantee data availability. Malicious actors can still opt not to submit all the transactional data, although the proof states that all taken state transitions are valid. Consequently, users remain uninformed about their balances and cannot perform state updates using the information in newly added blocks. The main advantage of off-chain storage is the decrease of on-chain posted data, and therefore, increasing network efficiency and scalability. Despite the benefits of storing information off-chain, it creates a complex challenge for data availability. However, there are proposed solutions to this long-standing problem, which will be highlighted later in this paper.

The data availability problem is relevant for scalability methods due to the obligation on the main chain to allow anyone to validate the computations performed off the main chain.

[38] Ethereum. Retrieved September 8, 2022, from: https://ethereum.org/en/developers/docs/data-availability/.

[39] Ibid

How do scalability methods handle data availability?

ZK- Rollups

The regular ZK rollup uses the main chain as a data layer, meaning that the rollup posts data on the main chain. On-chain data availability for Layer 2s is the highest security approach, considering any data posted on the main chain is immutable. However, the drawback of on-chain data availability is that it is impossible to avoid the call data gas cost. The data availability layer can become expensive because a block has finite space. This means that the overall transaction costs increase if the demand increases, considering the supply of transactions is finite, and assuming that gas costs remain the same. To avoid this problem, the data availability layer can be moved to off-chain; however, this leads to complex challenges and can lead to complex trust issues, often involving a centralised party.

Pure Validums

Validiums are scaling solutions that utilise off-chain availability and computation. Therefore, the data availability and execution and settlement layers are shifted off-chain. Validiums use validity proofs that are similar to ZK-proofs (i.e. ZK-SNARKs). This prevents invalid transitions and improves the security guarantees of the validium chain. Funds belonging to validium users are controlled by a smart contract on the main chain; therefore, the design is comparable with a ZK-rollup.

Consequently, near-instant withdrawals are possible as long the validity proof for the withdrawal request has been verified on the main chain. However, the main difference between ZK-rollups and Validiums is the way data is stored. Thus, putting the scalability solutions respectively in a different position on the data availability spectrum.

Pure Validiums utilise an off-chain data availability model, whereby transactional data is stored off-chain. Thus, no transactional data will be posted to the main chain. The storage of data off-chain leads to a significant centralisation issue. Users cannot withdraw funds from the on-chain smart contract if an operator acts maliciously by refusing to provide the necessary data to generate a Merkle proof. A Merkle proof validates the existence of the user's withdrawal transaction in a verified transaction batch. Thus, a Merkle proof is needed for a user to enable the on-chain smart contract to process a withdrawal.

Plasma

Plasma chains' users rely on the operator to provide block data if they need to create fraud proofs challenging invalid transactions. This system works as long the operator is honest and processes fraud proofs. However, if an operator has malicious intent, adding invalid transactions (e.g. hijacking funds) to the block can create a mass user exit of the child chain. Assuming users know about the dishonest operator, they exit the child chain as soon as possible to protect their funds. Users withdraw their funds from the child chain back to the parent chain; however, this is not as effortless as it might seem. Mass exits are caused by the lack of data availability on a Plasma chain, leading to complex issues.

---

[40] Ethereum. Retrieved September 8, 2022, from: https://ethereum.org/en/developers/docs/scaling/validium/.

[41] Ethereum. Retrieved September 8, 2022, from: https://ethereum.org/en/developers/docs/scaling/validium/#:~:text=Validium%20is%20a%20scaling%20solution,data%20on%20the%20Ethereum%20Mainnet.

A successful exit of the child chain can solely happen after the challenge period (e.g. a week). In that case, the users must prove that the operator has malicious intent. The user can provide proof by providing the last valid state of the child chain; however, this implies that the user has to acquire the whole child chain, essentially turning into a full node. This is unrealistic for most users, considering the required technical overhead and equipment to set up a full node. A side effect is that the entire valid state of the plasma chain (i.e. child chain) has to be posted on the parent chain, leading to significant congestion on the parent chain.

Summary

Data Availability is the theoretical guarantee that a block proposer publishes all transaction data for a block and makes that data available to other network participants. On-chain data availability is the most common approach, whereby block proposers are forced to publish all transactional data (i.e. block body) and have other nodes validating it. On-chain data availability is very apparent in monolithic blockchain platforms, whereby all core tasks are performed on one layer. However, scalability methods also utilise on-chain data availability, meaning that layer 2 does post transactional data on the layer 1. A particular set of scalability methods uses off-chain data availability; commonly, the data availability layer is offloaded by another layer. Therefore, block producers do not publish transactional data to the main chain but instead provide a cryptographic commitment to prove the availability of the data.

The regular ZK rollup uses the main chain as a data layer, meaning that the rollup posts data on the main chain. On-chain data availability for Layer 2s is the highest security approach, considering any data posted on the main chain is immutable. However, the drawback of on-chain data availability is that it is impossible to avoid the call data gas cost.

Validiums are scaling solutions that utilise off-chain availability and computation. Therefore, the data availability and execution and settlement layers are shifted off-chain. Validiums use validity proofs that are similar to ZK-proofs (i.e. ZK-SNARKs). This prevents invalid transitions and improves the security guarantees of the validium chain. Pure validiums utilise an off-chain data availability model, whereby transactional data is stored off-chain. Thus, no transactional data will be posted to the main chain. The storage of data off-chain leads to a significant centralisation issue. Users cannot withdraw funds from the on-chain smart contract if an operator acts maliciously by refusing to provide the necessary data to generate a Merkle proof.

Plasma chains' users rely on the operator to provide block data if they need to create fraud proofs challenging invalid transactions. This system works as long as the operator is honest and processes fraud proofs. A successful exit of the child chain can solely happen after the challenge period (e.g. a week). However, if an operator has malicious intent, adding invalid transactions (e.g. hijacking funds) to the block can create a mass user exit of the child chain. In that case, the users must prove that the operator has malicious intent. However, this is a relatively technical process which arguably cannot be done by the regular user.

Photo by Christopher Gower on Unsplash

Off-chain data availability has a few significant issues, but there are solutions available

Are these off-chain data availability problems realistic? Literature has shown that there multiple solutions available.

As blockchains get longer over time, the blockchain gets more resource intensive in terms of required storage and resources to participate (i.e. High-speed internet connection). However, it is possible to store the data off-chain, but as highlighted earlier in this research, this could lead to significant trust issues and consequently, data availability issues. Hence, to solve the problem of on-chain data availability being overly resource-inefficient, there are methods to verify high volumes of data without requiring a single participant to download the entire blockchain.

Randomly Sampled Committees (RAC) is a method that randomly splits a block into data chunks (i.e. blobs), and a random set of validators verifies these blobs. Each group of validators (i.e. committee) generates a signature attesting that they have validated the blob. Consequently, the network solely accepts the blob if there are signatures from the majority of the corresponding committee. It is nearly impossible for a malicious actor to disrupt the committee, considering the set of validators is randomly chosen. Data Availability Sampling adds another security layer on top of the committee and enhances data availability.

Data Availability Sampling is nearly identical to randomly sample committees; however, there is a subtle but significant difference. Instead of using a committee assigned to a blob, every node verifies a part of every blob instead of downloading one whole blob. This means that every node does participate in the validation of every blob, but instead of verifying the whole blob, they verify only a part of it. However, this method is not perfect, considering there is still a chance a malicious node that can hide a tiny amount of data, considering DAS only proves that most of the data is available, at least 50.1%.

Data Availability Proofs

A malicious block producer could prevent full nodes from generating proofs by withholding the required data and only releasing the block header to the network (i.e. data withholding attack). Additionally, the block producer posts the data long after the block has been published, making the block invalid. This would trigger a rollback of transactions, considering all blocks made on top of the invalid block are also determined invalid. To solve this issue, erasure coding is utilised.

Erasure coding enables any network participant to reconstruct and republish a blob with at least 50.1% of the data. Considering DAS does prove that at least 50.1% of the data is available, erasure coding provides a solution to the remaining issue whereby malicious nodes can hide a tiny amount of data.A simple mathematical analogy to put erasure coding in perspective is one that can recover a straight line with any given two distinct points. In the case of more complex polynomials, one would need half of the total data points to determine the missing data points (figure 5). Therefore, any blob can be evaluated, which improves the data availability.

In the case of a malicious actor, the node has to withhold at least 50.1% of the block to perform a data withholding attack, the data variant of selfish mining. However, in that case, the block would have been invalid in the first place, considering DAS determines blocks invalid if it cannot prove that at least 50.1% of the data is available.

[42] HackMD. Retrieved September 10, 2022, from: https://hackmd.io/@vbuterin/sharding_proposal#ELI5-data-availability-sampling

[43] Al-Bassam, M., Sonnino, A., & Buterin, V. (2018). Fraud and data availability proofs: Maximising light client security and scaling blockchains with dishonest majorities.

https://ethereum.org/en/developers/docs/scaling/validium/#:~:text=Validium%20is%20a%20scaling%20solution,data%20on%20the%20Ethereum%20Mainnet.

[44] Ibid.

In the case of a malicious actor, the node has to withhold at least 50.1% of the block to perform a data withholding attack, the data variant of selfish mining. However, in that case, the block would have been invalid in the first place, considering DAS determines blocks invalid if it cannot prove that at least 50.1% of the data is available.

Figure 6 Erasure coding overview

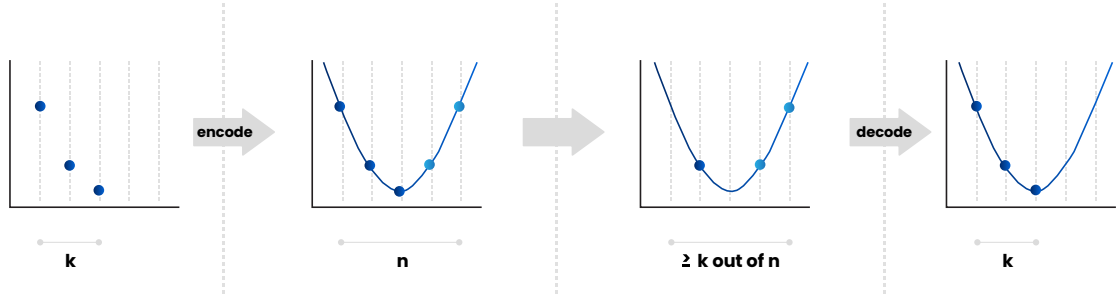

Data Availability Committee

Pure Validiums store all transaction data off-chain, meaning only the block producer can access it. This leads to complex challenges whereby the block producer can publish invalid transactions. Considering the block producer is not posting any transaction data on the main chain, the true state will be concealed. This issue is partly solved by a Data Availability Committee (DAC).

A DAC is a set of trusted parties that will store copies of off-chain data. The DAC will store the data off-chain but is required to make the data available in case there is a dispute. Additionally, in an emergency, the ASC will no longer accept new state updates and solely accept users to withdraw funds by providing a Merkle proof for the latest state. Users can access the Merkle path to their accounts through the publicised data and use that to retrieve the Application Smart Contract (ASC) funds.

The issue with a DAC is that it does rely on a particular set of trusted parties; therefore, the system is relatively centralised. This could be solved by adding nodes (i.e. trusted parties), but this would make the overall system less resource efficient.

To solve the problem of on-chain data availability being overly resource-inefficient, there are methods to verify high volumes of data without requiring a single participant to download the entire blockchain.

Summary

As blockchains get longer over time, the blockchain gets more resource intensive in terms of required storage and resources to participate. However, it is possible to store the data off-chain; this could lead to significant trust issues and, consequently, data availability issues. Hence, to solve the problem of on-chain data availability being overly resource-inefficient, there are methods to verify high volumes of data without requiring a single participant to download the entire blockchain.

Randomly Sampled Committees (RAC) is a method that randomly splits a block into blobs, and a random set of validators verifies these blobs. Each committee generates a signature attesting that they have validated the blob. Consequently, the network solely accepts the blob if there are signatures from the majority of the corresponding committee. Data Availability Sampling adds another security layer on top of the committee and enhances data availability.

Data Availability Sampling is nearly identical to randomly sampled committees; however, there is a subtle but significant difference. Instead of using a committee assigned to a blob, every node verifies a part of every blob instead of downloading one whole blob. However, this method is not perfect, considering there is still a chance a malicious node that can hide a tiny amount of data, considering DAS only proves that most of the data is available, at least 50.1%. To solve this issue, erasure coding is utilised.

Erasure coding enables any network participant to reconstruct and republish a blob with at least 50.1% of the data. Considering DAS does prove that at least 50.1% of the data is available, erasure coding provides a solution to the remaining issue whereby malicious nodes can hide a tiny amount of data.

Pure Validiums store all transaction data off-chain, meaning that only the block producer can access it. This leads to complex challenges whereby the block producer can publish invalid transactions. Considering the block producer is not posting any transaction data on the main chain, the true state will be concealed. This issue is partly solved by a Data Availability Committee (DAC).

A DAC is a set of trusted parties that will store copies of off-chain data. The DAC will store the data off-chain but is required to make the data available in case there is a dispute. Additionally, in an emergency, the ASC will no longer accept new state updates and solely accept users to withdraw funds by providing a Merkle proof for the latest state. Users can access the Merkle path to their accounts through the publicised data and use that to retrieve the Application Smart Contract (ASC) funds. The issue with a DAC is that it relies on a particular set of trusted parties; therefore, the system is relatively centralised. This could be solved by adding nodes (i.e. trusted parties), but this would make the overall system less resource efficient.

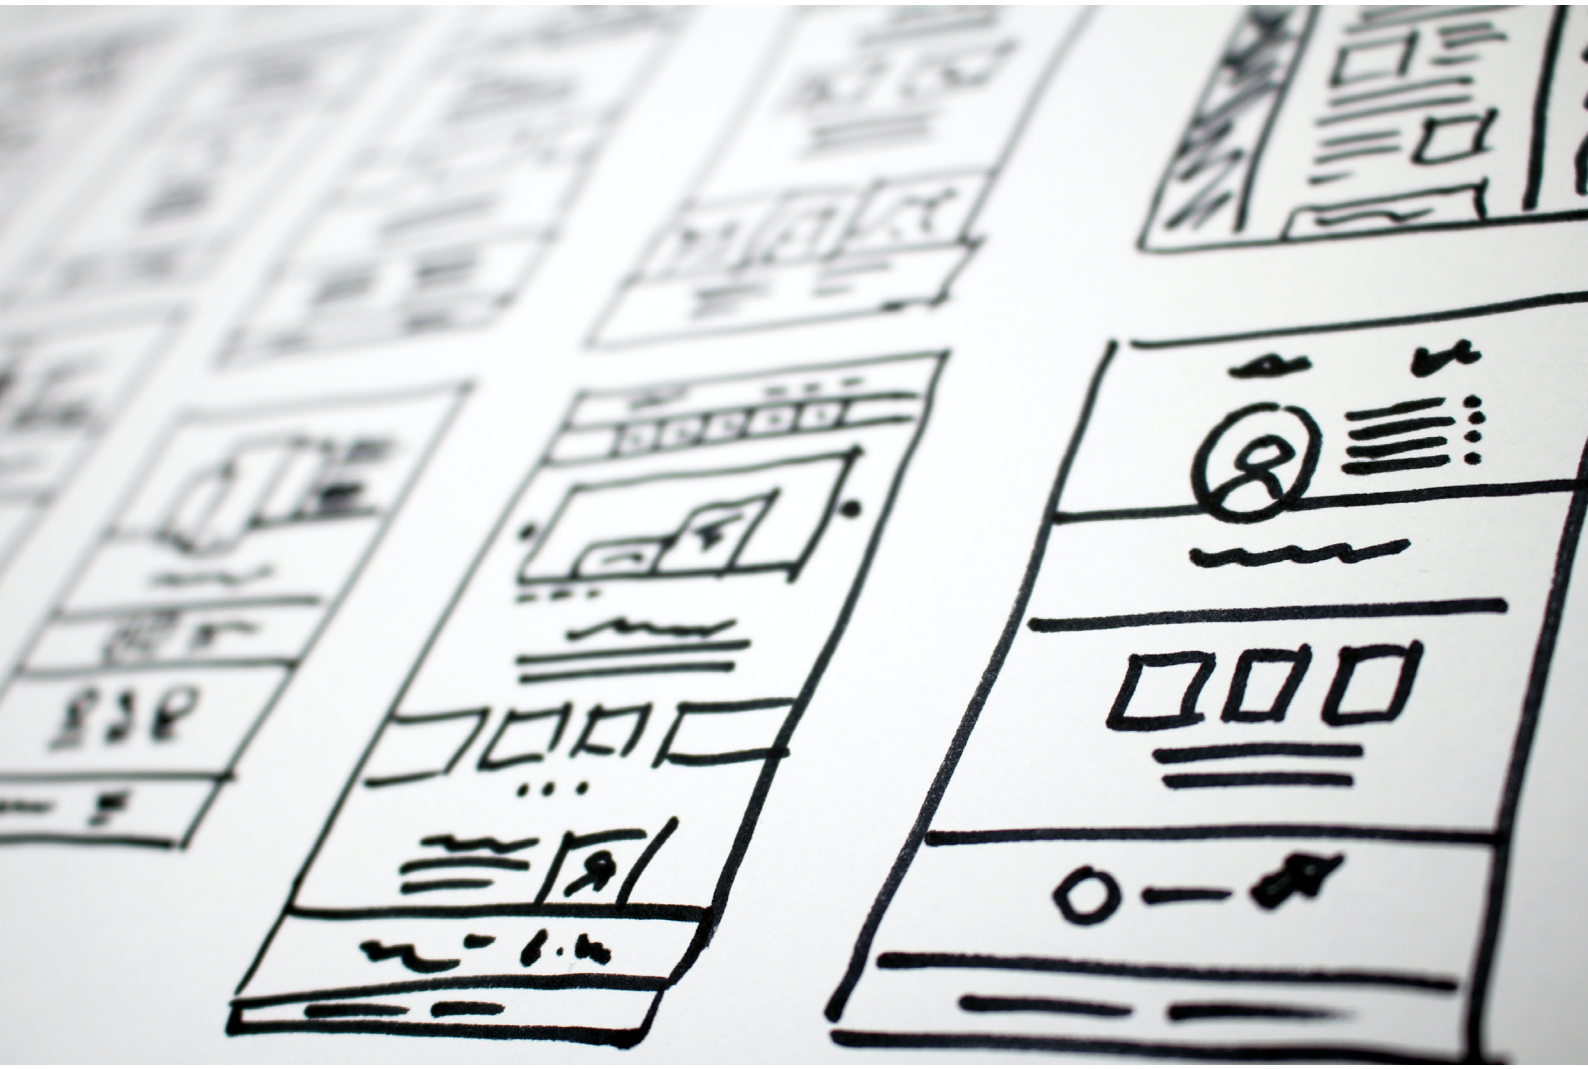

Photo by Hal Gatewood on Unsplash

Syscoin, among the first blockchain projects to implement roll-ups

Syscoin has a unique ecosystem design with merged mining, Z-DAG technology, and the to-be-implemented roll-ups.

Syscoin is a public decentralised high-performing blockchain network. The ecosystem solves the blockchain trilemma (i.e. the challenge of developing a secure, decentralised, and fast blockchain ecosystem) by having a four-layer tech stack. A tech stack is the collection of technologies an organisation uses to build an application. Syscoin is being used as the host layer, with Bitcoin as a consensus method providing an efficient foundation. Syscoin utilises Bitcoin as a consensus method by utilising the merged mining mechanism. Merged mining allows Syscoin to enhance the security of its network, applying Proof of Work (PoW) by recycling the Bitcoin network's energy. Miners can mine two or more cryptocurrencies simultaneously when they merge-mine without sacrificing any mining performance. Therefore, the miners can use computational power to mine blocks on multiple chains with the same algorithm (i.e. SHA-256 for Bitcoin).

On top of that, an EVM layer is being used as the operating layer, as Ethereum is widely adopted. An Ethereum virtual machine (EVM) is essentially a machine that mimics a physical computer. The computer is run by all the full nodes of the Ethereum network. The Ethereum Network is decentralised, meaning all the full nodes have to agree (i.e. come to a consensus) on how EVM behaves and how computations are made. Therefore, full nodes individually copy and verify transactions on the blockchain. Due to the individual calculations, every individual full node has its computation which in the best case is the same as all the other individual computations of every full node.

The third layer is a software development kit layer (SDK) that will allow Zero-Knowledge Roll-ups to increase the network's overall scalability. The last layer consists of an application layer. The application layer is an abstract layer that hides all the complex computations and overall technical details and serves as a general user interface for the network. Therefore, the application layer hides the system's operations to enhance user experience (UX). For example, a decentralised application (dApp) runs on the application layer with an intuitive user interface design, and therefore, consumers will not notice the underlying tech.

Figure 7 Syscoin Four-layered tech stack

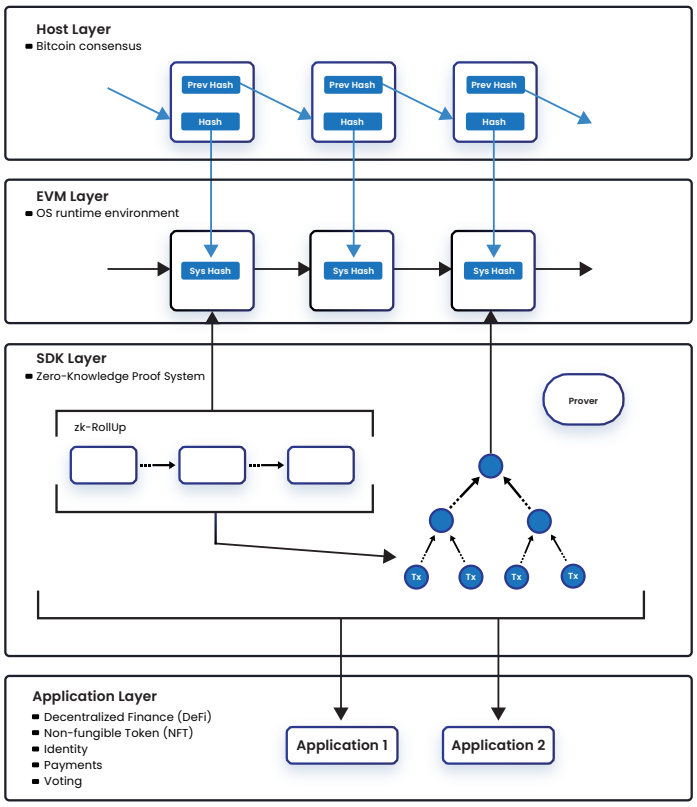

UTXO Model

Syscoin utilises the Transaction Output (UTXO) model of Bitcoin. The UTXO accounting model works similarly to cash. Whenever a user receives or spends Bitcoin, the transaction is recorded as a UTXO. Therefore, a user's wallet represents the net of all inputs and outputs of the combined Bitcoin UTXOs. The amount left is the amount that is 'unspent.' A bitcoin transaction has both an input and an output. The input is the address where the bitcoin is being sent from, and the output is the address where it is sent to. If an output has been spent, it is impossible to spend it again. However, a UTXO can be used or spent as an input in another transaction. A simplified example is required to put this abstract concept into perspective (Figure 8).

Syscoin and the Syscoin Platform Tokens (SPT) utilise the UTXO model; therefore, the Syscoin asset model is built on top of the Bitcoin UTXO model. If there are significant innovative breakthroughs in the UTXO model, Syscoin will benefit and capitalise on these innovations. The latest innovation on the Bitcoin network is the Taproot integration. Taproot increases transaction efficiency, privacy, and the potential for smart contracts that can be used to eliminate intermediaries.

What happens if you spend Bitcoin?

Assume you would like to send 5 Bitcoin over to a friend but you have an input UTXO of 12 BTC. You cannot simply spend these 5 BTC; and you have to spend the entire 12 BTC. Naturally, that is not something you would like to do as you would like to send only 5 BTC. What happens when you send 5 BTC to your friend is that two UTXOs are generated. The first UTXO of 5 BTC will be sent to your friend and the second UTXO is the difference between the input UTXO and the output UTXO, which is 7 BTC.

Figure 8 UTXO Model overview

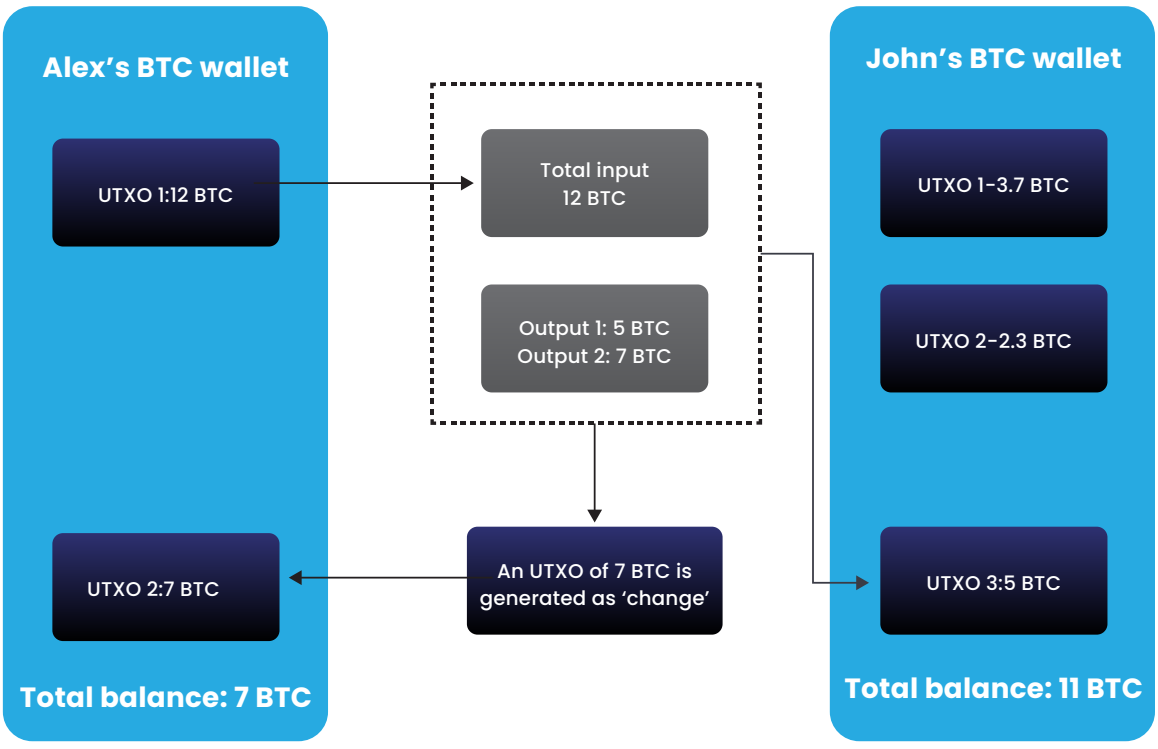

Scaling a blockchain without making significant trade-offs

Syscoin aims to solve the blockchain trilemma by utilising Bitcoin's decentralisation and security, the flexibility of Ethereum, and the scalability of rollups. Syscoin presents an exceptional value proposition whereby it uses the strengths of other projects to create synergy and network effects for holders and developers. The ecosystem is full EVM compatible through Syscoin's Network Enhanced Virtual Machine, allowing projects to migrate to Ethereum easily. The ecosystem is fully secure through merge mining, whereby currently, 30% of the total hash rate is merge mining Syscoin.

Syscoin currently utilises the Zero Confirmation Directed Acyclic Graph as a scalability method. Z-DAG allows Syscoin to scale with a theoretical exceptionally high TPS (I.e. 60,000 to 145,000) while retaining a high level of security. However, that is the maximum TPS (i.e. TPS burst) meaning the network cannot sustainably hold these levels for an extended period. Therefore, it is safer to assume that the Z-DAG technology will allow for 8,000 to 15,000 TPS. Z-DAG transactions can instantly settle as a result of the network being able to anticipate what transactions will be in the next block and how these will be ordered with confidence. Therefore, participants of the network can instantly settle payments without having to wait for block confirmations.

The parent chain miners order a list of transactions out of their memory pool which they sort by time; this is done after they have made a list based on transaction fees (e.g. supply and demand). To ensure there are no negative balances, a strict validation process by every peer is used when processing a block, comparable with the Bitcoin balance checks (i.e. UTXOs). The conflicting transactions and blocks will be rejected if the network cannot get an absolute consensus. The system utilises a 10-second delay to subsequent transfers made by the same asset holder to prevent double-spending and reduce the latency of ordering transactions over time. As stated before, the Z-DAG layer falls back on the merged Bitcoin layer; therefore, if there are any differences between the PoW block and the real-time state, they will be resolved upon the confirmation of a block, as Block $_{t}$ depends onBlock $_{t-1}$. The Z-DAG layer allows Syscoin to reduce transaction settlement times to near-instant while having a slight risk of double-spending over a minimum latency time.

Rollux

The ecosystem launched its rollup suite, Rollux, earlier this year. Rollux is a complete layer 2 solution suite that initially employs OP Rollups, followed by ZK-rollups once they realise maturity.The Rollups will utilise Syscoin as a gas token; however, considering Syscoin is a non-profit foundation, the overall user costs are allegedly lower than competitors. The Syscoin network will utilise a 64.000.000 bytes block, with two blobs of 32.000.000 bytes. As stated earlier, OP-rollups have a higher byte usage per transaction than ZK-rollups because all transactions need to include the sender's signature. The expected TPS for OP rollups is roughly 2000, according to Jagdeep Sidhu, core developer of Syscoin. However, it is expected to increase over time if transactional data can be stored off-chain. The main benefit of OP-rollups, compared to Z-DAG, is that they will allow for EVM executions.

---

[45] Syscoin. Retrieved September 18, 2022, from: https://syscoin.readme.io/docs/what-is-z-dag.

[46] Syscoin. Retrieved September 18, 2022, from: https://syscoin.org/news/introducing-rollux-syscoins-rollup-suite-ready-to-take-market-by-storm.

Proof of Data Availability, an intuitive way to provide data availability

Data availability is a Big Data problem which scales with the number of transactions. The main issue for most scalability methods is to guarantee data availability while preventing a significant data overhead. To put this in perspective, rollups are expected to add 2 megabytes (MB) a second of dedicated data space to the Ethereum main chain. Consequently, this will lead to approximately 63,000,000 MB to the mainchain. The enormous data overhead could lead to centralisation whereby individual users cannot afford to run nodes. However, off-chain data availability can significantly decrease the data on the main chain.

Off-chain data availability is still relatively complex, often resulting in complex trust and centralisation issues. Therefore, off-chain data availability used by pure Validiums is considered a suboptimal scalability method for solving the blockchain trilemma. The trade-off of high scalability but having low-grade security is, in most cases, not optimal. However, Syscoin aims to solve off-chain data availability issues by replicating and archiving the data by Proof of Data Availability (PoDA).

In the first stage, the data is replicated, meaning that any replicating node can copy data. This ensures that the data gets widespread through the network. The assumption is that if there are enough nodes, there is at least one honest node (i.e. Honest Minority Assumption). In that case, the data will always remain accessible, considering the honest node will be able to provide it. The main advantage of having an Honest Minority Assumption is that it has significantly lower overhead costs than an Honest Majority Assumption. Smart contracts require an Honest Majority Assumption, whereby at least 51% of the network has to form a consensus. By archiving the data instead, there is no need for consensus or spam protection. Therefore, the costs involved are significantly lower.

The replicated data is archived after a short period, meaning that the data is removed from the main chain and stored off-chain instead. By design, assuming Honest Minority Assumption, the data remains available even after it was archived. Considering network participants (i.e. archiving nodes) can request the data of the replicating nodes. Hence, the replication period must be long enough for archiving nodes to acquire the data to ensure decentralisation and increase the probability of the Honest Minority Assumption. Considering it is safe to assume that the digital infrastructure (i.e. internet connections) throughout the world differs, the replication period should be long enough to allow for participation by any node in any geographical location. Syscoin has increased its block time to solve this issue, reducing the replication period from 60 seconds to 150 seconds.

---

[47] Ethereum. Retrieved September 20, 2022, from: https://notes.ethereum.org/@vbuterin/data_sharding_roadmap#Who-would-store-historical-data-under-sharding.

[48] House of Chimera. Retrieved September 23, 2022, from: https://notes.ethereum.org/@vbuterin/data_sharding_roadmap#Who-would-store-historical-data-under-sharding.

Why is Proof of Data Availability needed?

ZK-rollups do have a two-step process to transition state on the main chain. Firstly, an aggregator provides validity proofs to the smart contract living on the main chain. These validity proofs consist of batched transactional data, where network participants (e.g. verifiers) can verify if all included transactions are valid. In the second stage, the aggregator calls the verifying contract again to enforce network validation of data liveness and correctness. The data posted on the main chain will significantly increase over time. The amount of data that is posted depends on the scalability method. OP rollups utilise more data per transaction than ZK-rollups, meaning that, generally, the finite storage of a block will be filled less efficiently (i.e. fewer transactions). Therefore, it would be ideal to have a data availability layer whereby each rollup archived its data off-chain while the guarantee of data retrievability remains valid. This means that a user can retrieve historical information from the blockchain; the importance of having access to historical data is to remain the whole blockchain intact. However, this assumes rollups in a shared cost model are equally efficient. To put this in perspective, an example will be utilised.

Let's assume that there is rollup X and Rollup Y, whereby Rollup X has a TPS of 500 while Rollup Y has a TPS of 1000. In a shared cost model, the involved costs will be equally split between these two rollups. The main issue here is that even costs are not split by effort and do not consider that Roll up X has a significantly lower real overhead cost considering its TPS is lower. The more transactions have to be stored and handled, the more overhead cost incur (e.g. electricity, hardware maintenance). Hence, the shared costs involved in PoDA depend on the throughput of the rollup. Consequently, the shared costs of Rollup X will be lower than Rollup Y, considering the overall throughput of Rollup X is lower than Rollup Y.

By offloading the transactional and settlement layer to a rollup (I.e. layer 2), the overall transaction cost of transferring assets significantly decreases. Additionally, considering the asset transactions are handled off-chain, the overall costs of interacting with smart contracts to interact with DeFi products on the main chain (i.e. layer 1) decreases. Due to that, the two gas markets, layer 1 and layer 2, are relatively independent of each other. The assumption is that layer 2 is significantly cheaper to transfer assets. Hence users will always prefer the layer 2 above the layer 1, ceteris paribus (i.e. all things equal). Consequently, the overall posted data on the main chain decreases, considering most of the transactions are handled off-chain, which eventually will lead to lower involved costs to interact with any smart contract living on the main chain without making any trade-offs in terms of decentralisation or security.

Proof of Data Availability is in theory more secure and resource efficient as Data Sampling. Data Sampling might be able to create more throughput, however, PoDA is significantly more secure considering data is more widespread through the whole network. Therefore, making it censorship resistant and making data withholding attack is much less probable and complex to pull-off. Additionally, it significantly lowers the chances of a targeted DDOS attack. Considering if data is not sufficiently being spread through the network, nodes can be targeted to disrupt the network. In the case of PoDA, these probabilities are fairly slim in comparison with Data Sampling.

[49] Medium. Retrieved September 25, 2022, from: https://jsidhu.medium.com/blockchain-idealisms-b61c5781ddc3.

Summary

Syscoin is a public decentralised high-performing blockchain network. Syscoin aims to solve the blockchain trilemma by utilising Bitcoin's decentralisation and security, the flexibility of Ethereum, and the scalability of rollups. Syscoin presents an exceptional value proposition whereby it uses the strengths of other projects to create synergy and network effects for holders and developers. To achieve this, Syscoin utilises a four-layer tech-stack, consistent out of a Host, EVM, SDK and application layer.

Syscoin is being used as the host layer, with Bitcoin as a consensus method providing an efficient foundation. Syscoin utilises Bitcoin as a consensus method by utilising the merged mining mechanism. Merged mining allows Syscoin to enhance the security of its network, applying Proof of Work (PoW) by recycling the Bitcoin network's energy. On top of that, an EVM layer is being used as the operating layer, as Ethereum is widely adopted. An Ethereum virtual machine (EVM) is essentially a machine that mimics a physical computer. The computer is run by all the full nodes of the Ethereum network. The Ethereum Network is decentralised, meaning all the full nodes must agree on how EVM behaves and how computations are made.

The third layer is a software development kit layer (SDK) that will allow Zero-Knowledge Roll-ups to increase the network's overall scalability. The last layer consists of an application layer. The application layer is an abstract layer that hides all the complex computations and overall technical details and serves as a general user interface for the network.

Syscoin currently utilises the Zero Confirmation Directed Acyclic Graph as a scalability method. Z-DAG allows Syscoin to scale with a theoretical exceptionally high TPS (I.e. 60,000 to 145,000) while retaining a high level of security. However, that is the maximum TPS (i.e. TPS burst), meaning the network cannot sustainably hold these levels for an extended period. Therefore, assuming that the Z-DAG technology will allow for 8000 to 15000 TPS is safer. To further increase the scalability of the network, Syscoin will implement roll-ups. The ecosystem launched its roll-up suite, Rollux, earlier this year. Rollux is a complete layer 2 solution suite that initially employs OP Rollups, followed by ZK-rollups once they realise maturity. The Rollups will utilise Syscoin as a gas token; however, considering Syscoin is a non-profit foundation, the overall user costs are allegedly lower than competitors.

Syscoin utilises its Proof of Data Availability (PoDA) to ensure Data availability. In the first stage, the data is replicated, meaning that any replicating node can copy data. This ensures that the data gets widespread through the network. The assumption is that if there are enough nodes, there is at least one honest node (i.e. Honest Minority Assumption). In that case, the data will always remain accessible, considering the honest node will be able to provide it. The replicated data is archived after a short period, meaning that the data is removed from the main chain and stored off-chain instead. By design, assuming Honest Minority Assumption, the data remains available even after it was archived. Network participants (i.e. archiving nodes) can request the data of the replicating nodes at any time. Therefore, PoDA is more secure as Data sampling the Honest Minority Assumption; data is always available and cannot be censored, making data withholding attacks complex.

September 2022

Designed by House of Chimera

HouseOfChimera.com