The Blockchain Trilemma

An Evaluation Framework

Giovanni Quattrocchi and Filippo Scaramuzza, Politecnico di Milano

Damian A. Tamburri, Eindhoven University of Technology

// We present a validated framework for the evaluation and comparison of three main third-generation blockchain categories, based on the three main nonfunctional aspects that discriminate their use for the design and orchestration of complex blockchain-oriented service applications, namely: scalability, decentralization, and security. //

©SHUTTERSTOCK.COM/YURCHANKA SIARHEI

WITH THE INCREASE of network traffic and load onto Bitcoin—which by definition is a so-called *proof-of-work* (PoW) network and is highly centralized—and Ethereum, respectively, the first and the second generations of blockchains, two main problems have arisen: scalability and interoperability,[1] namely, the interaction of multiple blockchain-oriented architectures within the same larger-scale system. To solve these two issues, several technologies have been introduced, leading onto a third generation of blockchains, more able to cope with the trend characteristics and nonfunctional requirements. We divide these such technological solutions into three main categories: 1) Layer 1 solutions, 2) Layer 2 solutions, so-called *rollups* (divided into optimistic and zero-knowledge approaches), and finally 3) side-chains. In this article, we present a framework for the evaluation and comparison of these approaches; our comparison features both a quantitative and qualitative lens of analysis rotating around the so-called blockchain *trilemma*,[2] namely, the conflicting relation binding together the three key properties that a blockchain strives to exhibit operationally. On the one hand, the trilemma is per se not a novel notion but, on the other hand, to the best of our knowledge we are the first to approach the problem over the technologies targeted by this article and at the architecture level, with the end-goal of providing an early insight into the tradeoffs entailed by the trilemma itself, and the design principles stemming from their evaluation. The trilemma leads to conclude that it is possible to have only two out of these three properties regardless of the architecture decisions made to approach all three at once. Blockchains establish trust through decentralized ledgers and consensus mechanisms, initially relying on energy-intensive PoW like Bitcoin and Ethereum, leading to scalability and centralization concerns. Innovations in blockchain architecture and consensus algorithms, such as proof-of-stake (PoS), have emerged to address these issues, exemplified by Ethereum's transition to PoS in Ethereum 2.0, enhancing efficiency and security. To improve blockchain scalability and efficiency, new solutions are being developed. This work

Digital Object Identifier 10.1109/MS.2024.3417341
Date of publication 26 June 2024; date of current version 10 October 2024.

focuses on five recently developed platforms that were elicited in an informal panel held with over 50+ blockchain architects as part of the NWO-MVI CHAIN project (https://chainresearch.eu/events/online-meeting-platform-for-responsible-innovation/).

Within the scope of the aforementioned project, a pilot study to identify with a simple card-sorting exercise[3] with over 40 practitioners of 3+ years of blockchain-oriented design experience, was enacted the result of which led us to a few select Layer 1 (L1) solutions, such as Cardano and Solana, which introduce alternative consensus algorithms and sharding to process transactions more effectively. In the scope of the same exercise, Layer 2 (L2) solutions emerged as competitive assets, with features including rollups and sidechains like Arbitrum, zkSync, and Polygon, operate on top of base layer blockchains to offload computation and enhance transaction throughput, maintaining security while increasing performance.

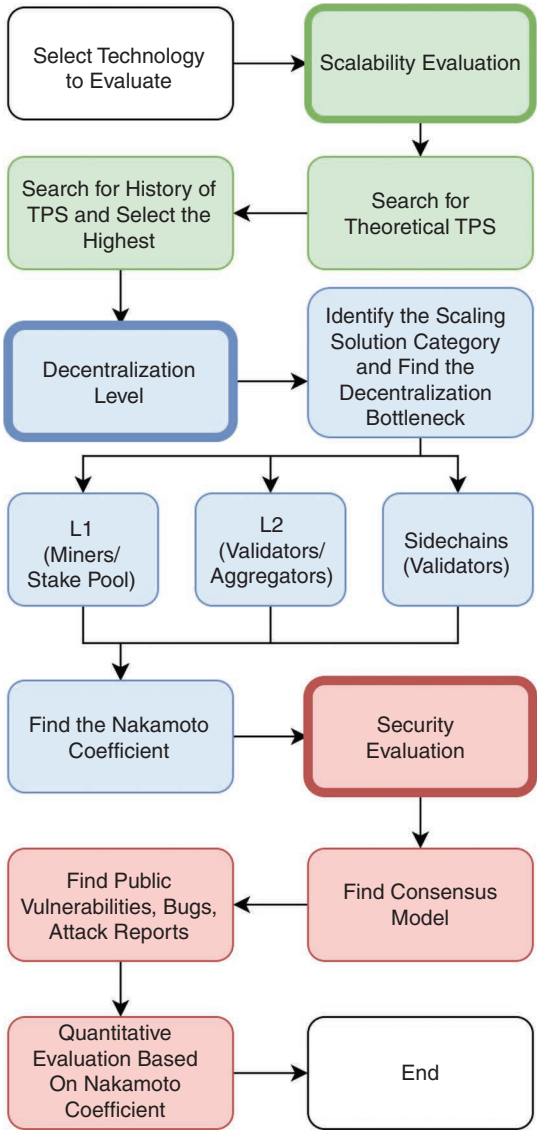

FIGURE 1. Process model diagram.

Evaluation Framework Definition

To find a way to evaluate the main properties of these scaling solutions, we took inspiration from the problem statement that shaped and inspired these solutions, i.e., the *Blockchain Trilemma*.

The scalability trilemma states that there are three properties that a blockchain tries to have and that it is possible to have only two out of these three, with an architecture tradeoff existing across all three architecture variability points. The three properties are as follows:

- *Scalability*: The chain can process more transactions than a single regular node can verify, e.g., a consumer laptop.
- *Decentralization*: The chain can run without any trust dependencies on a small group of large, centralized actors. This is typically interpreted to mean that there should not be any trust (or even honest-majority assumption) in a set of nodes that you cannot join with just a consumer laptop.
- *Security*: The chain can resist a large percentage of participating nodes trying to attack it (ideally 50%).

The intended process model is shown in Figure 1 while the replication package of our analysis is available at https://zenodo.org/records/10251476.

Scalability Evaluation

There is no actual research work on the definition of a metric for quantifying scaling in blockchain technologies. The first idea, and what in the end we ended up using, is the adoption of the so-called *transactions per second* (TPS), i.e., the number of transactions successfully added to a block and validated in one second. To be as objective as possible, we used two main types of TPS: The first is the one claimed to be the theoretical one, and the second one is the maximum value that has been measured over the history of transaction count per day. We excluded from the analysis a real-time TPS computation since it relies too much on the current load of the network and on other factors that are time-dependent.

Decentralization Evaluation

Despite the widely acknowledged importance of this property, most discussion on the topic lacks quantification. Sarada Prasad Gochhayat et al.[4] proposed a work that aims to measure decentralization using the *Gini coefficient*. There are striking similarities between the concepts of "too much inequality– and "too much centralization." Specifically, we can think of a nonuniform distribution of wealth as highly unequal and a nonuniform distribution of power as highly centralized.

The Gini coefficient is calculated using the Lorenz curve as depicted in Figure 2. We computed the cumulative distribution of held power starting from lowest values to highest, such as staked tokens, relative to the number of entities possessing that power (e.g., validators). This process generates a curve, i.e., the Lorenz curve, that visually encapsulates the cumulative distribution of decentralization in the analyzed platform. A Lorenz curve that closely aligns to linear function $(y = x)$ indicates a more even distribution of power, thus high decentralization, whereas a significant divergence from this line suggests inequality and greater centralization.

The equation for the Gini coefficient can be calculated from the area B under the Lorenz curve and the area A under the so-called *line of equality* as shown as follows:

$$\text{Gini coefficient} = \frac{A}{A+B}.$$

However, the Gini coefficient has an issue: While a high value tracks with our intuitive notion of a "more centralized" system, the fact that the Gini coefficient is restricted to a 0–1 scale means that it does not directly measure the number of individuals or entities required to compromise a system. An alternative approach is to define a similar metric based on the Lorenz curve from which the Gini coefficient is calculated, which is called the *Nakamoto coefficient*.[6] Given a subsystem s with K entities, let $p_1 > \cdots > p_K$ be the proportions of the subsystem controlled by each of the K participants such that $\Sigma_i^K p_i = 1$. Then we define the Nakamoto coefficient as

$$N_s := \min\left\{k \in [1,\ldots,K] : \sum_{i=1}^{k} p_i \geq \text{threshold}\right\}.$$

In other words, the Nakamoto coefficient (the higher the better) of a subsystem N_s is the minimum number of entities that holds a specific threshold of control. In the case of PoW networks, such a threshold is set to 51%, while for PoS networks is set to 33%.[7] This metric is the one we have chosen to conduct the research, where the individuals or entities are taken into account differ based on the underlying technology:

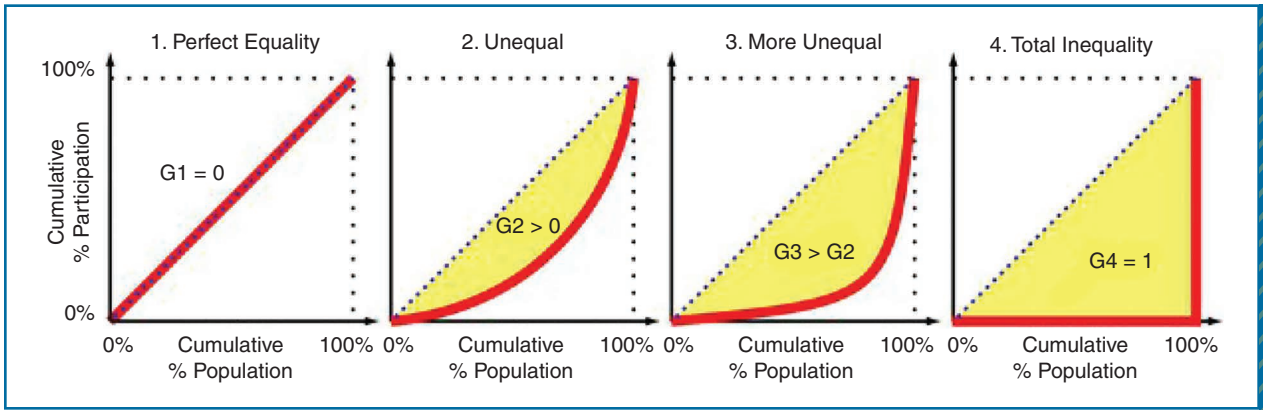

FIGURE 2. The Lorenz curve is shown in red. As the cumulative distribution diverges from a straight line, the Gini coefficient (G) increases from 0 to 1.[5]

- *L1 solutions and sidechains*: the individuals or entities are the blocks validator nodes
- *L2 solutions*: the individuals or entities are the aggregator nodes.

Security Evaluation

The best way to measure security quantitatively should be to count the number of vulnerabilities and exploits/attacks. Since this is practically impossible, we have tried to mention the well-known vulnerabilities in the consensus model adopted and to highlight some special efforts spent by the developers to solve specific security issues. Where possible, it would be appropriate to add a discussion about the known vulnerability and the reports of attacks or bugs.

Scalability Evaluation

Cardano

There is no actual claim about a theoretical number of TPS from the Cardano Team. So, to find this metric, we crafted a method based on the protocol parameters: Since the maximum size of a block (its body in bytes) is equal to **maxBlockBodySize: 65536** bytes. We then took the transaction average size using Cardano GraphQL (https://github.com/input-output-hk/cardano-graphql), obtaining a value of 613 b. This means that, on average, the number of transactions per block is equal to 107. Therefore, since a block is validated on average every 20 s, the theoretical TPS we obtained is 5.35. As for the maximum TPS Cardano has ever reached, from Messari,[8] we can see it got to 495,825 transactions in a day. Hence, the maximum throughput ever got is 5.74. The theoretical TPS we computed is slightly less than the maximum ever reached. This is due to our assumption that the only factor affecting the TPS metric is the block's body size. Some other efficiency improvements might be present in the protocol's implementation. Nevertheless, such considerations are beyond the scope of this conceptual work.

Solana

The theoretical value for TPS in Solana, as stated in its resident implementation documentation, can be computed as follows. On a 1-Gb/s network connection, the maximum number of transactions possible is 1 gigabit per second/176 bytes = 710k TPS. From the Solana explorer (https://analytics.solscan.io/public/dashboard/) it is possible to see the max count of transactions per day, counting only the meaningful ones, i.e., payments and instructions (this tool also takes into account special transactions, like votes, that the Solana protocol needs to operate). This count reaches its maximum at 152,330,000 transactions per day, i.e., 1763 TPS.

Arbitrum

The theoretical value of TPS we found is 40,000. This value was calculated if the maximum capacity of Ethereum was filled by the Arbitrum rollup batches and with no other L1 transactions. From the Arbitrum explorer (https://arbiscan.io/chart/tx) we got the highest number of transactions per day, which is 267,608, i.e., 3.09 TPS.

zkSync

The theoretical value of TPS as stated in the documentation is ~2000.[9] This value was evaluated with the current block gas limit of 12.5 M. From EthTPS (https://ethtps.info/Network?name=ZKSync), we can get the maximum value for TPS in the last two years, which is equal to 0.521.

Polygon

The theoretical value of TPS is 7200, as stated by the Polygon team. This value has been obtained in a Polygon local testnet. Anyway, the maximum value of transactions that have been reached in the mainnet is 9,177,310 (https://polygonscan.com/chart/tx). Hence Polygon reached ~101.97 TPS.

Decentralization Evaluation

Cardano

During each epoch, rewards are distributed among all stakeholders, who have delegated to a stake pool, either their stake pool or another pool. These rewards are autogenerated by the protocol itself and are not managed by the stake pools. From PoolTool (https://pooltool.io), we got the list of stake pools and the amount of ADA staked. As shown in Figure 3, the Nakamoto coefficient is 119.

Solana

Solana adopts the Delegated Proof of Stake consensus algorithm, where validating transactions (blocks) are nodes called *Validators*. We got the list of validators from Validators.app (https://www.validators.app/?locale=en&network=mainnet&order=stake) and the relative amount of SOL staked. As shown in Figure 3, the Nakamoto coefficient is 18.

Arbitrum

Arbitrum, as an optimistic rollup, batch inner transactions and sends them, grouped as one, to the Ethereum L1. To do so, Arbitrum takes advantage of three types of batcher nodes: *forwarders*, *aggregators*, and *sequencers*. Users can send their L2 transactions to any of these three nodes. Forwarder nodes forward any L2 transactions to a designated address

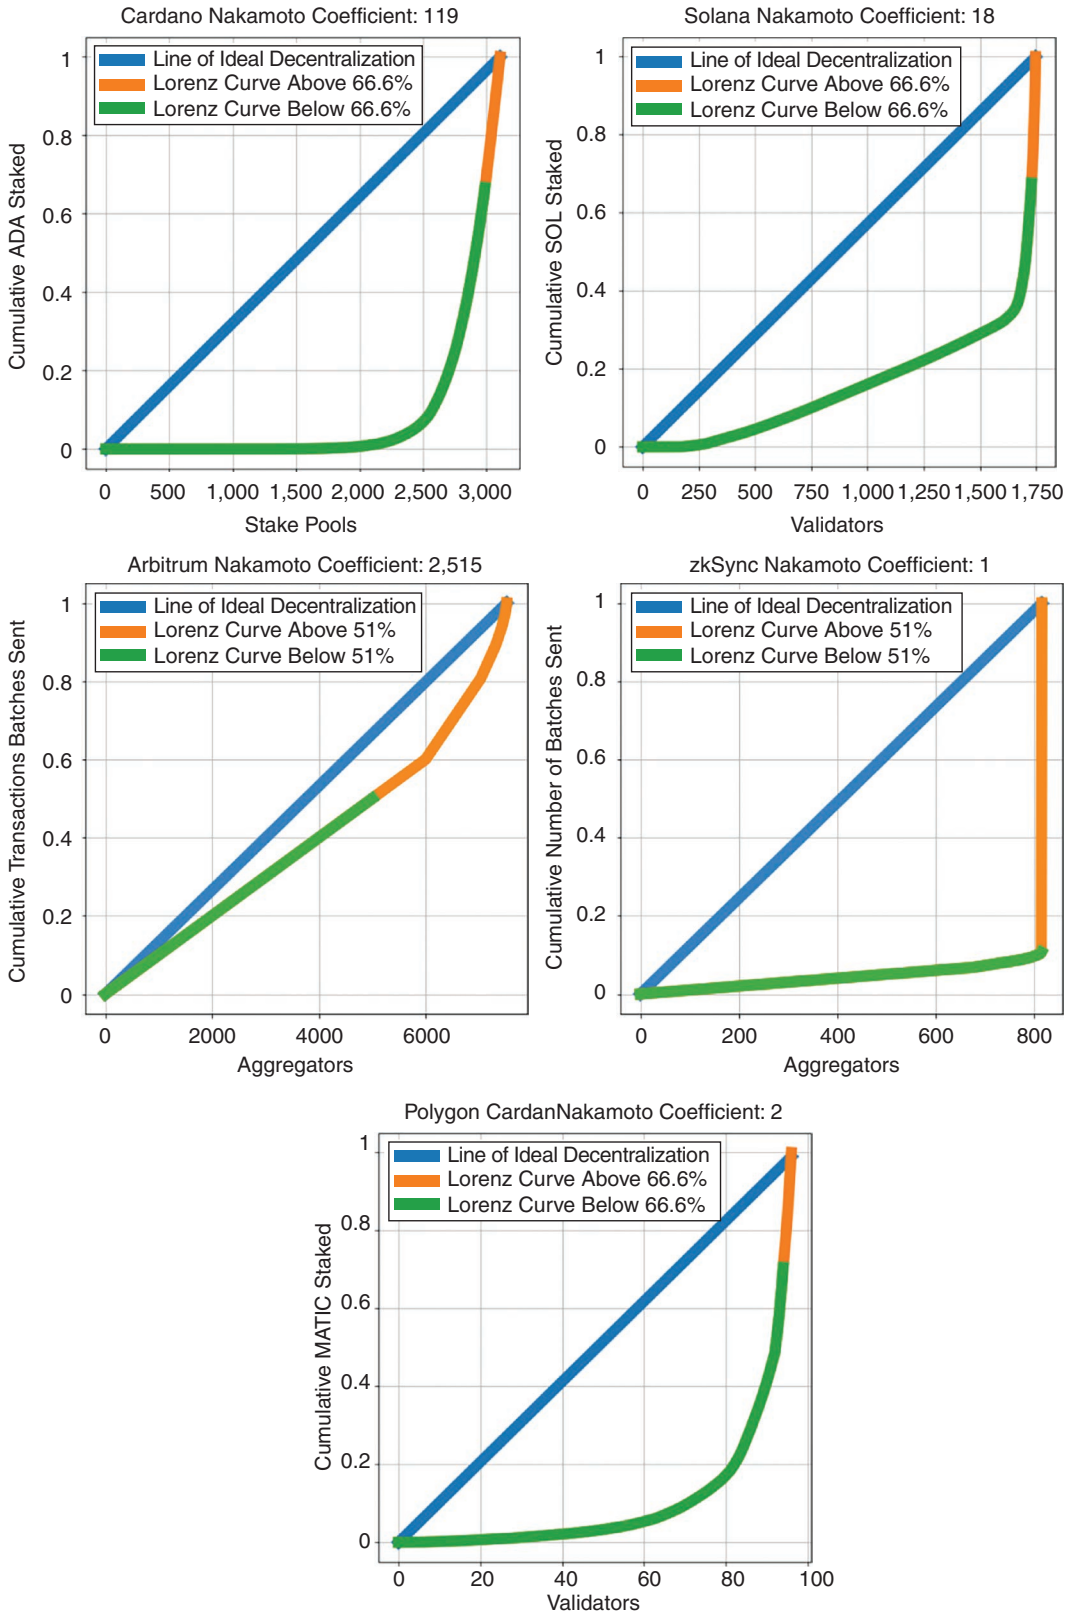

FIGURE 3. Nakamoto coefficients for decentralization in Cardano (119), Solana (18), Arbitrum (2,515), zkSync (1), and Polygon (2).

of another node. The designated node can be either a sequencer or an aggregator and is referred to as the aggregator address. Aggregator nodes will take a group of incoming L2 transactions and batch them into a single message to the delayed inbox (a smart contract on Ethereum). The sequencer node will also take a group of incoming L2 transactions and batch them into a single message, but it will send the batch message to the sequencer inbox instead. If the sequencer node stops adding transactions to the sequencer inbox, anyone can force the sequencer inbox to include transactions from the delayed inbox via a smart contract function call. This allows the Arbitrum network to always be available and resistant to a malicious sequencer.[10] In result, as shown in Figure 3 the Nakamoto coefficient is 2,515.

In Figure 3, we present the Lorenz curves concerning the aggregator nodes. Since an important role is played by the Arbitrum sequencer, too, we calculated, in the same way, the number of transactions in the same period, which is equal to 2,756. To have a comparison, we calculated the number of transactions of the top 2,515 aggregators, which is equal to 4,999.

zkSync

zkSync, as a ZK rollup, has a single smart contract called *delayed inbox*, to which aggregator nodes send batches of transactions. In this evaluation, our focus on the decentralization bottleneck rotates around the number of aggregators and the transaction batches sent to the delayed inbox. This point is crucial as more influential and widely adopted aggregators can manipulate transactions' order or introduce delays, potentially undermining the integrity of the entire network. This analysis has been conducted over 10,000 different transactions thanks to the Etherscan application programming interfaces. As shown in Figure 3, the Nakamoto coefficient is 1.

Polygon

Polygon basically acts as a PoS blockchain with a set of validators. The set of validators was taken from the Polygon website (https://wallet.polygon.technology/staking/validator). The analysis has been conducted by taking the number of validators and the respective amount of MATIC staked. As shown in Figure 3 the Nakamoto coefficient is 2.

Discussion

L1 solutions like Cardano and Solana exhibit similar trends, with Nakamoto coefficients of 119 and 18, respectively. However, these figures are modest when compared to the total number of validators and stake pools, which surpass 3,000 and 1,750, respectively. This indicates that control of the network can be achieved by gaining influence over 4% for Cardano and 1% for Solana of selected selected peers. Regarding L2 solutions like Arbitrum and zkSync, a similar methodology was employed to calculate the Nakamoto coefficient, leveraging the number of aggregators (instead of stakers or validators) and corresponding transaction volumes (instead of tokens staked).

Arbitrum provides two evaluation perspectives: One considers the single default sequencer managed by the platform creators, while the other evaluates user-operated sequencers as an alternative trust model Considering the single sequencer, Arbitrum would exhibit extreme centralization with a Nakamoto coefficient of 1. However, given the option for users to create their own sequencers, we chose to present the most decentralized setup, resulting in a Nakamoto coefficient of 2,515.

In contrast, zkSync does not offer such alternatives, leading to a Nakamoto coefficient of 1, as a single aggregator predominantly handles most L2-to-L1 transactions. Finally, Polygon, operating as a sidechain solution, exhibits a Nakamoto coefficient of 2, indicating a significant level of centralization.

Overall, third-generation blockchains do not appear to excel in decentralization, with the notable exception of Arbitrum, which achieves this when relying on user-operated aggregators.

Security Evaluation

Cardano

Cardano has a special implementation of the Proof of Stake consensus algorithm, named *Ouroboros*.[11] Ouroboros processes transaction blocks by dividing chains into epochs, which are further divided into time slots. A slot leader is elected for each time slot and is responsible for adding a block to the chain. To protect against adversarial attempts to subvert the protocol, each new slot leader is required to consider the last few blocks of the received chain as transient: Only the chain that precedes a prespecified number of transient blocks is considered settled. This is also referred to as the settlement delay and is the mechanism through which the ledger is securely passed between participants. Another interesting point in the security of Cardano is that it's written in Haskell, a functional programming language that encourages building a system using pure functions. This leads to a design where components are conveniently testable in isolation.

Solana

In PoW consensus, block times need to be very large (~10 min) to minimize the odds of multiple validators producing a new valid block at the same time. There's no such constraint in PoS consensus, but without reliable timestamps, a validator cannot determine the order of incoming blocks. The popular workaround is to tag each block with a timestamp. Because of clock drift and variance in network latencies, the timestamp is only accurate within an hour or two. To "workaround the workaround," these systems lengthen block times to provide reasonable certainty that the median timestamp on each block is always increasing. Solana takes a very different approach, which it calls *proof of history* or PoH. Leader nodes "timestamp" blocks with cryptographic proofs that some duration of time has passed since the last proof. All data hashed into the proof most certainly have occurred before the proof was generated. The node then shares the new block with validator nodes, which can verify those proofs. The blocks can arrive at validators in any order or even could be replayed years later. With such reliable synchronization guarantees, Solana can break blocks into smaller batches of transactions called *entries*. Entries are streamed to validators in real time before any notion of block consensus.[12] Another valuable point in Solana security is that smart contracts are written in Rust. This language puts its entire attention to detail and pattern recognition. Solidity surely has more well-known vulnerable patterns making it easier to identify weak code and fix (or exploit).

Arbitrum

It's worth noting that Arbitrum, being an Ethereum rollup, doesn't implement its consensus model, but instead, it relies on Ethereum's one. Anyway, Arbitrum has several features, such as a Rollup, to avoid the so-called *censorships attacks:* Once a forking attack occurs, it will inevitably be discovered by a certain challenger. Therefore since the attack will fail as long as there is a virtuous challenger, the attacker must jam all possible challengers. If there are many such challengers, the attack will already be difficult to accomplish.

zkSync

zkSync is built on zkRollup architectures.[13] This means every single transaction is verified by a smart contract on the Ethereum mainnet utilizing verifying the proof of the validity of the block. Thus, no validator can ever move the system into an incorrect state or take users' funds. In the ultimate emergency case of all validators being shut down or becoming unresponsive, the emergency exit mechanism ensures that users will keep control of their assets. It works as follows:

1. If transactions of a user are being ignored, for any reason, by the validators, an exit request can be submitted directly on the mainnet into the priority queue.
2. Validators are obliged to process priority queue requests within a short time window (≈1 week)
3. In case the validators fail to process the requests, the system enters exodus mode and every user can immediately exit all of their assets by making a direct transaction on the Ethereum mainnet.

This means every single transaction is verified by a smart contract on the Ethereum mainnet utilizing verifying the proof of the validity of the block.

Polygon

Rewards are distributed to all stakers proportional to their stake at every checkpoint with an exception being the proposer getting an additional bonus. User reward balance gets updated in the contract which is referred to while claiming rewards. Of course, validating does have some risks. Stakes are at risk of getting slashed in case the validator node commits a malicious act like double signing and validator downtime which also affects the linked delegators at that checkpoint. As long as two-thirds of the weighted stake of the validators is honest, the chain will progress accurately. Validators stake their MATIC tokens as collateral to work for the security of the network and, in exchange for their service, earn rewards.

Quantitative Comparison

Herein, we present an initial quantitative assessment of the selected blockchain technologies within the trilemma. The results, summarized in Table 1, facilitate a comparative

evaluation of each technology's performance in these key areas.

Scalability, indicated by TPS, varies significantly among the technologies, with Solana demonstrating high throughput, wherefore the only previous work that approached such a proposition is the review by Sanka et al.[14] Decentralization, reflected by the Nakamoto coefficient, highlights Arbitrum's high score, suggesting a robust decentralized network. Security, while difficult to quantify directly, has been computed as the cost in U.S. dollars to obtain control of the blockchain. In particular, we used the following metric:

$$\text{cost} := c * \sum_{i}^{n} s_i$$

where n is the total amount of validators, c is the token cost in U.S. dollars, and s_i is the amount of tokens staked by validator i. The higher this cost is the more difficult it is to take control of the blockchain.

For the security cost metric computation in PoS blockchains, we utilized available staking data specific to each blockchain. However, for rollups like Arbitrum and zkSync, the absence of a native staking mechanism necessitated a different method. Given their reliance on Ethereum for final transaction settlement, we used Ethereum's PoS security data as a proxy. This decision stems from the operational characteristics of rollups, where the security, especially in a decentralized scenario with multiple independent batcher nodes, is closely aligned with that of Ethereum's PoS model.

Obtained results highlight distinct tradeoffs among scalability, decentralization, and security. For example, Solana exhibits an impressive scalability with a TPS of 1,763, suggesting its capability to handle high transaction volumes efficiently. However, its Nakamoto coefficient is relatively low at 18, indicating a lesser degree of decentralization compared to other platforms such as Arbitrum, which boasts a Nakamoto coefficient of 2,515, reflecting a high level of network distribution and fault tolerance.

In terms of security, both zkSync and Arbitrum exhibit high values, with a cost of US$20.6 billion since they rely on the underlying secure Ethereum L1 network. Cardano and Polygon offer a balance between these metrics, with Cardano achieving a TPS of 5.74 and a Nakamoto coefficient of 119, and Polygon reaching a TPS of 101.97 with a Nakamoto coefficient of 2. This illustrates how each platform navigates the scalability-decentralization-security trilemma, optimizing certain aspects while compromising others to varying degrees.

It is crucial to acknowledge the influence of technology popularity on the maximum TPS in the scalability evaluation. A notable limitation of the TPS metric is its sensitivity to the adoption and popularity of a particular blockchain technology. For instance, a blockchain's widespread use and acceptance can impact its maximum TPS, as more participants and nodes contribute to the network, potentially leading to more transactions being processed. This phenomenon underscores the need to consider other approaches, like the theoretical TPS we included in the study.

This study underscores the complexity of optimizing blockchain technologies across scalability, decentralization, and security. The suitability of a blockchain for specific applications depends on the particular requirements and tradeoffs considered acceptable for the use case. This analysis provides a foundation for understanding these tradeoffs, aiding in the informed selection of blockchain technologies for third-generation applications.

TABLE 1. Comparison of blockchain approaches.

Blockchain	Scalability (TPS)	Decentralization (Nakamoto coefficient)	Security (cost in billions of U.S. dollars to obtain control of the blockchain)
Cardano	5.74	119	0.528
Solana	1763	18	9.11
Polygon	101.97	2	0.846
Arbitrum	3.09	2515	20.6
zkSync	0.521	1	20.6

The academic literature currently lacks explicit analysis of strategies to evaluate and compare third-generation blockchains due to the emerging nature of the research field. This work presents a framework to evaluate third-generation blockchains based on the Blockchain Trilemma.

In conclusion, further research is needed to properly ascertain the

ABOUT THE AUTHORS

GIOVANNI QUATTROCCHI is a junior assistant professor at Politecnico di Milano, 20133 Milan, Italy. His research interests include self-adaptive systems, software architectures, edge computing, and blockchain-based systems. Quattrocchi received his Ph.D. in computer engineering from Politecnico di Milano. Contact him at giovanni.quattrocchi@polimi.it.

FILIPPO SCARAMUZZA is a research assistant at Politecnico di Milano, 20133 Milan, Italy. His research interests include blockchain-oriented designs and their nonfunctional aspects. Scaramuzza received his M.Sc. in computer science and engineering from Politechnico di Milano. Contact him at filippo.scaramuzza@mail.polimi.it.

DAMIAN A. TAMBURRI is an associate professor at the Jheronimus Academy of Data Science, Eindhoven University of Technology, 5211 s' Hertogenbosch, The Netherlands. His research interests include data-intensive services DevOps/DataOps, social software engineering, and AI software engineering. Tamburri received his Ph.D. in information and software engineering from VU University Amsterdam. Contact him at d.a.tamburri@tue.nl.

nature and nurture of the blockchain trilemma, especially in sight of multimixes of third-generation blockchain-oriented computing. Furthermore, other paradigms, such as the blockchain quadrilemma,[15] merit exploration to propose novel insights, while also integrating the findings presented in this study. At the same time several key findings of our study act as design principles for blockchain-oriented design of the future. First, the assets on Table 1 provide valuable starting points to analysis blockchain-oriented designs featuring those specific technologies, for example, in the scope of ad-hoc architecture tradeoff analysis exercises. Second, the contents of Figure 3 induces a partial order between the select target technologies; such order is to be used when designing multichain solutions striving to achieve specific properties from the trilemma (e.g., decentralization). Third, finally, our approach itself serves as a general methodology to evaluate blockchain technology against the trilemma wherefore the design needs dictate it.

Acknowledgment

This research was supported in part by the Dutch Research Council (NWO) as part of the CHAIN project (https://chainresearch.eu/events/online-meeting-platform-for-responsible-innovation/) within the program "Responsible Innovation. Designing for Public Values in a Digital World" and by project 3A-Italy Circular and Sustainable Made in Italy—MICS (3A-ITALY) CUP D43C22003120001 (Grant PE00000004).

References

1. U. W. Chohan, *The Limits to Blockchain? Scaling vs. Decentralization*. Rochester, NY, USA: SSRN, Jan. 2019, doi: 10.2139/ssrn.3338560.
2. S. He, Y. Ning, H. Chen, C. Xing, and L.-J. Zhang, "Layered consensus mechanism in consortium blockchain for enterprise services," in *Blockchain—ICBC* (Lecture Notes in Computer Science). J. Joshi, S. Nepal, Q. Zhang, and L.-J. Zhang, Eds. Cham, Switzerland: Springer-Verlag, 2019, vol. 11521, pp. 49–64.
3. T. Zimmermann, "Card-sorting," in *Perspectives on Data Science for Software Engineering*, T. Menzies, L. A. Williams, and T. Zimmermann, Eds., New York, NY, USA: Academic Press, 2016, pp. 137–141.
4. S. P. Gochhayat, S. Shetty, R. Mukkamala, P. Foytik, G. A. Kamhoua, and L. Njilla, "Measuring decentrality in blockchain based systems," *IEEE Access*, vol. 8, pp. 178,372–178,390, 2020, doi: 10.1109/ACCESS.2020.3026577.
5. J. Hanson. "Can the Gini coefficient be negative?" Accessed: May 04, 2022. [Online]. Available: https://www.quora.com/Can-the-Gini-coefficient-be-negative
6. S. Chu and S. Wang, "The curses of blockchain decentralization," 2018, *arXiv:1810.02937*.
7. "Ethereum proof-of-stake attack and defense." Ethereum.Org. Accessed: Jul. 1, 2024. [Online]. Available: https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/attack-and-defense/
8. "Cardano (ADA) transaction count." Messari. Accessed: May 05, 2022.

[Online]. Available: https://messari.io/charts/cardano/txn-cnt
9. "zkEVM FAQ." zkSync. Accessed: May 08, 2022. [Online]. Available: https://docs.zksync.io/zkevm/#how-scalable-is-a-zk-rollup
10. Kyle Charbonnet. "A technical introduction to Arbitrum's optimistic rollup." Medium. Accessed: May 09, 2022. [Online]. Available: https://medium.com/privacy-scaling-explorations/a-technical-introduction-to-arbitrums-optimistic-rollup-860955ea5fec
11. A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-of-stake blockchain protocol," in *Advances in Cryptology – CRYPTO 2017*, J. Katz and H. Shacham, Eds., Cham, Switzerland: Springer-Verlag, 2017, pp. 357–388.
12. Anatoly Yakovenko, Solana Foundation. "Solana: A new architecture for a high performance blockchain." Accessed: Jul. 1, 2024. [Online]. Available: https://solana.com/solana-whitepaper.pdf. v0.8.13
13. Matter Labs. zkSync documentation. zkSync. Accessed: May 04, 2022. [Online]. Available: https://docs.zksync.io/userdocs/
14. A. I. Sanka and R. C. C. Cheung, "A systematic review of blockchain scalability: Issues, solutions, analysis and future research," *J. Netw. Comput. Appl.*, vol. 195, Dec. 2021, Art. no. 103232, doi: 10.1016/j.jnca.2021.103232.
15. F. Mogavero, I. Visconti, A. Vitaletti, and M. Zecchini, "The blockchain quadrilemma: When also computational effectiveness matters," in *Proc. IEEE Symp. Comput. Commun. (ISCC)*, 2021, pp. 1–6, doi: 10.1109/ISCC53001.2021.9631511.

Open Access funding provided by 'Politecnico di Milano' within the CRUI CARE Agreement

Digital Object Identifier 10.1109/MS.2024.3459148