INTERNATIONAL DEVELOPMENT PLANNING REVIEW
ISSN:1474-6743 | E-ISSN:1478-3401

BLOCKCHAIN SCALABILITY SOLUTIONS: LEVELS AND METHODS

Anna V. Aleshina[12*†], Andrey L. Bulgakov[12*†], Alexey D. Demidov[1*†]

1. Moscow Center for Fundamental and Applied Mathematics, Lomonosov Moscow State University, 119991, Russian Federation; bulgakoval@my.msu.ru (A.L.B.); aleshina@econ.msu.ru (A.V.A.); demidov.ad@bk.ru (A.D.D.);
2. Faculty of Economics, Lomonosov Moscow State University, 119991, Russian Federation; bulgakoval@my.msu.ru (A.L.B.); aleshina@econ.msu.ru (A.V.A.)

* Correspondence: aleshina@econ.msu.ru; bulgakoval@my.msu.ru; demidov.ad@bk.ru;
† These authors contributed equally to this work.

Abstract

This paper explores modern approaches to addressing key challenges of scalability, security, and performance in blockchain systems. The focus is on sharding technologies, directed acyclic graph (DAG) architectures, and consensus algorithms such as Proof of Work (PoW), Proof of Stake (PoS), and Practical Byzantine Fault Tolerance (PBFT). The study highlights successful implementations of sharding in systems like RepChain, OmniLedger, and Meepo, showcasing their efficiency in cross-shard communication, atomic transactions, and reputation mechanisms. The advantages and limitations of DAG-based architectures, applied in IOTA and Nano, are analyzed, emphasizing their potential for high scalability and energy efficiency. An evaluation of consensus algorithms underscores their impact on the balance between security, decentralization, and performance.

The findings confirm the potential of integrating these solutions to develop robust and scalable blockchain platforms suitable for real-world applications in sectors such as finance, logistics, and the Internet of Things (IoT). This paper provides a foundation for further research on advancing decentralization technologies and enhancing their adaptability to practical deployment scenarios.

Keywords: blockchain technologies, sharding, scalability,consensus algorithms, security, decentralisation,

Introduction

Blockchain, as one of the key technologies of the digital era, has demonstrated enormous potential in transforming many economic sectors over the past decades. Its applications go far beyond cryptocurrencies and include such areas as finance, logistics, healthcare, energy, and public administration. The technology is based on the principle of decentralization, which ensures transparency, immutability of data, and security of transactions. However, as the number of blockchain users increases and it scales, researchers and developers face new challenges related to network performance, throughput, and sustainability.

Key approaches to solving the scalability problem are sharding and the use of new architectures such as directed acyclic graphs (DAGs). Sharding allows dividing the network into

separate subgroups of nodes (shards) that process transactions in parallel, which significantly reduces the load on each node and increases throughput. DAG architectures, on the contrary, offer a departure from the traditional linear blockchain in favor of distributed transaction graphs, which helps to improve the efficiency of transaction processing. These solutions open up new horizons for the implementation of blockchain in large-scale projects, but also raise new issues of security and governance.

One of the most discussed topics remains the relationship between the technical characteristics of blockchain platform ecosystems and their practical capabilities for application in various industries. In particular, such promising models as RepChain, OmniLedger and Meepo demonstrate innovative approaches to ensuring inter-shard communication and security, but require further study under real-world loads. These systems implement reputation mechanisms, atomic inter-shard transactions and fault-tolerance mechanisms that ensure high performance and stability even under high loads.

The purpose of this paper is to study the scalability potential of blockchain platforms using modern solutions, including sharding and DAG. Particular attention is paid to the analysis of the performance of private Ethereum networks based on the Proof of Stake (PoS) consensus, which is of particular importance for industrial applications. The results of testing such networks provide valuable data for assessing their suitability for real conditions, and also serve as a basis for identifying promising areas for further development**.**

Sharding

Scalability and Security with Sharding: RepChain Experience

One of the most promising approaches to solving scalability and security problems in blockchain systems is the use of sharding. The sharding technology, effectively implemented in RepChain, allows to significantly increase the network throughput while maintaining its reliability and resistance to attacks.

RepChain implements a strategy of dividing nodes into groups (shards), each of which processes transactions in parallel. This approach reduces the load on individual validators and ensures a linear increase in system performance as the number of nodes increases. Results show that RepChain is capable of achieving a throughput of up to 6852 transactions per second, making it one of the most productive solutions on the market.

A key element of RepChain scalability is the balanced distribution of nodes between shards. The use of reputation scores prevents the concentration of weak or malicious nodes in one shard, which increases the overall level of system security. Shard leaders are selected based on their reputation, which ensures the ability to efficiently process significant transaction volumes.

The Reputation Chain plays a central role in ensuring security and performance. The system generates reputation scores based on the historical behavior of validators, including their activity and correctness in processing transactions. This data is used to select shard leaders, which helps to evenly balance the load between nodes. Moreover, a high level of reputation incentivizes validators to behave honestly and productively, as this increases their chances of receiving rewards.

RepChain's reputation mechanisms ensure network resilience even in the face of targeted attacks such as "self-promotion" or "slander". Strong consensus and warning mechanisms prevent attempts to disrupt the network. In addition, the protocol remains effective in the face of "masking" or "adaptive surveillance" attacks, demonstrating high resilience to internal and external threats.

Thus, RepChain is a shining example of the successful application of sharding to solve key problems of scalability and security of blockchain systems. Her experience highlights the importance of integrating reputation mechanisms into the architecture of modern decentralized platforms, which opens up prospects for their wider use in real-world scenarios.

OmniLedger: statistically representative shards and cross-shard transaction mechanism

Continuing to study modern approaches to the scalability of blockchain systems, it is worth highlighting OmniLedger, a protocol that offers a statistically sound method for forming shards and an innovative mechanism for processing cross-shard transactions. These technologies can improve both the security and performance of the network, ensuring its stability even under high loads.

Statistically representative shards

One of the key features of OmniLedger is the use of random distribution of validators across shards, which minimizes the likelihood of their compromise. This is achieved using the RandHound protocol, which ensures the creation of a bias-free distribution of random numbers. This approach allows for an even distribution of nodes, reducing the likelihood that a single shard will contain a significant proportion of malicious participants (no more than 25%).

To further enhance security, OmniLedger regularly mixes the shards, preventing them from being compromised in the long term. The use of large shards, combined with periodic node rotation, makes the system less vulnerable to targeted attacks.

Cross-shard transaction mechanism (Atomix)

To process transactions between shards, OmniLedger implements the Atomix mechanism, which guarantees their atomicity. The process consists of two main phases:

1. Lock phase:

The validators of the shards containing the transaction inputs check their correctness and fix the funds in the "locked" state. At this stage, a proof-of-acceptance is generated, ensuring that the shard is ready for further actions.

2. Unlock phase:

If all participating shards confirm the transaction, an "unlock-to-commit" is created, and the transaction is finalized. In the event of at least one shard failure, the funds are returned to the sender via an "unlock-to-abort". This approach eliminates the possibility of double spending and automates the return of funds to the user.

Atomix not only ensures strict transaction atomicity, but also involves clients in the confirmation process, which increases transparency and reduces the overhead of inter-shard communication.

Advantages of the OmniLedger approach

OmniLedger demonstrates a number of significant advantages:

• Scalability: Linear growth of the system throughput is achieved through the use of sharding. This makes the network suitable for processing huge amounts of data in large-scale applications.

• Security: Statistical distribution of validators and regular rotation of nodes minimize the risk of system compromise.

• Efficiency: Atomix guarantees reliable transaction execution with minimal overhead, maintaining high performance even with complex operations.

Thus, OmniLedger is an innovative solution that combines statistically sound share management methods and advanced inter-shard communication mechanisms. Its architecture provides the necessary balance between performance, security and ease of implementation, which makes it a universal tool for modern blockchain applications.

Meepo: Optimizing Inter-Shard Communication and Ensuring Transaction Atomicity

Improving inter-shard communication mechanisms is one of the most important tasks for increasing the performance and stability of blockchain systems. The Meepo protocol offers a comprehensive solution aimed at eliminating the main problems that arise in the process of interaction between shards and ensuring strict transaction atomicity.

In modern blockchain systems with sharding, cross-shard transactions face a number of limitations that reduce their efficiency and scalability. Several aspects stand out among the key issues. The efficiency of cross-shard transactions is hampered by frequent interactions between shares, which makes it impossible to use a fixed distribution of accounts, since the unpredictability of user behavior requires adaptive approaches. The limited flexibility of existing transaction models does not allow for effective interactions between more than two shares, which reduces the performance of complex smart contracts. The problems of strict transaction atomicity are caused by the fact that asynchronous methods do not guarantee immediate confirmation, and two-phase commit locks resources, limiting parallel execution of operations. In addition, maintaining synchronization between shares to ensure their availability increases the complexity of network management.

To overcome these challenges, Meepo implements a number of innovative approaches. Block completion and validation allows shards to exchange data packets containing inter-shard calls, which reduces latency and improves overall network performance due to optimizations across consortium local networks. The partial call merging strategy allows for simplified transaction execution by having shards issue calls with partial parameters, which are then merged

at the system level. In case of errors, transactions are removed from the block and the block is replayed taking into account the correction, which preserves the atomicity of transactions without affecting performance. To prevent failures, backup nodes are used that are activated when the primary node fails, preventing a complete system shutdown and maintaining its high availability.

The Meepo protocol has been tested on the Go-Ethereum platform using the Proof of Authority (PoA) consensus mechanism. Experimental results demonstrate its high efficiency, as the protocol copes with processing inter-shard transactions even under significant load. Replay-epoch effectively eliminates errors while maintaining transaction atomicity, and the use of backup nodes ensures high system availability, reducing the risk of failures.

The Meepo protocol is a promising solution for blockchain systems that require high performance and strict transaction atomicity. Its architecture meets modern scalability and reliability requirements, making it a suitable choice for consortium networks and other high-load applications.

Directed Acyclic Graph (DAG): An Alternative to the Linear Blockchain Structure

Directed Acyclic Graph (DAG) Basics

Directed Acyclic Graph (DAG) offers an alternative approach to organizing data in blockchain systems by replacing the linear chain of blocks with a decentralized graph structure. The main difference with DAG is that there is no need to combine transactions into blocks. Each new transaction is added to the graph, confirming one or more previous ones, which eliminates the need for a linear sequence. This solution provides more flexible and scalable data management.

The main principles of DAG include the absence of a single chain of truth, autonomous confirmation of transactions, and parallel processing of operations. In such a system, confirmations are distributed among all nodes, eliminating dependence on centralized validators. Due to this, each new transaction becomes an active participant in the process, confirming previous ones, which minimizes the risk of double spending and increases network throughput.

Examples of DAG-based systems

DAG applications can be seen in systems such as IOTA and Nano, each of which exhibits unique features of the technology.

IOTA implements a DAG structure called Tangle, which is aimed at the Internet of Things (IoT) market. Each transaction on the network confirms the previous two, creating a decentralized network without the need for mining. This approach reduces energy consumption, but makes the network vulnerable to low activity, increasing the likelihood of double spending attacks.

Nano uses the Block Lattice concept, where each account is represented by a separate blockchain. This allows for instant transactions and eliminates fees, making the system effective for payment transactions. However, Nano has the disadvantage of relying on a representative voting system, which reduces decentralization in the event of conflicts.

Advantages of DAG

One of the key advantages of DAG is high scalability. By processing transactions in parallel, the network can handle large loads, which is especially important for systems such as IoT. The elimination of mining makes DAG systems more energy-efficient and environmentally friendly, and the absence of fees increases their attractiveness to users. In addition, the absence of a linear chain reduces the likelihood of a single point of failure, which increases the resilience of the system.

DAG Limitations

Despite many advantages, DAG technologies face a number of limitations. One of the problems is vulnerability to attacks during low network activity, as is observed in IOTA. In addition, the maturity of such technologies remains at the development stage, which limits their use on a mass scale. The implementation of consensus in DAG requires more complex algorithms, which complicates development and integration. Additional mechanisms, such as the coordinator node in IOTA, can reduce the level of decentralization, introducing elements of centralization to protect against attacks.

Real-world DAG application scenarios

DAG-based systems find their application in various industries. For example, IOTA is actively used for microtransactions in the IoT, where minimal fees and high processing speed are important. Nano is focused on instant transfers with zero fees, making it attractive for financial transactions. Thus, DAG technologies continue to open new horizons for blockchain applications, providing flexible and scalable solutions.

Consensus Algorithms: The Basics of Decentralized Systems

Consensus algorithms are the foundation of blockchain systems, ensuring that data is consistent across distributed nodes. These algorithms ensure that network participants agree on the state of the blockchain, even in the absence of trust and the presence of potentially malicious nodes. The choice of a particular consensus algorithm directly affects the performance, security, energy efficiency, and degree of decentralization of the network.

Today, there are various approaches to implementing consensus, each with its own advantages and limitations. Among the most common are Proof of Work (PoW), Proof of Stake (PoS), and Practical Byzantine Fault Tolerance (PBFT). They are used depending on the goals and characteristics of the network, as well as the requirements for its performance and security. Below is a detailed analysis of these approaches.

Proof of Work (PoW): High Security, Low Performance

Proof of Work (PoW) is one of the first and most widely used consensus algorithms, which underpins blockchains such as Bitcoin and Ethereum (before switching to PoS). Its operating principle is that network participants (miners) perform complex calculations to solve a cryptographic problem. After verifying the solution, the network adds a new block to the chain.

One of the main advantages of PoW is its high security. The difficulty of performing a 51% attack makes control over the network extremely expensive, since an attacker must have more than half of the network's computing power. PoW is also well studied and has proven its reliability. However, this algorithm suffers from low performance: the limitation on the rate of block creation leads to low throughput, for example, Bitcoin processes about 7 transactions per second. In addition, mining requires significant energy costs, which has a negative impact on the environment. Despite its shortcomings, PoW continues to be used in systems where security is a priority.

Proof of Stake (PoS): Resource Saving, But Centralization Problems

Proof of Stake (PoS) is a more energy-efficient alternative to PoW. In this model, block validators are elected proportionally to the number of tokens they hold and stake. This approach reduces energy consumption, since there is no need to perform complex calculations. Thanks to this, PoS networks such as Ethereum (after switching to PoS), Cardano, and Solana are able to process more transactions in less time.

However, PoS faces the problem of centralization. Participants with a large number of tokens get more opportunities to create new blocks, which can lead to the concentration of control in the hands of a few validators. It also contributes to increasing imbalance, as validators with large stakes continue to increase their income, widening the gap with the rest of the network participants. Despite these drawbacks, PoS is actively used in projects focused on energy efficiency and high performance.

Practical Byzantine Fault Tolerance (PBFT): Fast, but Limited Scalability

Practical Byzantine Fault Tolerance (PBFT) is a consensus algorithm designed to operate in a partially trusted environment. In PBFT, nodes agree on the state of the system through messaging, with each transaction being confirmed by a majority of nodes. PBFT provides high transaction processing speeds because it requires a minimum number of communication rounds.

Fault tolerance makes the system resilient to malicious actions or errors if the proportion of such nodes does not exceed one-third. However, PBFT has limited scalability: as the number of nodes increases, the complexity of communication increases exponentially, making it less effective for large decentralized networks. The significant network load caused by messaging also limits its use. Despite these limitations, PBFT is actively used in enterprise blockchains such as Hyperledger Fabric and Tendermint, which operate with a limited number of trusted participants.

Conclusion

The conducted research has shown that the scalability, security and performance of blockchain systems can be significantly improved by using advanced technologies such as sharding, directed acyclic graphs (DAG) and various consensus algorithms. Each of the considered solutions makes an important contribution to the development of decentralized networks, but at the same time comes with specific limitations and challenges that require further study and improvement.

Sharding technologies, presented by the example of RepChain, OmniLedger and Meepo, have demonstrated high potential for increasing the throughput and resilience of blockchains. The use of reputation mechanisms, statistical distribution of validators and inter-shard transactions ensures the reliability of the network even under high load. These developments create the basis for the implementation of scalable blockchain systems capable of coping with the growth of the number of users and the volume of data.

DAG architectures such as IOTA and Nano offer a radically different approach to transaction processing, providing the ability to execute operations in parallel. This approach not only improves network performance, but also reduces energy costs and transaction fees. However, using DAG requires addressing issues such as vulnerability during low network activity and the complexity of consensus implementation.

Analysis of consensus algorithms has revealed their important role in ensuring a balance between security, performance, and energy efficiency. Proof of Work remains the benchmark for security, although its energy costs and low performance limit scalability. Proof of Stake provides higher performance and reduces energy consumption, but faces centralization challenges. Practical Byzantine Fault Tolerance, in turn, demonstrates excellent speed indicators in conditions of partial trust, which makes it suitable for enterprise blockchains.

Patens

Aleshina, A.V., Bulgakov, A.L., Smirnov, S.D., Demidov, A.D., Milyutin, M.A., Xin, Y., & Dmitrieva, L.V. Model for studying the scalability of the Ethereum network with load testing capabilities. Patent No. RU 2024684956. Federal Service for Intellectual Property, Russian Federation, October 22, 2024.

Author Contributions: These authors contributed equally to this work
Funding: The paper was published with the financial support of the Ministry of Education and Science of the Russian Federation as part of the program of the Moscow Center for Fundamental and Applied Mathematics under the agreement № 075-15-2022-284.
Institutional Review Board Statement: Not applicable.
Informed Consent Statement: Not applicable.
Data Availability Statement: Data sharing not applicable to this article as no datasets were generated or analyzed during the current study.The original contributions presented in this study are included in the article and supported by the following patent:
Aleshina, A.V., Bulgakov, A.L., Smirnov, S.D., Demidov, A.D., Milyutin, M.A., Xin, Y., & Dmitrieva, L.V. (2024). Model for studying the scalability of the Ethereum network with load testing capabilities (Patent No. RU 2024684956). Federal Service for Intellectual Property, Russian Federation.
In addition to the patents, the authors have made the source code for experiments and the private Ethereum network configuration publicly available. The code, created by the authors, can be accessed on GitHub at https://github.com/Ieef-msu/pos-ethereum-network-bench.
Acknowledgments: The authors thank the Moscow Center for Fundamental and Applied Mathematics for providing resources necessary for conducting the experiments.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Nasir, M. H., Arshad, J., Khan, M. M., Fatima, M., Salah, K., & Jayaraman, R. (2022). Scalable blockchains—A systematic review. *Future generation computer systems*, *126*, 136-162. https://doi.org/10.1016/j.future.2021.07.035

2. Sang, Y., Cali, U., Kuzlu, M., Pipattanasomporn, M., Lima, C., & Chen, S. (2020, August). IEEE SA blockchain in energy standardization framework: Grid and prosumer use cases. In *2020 IEEE Power & Energy Society General Meeting (PESGM)* (pp. 1-5). IEEE. https://doi.org/10.1109/PESGM41954.2020.9281709

3. Sanka, A. I., & Cheung, R. C. (2021). A systematic review of blockchain scalability: Issues, solutions, analysis and future research. *Journal of Network and Computer Applications*, *195*, 103232. https://doi.org/10.1016/j.jnca.2021.103232

4. Kokoris-Kogias, E., Jovanovic, P., Gasser, L., Gailly, N., Syta, E., & Ford, B. (2018, May). Omniledger: A secure, scale-out, decentralized ledger via sharding. In *2018 IEEE symposium on security and privacy (SP)* (pp. 583-598). IEEE. https://doi.org/10.1109/SP.2018.000-5

5. Huang, C., Wang, Z., Chen, H., Hu, Q., Zhang, Q., Wang, W., & Guan, X. (2020). Repchain: A reputation-based secure, fast, and high incentive blockchain system via sharding. *IEEE Internet of Things Journal*, *8*(6), 4291-4304. https://doi.org/10.1109/JIOT.2020.3028449

6. Bamakan, S. M. H., Motavali, A., & Bondarti, A. B. (2020). A survey of blockchain consensus algorithms performance evaluation criteria. *Expert Systems with Applications*, *154*, 113385. https://doi.org/10.1016/j.eswa.2020.113385