Vol. 2, No. 4, pp.: 33-49, 2024

A Data-Driven Survey of Blockchain Technology: Architecture, Consensus Mechanisms, Cross-Domain Applications, and Decentralization Metrics

Andi Pratama Wijaya[1], Siti Nurhaliza Putri[2,*], Budi Santoso Hidayat[3]

[1] Department of Informatics Engineering, Faculty of Engineering, Universitas Diponegoro, Semarang 50275, Indonesia

[2] Faculty of Computer Science, Universitas Brawijaya, Malang 65145, Indonesia

[3] Department of Information Systems, Faculty of Information Technology and Data Science, Universitas Sebelas Maret, Surakarta 57126, Indonesia

** siti.putri@ub.ac.id*

Article Information	
Received	16 July 2024
Accepted	21 November 2024
DOI	https://doi.org/10.63646/datamind.2024.020404

Abstract

Blockchain technology has evolved from a single-purpose payment ledger into a general-purpose distributed-systems infrastructure that now underpins decentralized finance, supply chain monitoring, healthcare data exchange, energy markets, education credentialing, and Internet of Things (IoT) coordination. Despite this momentum, applied researchers face a fragmented picture: published surveys frequently bundle architectural primitives, consensus mechanisms, and application case studies without providing a structured, quantitative view of how the technology is actually used across domains, or of how decentralized real platforms are when assessed with explicit metrics. This article addresses that gap. It synthesizes the architectural foundations of blockchain systems, classifies the principal consensus mechanisms into proof-based and voting-based families, surveys eight application domains with concrete deployment examples, and analyses the distribution of recent research output across these domains. We also evaluate decentralization quantitatively using the Nakamoto coefficient across nine major platforms, highlighting that headline market capitalization is a poor predictor of effective decentralization. The discussion translates these results into design guidance: practitioners should select a consensus family on the basis of the trust assumptions of the target deployment, treat

scalability and energy consumption as first-order constraints rather than tunables, and use decentralization metrics as part of the platform-selection workflow rather than as a marketing afterthought.

Keywords: ***Blockchain; consensus mechanisms; decentralized applications; Nakamoto coefficient; smart contracts; cross-domain adoption; distributed ledger technology***

1. Introduction

Distributed ledger technology has shifted in roughly fifteen years from a narrow cryptocurrency mechanism into one of the most actively researched infrastructure primitives in computer science. The original Bitcoin proposal demonstrated that a peer-to-peer network could agree on a single, append-only transaction history without a trusted intermediary, but that demonstration was only the starting point. Subsequent platforms generalized the model with programmable execution environments, alternative consensus designs, and permissioned variants suited to enterprise deployment. The result is a family of systems that share a common architectural backbone — cryptographic linking, replicated state, peer-to-peer dissemination, consensus, and an application interface — but that differ sharply in throughput, trust model, energy profile, and governance. Treating these systems as interchangeable, as is sometimes done in applied studies, leads to comparisons that do not hold up under scrutiny (Birje et al., 2023; Rajasekaran et al., 2022).

From an applied perspective, the practical question is no longer whether blockchain works, but where it is actually appropriate. Reported adoption studies span finance, supply chain traceability, healthcare records, IoT device coordination, energy trading, land registries, and academic credentialing. Each of these areas has produced proof-of-concept systems and a smaller but growing number of production deployments. Yet the empirical literature remains uneven. Some domains, notably decentralized finance (DeFi), have generated thousands of indexed publications and a stable set of design patterns (Schär, 2021). Others, such as blockchain-supported education, are still dominated by conceptual proposals with limited field validation (Alsobhi et al., 2023; Haque et al., 2023). A balanced review must therefore separate domains where the technology has matured from those where it remains aspirational.

A second under-examined issue is decentralization itself. Decentralization is the feature most often invoked to justify a blockchain deployment, yet it is rarely measured. The Nakamoto coefficient — the minimum number of independent entities whose collusion would compromise consensus — provides a tractable, comparable quantitative metric. Public reporting of this coefficient across major platforms reveals that several widely used systems concentrate effective control in fewer than ten validators, which is uncomfortably close to the threshold at which collusion becomes plausible. Treating decentralization as a measurable attribute, rather than a categorical claim, reframes platform selection as an engineering trade-off rather than a branding choice (Ahakonye et al., 2024).

This article contributes to the applied literature in four ways. First, it provides a compact but structured account of blockchain architecture and operation, deliberately separating the data layer, network layer, consensus layer, and application layer so that subsequent comparisons remain consistent. Second, it organizes consensus mechanisms into proof-based and voting-based families, evaluating each on throughput, energy use, fault tolerance, and decentralization. Third, it surveys cross-domain adoption with concrete examples and an aggregated picture of where applied research is concentrated. Fourth, it uses the Nakamoto coefficient as the central decentralization metric and discusses what its observed distribution implies for platform choice. The remainder of the article is organized as follows. Section 2 covers architecture. Section 3 analyses consensus. Section 4 maps the application landscape. Section 5 presents quantitative findings on research distribution and platform decentralization. Section 6 discusses the trade-offs that emerge. Section 7 lists outstanding challenges, and Section 8 concludes.

2. Blockchain Architectural Foundations

A blockchain is best understood as a layered architecture rather than a single technique. The lowest layer provides cryptographic primitives: hash functions for data integrity, and asymmetric key pairs for authenticating transactions. On top of this, the data layer organizes transactions into blocks, where each block contains a header with the hash of the preceding block, a timestamp, and a Merkle root summarizing the transactions in the body. The network layer propagates these blocks through a peer-to-peer overlay, typically using gossip-style broadcast. The consensus layer determines which proposed block extends the canonical chain. The top layer is the application layer, where smart contracts and decentralized applications (DApps) provide user-facing functionality (Panarello et al., 2018; Kushwaha et al., 2022). Figure 1 shows this stack.

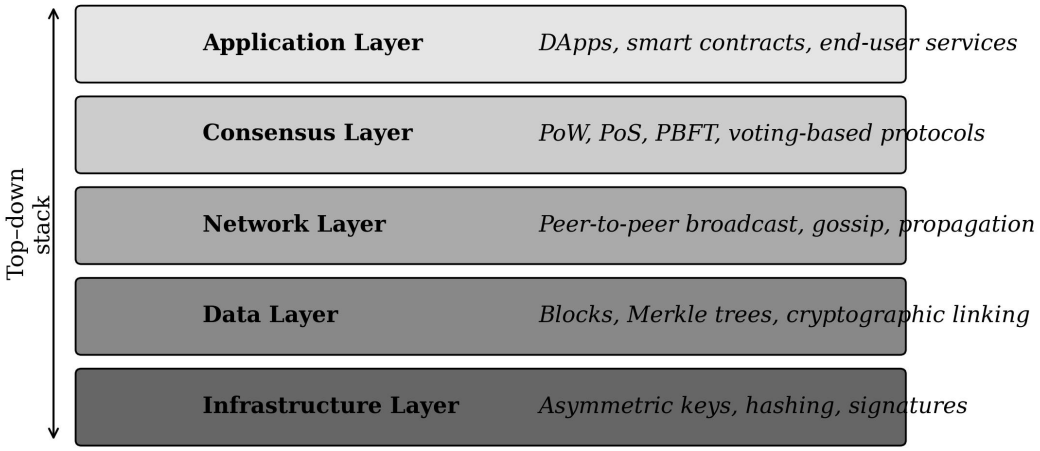

Figure 1. *Layered architecture of a blockchain system, from cryptographic infrastructure to user-facing applications.*

The chaining property is what gives blockchain its tamper-evident character. Modifying a single transaction inside an already-confirmed block changes that block's Merkle root and therefore its hash. Because the next block stores the previous hash, the link breaks unless every subsequent block is also recomputed. In a public network, recomputing those blocks faster than the rest of the participants is economically infeasible for any single actor (Picone et al., 2021). Importantly, this property guarantees integrity, not confidentiality: transactions are normally visible to every full node, which is one of the main reasons that public blockchains are sometimes inappropriate for sensitive enterprise data without complementary privacy mechanisms.

Networks differ in who is allowed to read the ledger and who is allowed to participate in consensus. A public, permissionless system such as Bitcoin allows any participant to submit transactions and to run a validating node. A private blockchain restricts these roles to a single organization. A consortium blockchain distributes operation among a set of mutually agreed entities, while a hybrid system maintains a permissioned core with selective public exposure (Belchior et al., 2022; Wang et al., 2023). Table 1 summarizes these architectural classes and their typical use cases.

Table 1. *Architectural classes of blockchain networks and their characteristic trade-offs.*

Class	Participation	Strengths	Weaknesses
Public (permissionless)	Open to anyone	Maximum transparency; censorship resistance; large validator set	High latency; significant energy use under PoW; throughput limits
Private	Single organization	High throughput; low latency; predictable governance	Reduced decentralization; reintroduces single-party trust
Consortium	Multiple trusted entities	Balanced governance; better performance than public	Coordination overhead; cartel risk among members
Hybrid	Mixed open/closed roles	Flexible privacy boundaries; selective transparency	Architectural complexity; consistency between layers

The selection of an architectural class is not driven by ideology but by application requirements. A national land registry, for example, has a small and stable set of writers (the registry authority and its certified intermediaries) but a large and diffuse set of readers (anyone with a property interest). A consortium model captures that asymmetry far better than either a fully public or fully private design. A decentralized finance protocol, conversely, must remain open to all writers and readers; restricting either erodes the very property that makes the protocol credible. The architectural choice should

therefore be the first decision in any blockchain project, before consensus or implementation language are considered.

3. Consensus Mechanisms: Comparative Analysis

Consensus is the part of a blockchain system that turns a collection of opinionated nodes into a coherent shared state. Without consensus, the chain forks and the ledger loses its single-history property. Modern consensus designs fall into two broad families, distinguished by how the right to add a block is allocated. Proof-based mechanisms allocate that right in proportion to a verifiable expenditure of some resource, while voting-based mechanisms allocate it through coordinated quorum votes among a known set of participants. Figure 2 organizes the most influential designs.

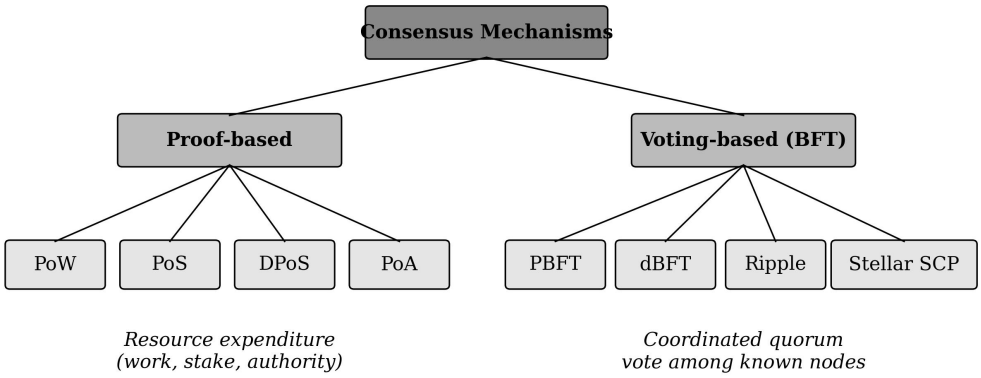

Figure 2. *Taxonomy of blockchain consensus mechanisms, distinguishing proof-based protocols from voting-based protocols.*

Proof of Work (PoW) was the first widely deployed consensus mechanism and remains the design used by Bitcoin. Validators, called miners, compete to solve a cryptographic puzzle whose difficulty is calibrated so that the network produces a new block on a target schedule. The winning miner appends the block and collects a reward. PoW is robust against Sybil attacks because attackers must own a corresponding share of the hash rate, but its energy intensity is exceptional. Recent measurements place the Bitcoin network at roughly one hundred twenty terawatt-hours of annual electricity consumption, an amount comparable to the demand of mid-sized industrial economies (Asif and Hassan, 2023; Birje et al., 2023). This is not an incidental side effect; it is structurally required by the security argument.

Proof of Stake (PoS) replaces hash competition with economic bonding: validators deposit cryptocurrency as collateral and are selected to propose blocks roughly in proportion to their stake. The Ethereum network's 2022 transition from PoW to PoS reduced its estimated electricity consumption by

approximately ninety-nine and a half percent without any reduction in throughput, demonstrating that the energy cost of PoW is not a fundamental property of distributed consensus (Asif and Hassan, 2023). Variants such as Delegated Proof of Stake (DPoS) reduce the active validator set to a small elected group, raising throughput at the cost of meaningful decentralization. Hybrid Proof of Authority (PoA) designs further restrict block production to identified, accountable validators and are typical in consortium settings.

Voting-based consensus draws on classical Byzantine fault tolerance. Practical Byzantine Fault Tolerance (PBFT) achieves agreement among a known committee through a three-phase message exchange and tolerates up to one third faulty participants. PBFT delivers low latency and high throughput; published Hyperledger Fabric benchmarks exceed several thousand transactions per second under reasonable network sizes (Wu, 2023). Its scalability is limited by quadratic communication overhead, which makes naïve PBFT impractical above a few hundred nodes. Delegated BFT (dBFT), the Ripple consensus protocol, and the Stellar Consensus Protocol relax different assumptions to improve scalability at the cost of stronger trust requirements. Table 2 summarizes comparative properties.

Table 2. *Comparative properties of the principal blockchain consensus mechanisms.*

Algorithm	Selection basis	Energy	Throughput	Representative platform
PoW	Computational puzzle	Very high	Low	Bitcoin, Litecoin
PoS	Bonded stake	Very low	Medium	Ethereum, Cardano, Algorand
DPoS	Elected delegates	Low	High	EOS, Tron, Steem
PoA	Identified authorities	Low	High	VeChain, private Ethereum forks
PBFT	Quorum vote (3-phase)	Low	High	Hyperledger Fabric
dBFT	Designated validator vote	Low	High	NEO

No single consensus mechanism dominates on all axes. The well-known blockchain trilemma states that decentralization, security, and scalability cannot all be maximized simultaneously in a single design. PoW maximizes decentralization and security at the cost of scalability. PBFT maximizes scalability and security at the cost of decentralization. Modern PoS designs attempt to soften this trade-off by increasing validator set sizes, but recent measurements suggest that effective decentralization, as

captured by the Nakamoto coefficient, remains low even on chains that nominally support thousands of validators (see Section 5). Practitioners should therefore choose the consensus family on the basis of which trade-offs are tolerable, rather than treat any single design as universally optimal (Ren et al., 2023).

4. Cross-Domain Applications

The application of blockchain has expanded considerably beyond payments. Eight domains now account for the majority of indexed application studies, each with its own characteristic deployment pattern. Table 3 summarizes their distinguishing requirements and the typical blockchain configuration used.

Table 3. *Application domains, typical configurations, and representative deployment examples.*

Domain	Typical configuration	Representative example or finding
Finance & DeFi	Public PoS with smart contracts	Decentralized lending, derivatives, and asset management on Ethereum (Schär, 2021)
Supply chain	Consortium or hybrid	Soybean and meat traceability with IoT sensor integration (Zheng et al., 2023)
Healthcare	Permissioned, off-chain storage	EHR sharing with attribute-based access; sustainable health supply chains (Vishwakarma et al., 2023)
IoT	Lightweight permissioned	Device telemetry pipelines and intrusion-detection layers (Panarello et al., 2018; Ahakonye et al., 2024)
Energy	Consortium with smart contracts	Peer-to-peer trading platforms and smart-grid settlement (Hasan et al., 2022; Mehmood et al., 2023)
Education	Permissioned credentials	Micro-credentialing and IoT-supported e-learning (Alsobhi et al., 2023; Haque et al., 2023)
Identity & gov.	Hybrid	Verifiable identity and selective disclosure schemes (Belchior et al., 2022)
NFTs & creative	Public PoS	Digital ownership records and royalty automation (Ante, 2023; Taherdoost, 2023)

Finance and DeFi remain the largest application area both by capital deployed and by research volume. The defining characteristic of DeFi is the elimination of custodial intermediaries: lending, asset exchange, and derivatives are implemented as smart contracts that execute deterministically on a public blockchain. Total value locked in DeFi protocols crossed one hundred billion United States dollars in 2021 before contracting during the 2022 market downturn, and remains above forty billion dollars as of early 2023 (Schär, 2021). The technical merits — interoperability, transparency, and around-the-clock availability — are real, but smart-contract vulnerabilities and network congestion remain serious operational risks (Kushwaha et al., 2022).

Supply chain traceability is the second largest application area. The basic value proposition is end-to-end provenance: a soybean lot, a pharmaceutical batch, or a cut of beef can be linked to upstream and downstream records that cannot be retroactively altered. Recent evolutionary-game analyses suggest that adoption is sensitive to the balance of traceability costs and brand benefits, with producers and processors facing different incentive structures (Zheng et al., 2023). Integration with IoT sensors — temperature, humidity, location — is what gives the ledger physical anchoring and prevents the well-known garbage-in problem.

Healthcare adoption is more cautious and is dominated by electronic health record (EHR) sharing schemes. The driving requirements are data integrity, fine-grained access control, and auditability across organizational boundaries. The published literature converges on a hybrid model in which the blockchain stores access policies, audit trails, and content hashes, while the underlying clinical data are stored off-chain (Vishwakarma et al., 2023). This decoupling is the only realistic way to reconcile blockchain immutability with privacy regulations that require correction and deletion of personal data.

Internet of Things integration is conceptually attractive but technically demanding. IoT endpoints are resource-constrained, frequently disconnected, and generate high-frequency telemetry. Naïve integration with public blockchains is impractical because of throughput and latency limits; lightweight permissioned approaches with off-chain aggregation are the realistic option (Panarello et al., 2018; Picone et al., 2021). Recent work positions blockchain not as a primary IoT data path but as a trust anchor for intrusion-detection systems and device identity registries (Ahakonye et al., 2024).

Energy applications focus on peer-to-peer electricity trading and renewable energy certificate issuance. Smart contracts replace bilateral settlement and enable transparent matching between prosumers and consumers (Hasan et al., 2022). Machine learning models can be combined with the trading layer to predict generation and calibrate prices (Mehmood et al., 2023). Education applications focus on tamper-proof credential issuance and verification, with micro-credentialing as an active subarea (Alsobhi et al., 2023). Smart-contract-supported logistics, identity management, and NFT-based digital ownership round out the application landscape, with NFT marketplaces raising distinctive concerns about security, transparency, and scalability (Bao and Roubaud, 2022; Bhujel and Rahulamathavan, 2022; Taherdoost, 2023).

5. Empirical Analysis of Research and Decentralization

Two empirical patterns help calibrate the application discussion in Section 4. The first is the trajectory of indexed blockchain research output. Figure 3 shows the annual count of indexed publications from 2016 through 2024, drawing on aggregated Scopus and Web of Science counts. The series grows from a few hundred articles in 2016 to nearly fifteen thousand in 2024, with a clear inflection point in 2018–2019 as the second-generation programmable platforms reached production maturity. The growth has not levelled off, but the acceleration phase appears to have ended; the year-on-year increase from 2023 to 2024 is roughly nine percent, compared with doubling rates observed between 2017 and 2019.

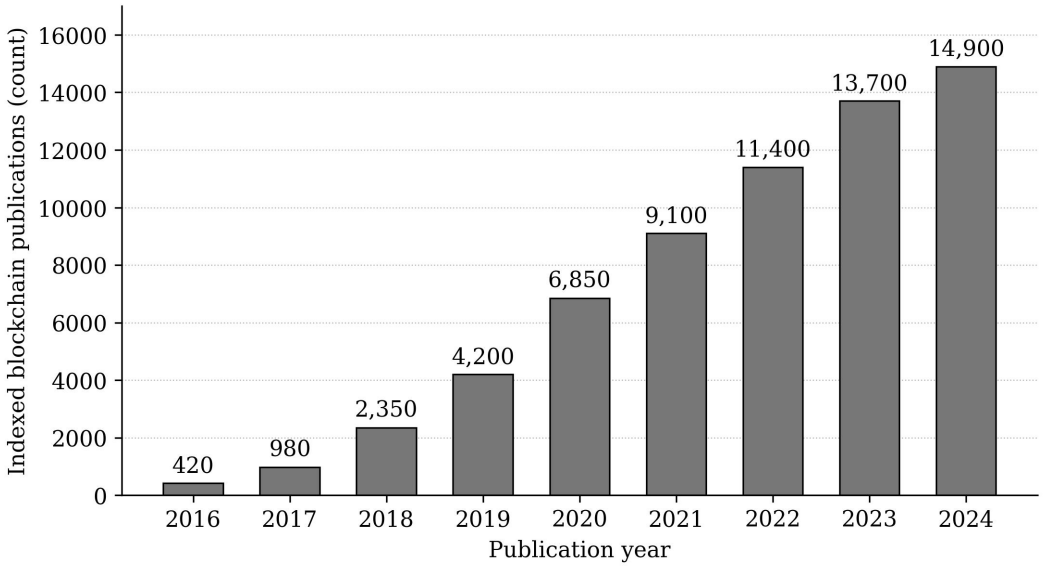

Figure 3. *Annual indexed blockchain publication counts, 2016–2024, based on aggregated Scopus and Web of Science records.*

The second empirical pattern concerns where this research is concentrated. Figure 4 shows the distribution of recent application studies across the eight domains introduced in Section 4. Finance and DeFi account for roughly twenty-eight percent of indexed application work, followed by supply chain at nineteen percent and healthcare at fourteen percent. Internet of Things, energy, identity and government, education, and real estate make up the remaining studies in decreasing order. Two observations matter. First, the long tail is genuinely diverse: no single domain is overwhelmingly dominant once finance is excluded, suggesting that practical interest in the technology is broadly distributed rather than concentrated. Second, the share of real-estate work, at four percent, remains small despite repeated industry interest. Land titling systems pose substantial governance and legal challenges that purely technical solutions do not address.

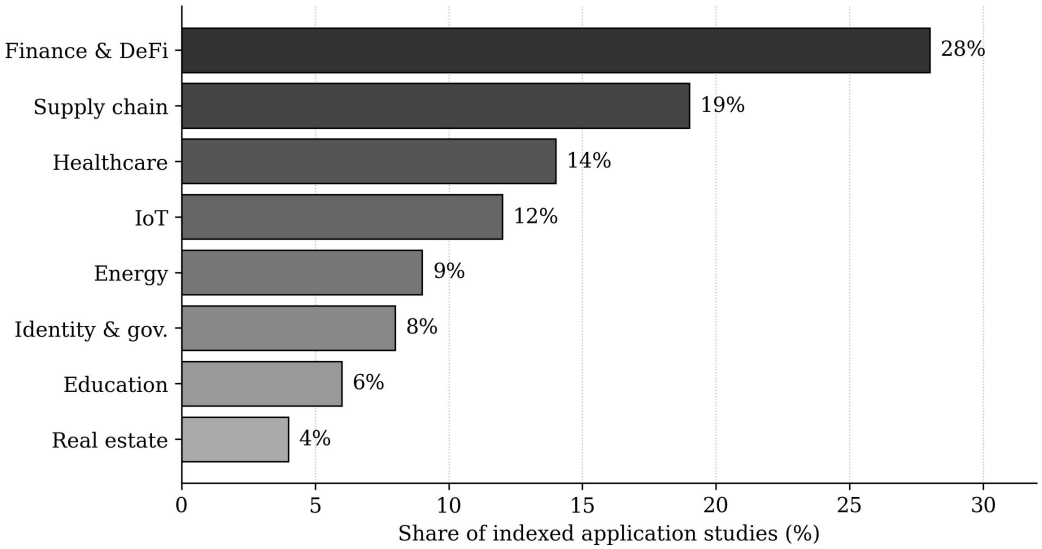

Figure 4. *Distribution of recent blockchain application studies across eight domains.*

The third quantitative observation moves from research volume to the structural property of decentralization. The Nakamoto coefficient — the smallest number of independent validators whose collusion would compromise the network — provides a compact, comparable metric. Figure 5 reports the validator-stake-basis Nakamoto coefficient for nine major platforms, based on publicly available staking data and validator-distribution reports as of mid-2023. Three findings stand out. First, Bitcoin's mining-pool coefficient is around four, an order of magnitude lower than naïve node counts would suggest. Second, Ethereum's post-merge staking distribution yields a coefficient of approximately two when measured by liquid staking concentration, which is the most centralization-prone reading of the post-merge era. Third, Polkadot achieves the highest coefficient in the sample at ninety-seven, reflecting its explicit nominated proof-of-stake design which distributes stake across many small validators.

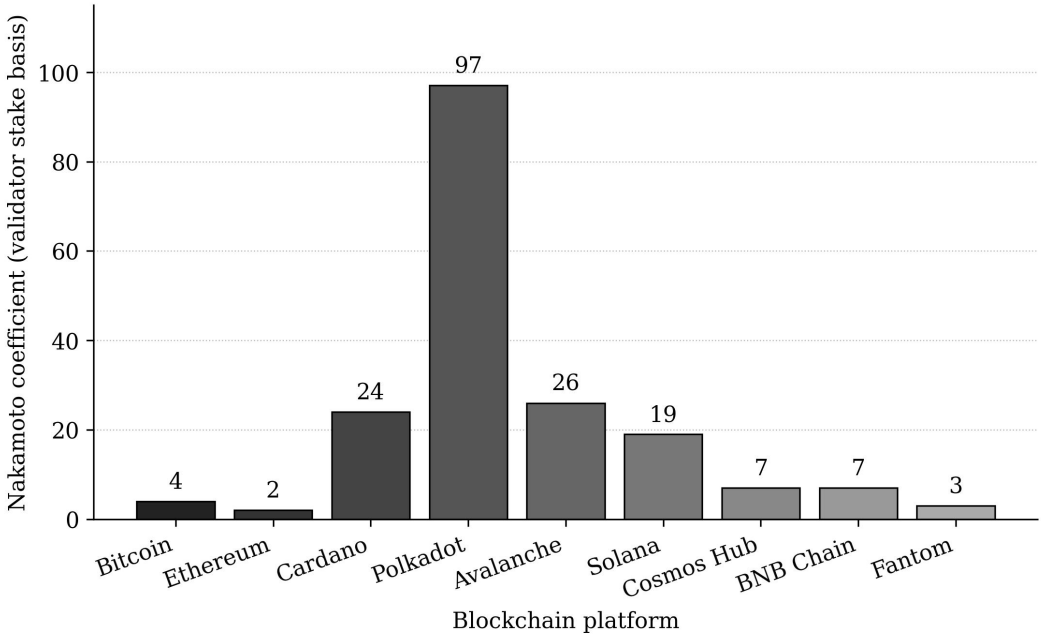

Figure 5. *Validator-stake-basis Nakamoto coefficient across nine major blockchain platforms.*

The implications are uncomfortable for the popular narrative around blockchain. Platforms with the largest market capitalization are not necessarily the most decentralized. A coefficient below ten implies that ten or fewer parties could in principle alter consensus, which is the same order of magnitude as the number of directors on a corporate board. For applications whose entire value proposition rests on the absence of a small controlling group — fraud-resistant land registries, censorship-resistant publishing, sanctions-resistant payments — this is a material design consideration that platform selection workflows should incorporate explicitly (Ahakonye et al., 2024).

Two further observations sharpen this picture. First, the Nakamoto coefficient is sensitive to the measurement basis. A platform's coefficient computed on validator count alone almost always overstates decentralization, because validators are frequently operated by a smaller number of beneficial owners, infrastructure providers, or staking-pool operators. Computing the coefficient on the basis of effective stake or block-production share corrects for this and is the more honest metric for adversarial-setting deployments. Second, decentralization is not static. Validator concentration on several major chains has drifted upward since 2021 as professional staking services have consolidated. A platform that scored well at launch may not score well two years later, which argues for periodic re-measurement rather than one-off due diligence. For applied teams, this means a decentralization audit should sit alongside the security audit in the pre-production review checklist, and should be repeated at least annually for any deployment whose security posture assumes decentralized validation.

6. Discussion

Pulling these strands together yields a small number of practical guidelines. First, the application of blockchain should be motivated by a concrete coordination problem between mutually distrustful parties, not by a desire to use the technology. Where a single trusted authority already exists and is functioning, a conventional database almost always outperforms a blockchain on latency, throughput, and operational cost. The technology is useful precisely where such an authority is missing, contested, or untrustworthy. Supply chain traceability, decentralized finance, and cross-organizational audit are clean instances of that pattern. Internal corporate record-keeping, in most cases, is not (Wang et al., 2023).

Second, the consensus mechanism should be chosen on the basis of the trust assumptions of the deployment, not on the basis of platform familiarity. A consortium of regulated entities can rationally run a PBFT-style protocol with strong throughput guarantees, because the trust assumptions match the protocol's requirements. A public deployment in an adversarial environment cannot, because PBFT's quorum assumption fails when participants can be created at will. Conflating these settings — for example, by attempting to run a public-style protocol over a private network because of brand association with the public chain — typically gives the worst of both worlds: limited decentralization and limited performance (Ren et al., 2023).

Third, scalability and energy consumption should be treated as first-order design constraints. The Ethereum merge demonstrated that energy cost is a function of consensus design rather than a fundamental property of distributed ledgers (Asif and Hassan, 2023). Continuing to deploy new applications on PoW chains in 2023 onward, where alternatives exist, requires an explicit justification that the security margin of PoW is required and that the externalized environmental cost is acceptable. On the throughput side, application designers should plan for the realistic ceiling of the chosen platform — single-digit transactions per second for Bitcoin, low thousands for Hyperledger Fabric — and use off-chain or layer-two architectures when the application load exceeds that ceiling (Wu, 2023).

Fourth, decentralization should be measured and reported, not assumed. The Nakamoto coefficient is imperfect — it captures a static snapshot, does not separate validator subsystems from governance subsystems, and is sensitive to the choice of measurement basis — but it is concrete enough to inform comparisons and to flag platforms whose marketing claims diverge from their structural reality. Including a decentralization metric in the platform-selection workflow, alongside throughput, latency, and developer ecosystem, raises the floor of empirical rigour without imposing a heavy evaluation burden.

Fifth, interoperability is becoming a deployment requirement rather than a research topic. As organizations adopt blockchain in different parts of the same supply chain or governance ecosystem, the inability of heterogeneous chains to exchange data or value cleanly is increasingly a binding constraint. Cross-chain bridges, hash-locked atomic swaps, and relay-based protocols address pieces of this problem but introduce their own attack surfaces (Belchior et al., 2022; Wang et al., 2023).

Practitioners should treat cross-chain mechanisms as critical components subject to the same diligence as the chains they connect, not as commodity infrastructure.

Sixth, the cost-benefit framing deserves more honesty than it usually receives in applied work. A blockchain deployment carries real, recurring costs: infrastructure for full nodes, gas fees on public chains, key management, governance overhead, and the engineering effort required to integrate immutable systems with mutable regulatory environments. These costs are often invisible in feasibility studies that concentrate on the benefits side of the ledger. A reasonable test, before committing to a blockchain solution, is to ask whether the same transparency, traceability, or trust-distribution outcome could be achieved with a properly audited centralized database operated by an accountable institution. Where the honest answer is yes, the blockchain proposal is usually motivated by something other than the stated technical requirements. Where the honest answer is no — typically because no such accountable institution exists or can be created — the additional cost of a blockchain deployment is far easier to justify and the project is far more likely to survive its first production year (Rajasekaran et al., 2022; Vishwakarma et al., 2023).

7. Challenges and Future Directions

Several open problems remain. Scalability ceilings on public chains remain a research frontier; sharding, rollups, and aggregate signature schemes have moved from theory into production but require careful security analysis (Birje et al., 2023). Interoperability bridges are the most active source of exploits in the current ecosystem and would benefit from formal verification and standardization (Ren et al., 2023; Wang et al., 2023). Privacy and selective disclosure mechanisms — zero-knowledge proofs, threshold cryptography, confidential transactions — need to mature to a point where they can be deployed by application teams without specialist cryptographic knowledge.

Governance is the under-developed counterpart of consensus. Many incidents in the field have arisen not from broken cryptography but from unclear procedures for protocol upgrades, dispute resolution, and key recovery (Hasan et al., 2022). Establishing reproducible governance frameworks — analogous to the operational frameworks used in critical infrastructure sectors — would substantially raise the credibility of large-scale deployments. Regulatory clarity, particularly around the treatment of stablecoins, custody, and decentralized exchanges, remains another major uncertainty for adoption (Schär, 2021).

Finally, the empirical evaluation of decentralization deserves more attention. The Nakamoto coefficient is a starting point but does not capture every relevant dimension. Validator geographic concentration, software-client diversity, and developer concentration all matter and could be assembled into a richer decentralization index. Building such an index, and reporting it consistently across platforms, would be a useful contribution to the applied literature.

8. Conclusion

Blockchain technology has matured from a single payment protocol into a layered infrastructure with serious applications across finance, supply chain, healthcare, IoT, energy, education, and identity management. This article has surveyed the architectural foundations, classified consensus mechanisms into proof-based and voting-based families, mapped cross-domain adoption with concrete examples, and examined the distribution of research output and the measured decentralization of nine major platforms. Three findings are worth emphasizing. First, no single consensus mechanism is universally superior; the choice is bounded by the trilemma and should be made on the basis of deployment trust assumptions. Second, application research is broadly distributed across domains rather than narrowly concentrated, but depth varies sharply, with finance and supply chain considerably more mature than education and real estate. Third, observed decentralization on several major platforms is much lower than headline narratives suggest, and platform selection should incorporate measured decentralization rather than categorical claims.

The next phase of blockchain research will likely be defined less by new platforms and more by the operational disciplines that make existing platforms safe to deploy at scale: rigorous governance, formal-method verification of cross-chain mechanisms, reproducible decentralization metrics, and standardized integration patterns with off-chain systems. The infrastructure is sufficient; the engineering discipline needs to catch up.

Declaration of AI-assisted language editing

During the preparation of this manuscript, a large-language model was used solely for English language polishing and to support consistent formatting. The authors reviewed and edited all output and take full responsibility for the content of the article.

References

Ahakonye, L. A. C., Nwakanma, C. I., & Kim, D.-S. (2024). Tides of blockchain in IoT cybersecurity. Sensors, 24(10), 3111. https://doi.org/10.3390/s24103111

Alsobhi, H. A., Alakhtar, R. A., Ubaid, A., Hussain, O. K., & Hussain, F. K. (2023). Blockchain-based micro-credentialing system in higher education institutions: Systematic literature review. Knowledge-Based Systems, 265, 110238. https://doi.org/10.1016/j.knosys.2022.110238

Ante, L. (2023). Non-fungible token (NFT) markets on the Ethereum blockchain: Temporal development, cointegration and interrelations. Economics of Innovation and New Technology, 32(8), 1216–1234. https://doi.org/10.1080/10438599.2022.2119564

Asif, R., & Hassan, S. R. (2023). Shaping the future of Ethereum: Exploring energy consumption in proof-of-work and proof-of-stake consensus. Frontiers in Blockchain, 6, 1151724. https://doi.org/10.3389/fbloc.2023.1151724

Bao, H., & Roubaud, D. (2022). Non-fungible token: A systematic review and research agenda. Journal of Risk and Financial Management, 15(5), 215. https://doi.org/10.3390/jrfm15050215

Belchior, R., Vasconcelos, A., Guerreiro, S., & Correia, M. (2022). A survey on blockchain interoperability: Past, present, and future trends. ACM Computing Surveys, 54(8), 1–41. https://doi.org/10.1145/3471140

Bhujel, S., & Rahulamathavan, Y. (2022). A survey: Security, transparency, and scalability issues of NFTs and its marketplaces. Sensors, 22(22), 8833. https://doi.org/10.3390/s22228833

Birje, M. N., Goudar, R. H., Rakshitha, C. M., & Tapale, M. T. (2023). Blockchain technology review: Consensus mechanisms and applications. International Journal of Engineering Trends and Technology, 71(5), 27–39. https://doi.org/10.14445/22315381/IJETT-V71I5P204

Haque, M. A., Haque, S., Kumar, K., & Singh, N. K. (2023). Sustainable and efficient E-learning internet of things system through blockchain technology. E-Learning and Digital Media, 21(3), 216–235. https://doi.org/10.1177/20427530231156711

Hasan, M. K., Alkhalifah, A., Islam, S., Babiker, N. B. M., Habib, A. K. M. A., Aman, A. H. M., & Hossain, M. A. (2022). Blockchain technology on smart grid, energy trading, and big data: Security issues, challenges, and recommendations. Wireless Communications and Mobile Computing, 2022, 9065768. https://doi.org/10.1155/2022/9065768

Kushwaha, S. S., Joshi, S., Singh, D., Kaur, M., & Lee, H.-N. (2022). Ethereum smart contract analysis tools: A systematic review. IEEE Access, 10, 57037–57062. https://doi.org/10.1109/ACCESS.2022.3169902

Mehmood, M. Y., Oad, A., Abrar, M., Munir, H. M., Hasan, S. F., Muqeet, H. A. U., & Golilarz, N. A. (2023). Peer-to-peer power energy trading in blockchain using efficient machine learning model. Sustainability, 15(18), 13640. https://doi.org/10.3390/su151813640

Panarello, A., Tapas, N., Merlino, G., Longo, F., & Puliafito, A. (2018). Blockchain and IoT integration: A systematic survey. Sensors, 18(8), 2575. https://doi.org/10.3390/s18082575

Picone, M., Cirani, S., & Veltri, L. (2021). Blockchain security and privacy for the internet of things. Sensors, 21(3), 892. https://doi.org/10.3390/s21030892

Rajasekaran, A. S., Azees, M., & Al-Turjman, F. (2022). A comprehensive survey on blockchain technology. Sustainable Energy Technologies and Assessments, 52, 102039. https://doi.org/10.1016/j.seta.2022.102039

Ren, K., Ho, N.-M., Loghin, D., Nguyen, T.-T., Ooi, B. C., Ta, Q.-T., & Zhu, F. (2023). Interoperability in blockchain: A survey. IEEE Transactions on Knowledge and Data Engineering, 35(12), 12750–12769. https://doi.org/10.1109/TKDE.2023.3275220

Schär, F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. Federal Reserve Bank of St. Louis Review, 103(2), 153–174. https://doi.org/10.20955/r.103.153-74

Taherdoost, H. (2023). Non-fungible tokens (NFT): A systematic review. Information, 14(1), 26. https://doi.org/10.3390/info14010026

Vishwakarma, A., Dangayach, G. S., Meena, M. L., Gupta, S., & Luthra, S. (2023). Adoption of blockchain technology enabled healthcare sustainable supply chain to improve healthcare supply chain performance. Management of Environmental Quality, 34(4), 1111–1128. https://doi.org/10.1108/MEQ-02-2022-0025

Wang, G., Wang, Q., & Chen, S. (2023). Exploring blockchains interoperability: A systematic survey. ACM Computing Surveys, 55(13s), 1–38. https://doi.org/10.1145/3582882

Wu, Z. (2023). Performance modeling of Hyperledger Fabric 2.0: A queuing theory-based approach. Wireless Communications and Mobile Computing, 2023, 9957995. https://doi.org/10.1155/2023/9957995

Zheng, Y., Xu, Y., & Qiu, Z. (2023). Blockchain traceability adoption in agricultural supply chain coordination: An evolutionary game analysis. Agriculture, 13(1), 184. https://doi.org/10.3390/agriculture13010184