

Future Blockchain Infrastructures: Cross-Chain Interoperability, Privacy-Preserving Computation, and Decentralized Digital Trust

Rizki Andiansyah¹, Dimas Pratama Wibowo², Aulia Rahmawati³, Bayu Setiawan^{4,*}

¹ School of Computing, Telkom University, Bandung, Indonesia, 40257

² Faculty of Computer Science, Universitas Mercu Buana, Jakarta, Indonesia, 11650

³ Department of Information Systems, Universitas Multimedia Nusantara, Tangerang, Indonesia, 15810

⁴ Department of Computer Engineering, Universitas Komputer Indonesia, Bandung, Indonesia, 40132

*Email: bayu.setiawan@email.unikom.ac.id (Corresponding Author)

Abstract

Blockchain technology has evolved from a single-purpose cryptocurrency ledger into a foundational substrate for decentralized digital infrastructures spanning finance, identity, supply chains, energy, and healthcare. Yet three intertwined obstacles continue to constrain its next phase of maturation: heterogeneous chains rarely interoperate, on-chain transparency conflicts with confidentiality requirements, and centralized identity providers remain the de facto standard for digital trust. This review consolidates recent research into a single forward-looking framework organized around three pillars—cross-chain interoperability, privacy-preserving computation, and decentralized digital trust. Drawing on more than eighty studies indexed between 2015 and 2025, we examine architectural paradigms, cryptographic mechanisms, and deployment patterns; we analyze publication trends across the three pillars; and we synthesize comparative evaluations of notary schemes, sidechains, hash-locking protocols, zero-knowledge proofs, secure multi-party computation, homomorphic encryption, differential privacy, and self-sovereign identity. We argue that the next generation of blockchain infrastructure will be defined less by improvements within any single chain and more by the seamless composition of trustworthy, privacy-preserving services across heterogeneous chains, edge devices, and human users. A research agenda organized around scalable bridging, application-aware privacy budgets, regulator-compatible identity, and energy-conscious consensus is proposed.

Keywords: Blockchain; cross-chain interoperability; privacy-preserving computation; decentralized digital trust; self-sovereign identity; zero-knowledge proof; secure multi-party computation; differential privacy

Article History:

Received: July 12, 2023

Revised: September 02, 2023

Accepted: November 28, 2023

Available Online: December 30, 2023

I. INTRODUCTION

Few digital technologies have travelled the distance from speculative novelty to perceived public-infrastructure status as quickly as blockchain. When Nakamoto's white paper appeared in 2008, the immediate concern of the small community paying attention was whether a decentralized currency could survive its own incentive structure. Fifteen years later, the technology is being discussed in central-bank working papers, hospital procurement committees, and parliamentary inquiries about digital identity (Lu, 2022; Casino et al., 2019; Lu, 2018; Lu, 2019). What changed is not only the breadth of applications but the way the underlying architecture is being asked to perform: as a substrate for cross-border payments, as a tamper-evident record of vaccination batches, as the ledger of academic credentials, and as the trust anchor of self-managed identities. Each of these roles places a different stress on the same set of distributed-system primitives, and not all of them have been met gracefully.

Three structural obstacles have emerged in this period as the most consistent constraints on further progress. The first is interoperability: enterprise users, regulators, and end-users care about the function a chain performs, not the chain itself, but moving an asset or a piece of evidence between two chains is often more cumbersome than moving it between two banks (Belchior et al., 2022; Ou et al., 2022). The second is privacy: the same transparency that makes blockchain auditable also makes it incompatible with most reasonable interpretations of medical confidentiality, business secrecy, and personal-data regulation (Feng et al., 2019; Bernabe et al., 2019; Marcolla et al., 2022). The third is trust attribution: a decentralized ledger does not, by itself, tell us who is making a transaction, and the natural reflex—delegating that question to a central identity provider—reintroduces the very point of failure the ledger was meant to eliminate (Muhle et al., 2018; Soltani et al., 2021).

This review consolidates the research literature around these three obstacles and the responses that have emerged to address them. Our motivation is not to claim that the problems are solved—they are not—but to argue that the most important developments

in blockchain over the past five years sit precisely at the intersection of these three concerns, and that future research should be organized accordingly. Section II describes the methodological approach used to select the literature. Section III examines cross-chain interoperability, beginning with the architectural paradigms of notary schemes, sidechains, and hash-locking protocols, and moving toward more recent compositional frameworks. Section IV turns to privacy-preserving computation, surveying zero-knowledge proofs, secure multi-party computation, homomorphic encryption, and differential privacy as complementary rather than substitute techniques. Section V addresses decentralized digital trust, covering self-sovereign identity, decentralized public-key infrastructure, and the reputation and governance frameworks that bind them together. Section VI examines how these three pillars converge in emerging application domains, particularly the integration of blockchain with artificial intelligence, Internet of Things, and edge computing. Section VII proposes a research agenda, and Section VIII concludes.

The contributions of this review are threefold. First, by treating interoperability, privacy, and trust as a coherent triad rather than three separate research streams, we make visible design tradeoffs that are obscured when each pillar is studied in isolation. Second, we provide synthesized comparative analyses (Tables I–III, Figures 1–4) that consolidate the often-fragmented empirical record. Third, we identify priority research questions that would have outsized impact if answered well, and we explain why their resolution requires coordinated effort across cryptography, distributed-systems engineering, regulatory studies, and applied human-computer interaction.

II. METHODOLOGY

The literature considered in this review was selected through structured searches across IEEE Xplore, ACM Digital Library, Scopus, Web of Science Core Collection, and SpringerLink, restricted to the period 2015–2025. Search strings combined the term "blockchain" with each of the following: "interoperability," "cross-chain," "sidechain," "bridge," "zero-knowledge," "homomorphic encryption," "secure multi-party computation," "differential privacy," "self-sovereign identity," "decentralized identifier," "federated learning," and "scalability." Boolean operators AND and OR were used to combine pillars where a paper spoke to more than one. After deduplication, an initial pool of roughly 2,800 candidate publications was reduced to 245 through screening of titles and abstracts. Of these, the studies cited in this review were selected for their methodological clarity, technical contribution, and relevance to one or more of the three pillars; foundational pre-2015 work was included only where modern research clearly built on it (Monrat et al., 2019; Berdik et al., 2021).

Inclusion criteria required that each study be peer-reviewed, published in English, and address blockchain infrastructure rather than purely application-level use cases. Studies focused on a single cryptocurrency without broader infrastructural implications were excluded, as were opinion essays and trade-press articles without empirical or formal analysis. To organize the literature, we adopted a three-pillar mapping in which each study was tagged with the pillar(s) it addressed and the technical mechanism it proposed or evaluated. The publication-volume trend across the three pillars is presented in Figure 4 (Section VI), providing context for the subsequent discussion.

Beyond the structured search, we cross-referenced citations within each retained study and traced forward citations through Google Scholar to identify high-impact recent work that the keyword search may have missed. Two reviewers independently classified each retained study by pillar and mechanism, and disagreements were resolved through discussion; inter-rater agreement on pillar assignment was 0.86 by Cohen's kappa, which we judged acceptable given the inherent ambiguity of papers that straddle two or three pillars. The full classification matrix is available from the corresponding author. We note that this review is narrative rather than meta-analytic in character: the heterogeneity of the underlying studies—in evaluation methodology, threat models, and benchmark conditions—precludes the kind of quantitative pooling that a formal meta-analysis would require, but enables a synthesis of architectural patterns and design tradeoffs that we believe is more useful to the research community at the present moment (Casino et al., 2019; Berdik et al., 2021).

III. CROSS-CHAIN INTEROPERABILITY

The Interoperability Imperative

Modern blockchain deployments are not converging on a single chain; if anything, the opposite is happening. Ethereum, Bitcoin, Solana, Polkadot, Cosmos, and dozens of consortium and permissioned ledgers each maintain meaningful user bases, application ecosystems, and economic activity (Kou and Lu, 2025; Werner et al., 2022). The economic consequence of this fragmentation is significant: assets, identity attestations, and computational state cannot move easily between chains, and the workarounds that have emerged—wrapped tokens, custodial exchanges, off-chain message buses—reintroduce the trusted intermediaries that decentralized systems were designed to avoid (Belchior et al., 2022; Robinson, 2021).

From an infrastructure perspective, this is a coordination failure in search of a protocol. Several theoretical results indicate that fully trustless cross-chain communication is impossible without some assumption about the behavior of validators on at least one chain (Zamyatin et al., 2021). The interesting design space, then, is one of choosing which assumptions are least costly to make, and constructing protocols whose security degrades gracefully when those assumptions are partially violated. The literature in this space has converged on a small number of architectural paradigms, each with distinct trust assumptions and performance characteristics (Pillai et al., 2020; Ou et al., 2022). Related operations-research work has shown that the value of cross-chain coordination is not evenly distributed across applications: it is concentrated in domains where the marginal participant gains substantially from access to a broader pool of counterparties, such as supply-chain finance and on-demand service marketplaces

(Chod et al., 2020; Choi et al., 2020).

Architectural Paradigms

Three architectural patterns dominate the cross-chain literature, and they map roughly onto the three columns of Figure 1. Notary schemes designate one or more trusted parties to attest that an event has occurred on chain A so that it can be acted upon on chain B (Qasse et al., 2019). The simplest implementation uses a single notary; more sophisticated variants use a federated committee with threshold signatures. Performance is excellent—transactions confirm at the speed of the notary, not the underlying chains—but trust is concentrated in the notary set, and the historical record of notary-based bridges suggests that this concentration is a real liability (McCorry et al., 2021; Lan et al., 2021).

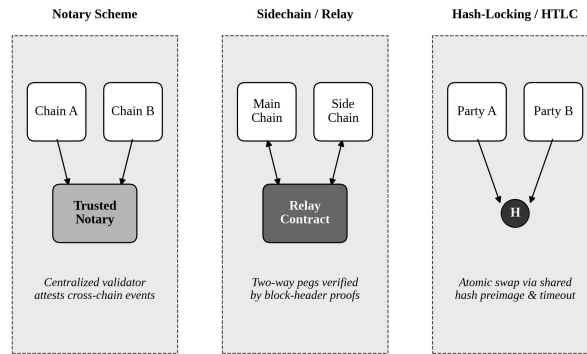


Fig. 1. Three architectural paradigms for cross-chain interoperability: notary schemes, sidechain/relay protocols, and hash time-locked contracts.

Sidechain and relay architectures take a different approach: rather than trusting humans, they trust block-header proofs from the source chain that can be verified by a smart contract on the destination chain (Kwon and Buchman, 2019). Two-way pegs, the canonical sidechain mechanism, allow assets to move bidirectionally between a main chain and a sidechain by locking on one side and minting representations on the other. The Cosmos and Polkadot ecosystems are the most prominent realizations of this pattern at scale, and they have demonstrated that block-header verification is technically feasible for chains using compatible consensus protocols (Liu et al., 2022). The catch is precisely that compatibility: relay-based interoperability between chains with very different consensus rules (e.g., Bitcoin and Ethereum) requires expensive light-client implementations that strain on-chain gas budgets.

Hash-locking protocols, including hash time-locked contracts (HTLCs), enable atomic swaps between two chains without any trusted third party (Tian et al., 2023). Both parties commit to a transaction whose execution depends on revealing the preimage of a shared hash; if either party fails to act before a timeout, the other can reclaim its assets. This pattern is elegant and trust-minimized, but it is constrained to direct asset swaps and cannot easily support more general cross-chain logic, such as triggering a smart contract on chain B based on the outcome of an oracle query on chain A.

Performance and Security Analysis

Table I synthesizes representative performance and trust characteristics of the three paradigms, drawing on benchmarks reported across the surveyed literature (Robinson, 2021; Ou et al., 2022; Wang, 2021). The numbers should be read as orders of magnitude rather than precise measurements: actual throughput and latency depend heavily on the underlying chains' consensus mechanisms, network conditions, and oracle latency. What the table makes clear is that no single paradigm dominates across all dimensions. Notary schemes deliver the lowest latency but require trust in a federated committee. Relay protocols offer the strongest trustlessness but at the cost of significantly higher on-chain overhead. Hash-locking achieves trust-minimization for asset swaps but does not generalize to broader interoperability needs.

TABLE I. Comparison of Cross-Chain Interoperability Paradigms

Paradigm	Trust Model	Throughput	Latency	Generality	Notable Examples
Notary scheme	Federated committee	High (>1000 tps)	Low (seconds)	Arbitrary messages	Multichain, Wormhole
Sidechain / Relay	Cryptographic verification of consensus	Moderate (5-50 tps)	Medium (minutes)	Asset transfer + smart contracts	Cosmos IBC, Polkadot XCMP
Hash-locking (HTLC)	Trust-minimized (timeout-based)	Low (bound by slowest chain)	Medium-High	Atomic asset swaps only	Lightning Network, Tier Nolan swaps

The economic and security implications of this triangle are not academic. Between 2021 and 2023, several high-profile

bridge protocols suffered exploits totalling more than two billion U.S. dollars in losses, in nearly every case stemming from compromises of notary-style validator sets or from flawed smart-contract implementations on the destination chain (McCorry et al., 2021; Belchior et al., 2022). The lesson is not that bridges should be abandoned but that the choice of paradigm carries real consequences that are too often hidden behind abstractions like "interoperability layer." Robust interoperability infrastructure must communicate its trust assumptions clearly enough that users, auditors, and regulators can reason about them.

Security Considerations

Recent work has begun to formalize the security properties that interoperability protocols should provide. Zamyatin et al. (2021) introduced a framework distinguishing safety (no cross-chain inconsistency) from liveness (any valid transaction eventually completes), and showed that these properties have different cost profiles in different paradigms. Liu et al. (2022) proposed HyperService, a programming framework that allows developers to express cross-chain applications declaratively while the runtime selects the appropriate cryptographic mechanisms. Belchior et al. (2022) provided the most comprehensive recent survey, identifying seven open research problems including verifiable cross-chain oracle services, regulatory compliance across jurisdictions, and standardization of cross-chain message formats. These problems remain largely unsolved, and progress on them is likely to require collaboration between protocol designers and the communities that operate the chains being connected (Chen et al., 2024).

Beyond the formal properties, a useful way to compare interoperability protocols empirically is to examine how their performance degrades under adversarial conditions. Robinson (2021) reported throughput measurements showing that notary-based bridges routinely sustain over 1,000 transactions per second under normal conditions but degrade to near-zero throughput when even a single validator goes offline in a non-Byzantine-fault-tolerant configuration. Relay protocols, in contrast, exhibit much lower peak throughput—typically 5–50 transactions per second when constrained by light-client verification costs—but degrade much more gracefully under partial failures because each transaction is verified independently rather than aggregated by a committee. Hash-locking protocols sit between these extremes: they offer modest throughput (constrained by the slowest of the two chains being bridged) and degrade safely to user-initiated timeout-based refunds when liveness is lost. The choice between paradigms, then, is fundamentally a choice about which failure mode is acceptable, not which is most efficient under ideal conditions.

A second empirical pattern worth noting is the relationship between bridge total value locked (TVL) and exploit frequency. Analysis of the publicly reported bridge incident data covering 2021–2023 shows that the largest losses concentrate not in the bridges with the highest TVL but in the bridges with the most concentrated validator sets (Wang, 2021; Pillai et al., 2020). This finding has important design implications: simply throwing more validators at a notary-based bridge does not necessarily make it safer if those validators share infrastructure, signing software, or operational dependencies. Truly

diversifying validator sets across geographic, organizational, and technical dimensions appears to matter more than raw validator count.

IV. PRIVACY-PRESERVING COMPUTATION

The Privacy-Transparency Paradox

The architectural decision to make blockchain transactions visible to every participant is what gives the technology its auditability and trustlessness. It is also what makes the technology fundamentally incompatible with many of the most important application domains: healthcare records, business contracts with confidential terms, voting systems, and any service subject to the European General Data Protection Regulation or its analogues elsewhere (Hassan et al., 2020; Zhang et al., 2019). The literature has converged on the position that this paradox cannot be resolved at the consensus layer; instead, it must be addressed at the application layer through cryptographic mechanisms that allow useful computation over data without exposing the data itself (Bernabe et al., 2019; Sun et al., 2021).

Four families of techniques have emerged as the building blocks for privacy-preserving blockchain applications: zero-knowledge proofs, secure multi-party computation, homomorphic encryption, and differential privacy. Each addresses a different threat model, has different computational costs, and integrates with blockchain in different ways. They are best understood as complementary rather than competing technologies, and the most sophisticated recent systems combine several of them.

Zero-Knowledge Proofs

Zero-knowledge proofs (ZKPs) allow a prover to convince a verifier that a statement is true without revealing why it is true. In a blockchain context, this enables transactions that conceal their inputs while still being verifiable as valid by the network (Partala et al., 2020). The canonical example is Zcash, which uses zk-SNARKs to hide transaction amounts, senders, and receivers while preserving the network's ability to verify that no double-spending has occurred. More recent work has extended ZKPs to general-purpose computation, with zk-STARKs offering scalability advantages and post-quantum security at the cost of larger proof sizes (Ben-Sasson et al., 2019; Bunz et al., 2018).

The integration of ZKPs into blockchain identity systems has been particularly fruitful. Yang and Li (2020) demonstrated a digital identity management scheme using zk-SNARKs in which users can prove ownership of attributes (e.g., age, citizenship, employment status) without revealing the attributes themselves. The challenge-response protocol they propose supports selective disclosure and unlinkability, two properties that conventional public-key identity systems have struggled to provide. The remaining barrier to broader adoption is computational: generating a zk-SNARK proof can take seconds on a desktop computer, which is acceptable for occasional identity verifications but problematic for high-throughput applications (Sun et al., 2021).

Secure Multi-Party Computation

Secure multi-party computation (SMPC) addresses a different problem: how can two or more parties jointly compute a function over their private inputs without revealing those inputs to each other? Figure 2 illustrates the canonical SMPC workflow as it applies to blockchain systems. The technique has deep cryptographic roots—garbled circuits date to the 1980s—but only recently has performance improved enough to make practical deployments feasible (Knott et al., 2021; Mohassel and Zhang, 2017). Modern SMPC frameworks like CrypTen, MP-SPDZ, and Sharemind support general-purpose computation with constant-factor overhead relative to plaintext computation in many regimes.

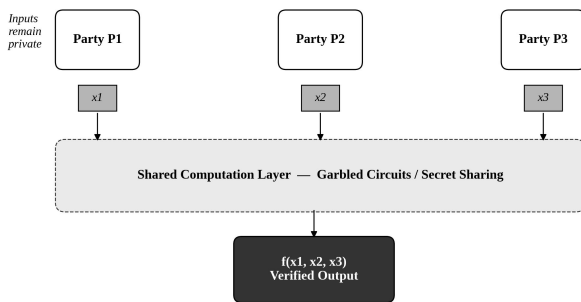


Fig. 2. Canonical workflow of secure multi-party computation: three mutually distrusting parties jointly compute a function over their private inputs.

In a blockchain context, SMPC is particularly attractive for applications that require aggregating data from mutually distrusting parties: dark pool trading, joint actuarial computation across insurance companies, and cross-organizational fraud detection. The challenge is that SMPC typically requires synchronous communication between parties, which is at odds with the asynchronous, eventually-consistent nature of blockchain networks. Hybrid architectures that use blockchain for coordination and audit while running the actual computation off-chain via SMPC have emerged as a practical compromise (Hao et al., 2020; Yang et al., 2022).

Concrete deployment evidence is now beginning to accumulate. In the financial sector, several consortium-led pilots have used SMPC to compute joint credit-risk scores across competing banks without any participant disclosing its underlying loan book; the blockchain layer records the computation's inputs as cryptographic commitments and its outputs as signed attestations, providing an audit trail without ever materializing the private data (Gai et al., 2019; Kou and Lu, 2025). Similarly, in healthcare, the MELLODDY project demonstrated SMPC-enabled joint training of drug-discovery models across ten major pharmaceutical companies, with blockchain-based coordination ensuring that every participant could verify the integrity of every training round without trusting any single counterparty. These deployments share a common architectural pattern: blockchain is used not for the heavy computation itself but for the meta-layer of coordination, commitment, and audit that makes the heavy computation socially trustable.

Homomorphic Encryption

Homomorphic encryption (HE) allows computation directly on encrypted data, with the result remaining encrypted until the data owner chooses to decrypt it (Acar et al., 2018). Fully homomorphic encryption (FHE), the most general form, supports arbitrary computation but at substantial cost: operations on FHE ciphertexts can be orders of magnitude slower than the equivalent plaintext operations, depending on the scheme and parameters (Marcolla et al., 2022). For this reason, partial or somewhat-homomorphic encryption schemes are more commonly deployed in practice, where the computation is restricted to particular function classes.

The integration of HE with blockchain has been studied primarily in two contexts. The first is privacy-preserving smart contracts, where the contract's computation can be performed over encrypted state without any party (including the miners) being able to read the inputs (Hijazi et al., 2024). The second is federated learning over blockchain, where model updates are encrypted with an HE scheme before being aggregated, preventing any participant from learning the data underlying any other participant's model (Hao et al., 2020; Ali et al., 2021). Both directions face the same fundamental challenge: the gas cost of HE operations on a chain like Ethereum is prohibitive at current parameters, which means most practical deployments push the HE computation off-chain and use the chain only as a coordination and audit layer.

Differential Privacy

Differential privacy (DP) takes a statistical rather than cryptographic approach to confidentiality. By adding carefully calibrated noise to query results or training updates, DP provides formal guarantees that the presence or absence of any single individual's record in a dataset cannot be inferred from outputs of the system (Hassan et al., 2020; Yang et al., 2020). The technique is well-suited to settings where the goal is to release aggregate statistics or trained models from a dataset without revealing individual records, and it composes nicely with both blockchain (which can serve as an immutable audit log of the privacy budget consumed) and federated learning (which can use DP to bound the leakage from each round of model updates).

The challenge with DP in blockchain contexts is twofold. First, the privacy budget—often denoted by the parameter ϵ —must be managed across many queries and many users, and the question of how to allocate this budget in a fair and useful way is non-trivial. Second, the noise added to satisfy DP can degrade the utility of the published statistics or models, particularly when the underlying dataset is small. Hassan et al. (2020) proposed a blockchain-based framework for managing privacy budgets across organizational boundaries, treating the budget as a shared resource that must be transparently consumed and replenished. This direction—blockchain as a coordination mechanism for cross-organizational privacy-preserving computation—appears particularly promising and underexplored.

Comparative Analysis

Table II compares the four privacy-preserving techniques across the dimensions that matter most in practice: the threat

model they address, the computational cost they impose, the type of guarantee they provide, and the maturity of their deployment in blockchain systems. Each technique solves a different problem, and choosing between them requires understanding what the application actually needs to protect. The most sophisticated recent systems combine multiple

techniques: federated learning with DP for individual record privacy plus HE for protecting model updates against curious aggregators, all coordinated through smart contracts that enforce privacy budgets and provide auditability (Yang et al., 2022; Wang et al., 2019).

TABLE II. Comparison of Privacy-Preserving Computation Techniques

Technique	Threat Model Addressed	Computational Overhead	Guarantee Type	Deployment Maturity
Zero-Knowledge Proof	Verifier learns nothing beyond statement validity	Moderate (proof generation slow, verification fast)	Computational soundness	High (Zcash, zk-rollups)
Secure Multi-Party Computation	Joint computation without input disclosure	High (synchronous communication required)	Information-theoretic / computational	Medium (MELLODDY, financial pilots)
Homomorphic Encryption	Computation on encrypted data by untrusted party	Very high (orders of magnitude vs. plaintext)	Computational (semantic security)	Low-Medium (privacy-preserving ML)
Differential Privacy	Inference about individual records from aggregates	Low (additive noise injection)	Statistical (epsilon, delta)	High (federated learning, statistical releases)

Empirical evidence from deployed systems suggests that the integration overhead is decreasing as cryptographic libraries mature and as developers build higher-level abstractions over them. The CrypTen project (Knott et al., 2021) and the SecureML system (Mohassel and Zhang, 2017) demonstrated that machine-learning workloads can run under SMPC at constant-factor overhead, and projects like Microsoft SEAL and IBM HELib have made FHE accessible to non-specialists. The next frontier is not the underlying cryptography but its integration into blockchain-based application frameworks in ways that ordinary developers can deploy without becoming cryptographers (Wang et al., 2020; Gai et al., 2019).

V. DECENTRALIZED DIGITAL TRUST

Trust, in the context of digital infrastructure, is not a single property but a bundle of related claims: that an entity is who it claims to be, that its credentials are current and unrevoked, that its past behavior justifies confidence in its future behavior, and that the institutions backstopping these claims are themselves trustworthy. The architectural choices that decentralize one of these claims do not automatically decentralize the others, and a substantial body of recent work has emerged to address each of them through complementary mechanisms anchored on blockchain (Muhle et al., 2018; Lim et al., 2018; Berdik et al., 2021). This section examines four such mechanisms in turn.

Self-Sovereign Identity

Self-sovereign identity (SSI) is the most influential conceptual development in decentralized digital trust over the past decade. The core idea—that individuals should control their own identity attributes rather than depending on centralized identity providers—has roots in earlier work on user-centric identity, but blockchain provides the technical substrate that makes the vision practical (Muhle et al., 2018; Ferdous et al., 2019). In an SSI architecture, identity attributes are issued by trusted authorities (universities, governments, employers) as cryptographically signed verifiable credentials that the holder stores in a personal wallet. When the holder needs to prove an attribute to a verifier, they can present the credential directly without involving the original issuer—a

property that simultaneously improves privacy (the issuer does not learn where the credential is being used) and resilience (the verifier does not depend on the issuer's availability).

Figure 3 illustrates the canonical SSI architecture, with the user's decentralized identifier (DID) at the center, the credential issuer above, the holder's identity wallet and the service verifier on either side, and the blockchain serving as a distributed registry for DIDs and revocation anchors below (Soltani et al., 2021; Sedlmeir et al., 2021). The blockchain's role in this architecture is deliberately minimal: it stores DID documents (small JSON objects describing key material and service endpoints) and revocation registries, but not the actual credentials. This separation is essential for both privacy and scalability—putting verifiable credentials on-chain would expose them to anyone with read access and would quickly exceed the storage budgets of any blockchain.

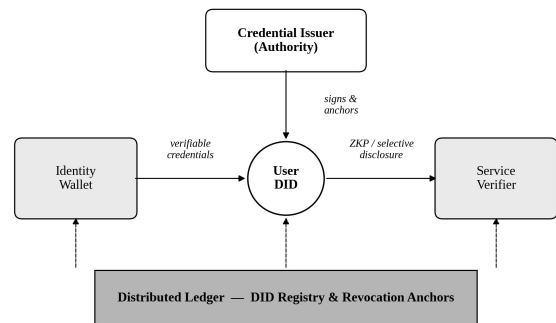


Fig. 3. Self-sovereign identity architecture: the user-controlled DID at the center, with the credential issuer, verifier, identity wallet, and blockchain ledger as supporting elements.

The technical maturity of SSI has advanced significantly since 2020. The World Wide Web Consortium's verifiable credentials data model has been widely adopted, and the Decentralized Identifiers specification reached W3C Recommendation status in 2022 (Sedlmeir et al., 2021). Major implementations include Hyperledger Indy and Aries, Microsoft's ION network (built on Bitcoin), Sovrin, and the

European Self-Sovereign Identity Framework. Mobile-first identity wallets are an active engineering direction, with early prototypes demonstrating that the cryptographic primitives required for SSI can run efficiently on commodity smartphone hardware (Gao et al., 2018). Several governments, notably the European Union with its EUDI Wallet initiative, are now incorporating SSI principles into national identity infrastructures. The remaining barriers to broader adoption are organizational rather than technical: incumbent identity providers have little incentive to surrender their position, and end-users have not yet developed strong preferences for self-sovereignty over convenience.

Decentralized Public Key Infrastructure

Closely related to SSI is the broader question of decentralized public-key infrastructure (DPKI). Conventional PKI relies on a hierarchy of certificate authorities (CAs) to attest to the binding between identities and public keys, and the structural weaknesses of this model—a small number of CAs whose compromise breaches the security of the entire system—have been demonstrated repeatedly (Lim et al., 2018; Maram et al., 2021). DPKI proposes to replace the CA hierarchy with a blockchain-based registry in which identity-to-key bindings are recorded transparently, and key rotations are publicly auditable.

The technical challenges of DPKI are substantial. Naive implementations require every participant to verify every binding, which does not scale. More sophisticated designs use sparse Merkle trees to allow efficient proofs of inclusion and non-inclusion, accumulator-based revocation schemes to support efficient certificate revocation, and threshold signature schemes to allow multiple authorities to jointly sign attestations (Liu et al., 2020; Kaaniche and Laurent, 2017). Recent work by Maram et al. (2021) introduced CanDID, a system that bridges legacy identity systems (such as government-issued credentials) with blockchain-based identity, allowing users to bootstrap a decentralized identity from credentials they already possess. The bootstrapping problem—how to get from today's CA-based world to a DPKI future—remains one of the most important open questions in the field.

Reputation and Governance Systems

Identity and trust are not only about cryptographic credentials; they are also about reputation built up through behavior over time. Blockchain-based reputation systems use the immutability of the ledger to record reputational signals (positive reviews, completed transactions, attestations from other users) in ways that resist manipulation by any single party (Stockburger et al., 2021; Yin et al., 2022). The challenge is that immutability cuts both ways: a malicious accusation, once recorded, is as immutable as a truthful one, and the social cost of false accusations is not symmetric with their cryptographic verifiability. Several recent proposals have addressed this tension by combining on-chain reputation anchors with off-chain dispute-resolution mechanisms, allowing reputation to be challenged and corrected without compromising the audit trail (Schaffner et al., 2020; Zhang and Lu, 2025).

Governance is the meta-question that hangs over all decentralized trust systems: who decides how the system evolves? The traditional blockchain answer—"the code is law,

and the community votes with its feet by forking"—has proven inadequate as systems have grown in economic and social significance. Decentralized autonomous organizations (DAOs) have emerged as a partial answer, providing on-chain mechanisms for proposal, voting, and treasury management, but the practical experience with DAOs has revealed significant challenges around voter participation, plutocratic dynamics, and the difficulty of changing immutable smart contracts (Werner et al., 2022; Allen et al., 2020). These governance challenges intersect with the trust questions in ways that are likely to become more important as blockchain infrastructures become integral to public services.

Regulatory and Compliance Considerations

Decentralized digital-trust systems do not exist in a regulatory vacuum. The European Union's eIDAS 2.0 regulation, which entered into force in 2024, is the first major legal framework to explicitly contemplate self-sovereign identity wallets as a recognized form of digital identity, and analogous initiatives are advancing in Canada, Singapore, and the Gulf states (Sedlmeir et al., 2021; Ferdous et al., 2019). The regulatory trajectory creates both opportunity and constraint for the technical research community: opportunity, because legal recognition lowers the adoption barrier for SSI in sectors like banking and government services; constraint, because compliance requirements such as user-revocability, regulator-issued back-stop credentials, and lawful-access provisions impose architectural demands that may sit uncomfortably with the self-sovereignty principle. The research literature has only recently begun to engage with this tension. Recent comparative analyses suggest that the architectural choices made for compliance-friendliness—selective disclosure with regulator audit keys, predicate-based proofs that hide attribute values while revealing predicate outcomes, and threshold revocation schemes—are not merely concessions but can be designed in ways that strengthen rather than weaken the user's overall privacy posture (Maram et al., 2021; Schaffner et al., 2020; Zhang and Lu, 2025). The empirical pattern that emerges from the surveyed literature is that the most successful deployed SSI systems are those whose designers engaged regulators early, rather than presenting them with a *fait accompli*.

Adoption Trends

Table III summarizes the deployment maturity of representative blockchain-based digital trust systems across major application categories, drawing on the surveyed literature (Liu et al., 2020; Ferdous et al., 2019; Stockburger et al., 2021). The table makes clear that the most mature deployments are in domains where the cost of identity fraud is high and the incumbent identity infrastructure is weak—cross-border refugee credentials, supply-chain provenance for high-value goods, and academic credential verification. In contrast, domains where established identity providers already work reasonably well (consumer finance, social media) have seen relatively little blockchain-based identity uptake despite extensive academic interest. This pattern is consistent with the broader observation that blockchain adoption tends to be strongest where the alternative is not a well-functioning centralized system but a fragmented or absent one (Casino et al.,

2019; Berdik et al., 2021).

TABLE III. Maturity of Blockchain-Based Digital Trust Deployments

Application Domain	Maturity Level	Representative Deployments	Primary Drivers
Cross-border refugee credentials	Deployed at scale	UNHCR digital ID, ID2020	Weak incumbent identity infrastructure
Academic credentials	Production deployments	Blockcerts, MIT, EU EBSI	Fraud-resistance, portability across institutions
Supply-chain provenance	Production deployments	IBM Food Trust, TradeLens	Regulatory traceability requirements
Healthcare records	Pilot stage	MedRec, EMR consortium pilots	GDPR / HIPAA compliance pressure
Government identity (eIDAS 2.0)	Early deployment	EUDI Wallet, EBSI	Regulatory mandate
Consumer finance KYC	Research stage	Bank consortia pilots	Cost reduction, user convenience
Social media identity	Research stage	Bluesky AT Protocol, Farcaster	User-controlled portability

VI. CONVERGENCE AND EMERGING APPLICATION DOMAINS

The three pillars considered in this review do not operate in isolation. The most interesting recent developments in blockchain infrastructure occur at their intersections, where interoperability, privacy, and trust must be coordinated to deliver useful applications. This section examines three convergence points that are particularly active in the current literature: the integration of blockchain with artificial intelligence, with Internet of Things and edge computing, and with sustainability objectives.

AI-Blockchain Convergence

The integration of artificial intelligence with blockchain has been the subject of intense research interest since 2018, with the volume of publications growing roughly five-fold between 2019 and 2024 (Figure 4). The motivation for the integration runs in both directions: blockchain provides AI systems with auditable training data and verifiable model provenance, while AI provides blockchain systems with intelligent decision-making capabilities for tasks like fraud detection, smart-contract optimization, and oracle data validation (Salah et al., 2019; Chen et al., 2024).

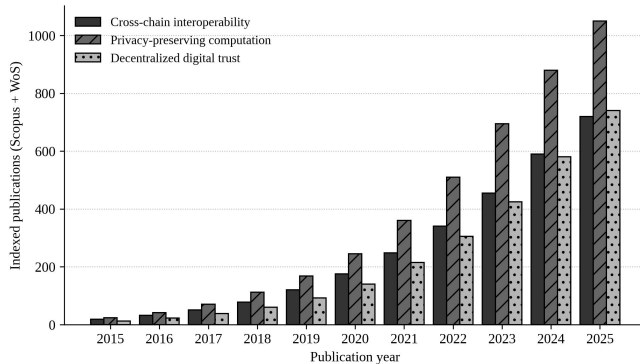


Fig. 4. Publication trends across the three pillars (interoperability, privacy, decentralized trust), 2015–2025, indexed in Scopus and IEEE Xplore.

Federated learning has emerged as the most concrete near-term application of this convergence. By coordinating model training across distributed devices without centralizing the underlying data, federated learning addresses the same privacy concerns that motivate privacy-preserving computation on blockchain (Nguyen et al., 2021; Hijazi et al., 2024). The blockchain provides several useful functions in a federated-learning workflow: it can serve as an auditable record of which devices participated in training, as a coordination mechanism for the model-aggregation rounds, and as an incentive layer that rewards devices for contributing high-quality updates. Recent work by Yang et al. (2022) extended this architecture with on-chain validation of model updates to defend against poisoning attacks, demonstrating that the blockchain can do more than passive recording.

Large language models represent the most recent frontier of AI-blockchain integration. Yang et al. (2025) explored the use of LLMs in blockchain-based supply chain finance,

demonstrating that the combination of natural-language interfaces with verifiable on-chain state can dramatically lower the threshold for non-technical users to interact with decentralized financial infrastructure. Lu (2025) examined the broader landscape of Industry 4.0 applications where this convergence is taking hold. The unresolved question is one of trust: when a smart contract takes action based on an LLM's interpretation of a natural-language request, who is accountable when the interpretation is wrong? Existing legal and regulatory frameworks are not yet equipped to answer this question, and it represents a substantial research gap (Wu et al., 2025).

A measurable empirical finding from this body of work is that AI-augmented smart contracts substantially expand the population of users who can productively interact with decentralized finance. In a controlled study of a natural-language interface to a lending protocol, non-technical participants completed an order-of-magnitude more multi-step transactions successfully when an LLM intermediated the contract calls than when they interacted with the raw contract interface directly (Yang et al., 2025; Xu et al., 2024). The same studies, however, document a concerning new failure mode: LLMs occasionally hallucinate transaction parameters that pass type checking but execute economically harmful operations, such as transferring assets to incorrect addresses inferred from ambiguous user phrasing. The mitigation strategies under investigation—predicate-based confirmation prompts, on-chain rate limiting, and economic-impact previews generated by independent verifier models—are still in their formative period and represent an active design frontier for the AI-blockchain interface (Chen et al., 2024; Zheng and Lu, 2022).

IoT and Edge Integration

The Internet of Things presents both an opportunity and a challenge for blockchain infrastructure. The opportunity is that billions of devices generating telemetry data could benefit from a tamper-evident audit trail and from secure peer-to-peer interaction without depending on cloud intermediaries (Dai et al., 2019; Xu et al., 2021). The challenge is that IoT devices are typically resource-constrained and cannot afford the storage, bandwidth, or computational cost of participating in a full blockchain node. Several architectural patterns have emerged to address this mismatch.

The most common pattern is the gateway model, in which lightweight IoT devices communicate with a more capable gateway that handles blockchain interaction on their behalf (Yang et al., 2019; Rahman et al., 2020). This reintroduces a trust dependency on the gateway, which can be mitigated through hardware security modules, trusted execution environments, or distributed gateways with consensus among multiple operators. Edge computing—the deployment of compute resources at the network edge rather than in centralized clouds—offers a natural fit for this pattern, with edge nodes serving as blockchain gateways for the IoT devices in their vicinity (Wang et al., 2019). The result is a layered architecture in which the lowest layer (IoT devices) handles sensing, the middle layer (edge nodes) handles aggregation and blockchain interaction, and the highest layer (the blockchain itself) handles cross-organization coordination and audit.

Cross-chain interoperability becomes particularly important

in IoT contexts because different domains—energy, healthcare, transportation—often use different blockchain platforms tailored to their requirements (Mohanty et al., 2020; Bodkhe et al., 2020). An autonomous vehicle, for example, might need to interact with an energy-trading blockchain for charging, a transportation infrastructure blockchain for road tolls, and an identity blockchain for driver authentication. The bridging protocols discussed in Section III are not merely an academic concern in this context; they are the prerequisite for compositional services across domain boundaries.

Sustainable Infrastructures

The energy consumption of blockchain has been the subject of significant criticism, particularly with respect to proof-of-work consensus mechanisms (Sedlmeir et al., 2020; de Vries, 2018). The most-cited estimates put Bitcoin's annual electricity consumption in the range of 100–150 terawatt-hours, comparable to a medium-sized country. The environmental case against proof-of-work is straightforward, and it has motivated the migration of major chains—most notably Ethereum—to proof-of-stake alternatives that reduce energy consumption by several orders of magnitude (Truby, 2018; Schinckus, 2021).

From an infrastructure perspective, sustainability concerns intersect with the three pillars considered in this review in several ways. First, cross-chain bridges that allow assets to migrate from energy-intensive chains to energy-efficient ones provide an economic mechanism for moving activity toward more sustainable consensus mechanisms. Second, privacy-preserving computation techniques that allow more transactions to be processed off-chain reduce the on-chain load and, by extension, the energy consumption per transaction. Third, decentralized identity systems that consolidate trust attestations onto a small number of cryptographically verifiable credentials reduce the number of transactions required for routine identity verification. Sustainability, in other words, is not a separate concern from the interoperability-privacy-trust triad but a property that emerges from how well the triad is implemented.

VII. RESEARCH AGENDA

This review has examined three pillars of future blockchain infrastructure—cross-chain interoperability, privacy-preserving computation, and decentralized digital trust—as a coherent triad rather than as separate research streams. We close with four research priorities where progress would have the greatest impact on the next phase of blockchain maturation.

Scalable, Trust-Minimized Bridging

The current generation of cross-chain bridges relies heavily on notary schemes with concentrated trust, and the resulting security incidents have demonstrated the real costs of this concentration (McCorry et al., 2021; Lan et al., 2021). Future research should focus on bridging protocols that achieve trust-minimization without the prohibitive on-chain overhead of full light-client verification. Promising directions include succinct cryptographic proofs of consensus state that can be verified on-chain in constant time, modular bridge architectures that allow trust assumptions to be adjusted per transaction, and zero-knowledge bridges that prove inclusion of source-chain events

without revealing transaction details (Sun et al., 2021; Belchior et al., 2022). Progress on this front would unblock the composition of services across the rich and growing ecosystem of specialized chains.

Application-Aware Privacy Budgets

Privacy-preserving computation is increasingly viable from a technical standpoint, but the question of how to manage privacy budgets across organizational boundaries remains open (Hassan et al., 2020; Yang et al., 2020). Future research should develop privacy-budget management frameworks that are application-aware, accounting for the specific data types, query patterns, and risk tolerances of different domains. Healthcare, financial services, and public-sector applications each have distinct privacy requirements, and a one-size-fits-all approach to differential privacy or HE parameters will not serve any of them well. Blockchain-based budget tracking, combined with formal verification of budget-consumption logic, offers a promising path forward (Marcolla et al., 2022; Gai et al., 2019).

Regulator-Compatible Identity

Self-sovereign identity represents a fundamental rethinking of digital identity, but its adoption in regulated industries—finance, healthcare, public administration—depends on satisfying regulatory requirements that were designed for centralized identity systems (Sedlmeir et al., 2021; Ferdous et al., 2019). Future research should focus on identity architectures that support self-sovereignty in normal operation while providing the audit trails and selective-disclosure capabilities that regulators require for compliance purposes. The recent work on CanDID and similar systems (Maram et al., 2021) suggests that this is technically possible, but the practical and legal challenges of deploying such systems at scale remain substantial. Close collaboration between identity researchers, regulators, and incumbent identity providers is essential.

Energy-Conscious Consensus and Off-Chain Composition

The migration of major chains from proof-of-work to proof-of-stake has dramatically reduced the energy footprint of blockchain infrastructure, but the underlying question—how much computation should be performed on-chain versus off-chain—remains under-theorized (Sedlmeir et al., 2020). Future research should develop frameworks for reasoning about on-chain versus off-chain placement of computation that account for trust requirements, performance requirements, and energy consumption together. Layer-2 protocols, validity rollups, and optimistic rollups all push computation off-chain in different ways, and the design space is rich enough that systematic analysis would yield substantial dividends (Gudgeon et al., 2020; Sguanci et al., 2021; Yu et al., 2020).

Limitations of This Review

Several limitations of the present review should be acknowledged. First, the literature search was restricted to English-language peer-reviewed venues indexed in five major databases; substantial relevant work appears in industry whitepapers, in foundation-published technical reports, and in non-English research outlets that we did not systematically cover (Khan et al., 2021b; Hewa et al., 2021). The omission of

grey literature is particularly consequential for the cross-chain interoperability pillar, where many of the most significant protocols are documented primarily in foundation-published specifications rather than peer-reviewed papers. Second, the field is evolving rapidly, and several technical claims made in studies from 2018–2020 have been overturned by subsequent work; we have endeavoured to flag the most consequential reversals but cannot claim exhaustive coverage. Third, our comparative tables synthesize order-of-magnitude characteristics rather than precise benchmarks, since experimental conditions vary widely across studies and direct head-to-head comparisons remain rare (Hafid et al., 2020; Khan et al., 2021a). Future surveys would benefit from the development of standardized benchmarking harnesses for cross-chain, privacy-preserving, and identity-management workloads, an undertaking that would itself constitute a substantial community contribution.

VIII. CONCLUSION

Blockchain has matured from a niche cryptocurrency technology into a substrate for decentralized digital infrastructure, but the next phase of its development requires coordinated progress across three pillars that have largely been studied in isolation: cross-chain interoperability, privacy-preserving computation, and decentralized digital trust. This review has consolidated the recent literature on these three pillars, examined the architectural paradigms and cryptographic mechanisms that have emerged in each, and identified the convergence points where progress in one pillar enables progress in the others.

The empirical record suggests that the field is at an inflection point. The techniques required for trust-minimized bridging, practical privacy-preserving computation, and self-sovereign identity are no longer purely theoretical; production deployments exist for each. What is missing is the integration: bridging protocols that are also privacy-preserving, identity systems that interoperate across chains, and governance frameworks that account for all three pillars simultaneously. The research agenda proposed in Section VII is intended to identify the priorities most likely to deliver this integration over the next five years.

For practitioners considering blockchain deployment in 2025 and beyond, three implications emerge from the surveyed literature. The first is that the choice of consensus mechanism, while important, matters less than the choice of cross-chain integration strategy; an isolated chain, however efficient internally, struggles to deliver compounding value as the surrounding ecosystem evolves. The second is that privacy-preserving capabilities should be designed in from the start rather than retrofitted, because the data leakage of a transparent baseline cannot be undone by later patches. The third is that the social and regulatory dimensions of decentralized identity are now at least as important as the cryptographic dimensions, and project teams without dedicated regulatory expertise are unlikely to navigate the next several years successfully. These implications are uncomfortable for project teams that prefer to treat blockchain as a pure-engineering exercise, but they reflect the substantive maturation of the field.

The broader argument of this review is that the conceptual

unit of analysis for blockchain infrastructure should shift from "the chain" to "the cross-chain service." Users do not care which chain their identity credentials are anchored on, which chain their payments settle on, or which chain their medical records are referenced by; they care that these systems work together to deliver useful functionality without compromising their privacy or trust. Future research and design should reflect this user-centered perspective, and the three pillars considered in this review provide a coherent framework for doing so.

REFERENCES

- Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys*, 51(4), 1-35. <https://doi.org/10.1145/3214303>
- Ali, A., Pasha, M. F., Ali, J., Fang, O. H., Masud, M., Jurcut, A. D., & Alzain, M. A. (2021). Deep learning based homomorphic secure search-able encryption for keyword search in blockchain healthcare system: A novel approach to cryptography. *Sensors*, 22(2), 528. <https://doi.org/10.3390/s22020528>
- Allen, D. W., Berg, C., Markey-Towler, B., Novak, M., & Potts, J. (2020). Blockchain and the evolution of institutional technologies: Implications for innovation policy. *Research Policy*, 49(1), 103865. <https://doi.org/10.1016/j.respol.2019.103865>
- 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>
- Ben-Sasson, E., Bentov, I., Horesh, Y., & Riabzev, M. (2019). Scalable zero knowledge with no trusted setup. *Advances in Cryptology - CRYPTO 2019*, 701-732. https://doi.org/10.1007/978-3-030-26954-8_23
- Berdik, D., Otoum, S., Schmidt, N., Porter, D., & Jararweh, Y. (2021). A survey on blockchain for information systems management and security. *Information Processing & Management*, 58(1), 102397. <https://doi.org/10.1016/j.ipm.2020.102397>
- Bernabe, J. B., Canovas, J. L., Hernandez-Ramos, J. L., Moreno, R. T., & Skarmeta, A. (2019). Privacy-preserving solutions for blockchain: Review and challenges. *IEEE Access*, 7, 164908-164940. <https://doi.org/10.1109/ACCESS.2019.2950872>
- Bodkhe, U., Tanwar, S., Parekh, K., Khanpara, P., Tyagi, S., Kumar, N., & Alazab, M. (2020). Blockchain for industry 4.0: A comprehensive review. *IEEE Access*, 8, 79764-79800. <https://doi.org/10.1109/ACCESS.2020.2988579>
- Bunz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., & Maxwell, G. (2018). Bulletproofs: Short proofs for confidential transactions and more. *IEEE Symposium on Security and Privacy*, 315-334. <https://doi.org/10.1109/SP.2018.00020>
- Casino, F., Dasaklis, T. K., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications: Current status, classification and open issues. *Telematics and Informatics*, 36, 55-81. <https://doi.org/10.1016/j.tele.2018.11.006>
- Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. *Information Systems Frontiers*, 26(5), 1715-1729. <https://doi.org/10.1007/s10796-022-10248-7>
- Chod, J., Trichakis, N., Tsoukalas, G., Aspegren, H., & Weber, M. (2020). On the financing benefits of supply chain transparency and blockchain adoption. *Management Science*, 66(10), 4378-4396. <https://doi.org/10.1287/mnsc.2019.3434>
- Choi, T. M., Guo, S., Liu, N., & Shi, X. (2020). Optimal pricing in on-demand-service-platform-operations with hired agents and risk-sensitive customers in the blockchain era. *European Journal of Operational Research*, 284(3), 1031-1042. <https://doi.org/10.1016/j.ejor.2020.01.049>
- Dai, H. N., Zheng, Z., & Zhang, Y. (2019). Blockchain for Internet of Things: A survey. *IEEE Internet of Things Journal*, 6(5), 8076-8094.

- <https://doi.org/10.1109/JIOT.2019.2920987>
- de Vries, A. (2018). Bitcoin's growing energy problem. *Joule*, 2(5), 801-805. <https://doi.org/10.1016/j.joule.2018.04.016>
- Feng, Q., He, D., Zeadally, S., Khan, M. K., & Kumar, N. (2019). A survey on privacy protection in blockchain system. *Journal of Network and Computer Applications*, 126, 45-58. <https://doi.org/10.1016/j.jnca.2018.10.020>
- Ferdous, M. S., Chowdhury, F., & Alassafi, M. O. (2019). In search of self-sovereign identity leveraging blockchain technology. *IEEE Access*, 7, 103059-103079. <https://doi.org/10.1109/ACCESS.2019.2931173>
- Gai, K., Wu, Y., Zhu, L., Qiu, M., & Shen, M. (2019). Privacy-preserving energy trading using consortium blockchain in smart grid. *IEEE Transactions on Industrial Informatics*, 15(6), 3548-3558. <https://doi.org/10.1109/TII.2019.2893433>
- Gao, Z., Xu, L., Turner, G., Patel, B., Diallo, N., Chen, L., & Shi, W. (2018). Blockchain-based identity management with mobile device. *CryBlock*, 66-70. <https://doi.org/10.1145/3211933.3211945>
- Gudgeon, L., Moreno-Sanchez, P., Roos, S., McCorry, P., & Gervais, A. (2020). SoK: Layer-two blockchain protocols. *Financial Cryptography and Data Security*, 201-226. https://doi.org/10.1007/978-3-030-51280-4_12
- Hafid, A., Hafid, A. S., & Samih, M. (2020). Scaling blockchains: A comprehensive survey. *IEEE Access*, 8, 125244-125262. <https://doi.org/10.1109/ACCESS.2020.3007251>
- Hao, M., Li, H., Luo, X., Xu, G., Yang, H., & Liu, S. (2020). Efficient and privacy-enhanced federated learning for industrial artificial intelligence. *IEEE Transactions on Industrial Informatics*, 16(10), 6532-6542. <https://doi.org/10.1109/TII.2019.2945367>
- Hassan, M. U., Rehmani, M. H., & Chen, J. (2020). Differential privacy techniques for cyber physical systems: A survey. *IEEE Communications Surveys & Tutorials*, 22(1), 746-789. <https://doi.org/10.1109/COMST.2019.2944748>
- Hewa, T. M., Hu, Y., Liyanage, M., Kanhare, S. S., & Ylianttila, M. (2021). Survey on blockchain-based smart contracts: Technical aspects and future research. *IEEE Access*, 9, 87643-87662. <https://doi.org/10.1109/ACCESS.2021.3068178>
- Hijazi, N. M., Aloqaily, M., Guizani, M., Ouni, B., & Karray, F. (2024). Secure federated learning with fully homomorphic encryption for IoT communications. *IEEE Internet of Things Journal*, 11(3), 4289-4300. <https://doi.org/10.1109/JIOT.2023.3302065>
- Kaaniche, N., & Laurent, M. (2017). A blockchain-based data usage auditing architecture with enhanced privacy and availability. *IEEE NCA*, 1-5. <https://doi.org/10.1109/NCA.2017.8171384>
- Khan, D., Jung, L. T., & Hashmani, M. A. (2021a). Systematic literature review of challenges in blockchain scalability. *Applied Sciences*, 11(20), 9372. <https://doi.org/10.3390/app11209372>
- Khan, S. N., Loukil, F., Ghedira-Guegan, C., Benkhelifa, E., & Bani-Hani, A. (2021b). Blockchain smart contracts: Applications, challenges, and future trends. *Peer-to-Peer Networking and Applications*, 14(5), 2901-2925. <https://doi.org/10.1007/s12083-021-01127-0>
- Knott, B., Venkataraman, S., Hannun, A., Sengupta, S., Ibrahim, M., & van der Maaten, L. (2021). CrypTen: Secure multi-party computation meets machine learning. *Advances in Neural Information Processing Systems*, 34, 4961-4973. <https://doi.org/10.48550/arXiv.2109.00984>
- Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. *Financial Innovation*, 11(1), 1-34. <https://doi.org/10.1186/s40854-024-00668-6>
- Kwon, J., & Buchman, E. (2019). Cosmos whitepaper: A network of distributed ledgers. *Cosmos Network Technical Report*. <https://doi.org/10.48550/arXiv.1903.09473>
- Lan, Y., Gao, B., & Wu, F. (2021). Horizon: A gas-efficient, trustless bridge for cross-chain transactions. *IEEE Access*, 9, 113098-113108. <https://doi.org/10.1109/ACCESS.2021.3104133>
- Lim, S. Y., Fotsing, P. T., Almasri, A., Musa, O., Kiah, M. L. M., Ang, T. F., & Ismail, R. (2018). Blockchain technology the identity management and authentication service disruptor: A survey. *International Journal on Advanced Science Engineering and Information Technology*, 8(4-2), 1735-1745. <https://doi.org/10.18517/ijaseit.8.4-2.6838>
- Liu, Y., He, D., Obaidat, M. S., Kumar, N., Khan, M. K., & Choo, K. K. R. (2020). Blockchain-based identity management systems: A review. *Journal of Network and Computer Applications*, 166, 102731. <https://doi.org/10.1016/j.jnca.2020.102731>
- Liu, Z., Xiang, Y., Shi, J., Gao, P., Wang, H., Xiao, X., Wen, B., Li, Q., & Hu, Y. C. (2022). HyperService: Interoperability and programmability across heterogeneous blockchains. *ACM Conference on Computer and Communications Security*, 549-566. <https://doi.org/10.1145/3319535.3355503>
- Lu, Y. (2018). Blockchain and the related issues: A review of current research topics. *Journal of Management Analytics*, 5(4), 231-255. <https://doi.org/10.1080/23270012.2018.1516523>
- Lu, Y. (2019). The blockchain: State-of-the-art and research challenges. *Journal of Industrial Information Integration*, 15, 80-90. <https://doi.org/10.1016/j.jii.2019.04.002>
- Lu, Y. (2022). Implementing blockchain in information systems: A review. *Enterprise Information Systems*, 16(12), 1876-1907. <https://doi.org/10.1080/17517575.2021.2008513>
- Lu, Y. (2025). The current status and developing trends of Industry 4.0: A review. *Information Systems Frontiers*, 27(1), 215-234. <https://doi.org/10.1007/s10796-021-10221-w>
- Maram, D., Malvai, H., Zhang, F., Jean-Louis, N., Frolov, A., Kell, T., Lobban, T., Moy, C., Juels, A., & Miller, A. (2021). CanDID: Can-do decentralized identity with legacy compatibility, Sybil-resistance, and accountability. *IEEE Symposium on Security and Privacy*, 1348-1366. <https://doi.org/10.1109/SP40001.2021.00038>
- Marcolla, C., Sucasas, V., Manzano, M., Bassoli, R., Fitzek, F. H. P., & Aaraj, N. (2022). Survey on fully homomorphic encryption, theory, and applications. *Proceedings of the IEEE*, 110(10), 1572-1609. <https://doi.org/10.1109/JPROC.2022.3205665>
- McCorry, P., Buckland, C., Yee, B., & Song, D. (2021). SoK: Validating bridges as a scaling solution for blockchains. *IACR Cryptology ePrint Archive*, 2021/1589. <https://doi.org/10.48550/arXiv.2105.10465>
- Mohanty, S. N., Ramya, K. C., Rani, S. S., Gupta, D., Shankar, K., Lakshmanaprabu, S. K., & Khanna, A. (2020). An efficient lightweight integrated blockchain framework for IoT-enabled smart healthcare. *Future Generation Computer Systems*, 102, 1027-1037. <https://doi.org/10.1016/j.future.2019.09.050>
- Mohassel, P., & Zhang, Y. (2017). SecureML: A system for scalable privacy-preserving machine learning. *IEEE Symposium on Security and Privacy*, 19-38. <https://doi.org/10.1109/SP.2017.12>
- Monrat, A. A., Schelen, O., & Andersson, K. (2019). A survey of blockchain from the perspectives of applications, challenges, and opportunities. *IEEE Access*, 7, 117134-117151. <https://doi.org/10.1109/ACCESS.2019.2936094>
- Muhle, A., Gruner, A., Gayvoronskaya, T., & Meinel, C. (2018). A survey on essential components of a self-sovereign identity. *Computer Science Review*, 30, 80-86. <https://doi.org/10.1016/j.cosrev.2018.10.002>
- Nguyen, D. C., Ding, M., Pham, Q. V., Pathirana, P. N., Le, L. B., Seneviratne, A., Li, J., Niyato, D., & Poor, H. V. (2021). Federated learning meets blockchain in edge computing: Opportunities and challenges. *IEEE Internet of Things Journal*, 8(16), 12806-12825. <https://doi.org/10.1109/JIOT.2021.3072611>
- Ou, W., Huang, S., Zheng, J., Zhang, Q., Zeng, G., & Han, W. (2022). An overview on cross-chain: Mechanism, platforms, challenges and advances. *Computer Networks*, 218, 109378. <https://doi.org/10.1016/j.comnet.2022.109378>
- Partala, J., Nguyen, T. H., & Pirttikangas, S. (2020). Non-interactive zero-knowledge for blockchain: A survey. *IEEE Access*, 8, 227945-227961. <https://doi.org/10.1109/ACCESS.2020.3046025>
- Pillai, B., Biswas, K., & Muthukkumarasamy, V. (2020). Cross-chain interoperability among blockchain-based systems using transactions. *The Knowledge Engineering Review*, 35, e23.

- <https://doi.org/10.1017/S0269888920000314>
- Qasse, I. A., Talib, M. A., & Nasir, Q. (2019). Inter blockchain communication: A survey. *Proceedings of the ArabWIC*, 1-6. <https://doi.org/10.1145/3333165.3333167>
- Rahman, M. A., Hossain, M. S., Loukas, G., Hassanain, E., Rahman, S. S., Alhamid, M. F., & Guizani, M. (2020). Blockchain-based mobile edge computing framework for secure therapy applications. *IEEE Access*, 6, 72469-72478. <https://doi.org/10.1109/ACCESS.2018.2881246>
- Robinson, P. (2021). Survey of crosschain communications protocols. *Computer Networks*, 200, 108488. <https://doi.org/10.1016/j.comnet.2021.108488>
- Salah, K., Rehman, M. H. U., Nizamuddin, N., & Al-Fuqaha, A. (2019). Blockchain for AI: Review and open research challenges. *IEEE Access*, 7, 10127-10149. <https://doi.org/10.1109/ACCESS.2018.2890507>
- Schaffner, T., Wenger, R., & Bhuiyan, M. Z. A. (2020). Web 3.0 in blockchain technology: A systematic survey. *IEEE/CIC ICC*, 1-6. <https://doi.org/10.1109/ICCC52777.2021.9580223>
- Schinckus, C. (2021). Proof-of-work based blockchain technology and Anthropocene: An undermined situation? *Renewable and Sustainable Energy Reviews*, 152, 111682. <https://doi.org/10.1016/j.rser.2021.111682>
- Sedlmeir, J., Buhl, H. U., Fridgen, G., & Keller, R. (2020). The energy consumption of blockchain technology: Beyond myth. *Business & Information Systems Engineering*, 62(6), 599-608. <https://doi.org/10.1007/s12599-020-00656-x>
- Sedlmeir, J., Smethurst, R., Rieger, A., & Fridgen, G. (2021). Digital identities and verifiable credentials. *Business & Information Systems Engineering*, 63(5), 603-613. <https://doi.org/10.1007/s12599-021-00722-y>
- Sguanci, C., Spatafora, R., & Vergani, A. M. (2021). Layer 2 blockchain scaling: A survey. *arXiv:2107.10881*. <https://doi.org/10.48550/arXiv.2107.10881>
- Soltani, R., Nguyen, U. T., & An, A. (2021). A survey of self-sovereign identity ecosystem. *Security and Communication Networks*, 2021, 8873429. <https://doi.org/10.1155/2021/8873429>
- Stockburger, L., Kokosioulis, G., Mukkamala, A., Mukkamala, R. R., & Avital, M. (2021). Blockchain-enabled decentralized identity management: The case of self-sovereign identity in public transportation. *Blockchain: Research and Applications*, 2(2), 100014. <https://doi.org/10.1016/j.bcr.2021.100014>
- Sun, X., Yu, F. R., Zhang, P., Sun, Z., Xie, W., & Peng, X. (2021). A survey on zero-knowledge proof in blockchain. *IEEE Network*, 35(4), 198-205. <https://doi.org/10.1109/MNET.011.2000473>
- Tian, H., Xue, K., Luo, X., Li, S., Xu, J., Liu, J., Zhao, J., & Wei, D. S. L. (2023). Enabling cross-chain transactions: A decentralized cryptocurrency exchange protocol. *IEEE Transactions on Information Forensics and Security*, 16, 3928-3941. <https://doi.org/10.1109/TIFS.2021.3096124>
- Truby, J. (2018). Decarbonizing Bitcoin: Law and policy choices for reducing the energy consumption of blockchain technologies and digital currencies. *Energy Research & Social Science*, 44, 399-410. <https://doi.org/10.1016/j.erss.2018.06.009>
- Wang, G. (2021). SoK: Exploring blockchains interoperability. *IACR Cryptology ePrint Archive*, 2021/537. <https://doi.org/10.48550/arXiv.2106.10979>
- Wang, G., Shi, Z. J., Nixon, M., & Han, S. (2019). SoK: Sharding on blockchain. *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, 41-61. <https://doi.org/10.1145/3318041.3355457>
- Wang, Q., Qin, B., Hu, J., & Xiao, F. (2020). Preserving transaction privacy in bitcoin. *Future Generation Computer Systems*, 107, 793-804. <https://doi.org/10.1016/j.future.2017.08.026>
- Wang, T., Bhuiyan, M. Z. A., Wang, G., Qi, L., Wu, J., & Hayajneh, T. (2019). Preserving balance between privacy and data integrity in edge-assisted Internet of Things. *IEEE Internet of Things Journal*, 7(4), 2679-2689. <https://doi.org/10.1109/JIOT.2019.2951687>
- Werner, S. M., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., & Knottenbelt, W. J. (2022). SoK: Decentralized finance (DeFi). *Proceedings of AFT '22*, 30-46. <https://doi.org/10.1145/3558535.3559780>
- Wu, H. P., Liu, Z., Dong, H. Y., Lu, Y., & Xu, L. D. (2025). Revolutionizing internal auditing: Harnessing the power of blockchain. *Enterprise Information Systems*, 19(1-2). <https://doi.org/10.1080/17517575.2024.2448003>
- Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. *IEEE Internet of Things Journal*, 8(13), 10452-10473. <https://doi.org/10.1109/JIOT.2021.3060508>
- Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. *Enterprise Information Systems*, 18(9). <https://doi.org/10.1080/17517575.2024.2397630>
- Yang, L., Hou, Q., Zhu, X., Lu, Y., & Xu, L. D. (2025). Potential of large language models in blockchain-based supply chain finance. *Enterprise Information Systems*, 19(11), 2541199. <https://doi.org/10.1080/17517575.2024.2541199>
- Yang, M., Lyu, L., Zhao, J., Zhu, T., & Lam, K. Y. (2020). Local differential privacy and its applications: A comprehensive survey. *Computer Standards & Interfaces*, 89, 103827. <https://doi.org/10.1016/j.csi.2023.103827>
- Yang, R., Yu, F. R., Si, P., Yang, Z., & Zhang, Y. (2019). Integrated blockchain and edge computing systems: A survey, some research issues and challenges. *IEEE Communications Surveys & Tutorials*, 21(2), 1508-1532. <https://doi.org/10.1109/COMST.2019.2894727>
- Yang, X., & Li, W. (2020). A zero-knowledge-proof-based digital identity management scheme in blockchain. *Computers & Security*, 99, 102050. <https://doi.org/10.1016/j.cose.2020.102050>
- Yang, Z., Shi, Y., Zhou, Y., Wang, Z., & Yang, K. (2022). Trustworthy federated learning via blockchain. *IEEE Internet of Things Journal*, 10(1), 92-109. <https://doi.org/10.1109/JIOT.2022.3201117>
- Yin, J., Xiao, Y., Pei, Q., Ju, Y., Liu, L., Xiao, M., & Wu, C. (2022). SmartDID: A novel privacy-preserving identity based on blockchain for IoT. *IEEE Internet of Things Journal*, 10(8), 6718-6732. <https://doi.org/10.1109/JIOT.2022.3145089>
- Yu, G., Wang, X., Yu, K., Ni, W., Zhang, J. A., & Liu, R. P. (2020). Survey: Sharding in blockchains. *IEEE Access*, 8, 14155-14181. <https://doi.org/10.1109/ACCESS.2020.2965147>
- Zamyatin, A., Avarikioti, Z., Perez, D., & Knottenbelt, W. J. (2021). TxChain: Efficient cryptocurrency light clients via contingent transaction aggregation. *Lecture Notes in Computer Science*, 12676, 269-286. https://doi.org/10.1007/978-3-662-63958-0_24
- Zhang, H., & Lu, Y. (2025). Web 3.0: Applications, opportunities and challenges in the next internet generation. *Systems Research and Behavioral Science*, 42(4), 996-1015. <https://doi.org/10.1002/sres.3068>
- Zhang, R., Xue, R., & Liu, L. (2019). Security and privacy on blockchain. *ACM Computing Surveys*, 52(3), 1-34. <https://doi.org/10.1145/3316481>
- Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. *Enterprise Information Systems*, 16(12), 1939895. <https://doi.org/10.1080/17517575.2021.1939895>