

Programmable Trust Fabrics for Cross-Domain Digital Value Exchange

Ahmad Rizki Pratama^{1,*}, Dewi Anggraini Hidayat², Bambang Setiawan Wibowo³

¹ Department of Informatics Engineering, Faculty of Science and Technology, Universitas Jambi, Jambi, Indonesia

² Department of Computer Science, Faculty of Engineering, Universitas Bengkulu, Bengkulu, Indonesia

³ Department of Information Systems, Faculty of Engineering, Universitas Mataram, Mataram, Indonesia

* Corresponding author: ahmad.pratama@unja.ac.id

Abstract

Distributed ledger technologies have matured into a heterogeneous landscape of public, permissioned, and hybrid ledgers, each engineered for a distinct trust regime. The practical consequence is that value, once committed to one ledger, is hard to move to another without sacrificing the very guarantees that motivated its commitment. This paper develops the concept of a programmable trust fabric, an explicitly layered abstraction in which cryptographic substrates, consensus and settlement, interoperability mechanisms, programmability, and application logic are treated as independently engineered yet composable layers. We argue that this framing dissolves a number of debates that have proceeded in parallel within the blockchain literature — public versus permissioned, scalability versus decentralisation, on-chain versus off-chain — by relocating them as design choices within layers rather than properties of whole systems. We synthesise published benchmarks and field reports to characterise the performance, security, and decentralisation trade-offs that arise at each layer, and we evaluate five families of cross-domain interoperability mechanisms — notary schemes, sidechains, hash-time-locked contracts, light-client relays, and federated bridges — along common axes. Application evidence drawn from cross-border finance, healthcare, supply chain, energy, and digital identity confirms that the fabric paradigm is operationally relevant rather than merely theoretical. We close by identifying open research questions on bridge security, scalability, privacy, regulation, and post-quantum migration that will define the next phase of the technology.

Keywords: Programmable trust; distributed ledger; cross-chain interoperability; smart contracts; atomic swaps; decentralised identity; layered architecture; digital value exchange

Article History:

Received: January 08, 2025

Revised: March 22, 2025

Accepted: May 11, 2025

Available Online: June 30, 2025

1. Introduction

Digital value exchange has been steadily reshaped by the proliferation of distributed ledger technologies (DLTs) over the past decade. What began as an experiment in peer-to-peer electronic cash has matured into a heterogeneous ecosystem of public, permissioned, and hybrid ledgers, each optimised for distinct trust requirements, throughput regimes, and regulatory contexts (Belchior et al., 2021; Yli-Huumo et al., 2016). As organisations transition from isolated proofs-of-concept to production deployments, a structural deficit has become apparent: each ledger constitutes a sovereign trust island, and the very properties that secure assets within a domain — finality rules, validator sets, cryptographic primitives — become barriers to value exchange across domains (Robinson, 2021; Herlihy, 2018). The notion of a single, universal ledger has receded; the realistic horizon is a world of many ledgers that must interoperate.

This paper develops the concept of a programmable trust fabric, a layered abstraction that explicitly separates the substrates that produce verifiable state, the mechanisms that translate state across administrative domains, and the application logic that consumes that state. The phrase “trust fabric” foregrounds a property that has been implicit in much of the blockchain literature but rarely treated as a first-class object of design: trust is not a binary attribute but a programmable resource, configurable along dimensions of who can read,

who can write, who can verify, and under what economic and legal constraints (Werbach, 2018; Lumineau et al., 2021). When a fabric is programmable, the conditions under which value moves between domains can themselves be encoded, verified, and audited, rather than relying on bilateral agreements, off-chain custodians, or legal escrows.

Three convergent forces motivate this reframing. First, the rise of decentralised finance (DeFi) has demonstrated that composability — the ability to chain together independently developed contracts into novel financial primitives — produces network effects that single-chain ecosystems cannot match (Werner et al., 2022; Schär, 2021). Composability across chains, however, remains immature, with cross-chain bridges accounting for a disproportionate share of catastrophic exploits in the period 2020–2024. Second, regulated industries such as healthcare, supply chain, and energy increasingly require interoperability not only among ledgers but between ledgers and traditional information systems (Mengelkamp et al., 2018; Esposito et al., 2018; Saberi et al., 2019). Third, sovereign digital currencies and tokenised securities are being issued on dedicated permissioned infrastructures whose value must nonetheless be reachable from open public networks (Schär, 2021; Hassani et al., 2018).

Despite considerable research on individual ledger architectures, on consensus design, and on specific cross-chain mechanisms, the literature has not converged on a coherent framework for reasoning about programmable trust across heterogeneous environments. Surveys of interoperability protocols catalogue mechanisms (notary schemes, sidechains, relays, hash-time-locked contracts, light-client bridges) but typically treat them as competing solutions rather than as composable building blocks within a unified stack (Belchior et al., 2021; Robinson, 2021). Studies of trust in blockchain either celebrate the “trustless” framing or critique it, but rarely connect the technical mechanics to the governance choices that operationalise trust in practice (Hawlitshchek et al., 2018; Beck et al., 2018). The present work attempts to bridge these strands.

Our analysis builds on a substantial body of foundational scholarship. Early systematic reviews mapped the breadth of blockchain research and identified maturity gaps that persist today (Yli-Huumo et al., 2016; Casino et al., 2019; Zheng et al., 2018). Concept papers articulated the programmability premise that underlies our fabric framing (Christidis & Devetsikiotis, 2016; Pilkington, 2016). Subsequent reviews placed blockchain within a wider technology-management agenda, examining drivers, barriers, and adoption pathways across industries (Risius & Spohrer, 2017; Frizzo-Barker et al., 2020; Hughes et al., 2019). Our contribution is to organise these strands around the cross-domain question, which earlier syntheses treated as one issue among many but which we contend is the defining question for the next decade of deployment.

Our contributions are fourfold. We first articulate a five-layer reference architecture for programmable trust fabrics that subsumes existing single-chain models as special cases. Second, we analyse the cross-domain interoperability mechanisms that connect these fabrics, providing a comparative evaluation of their security, latency, and decentralisation trade-offs. Third, we present an empirical synthesis drawn from published benchmarks, illustrating how design choices translate into measurable system behaviour. Fourth, we examine application domains where programmable trust fabrics have moved beyond conceptual proposals into operational deployments, and we identify the open research questions that remain. The remainder of the paper is organised as follows. Section 2 develops the conceptual foundations and introduces the layered model. Section 3 examines the architectural components that populate each layer. Section 4 surveys cross-domain interoperability mechanisms. Section 5 reports a comparative analysis. Section 6 discusses application scenarios. Section 7 enumerates challenges and future directions, and Section 8 concludes.

2. Conceptual Foundations: From Distributed Ledgers to Trust Fabrics

2.1 Programmable Trust as a Design Objective

Trust in distributed systems has historically been an organisational property: counterparties trusted each other, or they trusted a common intermediary, or they relied on legal recourse to compensate for the absence of trust. The decade since Nakamoto's proposal has demonstrated that trust can also be engineered through cryptographic primitives, replicated computation, and economic incentives (Bonneau et al., 2015; Tschorsch & Scheuermann, 2016). A programmable trust fabric inherits this engineering posture but generalises it: rather than fixing the trust assumptions of a system at design time, the fabric exposes them as configurable parameters that can be tightened or relaxed for specific transactions, domains, or counterparties.

Three properties distinguish programmable trust from earlier notions. First, verifiability: any party with appropriate cryptographic material can independently confirm that a stated transition occurred and was authorised by the rules in force (Kosba et al., 2016). Second, composability: the outputs of one trust-bearing operation can serve as inputs to another, even across domains, without re-establishing trust assumptions from scratch (Werner et al., 2022). Third, accountability: when something goes wrong, the fabric supports forensic reconstruction of the relevant state and decision history, a property increasingly demanded by regulators (Beck et al., 2018; Lumineau et al., 2021). These properties together transform trust from a static gating condition into a runtime resource that applications can request, combine, and reason about.

2.2 The Trust Fabric as a Layered Architecture

We propose a five-layer reference model for the programmable trust fabric, illustrated in Figure 1. From the bottom upward, the cryptographic substrate provides the primitives — hash functions, digital signatures, zero-knowledge proofs, and key-management facilities — on which all higher layers depend (Gervais et al., 2016). The consensus and settlement layer aggregates individual transitions into an ordered history with well-defined finality semantics, using proof-based or voting-based protocols depending on participant assumptions (Bamakan et al., 2020). The interoperability layer bridges domains, translating state and value across consensus boundaries through notary schemes, relays, hash-time-locked contracts, and atomic swaps. The programmability layer exposes the fabric to application developers via smart contracts, tokenisation primitives, and oracle interfaces (Atzei et al., 2017; Mendling et al., 2018). The application layer, finally, hosts decentralised applications, wallets, and cross-domain services that consume the lower layers' guarantees.

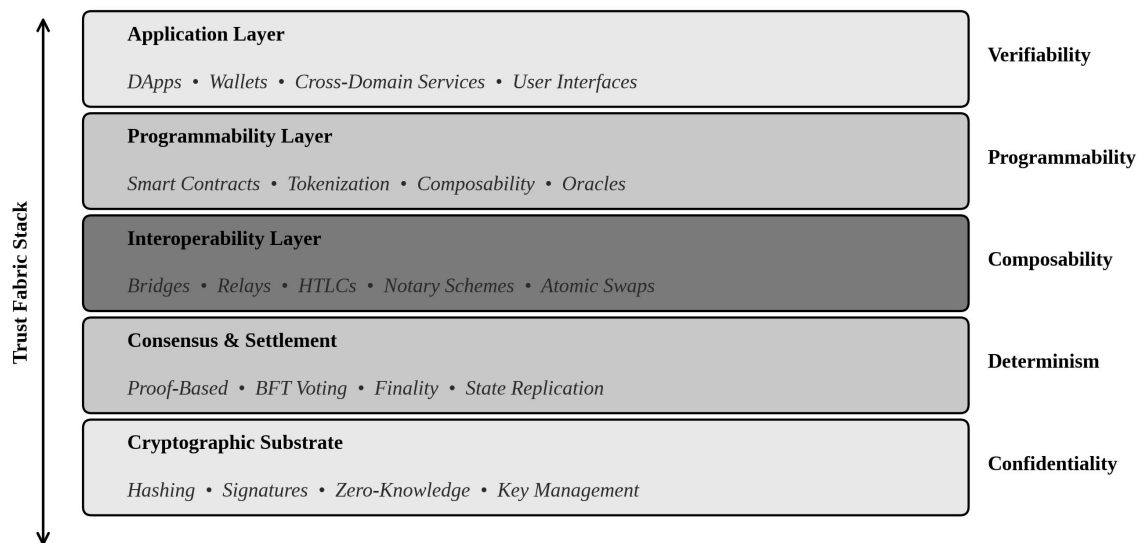


Figure 1. Five-layer reference architecture for a programmable trust fabric.

Several features of this model deserve emphasis. The layers are not strictly hierarchical in the sense of OSI-style encapsulation; in particular, the interoperability layer can interact directly with both the consensus layer (through light-client proofs) and the programmability layer (through contract-mediated bridges). The model is also fractal: each domain may itself instantiate the full five-layer stack, and a trust fabric composed of multiple domains is itself a trust fabric at a higher level of abstraction (Kannengießer et al., 2020). This fractality is what enables the fabric metaphor: trust is woven rather than stacked, with threads of varying provenance and strength interlacing to produce a composite material whose properties exceed those of any single ledger.

2.3 Properties and Trade-Offs

The properties of a trust fabric are subject to the familiar blockchain trilemma — security, scalability, and decentralisation — but the trilemma manifests differently across layers. The cryptographic substrate is largely immune to the trilemma: stronger primitives improve security without affecting scalability or decentralisation, though they may increase computational cost (Bonneau et al., 2015). The consensus layer is where the trilemma bites hardest, with proof-based protocols sacrificing throughput for decentralisation and voting-based protocols inverting that trade. The interoperability layer introduces a fourth dimension, trust translation cost, which captures the overhead of converting a state proof from one domain's format and assumptions to another's (Zamyatin et al., 2021; Robinson, 2021). The programmability layer faces a distinct expressiveness-versus-determinism trade-off, since richer contract languages enable more sophisticated logic but also broaden the attack surface (Atzei et al., 2017).

Table 1 summarises how three archetypal fabric configurations distribute these trade-offs. The public permissionless configuration maximises decentralisation and verifiability at the cost of throughput and regulatory alignment. The permissioned consortium configuration trades decentralisation for performance and regulatory tractability. The hybrid fabric, which embeds permissioned execution within a public substrate or vice versa, occupies a middle ground that has proven attractive for industries balancing public auditability with operational confidentiality (Esposito et al., 2018; Saberi et al., 2019).

Table 1. Comparative properties of three archetypal trust fabric configurations.

Property	Public permissionless	Permissioned consortium	Hybrid
Validator set	Open, dynamic	Closed, vetted	Tiered (open + vetted)
Throughput (typical)	10–100 TPS	1,000–10,000 TPS	100–5,000 TPS
Finality latency	Minutes (probabilistic)	Sub-second (deterministic)	Seconds (configurable)
Decentralisation	High	Low to moderate	Moderate
Regulatory tractability	Difficult	Straightforward	Negotiable per layer
Typical exemplar	Bitcoin, Ethereum L1	Hyperledger Fabric, Corda	Quorum, Polygon Edge

3. Architectural Components and Mechanisms

3.1 The Cryptographic Substrate

Every trust fabric rests on cryptographic primitives whose security properties cap the assurances higher layers can offer. Hash functions, particularly the SHA-2 and Keccak families, provide the collision and preimage resistance required to chain blocks immutably and to construct succinct commitments to large datasets via Merkle trees (Gervais et al., 2016; Garay et al., 2015). Digital signatures, predominantly elliptic-curve constructions such as ECDSA on the secp256k1 curve, provide authenticity and non-repudiation for transactions. Increasingly, fabrics incorporate advanced primitives — threshold signatures, BLS aggregation, and zero-knowledge proofs — to compress verification costs and to enable confidentiality features absent from earlier designs (Kosba et al., 2016).

The role of the substrate in cross-domain settings is subtle. Two domains may use incompatible curves or hash functions, in which case state proofs from one cannot be verified directly by the other without an adapter contract or a trusted relayer (Belchior et al., 2021). Standardisation efforts have produced common formats for proof transport, but the underlying assumption — that a domain's substrate is well-implemented and free of vulnerabilities — remains the foundation on which all cross-domain trust translation depends (Bonneau et al., 2015).

3.2 Consensus and Settlement

Above the substrate, consensus protocols transform individual transactions into a totally ordered history. Proof-of-work, the original consensus mechanism, achieves Sybil resistance by tying block production to verifiable computation but consumes substantial energy and offers probabilistic rather than deterministic finality (Eyal & Sirer, 2018; Gervais et al., 2016). Proof-of-stake variants substitute economic stake for computation and achieve faster, often deterministic finality, at the cost of more elaborate cryptoeconomic assumptions (Bamakan et al., 2020). Byzantine fault-tolerant protocols, including PBFT and its many descendants, provide immediate finality but require known validator sets and quadratic communication overhead, restricting their use to permissioned settings (Mendling et al., 2018; Pongnumkul et al., 2017).

From the trust fabric perspective, the choice of consensus protocol determines two cross-layer properties. The first is finality semantics: whether a transaction can be considered settled after a fixed time or only probabilistically, and how this propagates to applications that consume the chain's state. The second is validator transparency: whether the set of nodes responsible for state transitions is known and accountable, which directly influences regulatory alignment. Hybrid trust fabrics often deploy different consensus protocols in different sub-domains and reconcile finality semantics at the interoperability layer (Kannengießler et al., 2020; Belchior et al., 2021).

3.3 The Programmability Layer

Smart contracts elevate a ledger from a passive record-keeper to an active execution environment, enabling the encoding of arbitrary business logic that the network enforces (Atzei et al., 2017; Mendling et al., 2018). Ethereum's Virtual Machine and its Solidity language remain dominant, but the broader ecosystem includes WebAssembly-based runtimes, Move-language environments, and domain-specific languages targeting specific verticals (Mendling et al., 2018; Bamakan et al., 2020). The expressiveness of these environments is the source of their power and of their risk. Programmability enables tokenisation, automated market making, decentralised lending, and identity management; it also exposes the fabric to a class of vulnerabilities — reentrancy, integer overflow, oracle manipulation, governance attacks — that have produced the majority of high-profile losses in DeFi (Atzei et al., 2017; Werner et al., 2022; Khan & Salah, 2018).

In a programmable trust fabric, contracts are not confined to a single domain. Cross-domain contracts coordinate state across ledgers, often through bridge or relay infrastructure that itself runs under contract control. The composability that DeFi exploits within a single chain extends, in principle, across the entire fabric, though in practice it remains limited by the immaturity of standardised cross-chain messaging protocols and by the security profile of existing bridge implementations (Robinson, 2021; Belchior et al., 2021).

3.4 Identity, Authentication, and Privacy

A frequently overlooked component of trust fabrics is the identity layer that links cryptographic keys to off-chain entities. Pure pseudonymity, as in early Bitcoin, suffices for some applications but precludes others that require accountability, regulatory compliance, or selective disclosure (Kshetri, 2017; Salman et al., 2018). Self-sovereign identity frameworks, decentralised identifiers, and verifiable credentials provide a vocabulary for binding identity to keys without ceding control to a central registry (Liang et al., 2017). Privacy-preserving techniques — zero-knowledge proofs, ring signatures, confidential transactions — coexist with identity mechanisms, allowing fabrics to expose claims about an entity without disclosing the underlying data (Kosba et al., 2016).

In cross-domain settings, identity and privacy interact in non-trivial ways. A claim about an entity may be valid in one domain but require translation to be useful in another, particularly when the two domains operate under different regulatory frameworks. The fabric must therefore support not only cross-domain value transfer but cross-domain claim transfer, with auditable evidence of which claims were used to authorise which transitions (Ouaddah et al., 2016; Salman et al., 2018).

Figure 2 illustrates how these components interact in a representative cross-domain value transfer, with a programmable mediation layer translating between two heterogeneous domains. The mediation layer does not replicate the full state of either domain; rather, it consumes verifiable proofs of selected state transitions

and emits new transitions in the destination domain conditioned on those proofs.

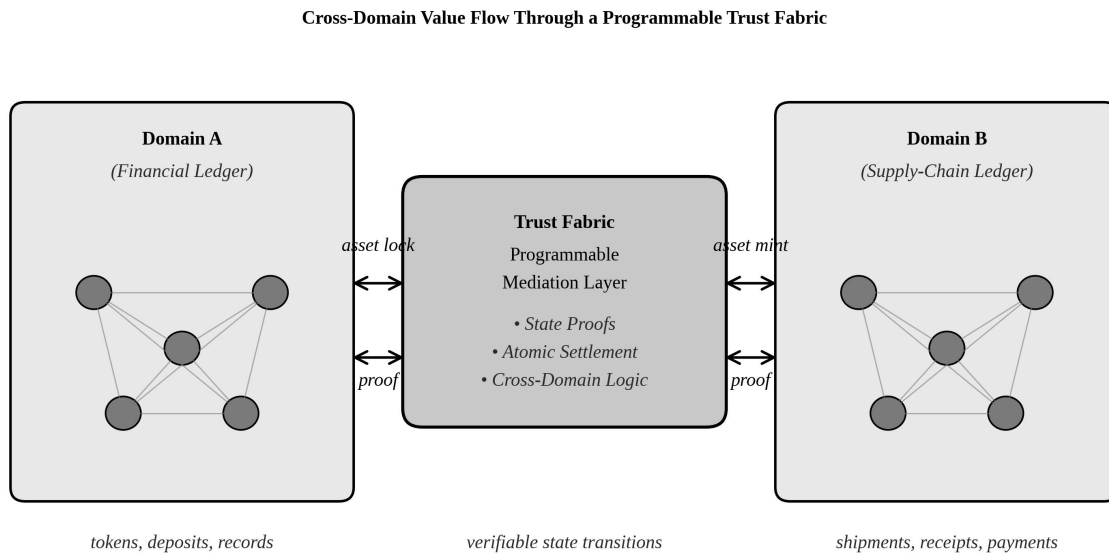


Figure 2. Cross-domain value flow mediated by a programmable trust fabric.

3.5 Substrate Patterns from the Internet-of-Things Literature

Cross-domain considerations are particularly acute when one of the domains is an Internet-of-Things (IoT) environment, where devices originate transactions but have limited capacity to verify chain state directly. A substantial literature has explored architectural patterns for integrating constrained devices with distributed ledgers (Christidis & Devetsikiotis, 2016; Conoscenti et al., 2016; Ali et al., 2018). Surveys covering edge-blockchain integration, gateway-mediated participation, and lightweight client protocols document recurring design points that recur in our fabric model (Dai et al., 2019; Yang et al., 2019; Fernández-Caramés & Fraga-Lamas, 2018; Lo et al., 2017). Specific protocol-level work on BLE-anchored gateways and constrained-device authentication extends these patterns to the device tier (Cha et al., 2018; Atlam et al., 2018). The recurring lesson is that the fabric's lower layers must accommodate participants that cannot run a full node, and that this accommodation is itself a cross-domain interoperability concern, since gateway-mediated participation effectively bridges an IoT sub-domain to a more capable ledger domain. The pattern language emerging from this work — relay, oracle, escrow, registry, witness — anticipates the abstractions we develop in Section 4 (Xu et al., 2018).

4. Cross-Domain Interoperability Mechanisms

Translating trust across domain boundaries is the defining engineering challenge of the trust fabric paradigm. Five mechanism families have emerged, each with distinctive trust assumptions and operational profiles: notary schemes, sidechain pegs, hash-time-locked contracts, relay-based light-client bridges, and federated validator bridges. We examine each in turn before discussing how they combine into composite designs.

4.1 Notary Schemes

Notary schemes designate a trusted party or quorum that observes events on one ledger and asserts their

occurrence on another. The notary need not understand the cryptographic semantics of either ledger; it only needs to be trusted to report faithfully (Belchior et al., 2021; Robinson, 2021). Centralised exchanges effectively act as notaries when they credit deposits and authorise withdrawals across ledgers. Multi-signature notary schemes distribute the trust across several entities, reducing single-point-of-failure risk but not eliminating the fundamental trust assumption.

Notary schemes are the simplest interoperability mechanism to deploy and remain the dominant pattern in production, particularly in custodial DeFi and in regulated cross-border payments (Schär, 2021; Hughes et al., 2019). Their weakness is precisely the trust they require: a compromised or coerced notary can fabricate transactions, and historical losses attributable to bridge compromises have largely involved notary-style designs (Belchior et al., 2021).

4.2 Sidechains and Pegged Assets

Sidechains are auxiliary ledgers connected to a primary chain through a two-way peg, enabling assets to be locked on one side and a corresponding representation issued on the other (Robinson, 2021). The economic peg can be enforced cryptographically through light-client verification, through a federation of validators that custody the locked assets, or through a hybrid arrangement that combines both. Sidechains were originally proposed as a scalability mechanism but have evolved into a general-purpose interoperability primitive, since the same machinery that pegs assets to a parent chain can peg them to any chain that supports the requisite verification logic.

From the fabric perspective, sidechains provide a useful decoupling: each sidechain can specialise in a particular workload — high throughput, privacy preservation, regulatory compliance — without imposing those constraints on the parent. Their drawback is the introduction of additional consensus surfaces that must be secured, and the dependency on the peg mechanism, whose failure can produce systemic losses across all assets bridged through it (Belchior et al., 2021).

4.3 Hash-Time-Locked Contracts and Atomic Swaps

Hash-time-locked contracts (HTLCs) implement bilateral exchanges between domains without requiring either party to trust the other or any third party (Herlihy, 2018). The mechanism leverages a shared cryptographic puzzle: one party publishes a hash commitment, both parties lock assets contingent on the puzzle being solved within a time window, and the secret is revealed in the act of claiming, simultaneously unlocking both sides of the trade. Figure 3 illustrates the canonical four-step exchange.

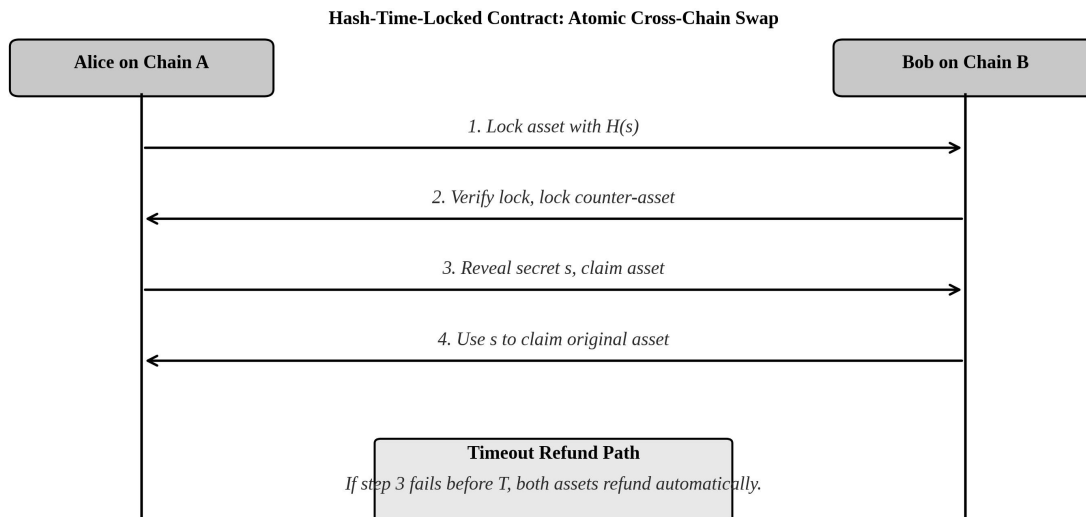


Figure 3. Hash-time-locked contract sequence for an atomic cross-chain swap.

Atomic swaps achieve genuine trustlessness — the protocol either completes correctly or refunds both parties — but they pay for this property with rigidity (Herlihy, 2018). HTLCs require both domains to support the same hash function and to provide adequate timelock primitives. They are inherently bilateral, scaling poorly to many-party exchanges. And the time windows must be calibrated to the slowest domain's finality semantics, which in cross-public-chain settings introduces minutes-to-hours of latency. HTLCs nonetheless remain a foundational building block, particularly in privacy-preserving exchanges and in payment-channel networks layered above public chains (Robinson, 2021).

4.4 Relay-Based Light-Client Bridges

Relay bridges generalise the verification approach of sidechain pegs to settings where the destination chain runs a light-client implementation of the source chain's consensus protocol (Belchior et al., 2021; Robinson, 2021). Block headers from the source chain are submitted to the destination, where a contract verifies their validity according to the source's rules. Once a header is accepted, Merkle proofs against it can authenticate arbitrary source-chain state to destination-chain contracts.

Relays push the trust assumptions of cross-domain communication as close as possible to the trust assumptions of the underlying chains. If the source chain's consensus is secure and the relay contract's verification logic is correct, the relay is secure. The catch is the computational cost of light-client verification: validating a chain of headers, especially for proof-of-work chains with non-trivial difficulty, can be prohibitively expensive in destination-chain gas units. Optimised constructions reduce this overhead through succinct proofs, but at the cost of additional cryptographic complexity (Belchior et al., 2021; Robinson, 2021).

4.5 Federated Bridges

Federated bridges occupy a middle ground between notary schemes and pure relays. A defined committee of validators observes the source chain and signs attestations that the destination chain accepts (Robinson, 2021). The committee can be made auditable, accountable, and economically incentivised through staking, reducing — though not eliminating — the trust required. Most production cross-chain bridges in operation today are federated, in part because the engineering effort to implement a full relay across heterogeneous consensus protocols remains high (Belchior et al., 2021).

4.6 Comparative Trade-Offs

Table 2 summarises the five families along security, performance, and deployability dimensions. No single mechanism dominates; the right choice depends on the asset values at risk, the latency tolerance of the application, and the regulatory context in which both domains operate. In practice, mature trust fabrics combine mechanisms: a federated bridge may carry routine traffic, while high-value cross-domain settlements use HTLCs or light-client relays. This composability — and the meta-design question of how to allocate traffic among mechanisms — has emerged as an active research area (Robinson, 2021).

Table 2. Cross-domain interoperability mechanisms compared along security, performance, and deployability dimensions.

Mechanism	Trust assumption	Latency	Throughput	Deployability
Notary scheme	Trusted operator	Low (sec)	High	Easy
Sidechain / 2-way peg	Federation or merge-mined	Moderate	Moderate to high	Moderate
Hash-time-locked contract (HTLC)	None (cryptographic)	Moderate to high	Low (per-pair)	Pairwise only
Light-client relay	Source-chain consensus	Moderate	Moderate	Engineering-intensive
Federated bridge	Bonded validator set	Low to moderate	High	Most common in practice

5. Comparative Analysis and Performance Evaluation

Having mapped the architectural space, we now turn to a comparative analysis that draws together quantitative observations from the published literature. Our aim is not to declare a single winner among mechanisms but to characterise the design envelope within which programmable trust fabrics must operate. The empirical landscape is fragmented — benchmarks use different workloads, network conditions, and hardware — but several robust patterns emerge from synthesising studies that share methodology (Pongnumkul et al., 2017; Bamakan et al., 2020; Gervais et al., 2016; Sankar, Sindhu, & Sethumadhavan, 2017). Broader technology-management surveys have likewise foregrounded the importance of cross-cutting evaluation as the technology moves from prototype to production (Ahram et al., 2017).

5.1 Throughput and Finality Latency

Figure 4 plots throughput against finality latency for six representative configurations spanning the design space. The figures are drawn from published benchmarks adjusted for comparability where possible, and they should be read as order-of-magnitude indicators rather than precise measurements. Three regimes are apparent. Public proof-of-work configurations, exemplified by PoW relay bridges between two public chains, exhibit the lowest throughput (single digits per second) and the longest finality latency (multiple block confirmations on both sides, totalling roughly ten seconds in our representative case). Permissioned BFT notary configurations achieve three orders of magnitude higher throughput at sub-two-second latency. Hybrid

arrangements occupy intermediate positions, with relay-based light-client bridges yielding several hundred transactions per second at finality latency of seconds, and federated validator-set bridges performing similarly.




Figure 4. Throughput and finality latency across representative trust fabric configurations.

HTLC-based atomic swaps occupy a distinctive position: while their throughput per channel is modest, their finality latency is dominated by the slower of the two underlying chains, which can introduce minute-scale delays even when the swap logic itself is simple. This explains why HTLCs are typically deployed for high-value, low-frequency settlements rather than for retail-grade payments (Herlihy, 2018; Robinson, 2021).

To make the comparison concrete, we summarise three quantitative observations that recur across the published benchmark studies. First, the throughput ratio between the slowest public configuration and the fastest permissioned configuration is consistently in the range of two to four orders of magnitude, with most studies clustering around a factor of one thousand. Second, finality latency exhibits bimodality: configurations either achieve sub-two-second finality (BFT and notary schemes) or fall into the multi-minute regime (relays anchored to PoW chains), with intermediate values being rare in production deployments and confined to specific layer-2 configurations. Third, the variance of measured throughput within a single configuration is larger than the variance between adjacent configurations, suggesting that engineering quality and operational tuning often dominate the choice of consensus protocol when both are competently implemented (Bamakan et al., 2020; Pongnumkul et al., 2017; Gervais et al., 2016). These observations have practical consequences: practitioners choosing among configurations for a given workload should weight expected operational variance heavily, and architects of cross-domain systems should design for the worst-case latency of the slowest involved domain rather than the average.

Sequence-level reasoning about cross-chain atomic exchanges further sharpens the latency analysis. The HTLC sequence depicted earlier in Figure 3 shows that the wall-clock latency of the four steps is the sum of finality times on the two domains, plus network propagation, plus the strategically chosen timeout buffers that ensure neither party can be defrauded by chain reorganisation. The composition of independent latency distributions is fundamental: pairing a fast permissioned domain with a slow public domain yields a swap whose effective latency tracks the slower side, eroding the speed advantage of the faster side. This insight, simple in retrospect, has substantial implications for how heterogeneous trust fabrics should be designed when latency budgets are tight, and it explains the empirical observation that high-frequency cross-domain

settlement is rarely deployed across consensus regimes that differ by more than an order of magnitude in finality time (Herlihy, 2018; Belchior et al., 2021).

5.2 Decentralisation Profiles

Performance metrics alone underspecify a trust fabric. A second axis is decentralisation, which captures the difficulty of corrupting the fabric by capturing a critical subset of its participants. The Nakamoto coefficient — the minimum number of independent entities whose collusion would compromise the fabric — has emerged as a convenient single-number proxy, but it averages over substantial heterogeneity (Gervais et al., 2016; Kannengießer et al., 2020). Public proof-of-work fabrics typically achieve high Nakamoto coefficients for hashpower (in the thousands), but lower coefficients for client implementations or for stake distribution. Permissioned fabrics, by design, have low coefficients (often single digits) but compensate with legal and contractual accountability of validators.

In cross-domain settings the relevant Nakamoto coefficient is not that of any single domain but that of the trust path traversed by a particular transaction. A swap mediated by a federated bridge with seven signers inherits the bridge's coefficient of seven, even if both endpoint chains have coefficients in the thousands (Belchior et al., 2021; Robinson, 2021). This observation reframes a common debate: the much-publicised security of public chains is often invoked to justify bridging assets to those chains, but the assets in transit are only as secure as the bridge itself. Programmable trust fabrics that expose path-level decentralisation metrics to applications would represent a substantial step forward in user-facing risk transparency.

5.3 Multi-Dimensional Comparison

Figure 5 collapses six dimensions — scalability, decentralisation, security, interoperability, programmability, and regulatory alignment — into a radar visualisation for three archetypal fabric profiles. The public permissionless profile achieves high decentralisation and respectable security but lags on scalability and regulatory alignment. The permissioned consortium profile inverts this: strong on scalability and regulatory alignment, weak on decentralisation. The hybrid trust fabric, which is the explicit aim of much current work, attempts a balanced profile, sacrificing some decentralisation in exchange for substantially improved interoperability and regulatory tractability.




Figure 5. Comparative profiles of three archetypal trust fabric configurations.

The figure illustrates a recurring observation in the literature: no profile dominates, and the appropriate choice is workload-dependent (Kannengießler et al., 2020; Hughes et al., 2019). Programmable trust fabrics — by exposing the choice points as parameters rather than fixing them at deployment — make it possible to allocate different traffic to different profiles within a single integrated system. A high-value institutional settlement might route through a permissioned consortium with formal regulatory oversight, while a retail micro-payment of the same notional asset might traverse a public sidechain with hash-locked exits.

5.4 Security Analysis Across Mechanisms

Table 3 summarises the principal threat vectors and defensive properties of the five interoperability mechanism families. The diversity of failure modes is itself a finding: there is no single class of attack that all mechanisms must defend against, and consequently there is no single defensive primitive that all mechanisms can adopt. Notary schemes are vulnerable to operator compromise and to bribery; HTLCs are vulnerable to chain-reorganisation attacks during the timelock window; relays inherit the security of the source chain but introduce additional verification surface; federated bridges are vulnerable to validator collusion proportional to the threshold required for attestation. Hybrid designs that combine mechanisms can dilute these vulnerabilities but at the cost of increased operational complexity and a larger total attack surface (Belchior et al., 2021; Karame & Capkun, 2018).

Table 3. Representative attack surfaces and known mitigations for cross-domain interoperability mechanisms.

Mechanism	Primary attack surface	Documented incident class	Mitigation pattern
Notary scheme	Notary compromise; insider risk	Custodial loss; key theft	HSM-backed keys; auditing; insurance

Sidechain	Peg-contract bugs; weak federation	Bridge exploit; double-spending	Formal verification; staked bond
HTLC	Timeout grief; chain congestion	Refund denial-of-service	Adaptive timeouts; submarine swaps
Light-client relay	Consensus forgery; data unavailability	Eclipse and long-range attacks	Checkpointing; weak subjectivity
Federated bridge	Validator collusion; key compromise	Multi-sig theft; oracle manipulation	Threshold signatures; rotation; bonds

6. Application Domains and Case Evidence

Programmable trust fabrics are no longer exclusively a research topic. A growing body of operational evidence demonstrates their use in domains where cross-domain value exchange has historically been blocked by trust deficits or by infrastructural fragmentation. We survey five domains in which deployments are sufficiently mature to permit empirical observations.

6.1 Cross-Border Finance and Settlement

Cross-border payments and securities settlement have been among the earliest beneficiaries of trust fabric architectures. Traditional correspondent banking relies on a chain of bilateral relationships, each of which absorbs latency, cost, and counterparty risk. Blockchain-based payment networks compress this chain by allowing institutions to settle on a shared ledger, with bridge mechanisms translating value into and out of local currencies and securities depositories (Schär, 2021; Hassani et al., 2018; Cocco, Pinna, & Marchesi, 2017). Real-world implementations have demonstrated settlement times measured in seconds rather than days, with documented cost reductions for participating institutions (Hughes et al., 2019). Simulation studies of cryptoasset markets have also illuminated the price-formation dynamics that emerge once tokenised instruments trade across multiple venues (Cocco, Concas, & Marchesi, 2017), and broader management research has begun to integrate blockchain into mainstream information-systems scholarship (Beck et al., 2017).

The trust fabric perspective clarifies why these deployments have been slower to spread than the technology might suggest. Settlement requires not only technical interoperability but the legal recognition of digital tokens as binding instruments, which in turn requires regulatory frameworks that the technology does not itself supply (Werbach, 2018; Lumineau et al., 2021). Fabrics that integrate auditable identity layers and that produce regulator-readable proofs of compliance are progressing more quickly than purely public configurations.

6.2 Healthcare Data Exchange

Healthcare presents a paradigmatic cross-domain challenge: patient records reside in heterogeneous institutional systems, each with strong incentives for retention and weak incentives for sharing, yet care quality and research utility depend on integrated access (Esposito et al., 2018; Salman et al., 2018). Programmable trust fabrics enable patients (or their delegates) to grant fine-grained, auditable access to records held in distinct institutional ledgers, with consent revocations propagating across the fabric and access

trails preserved for audit (Wang et al., 2018; Esposito et al., 2018). Pilot deployments in cancer care, organ transplantation, and clinical trials have reported measurable improvements in data-sharing velocity, though scale has remained limited by integration costs with legacy electronic health record systems.

6.3 Supply Chain Provenance

Supply chain applications are perhaps the most-discussed trust fabric use case, with extensive published literature and a track record of pilots in agriculture, pharmaceuticals, luxury goods, and electronics (Saberi et al., 2019; Kouhizadeh et al., 2021; Queiroz et al., 2020; Kamble et al., 2020; Wang, Han, & Beynon-Davies, 2019). The cross-domain dimension is fundamental: a single product passes through producer, processor, distributor, retailer, and consumer, each of whom maintains its own systems and trusts the others only conditionally. A programmable trust fabric anchors product histories to events recorded by IoT devices and verified by participants, producing tamper-evident provenance trails that can be selectively disclosed (Tian, 2017; Lin et al., 2018; Korpela, Hallikas, & Dahlberg, 2017). Domain studies have examined pharmaceutical traceability under cold-chain constraints (Lim et al., 2021), authentication in luxury supply chains (Choi, 2019), and the wider operations-management implications of fabric-mediated transparency (Tönnissen & Teuteberg, 2020; Pournader et al., 2020). Studies have documented reductions in counterfeiting risk and in dispute-resolution costs, while also identifying organisational barriers — data quality, intra-firm coordination, supplier readiness — that the technology alone cannot resolve (Min, 2019; Kouhizadeh et al., 2021; Treiblmaier, 2018). Life-cycle assessment frameworks have begun integrating fabric-anchored provenance data, opening a sustainability-accounting use case that crosses both organisational and environmental measurement boundaries (Zhang et al., 2020).

6.4 Energy Trading and Microgrids

Energy applications have evolved from conceptual demonstrations toward operational pilots, particularly in peer-to-peer local energy markets, electric-vehicle charging settlement, and renewable energy certificate tracking (Mengelkamp et al., 2018; Pop et al., 2018; Andoni et al., 2019). The trust fabric perspective is especially relevant because energy data crosses regulatory boundaries (national, regional, municipal), technical boundaries (grid operators, distributors, consumers), and economic boundaries (wholesale, retail, derivative markets). Smart contracts mediate auctions, settle micro-transactions, and enforce contractual constraints with timing precision unavailable in conventional billing systems (Mengelkamp et al., 2018; Andoni et al., 2019).

6.5 Digital Identity and Credentials

Decentralised identity infrastructures use trust fabrics as anchors for verifiable credentials that can be presented across domains without disclosing more than necessary (Liang et al., 2017; Ouaddah et al., 2016). Education credentials, professional licences, and right-to-work attestations are progressively being issued in formats that recipients can verify directly against an issuer's signature, while presenting only selected claims to relying parties. The fabric layer carries the issuer's public keys and revocation status; the credentials themselves typically remain in user-controlled wallets. The cross-domain property — that a credential issued in one jurisdiction can be verified in another without bilateral integration — is the principal source of value, though the cross-jurisdictional recognition of the credentials themselves remains a legal rather than technical question (Werbach, 2018; Beck et al., 2018). Adjacent work has examined property and registry use cases — real-estate transfer, land titling, asset-ownership records — that share the cross-jurisdictional verification challenge with identity systems (Notheisen, Cholewa, & Shanmugam, 2017; Maesa & Mori, 2020).

7. Challenges and Future Research Directions

Notwithstanding the encouraging trajectory, several issues constrain the broader adoption of programmable trust fabrics. We close with a non-exhaustive enumeration of the most pressing research directions, drawing on the literature surveyed throughout this paper as well as on patterns observed in deployment.

7.1 Bridge Security

Cross-chain bridges have been the locus of catastrophic losses in the period 2020–2024, with several incidents exceeding hundreds of millions of dollars in compromised value (Werner et al., 2022). The root causes are diverse — key management failures, smart contract bugs, validator collusion, oracle manipulation — but they share a common structure: the bridge concentrates value while distributing trust thinly, creating a target whose compromise unlocks disproportionate gains (Belchior et al., 2021; Atzei et al., 2017). These failure modes echo earlier analyses of blockchain-cloud integration risk, where shared infrastructure produced similar concentration effects (Park & Park, 2017). Future work must develop bridge architectures whose security degrades gracefully under partial compromise, perhaps by capping the value in transit through any single bridge path or by requiring multi-mechanism corroboration for high-value transitions (Robinson, 2021).

7.2 Scalability and Layer-2 Integration

The scalability gap between permissioned and public fabrics has narrowed but not closed. Layer-2 constructions — rollups, payment channels, sidechains — push throughput substantially beyond their base layers but introduce their own interoperability complications, since assets on different rollups face the same translation challenges that motivated this work (Robinson, 2021). Open research questions include the design of native cross-rollup messaging, the security model under which a layer-1 fabric can vouch for cross-layer-2 transactions, and the user-experience implications of multi-layer state.

7.3 Privacy, Confidentiality, and Selective Disclosure

Programmable trust fabrics must reconcile transparency — required for verification and audit — with confidentiality, required for commercial sensitivity and personal privacy (Kosba et al., 2016; Wang et al., 2018). Zero-knowledge proofs and confidential-transaction schemes are maturing rapidly, but their integration with cross-domain protocols remains immature. A particularly active sub-area is the design of cross-chain proofs that reveal only the minimal information necessary for the destination domain to act, without exposing the broader source-chain state (Belchior et al., 2021).

7.4 Regulatory Alignment and Governance

As trust fabrics carry larger fractions of the global value flow, regulatory engagement intensifies. Anti-money-laundering rules, securities law, data-protection regimes, and tax reporting all impose constraints that the fabric architecture must accommodate (Werbach, 2018; Beck et al., 2018; Lumineau et al., 2021). The governance of fabric protocols themselves — who can propose changes, who must ratify them, how forks are managed — remains an unsettled question, particularly when participants are distributed across jurisdictions with conflicting rules.

7.5 Quantum-Resistant Substrates

The cryptographic substrate on which all higher layers depend is vulnerable to large-scale quantum computers, particularly through Shor's algorithm against elliptic-curve signature schemes. While the timeline for cryptographically relevant quantum computers remains uncertain, the long-tail nature of historical chain data implies that a transition to post-quantum primitives will be required well before the threat materialises, lest historical transactions become forgeable in retrospect. Standardised post-quantum signature schemes are emerging, but their efficiency profile — particularly signature size — poses non-trivial challenges for high-throughput fabrics.

7.6 Toward Empirical Maturity

Finally, the field would benefit from greater methodological discipline in empirical evaluation. Benchmarks vary widely in workload, hardware, and instrumentation, making cross-study comparison difficult. Standardised benchmarks for interoperability mechanisms, common datasets for evaluating bridge security, and shared metrics for path-level decentralisation would accelerate the maturation of the field (Pongnumkul et al., 2017; Bamakan et al., 2020).

8. Conclusion

Programmable trust fabrics represent a conceptual shift in how distributed ledger technologies are deployed in practice. Rather than treating each ledger as an autonomous trust island, the fabric perspective emphasises the engineering of trust as a configurable resource that crosses administrative, technical, and legal boundaries. The five-layer reference architecture introduced in this paper — cryptographic substrate, consensus and settlement, interoperability, programmability, and application — provides a vocabulary in which existing single-chain models and emerging multi-chain compositions can be expressed and compared on common terms.

Cross-domain interoperability remains the defining challenge. We have surveyed five mechanism families — notary schemes, sidechains, hash-time-locked contracts, light-client relays, and federated bridges — and shown that each family occupies a distinct point in the design space defined by trust assumptions, performance, and deployability. No single mechanism dominates; mature trust fabrics will increasingly compose mechanisms, routing different traffic through different paths according to value, urgency, and regulatory context. The comparative evaluation presented in Section 5 illustrates how performance, decentralisation, and security trade off against one another, and how the trade-offs differ at the level of individual chains versus the level of cross-domain trust paths.

Application evidence from cross-border finance, healthcare, supply chain, energy, and digital identity confirms that the fabric paradigm is operationally relevant rather than merely theoretical. At the same time, several open challenges — bridge security, scalability, privacy, regulatory alignment, and quantum resistance — temper enthusiasm and define a substantial research agenda. The next phase of the technology will be shaped by how the community addresses these challenges, and by whether it can develop a shared methodological vocabulary for evaluating progress. The trajectory we have observed in adjacent computing infrastructures — the slow consolidation around standards, the gradual professionalisation of operational practice, the eventual disappearance of the underlying complexity from end-user experience — provides a reasonable template. Programmable trust fabrics are, in the relevant sense, infrastructure; their success will be measured not by the visibility of the technology but by its invisibility, by the moment at which value moves across domains as freely and reliably as packets move across networks today.

Declarations

Funding. The authors received no specific financial support for the research, authorship, or publication of this article.

Competing interests. The authors declare that they have no competing interests.

Author contributions. A.R.P. conceived the study, designed the layered architecture, and drafted the manuscript. D.A.H. conducted the interoperability mechanism survey and contributed to the comparative analysis. B.S.W. compiled the benchmark synthesis, prepared the figures, and reviewed the manuscript. All authors read and approved the final version.

Data availability. All data discussed in this article are drawn from publicly available sources cited in the reference list; no new datasets were generated.

Ethics. Not applicable; the study did not involve human subjects, animals, or personal data.

References

- Ahram, T., Sargolzaei, A., Sargolzaei, S., Daniels, J., & Amaba, B. (2017). Blockchain technology innovations. 2017 IEEE Technology & Engineering Management Conference (TEMSCON), 137–141. <https://doi.org/10.1109/TEMSCON.2017.7998367>
- Ali, M. S., Vecchio, M., Pincheira, M., Dolui, K., Antonelli, F., & Rehmani, M. H. (2018). Applications of blockchains in the Internet of Things: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 21(2), 1676–1717. <https://doi.org/10.1109/COMST.2018.2886932>
- Andoni, M., Robu, V., Flynn, D., Abram, S., Geach, D., Jenkins, D., McCallum, P., & Peacock, A. (2019). Blockchain technology in the energy sector: A systematic review of challenges and opportunities. *Renewable and Sustainable Energy Reviews*, 100, 143–174. <https://doi.org/10.1016/j.rser.2018.10.014>
- Atlam, H. F., Alenezi, A., Alassafi, M. O., & Wills, G. B. (2018). Blockchain with Internet of Things: Benefits, challenges, and future directions. *International Journal of Intelligent Systems and Applications*, 10(6), 40–48. <https://doi.org/10.5815/ijisa.2018.06.05>
- Atzei, N., Bartoletti, M., & Cimoli, T. (2017). A survey of attacks on Ethereum smart contracts (SoK). In M. Maffei & M. Ryan (Eds.), *Principles of Security and Trust* (pp. 164–186). Springer. https://doi.org/10.1007/978-3-662-54455-6_8
- Bamakan, S. M. H., Motavali, A., & Babaei Bondarti, A. (2020). A survey of blockchain consensus algorithms performance evaluation criteria. *Expert Systems with Applications*, 154, 113385. <https://doi.org/10.1016/j.eswa.2020.113385>
- Beck, R., Avital, M., Rossi, M., & Thatcher, J. B. (2017). Blockchain technology in business and information systems research. *Business & Information Systems Engineering*, 59(6), 381–384. <https://doi.org/10.1007/s12599-017-0505-1>
- Beck, R., Müller-Bloch, C., & King, J. L. (2018). Governance in the blockchain economy: A framework and research agenda. *Journal of the Association for Information Systems*, 19(10), 1020–1034. <https://doi.org/10.17705/1jais.00518>
- Belchior, R., Vasconcelos, A., Guerreiro, S., & Correia, M. (2021). A survey on blockchain interoperability: Past, present, and future trends. *ACM Computing Surveys*, 54(8), 1–41. <https://doi.org/10.1145/3471140>
- Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., & Felten, E. W. (2015). SoK: Research perspectives and challenges for Bitcoin and cryptocurrencies. 2015 IEEE Symposium on Security and Privacy, 104–121. <https://doi.org/10.1109/SP.2015.14>
- 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>
- Cha, S. C., Chen, J. F., Su, C., & Yeh, K. H. (2018). A blockchain connected gateway for BLE-based devices in the Internet of Things. *IEEE Access*, 6, 24639–24649. <https://doi.org/10.1109/ACCESS.2018.2799942>
- Choi, T. M. (2019). Blockchain-technology-supported platforms for diamond authentication and certification in luxury supply chains. *Transportation Research Part E: Logistics and Transportation Review*, 128, 17–29. <https://doi.org/10.1016/j.tre.2019.05.011>
- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the Internet of Things. *IEEE Access*, 4, 2292–2303. <https://doi.org/10.1109/ACCESS.2016.2566339>
- Cocco, L., Concas, G., & Marchesi, M. (2017). Using an artificial financial market for studying a cryptocurrency market. *Journal of Economic Interaction and Coordination*, 12(2), 345–365. <https://doi.org/10.1007/s11403-015-0168-2>
- Cocco, L., Pinna, A., & Marchesi, M. (2017). Banking on blockchain: Costs savings thanks to the blockchain technology. *Future Internet*, 9(3), 25. <https://doi.org/10.3390/fi9030025>

- Conoscenti, M., Vetrò, A., & De Martin, J. C. (2016). Blockchain for the Internet of Things: A systematic literature review. 2016 IEEE/ACS 13th International Conference of Computer Systems and Applications, 1–6. <https://doi.org/10.1109/AICCSA.2016.7945805>
- 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>
- Esposito, C., De Santis, A., Tortora, G., Chang, H., & Choo, K. K. R. (2018). Blockchain: A panacea for healthcare cloud-based data security and privacy? *IEEE Cloud Computing*, 5(1), 31–37. <https://doi.org/10.1109/MCC.2018.011791712>
- Eyal, I., & Sirer, E. G. (2018). Majority is not enough: Bitcoin mining is vulnerable. *Communications of the ACM*, 61(7), 95–102. <https://doi.org/10.1145/3212998>
- Fernández-Caramés, T. M., & Fraga-Lamas, P. (2018). A review on the use of blockchain for the Internet of Things. *IEEE Access*, 6, 32979–33001. <https://doi.org/10.1109/ACCESS.2018.2842685>
- Frizzo-Barker, J., Chow-White, P. A., Adams, P. R., Mentanko, J., Ha, D., & Green, S. (2020). Blockchain as a disruptive technology for business: A systematic review. *International Journal of Information Management*, 51, 102029. <https://doi.org/10.1016/j.ijinfomgt.2019.10.014>
- Garay, J., Kiayias, A., & Leonardos, N. (2015). The Bitcoin backbone protocol: Analysis and applications. In E. Oswald & M. Fischlin (Eds.), *Advances in Cryptology — EUROCRYPT 2015* (pp. 281–310). Springer. https://doi.org/10.1007/978-3-662-46803-6_10
- Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., & Capkun, S. (2016). On the security and performance of proof of work blockchains. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 3–16. <https://doi.org/10.1145/2976749.2978341>
- Hassani, H., Huang, X., & Silva, E. (2018). Banking with blockchain-ed big data. *Journal of Management Analytics*, 5(4), 256–275. <https://doi.org/10.1080/23270012.2018.1528900>
- Hawlicschek, F., Notheisen, B., & Teubner, T. (2018). The limits of trust-free systems: A literature review on blockchain technology and trust in the sharing economy. *Electronic Commerce Research and Applications*, 29, 50–63. <https://doi.org/10.1016/j.elerap.2018.03.005>
- Herlihy, M. (2018). Atomic cross-chain swaps. *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*, 245–254. <https://doi.org/10.1145/3212734.3212736>
- Hughes, A., Park, A., Kietzmann, J., & Archer-Brown, C. (2019). Beyond Bitcoin: What blockchain and distributed ledger technologies mean for firms. *Business Horizons*, 62(3), 273–281. <https://doi.org/10.1016/j.bushor.2019.01.002>
- Kamble, S. S., Gunasekaran, A., & Sharma, R. (2020). Modeling the blockchain enabled traceability in agriculture supply chain. *International Journal of Information Management*, 52, 101967. <https://doi.org/10.1016/j.ijinfomgt.2019.05.023>
- Kannengießer, N., Lins, S., Dehling, T., & Sunyaev, A. (2020). Trade-offs between distributed ledger technology characteristics. *ACM Computing Surveys*, 53(2), 1–37. <https://doi.org/10.1145/3379463>
- Karame, G., & Capkun, S. (2018). Blockchain security and privacy. *IEEE Security & Privacy*, 16(4), 11–12. <https://doi.org/10.1109/MSP.2018.3111241>
- Khan, M. A., & Salah, K. (2018). IoT security: Review, blockchain solutions, and open challenges. *Future Generation Computer Systems*, 82, 395–411. <https://doi.org/10.1016/j.future.2017.11.022>
- Korpela, K., Hallikas, J., & Dahlberg, T. (2017). Digital supply chain transformation toward blockchain integration. *Proceedings of the 50th Hawaii International Conference on System Sciences*, 4182–4191. <https://doi.org/10.24251/HICSS.2017.506>
- Kosba, A., Miller, A., Shi, E., Wen, Z., & Papamanthou, C. (2016). Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. 2016 IEEE Symposium on Security and Privacy, 839–858.

<https://doi.org/10.1109/SP.2016.55>

- Kouhizadeh, M., Saberi, S., & Sarkis, J. (2021). Blockchain technology and the sustainable supply chain: Theoretically exploring adoption barriers. *International Journal of Production Economics*, 231, 107831. <https://doi.org/10.1016/j.ijpe.2020.107831>
- Kshetri, N. (2017). Blockchain's roles in strengthening cybersecurity and protecting privacy. *Telecommunications Policy*, 41(10), 1027–1038. <https://doi.org/10.1016/j.telpol.2017.09.003>
- Liang, X., Shetty, S., Tosh, D., Kamhoua, C., Kwiat, K., & Njilla, L. (2017). ProvChain: A blockchain-based data provenance architecture in cloud environment with enhanced privacy and availability. 2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, 468–477. <https://doi.org/10.1109/CCGRID.2017.8>
- Lim, M. K., Li, Y., Wang, C., & Tseng, M. L. (2021). A literature review of blockchain technology applications in supply chains: A comprehensive analysis of themes, methodologies and industries. *Computers & Industrial Engineering*, 154, 107133. <https://doi.org/10.1016/j.cie.2021.107133>
- Lin, J., Shen, Z., Zhang, A., & Chai, Y. (2018). Blockchain and IoT based food traceability for smart agriculture. *Proceedings of the 3rd International Conference on Crowd Science and Engineering*, 1–6. <https://doi.org/10.1145/3265689.3265692>
- Lo, S. K., Xu, X., Chiam, Y. K., & Lu, Q. (2017). Evaluating suitability of applying blockchain. 2017 22nd International Conference on Engineering of Complex Computer Systems, 158–161. <https://doi.org/10.1109/ICECCS.2017.26>
- Lumineau, F., Wang, W., & Schilke, O. (2021). Blockchain governance—A new way of organizing collaborations? *Organization Science*, 32(2), 500–521. <https://doi.org/10.1287/orsc.2020.1379>
- Maesa, D. D. F., & Mori, P. (2020). Blockchain 3.0 applications survey. *Journal of Parallel and Distributed Computing*, 138, 99–114. <https://doi.org/10.1016/j.jpdc.2019.12.019>
- Mendling, J., Weber, I., van der Aalst, W., vom Brocke, J., Cabanillas, C., Daniel, F., Debois, S., Di Ciccio, C., Dumas, M., Dustdar, S., Gal, A., García-Bañuelos, L., Governatori, G., Hull, R., La Rosa, M., Leopold, H., Leymann, F., Recker, J., Reichert, M., ... Zhu, L. (2018). Blockchains for business process management — Challenges and opportunities. *ACM Transactions on Management Information Systems*, 9(1), 1–16. <https://doi.org/10.1145/3183367>
- Mengelkamp, E., Notheisen, B., Beer, C., Dauer, D., & Weinhardt, C. (2018). A blockchain-based smart grid: Towards sustainable local energy markets. *Computer Science — Research and Development*, 33(1–2), 207–214. <https://doi.org/10.1007/s00450-017-0360-9>
- Min, H. (2019). Blockchain technology for enhancing supply chain resilience. *Business Horizons*, 62(1), 35–45. <https://doi.org/10.1016/j.bushor.2018.08.012>
- Notheisen, B., Cholewa, J. B., & Shanmugam, A. P. (2017). Trading real-world assets on blockchain: An application of trust-free transaction systems in the market for lemons. *Business & Information Systems Engineering*, 59(6), 425–440. <https://doi.org/10.1007/s12599-017-0499-8>
- Ouaddah, A., Abou Elkalam, A., & Ait Ouahman, A. (2016). FairAccess: A new blockchain-based access control framework for the Internet of Things. *Security and Communication Networks*, 9(18), 5943–5964. <https://doi.org/10.1002/sec.1748>
- Park, J. H., & Park, J. H. (2017). Blockchain security in cloud computing: Use cases, challenges, and solutions. *Symmetry*, 9(8), 164. <https://doi.org/10.3390/sym9080164>
- Pilkington, M. (2016). Blockchain technology: Principles and applications. In F. X. Olleros & M. Zhegu (Eds.), *Research Handbook on Digital Transformations* (pp. 225–253). Edward Elgar Publishing. <https://doi.org/10.4337/9781784717766.00019>
- Pongnumkul, S., Siripanpornchana, C., & Thajchayapong, S. (2017). Performance analysis of private blockchain platforms in varying workloads. 2017 26th International Conference on Computer Communication and Networks,

- 1–6. <https://doi.org/10.1109/ICCCN.2017.8038517>
- Pop, C., Cioara, T., Antal, M., Anghel, I., Salomie, I., & Bertocini, M. (2018). Blockchain-based decentralized management of demand response programs in smart energy grids. *Sensors*, 18(1), 162. <https://doi.org/10.3390/s18010162>
- Pournader, M., Shi, Y., Seuring, S., & Koh, S. L. (2020). Blockchain applications in supply chains, transport and logistics: A systematic review of the literature. *International Journal of Production Research*, 58(7), 2063–2081. <https://doi.org/10.1080/00207543.2019.1650976>
- Queiroz, M. M., Telles, R., & Bonilla, S. H. (2020). Blockchain and supply chain management integration: A systematic review of the literature. *Supply Chain Management: An International Journal*, 25(2), 241–254. <https://doi.org/10.1108/SCM-03-2018-0143>
- Risius, M., & Spohrer, K. (2017). A blockchain research framework: What we (don't) know, where we go from here, and how we will get there. *Business & Information Systems Engineering*, 59(6), 385–409. <https://doi.org/10.1007/s12599-017-0506-0>
- Robinson, P. (2021). Survey of crosschain communications protocols. *Computer Networks*, 200, 108488. <https://doi.org/10.1016/j.comnet.2021.108488>
- Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Research*, 57(7), 2117–2135. <https://doi.org/10.1080/00207543.2018.1533261>
- Salman, T., Zolanvari, M., Erbad, A., Jain, R., & Samaka, M. (2018). Security services using blockchains: A state of the art survey. *IEEE Communications Surveys & Tutorials*, 21(1), 858–880. <https://doi.org/10.1109/COMST.2018.2863956>
- Sankar, L. S., Sindhu, M., & Sethumadhavan, M. (2017). Survey of consensus protocols on blockchain applications. 2017 4th International Conference on Advanced Computing and Communication Systems, 1–5. <https://doi.org/10.1109/ICACCS.2017.8014672>
- 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>
- Tian, F. (2017). A supply chain traceability system for food safety based on HACCP, blockchain & Internet of Things. 2017 International Conference on Service Systems and Service Management, 1–6. <https://doi.org/10.1109/ICSSSM.2017.7996119>
- Tönnessen, S., & Teuteberg, F. (2020). Analysing the impact of blockchain-technology for operations and supply chain management. *International Journal of Information Management*, 52, 101953. <https://doi.org/10.1016/j.ijinfomgt.2019.05.009>
- Treiblmaier, H. (2018). The impact of the blockchain on the supply chain: A theory-based research framework and a call for action. *Supply Chain Management: An International Journal*, 23(6), 545–559. <https://doi.org/10.1108/SCM-01-2018-0029>
- Tschorsch, F., & Scheuermann, B. (2016). Bitcoin and beyond: A technical survey on decentralized digital currencies. *IEEE Communications Surveys & Tutorials*, 18(3), 2084–2123. <https://doi.org/10.1109/COMST.2016.2535718>
- Wang, S., Zhang, Y., & Zhang, Y. (2018). A blockchain-based framework for data sharing with fine-grained access control in decentralized storage systems. *IEEE Access*, 6, 38437–38450. <https://doi.org/10.1109/ACCESS.2018.2851611>
- Wang, Y., Han, J. H., & Beynon-Davies, P. (2019). Understanding blockchain technology for future supply chains: A systematic literature review and research agenda. *Supply Chain Management: An International Journal*, 24(1), 62–84. <https://doi.org/10.1108/SCM-03-2018-0148>
- Werbach, K. (2018). Trust, but verify: Why the blockchain needs the law. *Berkeley Technology Law Journal*, 33(2), 487–550. <https://doi.org/10.15779/Z38H41JM9N>

- Werner, S., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., & Knottenbelt, W. (2022). SoK: Decentralized finance (DeFi). *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, 30–46. <https://doi.org/10.1145/3558535.3559780>
- Xu, X., Pautasso, C., Zhu, L., Lu, Q., & Weber, I. (2018). A pattern collection for blockchain-based applications. *Proceedings of the 23rd European Conference on Pattern Languages of Programs*, 1–20. <https://doi.org/10.1145/3282308.3282312>
- 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>
- Yli-Huumo, J., Ko, D., Choi, S., Park, S., & Smolander, K. (2016). Where is current research on blockchain technology? — A systematic review. *PLoS ONE*, 11(10), e0163477. <https://doi.org/10.1371/journal.pone.0163477>
- Zamyatin, A., Al-Bassam, M., Zindros, D., Kokoris-Kogias, E., Moreno-Sanchez, P., Kiayias, A., & Knottenbelt, W. J. (2021). SoK: Communication across distributed ledgers. In N. Borisov & C. Diaz (Eds.), *Financial Cryptography and Data Security* (pp. 3–36). Springer. https://doi.org/10.1007/978-3-662-64331-0_1
- Zhang, A., Zhong, R. Y., Farooque, M., Kang, K., & Venkatesh, V. G. (2020). Blockchain-based life cycle assessment: An implementation framework and system architecture. *Resources, Conservation and Recycling*, 152, 104512. <https://doi.org/10.1016/j.resconrec.2019.104512>
- Zheng, Z., Xie, S., Dai, H. N., Chen, X., & Wang, H. (2018). Blockchain challenges and opportunities: A survey. *International Journal of Web and Grid Services*, 14(4), 352–375. <https://doi.org/10.1504/IJWGS.2018.095647>