Network-Centric Auditing and Regulation for Cryptocurrency Systems: A Tutorial

Qingyang Ding, Qinnan Zhang, *Member, IEEE*, Xiaofei Yue, *Graduate Student Member, IEEE*, Jianbin Chen, Ziming Zhao, Zhaoxuan Li, Xiaochen Hao, Xinyu Zhang, and Nan Wang

***Abstract*—Public blockchains enable cryptocurrency and cross-chain asset flows, but anonymity and near-real-time decision-making complicate regulation. This paper adopts a network-centric perspective to study cryptocurrency regulation within a multi-layered network. We formalize sidechain interoperability and commitment-based anchoring as structural coupling and systematize prior research across three dimensions: (i) cross-layer governance, (ii) data-driven anomaly detection, and (iii) privacy as restricted information flow. Based on this, we derive design principles for verifiable anchoring, ex-ante checks, and low message complexity. As a case study, we present `AuditDC`, a dual-layer auditing system whose consortium smart contracts enforce ex-ante checks, maintain differentially private reputation, and emit verifiable cross-layer evidence. Finally, we summarize open questions in the network-centric regulatory space and propose a research agenda.**

***Index Terms*—Cryptocurrency regulation, multi-layer networks, cross-chain interoperability, auditing smart contracts.**

I. Introduction

DIGITAL asset transactions between users around the world have been enabled by public blockchains such as Bitcoin [1] and Ethereum [2]. However, due to the pseudonymous nature of transactions, cross-chain activity, and regulatory decision time frames, regulation requires special consideration [3], [4]. Empirical studies show how quickly simple payment methods can evolve into complex, interconnected transaction structures through address reuse [5], [6]. This creates challenges for regulators, exchanges, and infrastructure operators in auditing transactions ex-post and raises coordination issues among these parties. Platform providers increasingly offer interfaces for modular runtimes [7], enabling supervisory controls within the execution fabric rather than as an external element applied after transactions. This paper provides a regulatory framework for multi-layer networks and explains how cross-layer structure affects what regulators can verify and control.

In fact, modern cryptocurrency systems can be modeled as multiple interacting graphs with distinct semantics. Each type of consensus protocol generates a network of validators or peers from BFT consensus protocols [8], [9] to longest-chain consensus protocols [10]. Payment transactions generate a network of transactions on a public blockchain, while the supervisor can operate on a logically separate consortium blockchain of record-keeping on top of a public chain [7]. Cross-chain arrangements induce explicit networks of interactions in which transactions made in one blockchain can affect another [11], [12], [13], creating explicit links between layers that differ in their zones of potential failure and in how transactions can be confirmed. Anchoring, batching, and the selection of bridge validators are actions that redesign the global network to increase the difficulty of evasion or rollback while building trust at the bridge level, according to XCLAIM [14]. For example, treating the cryptocurrency system as a multi-layer network [15], [16] enables us to leverage concepts of robustness and control from the network science literature. Interdependent networks experience failure modes and phase transitions that do not appear on a single layer [17], [18], allowing us to see how constraints move from one layer to another in order to understand how actions taken by one party locally impact the entire network, such as throttling a bridge or freezing outputs onto that bridge. Treating regulation as an intervention in a network allows us to articulate the design objectives we want to achieve, enabling us to reshape the reachability or propagation of risk with minimal disruption. Concepts that derive from structural controllability and observability provide insight into what information should be anchored and what actions will have a macroscopic effect [19], [20]. Cascading failure analysis suggests explicit rollback limits and recovery methods for addressing the disruption caused by deteriorating cross-layer edges and when re-organizations occur [21], [22].

All effective supervisory designs must consider modern blockchain systems [23]. While permissioned and permissionless environments have different consensus mechanisms and ultimately achieve different latency and throughput trade-offs, this can be measured by BFT-style finality in consortium blockchains [8], [9] and PoS-style committee sampling, which facilitates both reduced confirmation time and simplified

Received 19 November 2025; revised 16 February 2026; accepted 25 February 2026. Date of publication 9 March 2026; date of current version 23 March 2026. The work of Qingyang Ding was supported by the Beijing Social Science Foundation under Grant 23GLC037. Recommended for acceptance by Dr. Liehuang Zhu. *(Corresponding authors: Qinnan Zhang; Xiaofei Yue.)*

Qingyang Ding, Jianbin Chen, Xiaochen Hao, Xinyu Zhang, and Nan Wang are with the School of Management, Beijing Union University, Beijing 100101, China (e-mail: gltqingyang@buu.edu.cn; jianbin.chen@buu.edu.cn).

Qinnan Zhang is with the Institute of Artificial Intelligence, Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing, Beihang University, Beijing 100191, China (e-mail: zhangqn@buaa.edu.cn).

Xiaofei Yue is with the School of Computer Science and Technology, Beijing Institute of Technology, Beijing 100081, China (e-mail: xfyue@bit.edu.cn).

Ziming Zhao is with the School of Software Technology, Zhejiang University, Ningbo 315100, China (e-mail: zhaoziming@zju.edu.cn).

Zhaoxuan Li is with the State Key Laboratory of Cyberspace Security Defense, Institute of Information Engineering, Chinese Academy of Sciences, Beijing 100093, China (e-mail: lizhaoxuan@iie.ac.cn).

Digital Object Identifier 10.1109/TNSE.2026.3672038

security arguments [24], [25]. In a permissioned environment, the Fabric project separates execution from ordering and validation while also providing contract interfaces that support auditability [7]. Thus, a network-centric view will place fast checks along relatively short critical paths, and anchor verification evidence to make verifying the correctness of a transaction inexpensive and make forgery very difficult. Ex-ante supervision requires cross-layer evidence that can be consumed by on-chain mechanisms to gate settlement actions.

Detection and privacy are vital components of transaction processing, but they are most useful when viewed as elements in an end-to-end network process rather than isolated analyses. The methods for performing transaction tracing and attribution have evolved from heuristic clustering [5] to machine learning applied to transaction graphs and temporal flows. However, they remain ex-post and thus do not provide a direct path to enforceable constraints. Implementations of cryptographic privacy technology (i.e., zero-knowledge protocols [26], [27] and short-range proofs [28]), in conjunction with differential privacy technology, provide primitives for releasing statistical results and ensuring that the privacy of individuals with respect to their reputations is maintained with formal guarantees [29], [30], [31], [32], [33], [34]. Within a network-centric framework, these technologies will provide the basic building blocks for constrained flow of information via the use of structural interventions and cross-layer anchoring.

Although much progress has been made on detecting historical transactions, there remains a significant gap between implementing actionable, ex-ante regulations. For example, many existing detection systems [5], [6], [35] analyze recorded transaction data and provide published labels. However, they do not generally implement cross-layer constraints that could prevent, slow, or delay the flow of high-risk transactions prior to settlement. From consensus protocols, they have typically optimized for safety and liveness. In bridge designs, they have generally focused on functionality and throughput, with little consideration for verifiable audit points. Therefore, we propose a multi-layer network perspective in which the public transaction layer, the consortium auditing layer, and the anchoring edges connecting them form an interdependent system. Using this new perspective, we survey architectures for regulation, learning-based detection systems, and privacy technologies. Unlike most previous surveys [3], [36] that consider security, consensus, and privacy in isolation, we demonstrate how they can work together across layers in a unified multi-layer network model. We then present a case study of AuditDC, a dual-layer auditing design that audits smart contracts to enable ex-ante regulation. It provides reliable results via differential privacy, while anchoring them to the transaction layer for enforcement verification via a cross-layer authorization interface.

To sum up, this paper makes the following contributions:

- We present a network-centric model that casts cryptocurrency regulation as a multi-layer, interdependent network with explicit cross-layer edges.
- We systematize prior work along three axes aligned with network-centric regulation and distill design principles and open challenges.
- We conduct a case study of AuditDC, an automatic and decentralized auditing system for cryptocurrency regulation on a dual-layer blockchain.

The remainder of this paper is organized as follows. Section II introduces structural background for cross-chain and network-centric regulation. Section III surveys related work and distills design principles and open challenges. Section IV presents a case study of AuditDC, and Section V analyzes its security and evaluates performance, scalability, and robustness. Section VI discusses open problems, with the paper concluded in Section VII.

II. Preliminaries

This section outlines a network-centric regulation model. We first define the "Regulating Blockchain with Blockchain" paradigm. Then, we leverage cross-layer structures to map transaction flows and interoperability as carriers of verifiable evidence. Finally, we summarize the main regulation mechanisms.

A. Blockchain Regulation Paradigm

Public blockchain networks process transactions globally and can achieve Byzantine fault-tolerant consistency within partially synchronous network models. Current platforms implement consensus protocols with (i) probabilistic finality and longest-chain consensus (e.g., Bitcoin [1]), (ii) deterministic PBFT-style protocols (e.g., PBFT [8] and HotStuff [9]), (iii) Hybrid designs with optimistic fast lanes and conservative delayed backends (e.g., Thunderella [37]) and (iv) Committee sampling-based proof-of-stake protocols (e.g., Algorand [24], Ouroboros [25] and Praos [38]). In permissioned environments, decoupling execution, ordering, and validation into clear, programmable interfaces enables governance logic to be embedded into transactions [39]. Hyperledger Fabric [7] demonstrates this by isolating endorsement execution from sequencing, enabling later verification via batched evidence. We introduce the Regulating blockchain with blockchain: deploying audit contracts at a consortium audit layer to create verifiable auditing constraints for public blockchain stakeholders. The cryptocurrency system is modeled as a multi-layer network comprising a public transaction layer, a consortium audit layer, and cross-layer evidence edges. Regulation involves observation and intervention to change connectivity and flow, while ensuring safety, liveness, and privacy [15], [40].

The dedicated audit layer has two uses. First, it reduces the critical verification path by placing the costliest or most policy-heavy computations as close to their policy execution point as possible. PBFT [8] and HotStuff [9] can provide safety and liveness of decisions made in the audit layer under partial synchrony. The audit layer allows ex-ante constraints to be imposed on a multi-layer network through structural operations to limit risk diffusion. The resilience of these designs relates to how cascading failures might occur in interdependent networks [17], [18] and how interventions in multi-layer networks link regtech goals through structural controllability and observability [19], [20].

Building these systems requires designing for adversarial reorganizations, liveness under noise, and anchoring cost. The

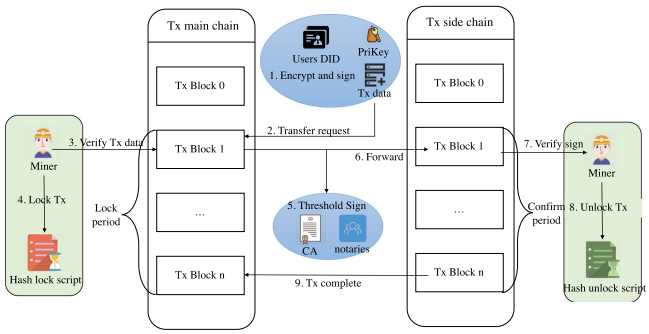

Fig. 1. Sidechain-based transfer across multiple blockchains.

choices made in selecting the substrate will determine the critical path from alert to enforcement, potential queueing, and potential failure modes within the auditor interaction network [41]. HotStuff [9] provides a simple chained pipeline using linear communication with stable leader rotations. PBFT [8] and Tendermint [42] provide predictable commit routes under partial synchrony through explicit view changes and quorum certificates. In proof-of-stake settings, Algorand [24] and Praos [38] provide more stable confirmation horizons by ensuring that each slot is verifiably randomly determined and that the effect of a malicious party is bounded.

B. Cross-Chain Interoperability

In permissioned environments, sidechain-enabled interoperability mechanisms improve interoperability and scalability between blockchains [43], [44], [45]. The traditional architecture employs two-way anchoring mechanisms, in which assets are locked to a primary, outward-facing ledger while a corresponding representation is maintained on a sidechain. We survey current research on cross-chain interoperability, including patterns and threat models that affect the global transaction network.

The above processes produce cross-layer evidence edges: commitments made to the state of one ledger can be validated against the state of another ledger; thus, regulators can enforce regulations across jurisdictions. Fig. 1 shows a common anchoring arrangement based on an m-of-n multisignature script signed by a subset of a permissioned consortium blockchain of validators acting as notary signers. Off-chain complementary protocols can shorten settlement time between connected ledgers and reduce contention. Fast bilateral micropayment channels [46], Sprites (reducing state channel collateral lock-time) [47], and other forms of access-controlled off-chain interaction expand interoperable deployment patterns in recent literature [48]. The reliability and discoverability of cross-layer evidence edges depend on the design and structure of their validator set and committed signature set, including parameters defining threshold authorization [49] and remote state validation. FlyClient [50] produces logarithmic-size proofs of existence for lightweight clients. In high-throughput and sharded environments, transactional topology manifests over time, changing how evidence edges can and should be placed and maintained [51], [52], [53], [54].

The existence of cross-layer evidence edges creates risks and opportunities. Validator collusion may create fraudulent signatures at the boundary between two blockchains, deceive light clients, or exploit weak validation assumptions. XCLAIM [14] provides formal models of trust assumptions for existing protocols enabling cross-border blockchain assets. NIPoPoWs [55], [56] provides a concise model of client acceptance of transactions, minimizing the risk of exploitation by full nodes. Introducing distributed cryptographic-strength metrics will support validator auditing and enable concise cumulative evidence via polynomial commitments to spline curves [57], with validator signatures threshold controlled via threshold ECDSA [58], [59]. Sampling validators via VRFs mitigates attackers targeting specific nodes, reducing the risk of cascading attacks. In multi-layer networks, reinforced honest edges or nodes will mitigate the potential for attacks to trigger a cascade, consistent with interdependent network theory [17], [18].

C. Network-Centric Regulation Mechanisms

A network-centric regulatory framework comprises several core components: consensus, contracts, privacy, and security techniques. The four components define what can be seen, what can be proven, and where intervention can be made. The configuration of ordering and quorum protocols defines the length of the critical path and the queuing of audit actions. The PBFT [8] algorithm tolerates up to f faulty replicas using a three-phase commit with $3f + 1$ total replicas. HotStuff [9] uses quorum-certificate chaining to achieve low communication overhead. To meet the demands of BFT variants (e.g., SBFT [60] and Mir BFT [61]), throughput improves through signature aggregation and batching [62]. This improvement enables audit batching and reduces auditing delay. By implementing a Fabric-like architecture [7], the segregation of endorsement, ordering, and validation enables the audit committee to function independently while providing a consistent view of cross-layer audit information. Both parties can interpret regulatory standards into mathematical models that enable the implementation of specific transactions in the transaction network. A regulator can use rate-limit predicates to restrict suspicious transaction flows by limiting transactions per second. Once this threshold is exceeded, regulatory rules can be applied through flow-obstruction mechanisms (e.g., freezing or isolation). Regulatory authorities may use blacklists and whitelists to adjust the granularity of interventions, minimizing collateral impact. To minimize enforcement latency after an alert, enforcement must limit critical-path latency, eliminate contract vulnerabilities (e.g., re-entrancy), and ensure a verifiable connection between evidence of enforcement and the enforcement decision. Replay resistance increases the likelihood of verification [50], [55], [63]. Audit records document a regulatory decision, and if they contain sensitive data, that data must be kept confidential until disclosure under a DP budget [29]. Regulatory authorities may use contracts as control points to modify edges and weights of identified high-risk subgraphs within a multi-layer network while maintaining an auditable record of interventions.

Using cryptographic hardening and topology, interoperable components may be operated by distinct operators without a

TABLE I
REPRESENTATIVE WORKS ON CRYPTOCURRENCY REGULATION, ANOMALY DETECTION, AND PRIVACY FROM A NETWORK-CENTRIC VIEW

Method	Layers	Characteristics	Limitations
Regulation architectures [66]	Multi	Separates audit from public settlement; structural traceability	Coordination overhead; weaker decentralization; limited pre-settlement control
Interoperability primitives [11]–[13]	Multi	Atomic swaps and channels move assets and proofs; create coupling points for enforcement	Bridge trust and liveness assumptions; operational overhead
Succinct light clients [50], [55]	Dual	Efficient remote state verification; routine anchoring and evidence transport	Requires remote-chain support; honest-majority-style assumptions
Sharded and deconstructed ledgers [51], [54], [65]	Multi	Structural parallelism and collective signing; higher throughput; shorter confirmation paths	Complex pipelines and tuning; fault handling more involved
Consensus for fast finality [24], [25], [41]	Single	Lower confirmation variance; committee sampling and BFT-style finality; clearer audit budgets	Protocol and network assumptions; fairness and incentives need care
Contract security and verification [67], [68]	Single	Static and formal analyses harden regulatory contracts; reduce runtime risk	Coverage gaps and false positives; cost for complex policies
Attested oracles and TEEs [69]	Single	Authenticated feeds for policy inputs with enclave attestations	TEE trust and side channels; availability risks
Transaction graphs and anomaly surveys [5], [70]	Single	Foundational graph measurements and anomaly primitives; basis for subgraph or edge patterns	Often post hoc and label driven; limited direct actuation
Graph learning and streaming detection [71]–[73]	Single	GNN and self-supervision improve accuracy; online pipelines for near real-time streams	Explanations and calibration lacking; brittle under shift; hard to operationalize
Illicit ecosystems and MEV [74]–[77]	Single	Empirical pipelines reveal ordering effects and cash-out paths; identify enforcement choke points	Cross-layer hops evade single-chain analytics; policy gaps persist
Cryptographic privacy [27], [28], [78], [79]	Dual	Mixing, commitments, zk proofs, and private contracts enable strong on-chain privacy	Prover or verifier overhead; integration friction with public regulation
Network and DP privacy [29], [80], [81]	Single	Metadata privacy in relay; DP-based telemetry and reputation that compose with audits	Deployment friction; DP utility and budgeting trade-offs
Privacy-preserving auditing frameworks [82], [83]	Dual	Encrypted or zk-backed auditing and selective disclosure in consortium blockchain settings	Communication and proof costs; limited integration with public-chain regulation

single point of failure. Cross-domain gateways must implement thresholding techniques and distributed key generation to ensure no single operator can produce imprecise cross-layer evidence. Two-party and multiparty ECDSA [58], [59] and polynomial commitments [57] can be defined as components that create global trust while minimizing the size of audit records needed for verification.

Off-chain technologies and deconstructed ledger technologies also define how regulators will interact with fund flows in the network. Atomic swap contracts with hash-time locks allow parties to reduce counterparty risk by enabling a controlled exit prior to failure. Herlihy [12] defines parameters for achieving atomicity of an atomic swap using hash-time locks. Lightning [13] processes most transactions off-chain with on-chain settlement guarantees. Perun [64] defines a framework for virtual channels and hubs based on a limited number of on-chain contracts. In a sharded or decomposed ledger, edges associated with each shard can maintain a global root blockchain or validate shard edges, changing how failures occur across partitions. ByzCoin [65], OmniLedger [51], RapidChain [53], and Prism [54] show how parallelism and collective authorization affect confirmation time and the placement of audit records relative to events. These design considerations identify where to anchor or verify cross-layer audit records so that ex-ante regulatory actions remain effective in the presence of partitions or confirmation failures.

III. NETWORK-CENTRIC REGULATION

In this section, we survey the literature along three axes: (i) network-centric regulation that treats supervision as ex-ante interventions on a multi-layer network, (ii) data-driven detection that discovers subgraphs and edge patterns indicative of risk, and (iii) privacy protection that constrains information flow without severing evidence paths. In addition, we compare representative works in Table I.

A. Cryptocurrency and Regulation

In Section. II, we discuss the chain-governing-chain abstraction, cross-layer evidence edges, and a short verifiable path from alert to enforcement. We now discuss how prior research implements these primitives in supervision architectures.

A representative example of regulatory architecture is separating public settlement from permissioned execution or audit, which promotes verifiable trails for traceability and policy enforcement [66]. The primary design question is how to attach ex-ante interventions to short critical paths so that intervention latency occurs before finality and evidence can be verified across domains. Network risks remain, and supervisors should be shielded from them. Propagation delays [84] impact fork rate and observation coverage. Deanonymization at the peering layer [85], eclipse attacks [86], and BGP hijacks [87] can affect what a supervisor can observe or reach. Relays that modify Erlay's bandwidth and diffusion dynamics [88] also change observability and the ability to act in near real time.

Ecosystems have become increasingly diverse, making interoperability essential for supervision. Sidechains [11], hashed time-lock contracts for atomic swaps [12], and payment-channel networks (e.g., Lightning [13]) connect administrative domains, where off-chain payment flows can be settled on-chain with certainty. In these instances, succinct light clients and proofs (e.g., FlyClient [50] and NIPoPoWs [55]) enable practical remote validation of evidence at reduced costs of transporting evidence between layers. Sharding and deconstructed ledgers

bring supervisory mechanisms closer to execution through new anchoring points and fault semantics [51], [53], [54], [65]. In this multi-layer network, permissioned execution layers (e.g., Fabric [7]) provide an ideal home for audit contracts and evidentiary documentation. Anchoring-depth determination may be supported via backbone-style rollback-risk analyses [89] and tie-breaking rules [90]. Research examining the decentralization of mining [91] provides insights into where cross-layer evidence edges should terminate and where robust evidence exists.

The literature suggests a continuing movement from single-chain settlement toward multi-domain value flows and permissioned execution layers. However, most current architectures lack a unified methodology for (i) positioning verifiable cross-layer evidence edges with explicit trust and public membership and (ii) associating ex-ante interventions with short, auditable alert-to-enforcement paths.

B. Data-Driven Abnormal Transaction Detection

Graph-centric analysis of on-chain activity has raised interest in analytic methods for these graphs. Initial work uses heuristics to cluster transactions and identify addresses [5]. Later work uses heuristics to identify wealth transfers and characterize addresses based on behavior [6]. Methods for identifying anomalies via statistical techniques, proximity methods, clustering, and graph techniques have been developed and used [70], [92], [93]. An example is OddBall [94], which identifies ego-networks that deviate from power-law distributions. Recent surveys synthesize anomaly-detection methods, data sources, and evaluation metrics for blockchain technology [95]. A highlighted point is the trade-off between recall and operational cost.

Supervised and semi-supervised methods for detecting illicit and fraudulent activity in cryptocurrency have been developed. Examples include webs of Bitcoin transactions and classification by GCNs [71], dissection of Ponzi schemes using features and heuristics on Ethereum [96], and indicators of contracts and behavior sequences related to schemes [97]. These methods motivate architectures for fraud detection that combine Deep Graph Infomax [72], GINs, and ensemble classification [98]. MDST-GNN leverages spatiotemporal characteristics of transaction data to identify local and long-term dependencies [99]. Methods for near real-time identification of subgraph anomalies include FAST [73], which uses synopsis and dynamic states, and SedanSpot [100], which identifies dense subgraphs from temporal anomalies. Several recent surveys characterize the spectrum of GNNs used to analyze blockchain-related data.

For regulatory enforcement, ex post classifiers achieve high accuracy in identifying illicit activity, but delays limit their leverage for settlement. Streaming detection systems trade some accuracy for speed. Cross-ledger relationships enable cross-self-consistent detection of abnormal transaction activity across multiple blockchains [101]. Key components for operationalizing regulatory enforcement include processing time, calibration thresholds, and explanation methods. It is important to calibrate models to address false positives, as thresholds define business rules for quarantining activity. The asymmetric rule structure for data collection impacts network visualization because the collected structure determines analysis visibility [84], [88]. BASIA enhances dataset similarity using a blockchain-generated index for similarity search [102]. These detection models can establish conditions under which alerts map to quarantine, settlement delay, or rate limits on subgraphs [103]. Cross-layer commitments that ensure submission reliability shorten the alert-to-enforcement window. Due to the limited explanatory capability of many GNNs, tools such as GNNExplainer [104] and PGExplainer [105] can provide subgraph evidence and confidence for selected actions.

Overall, GNN-based and streaming detection advances beyond early heuristic attribution, but linking detection to ex-ante enforcement remains challenging due to operational time frames and regulatory complexity.

C. Privacy Protection in Cryptocurrency

Mixers, confidential transaction systems, ring signatures, zero-knowledge proofs, and network-layer anonymity are different components of privacy technology. Mixcoin [78] establishes accountable mixing via escrowed transactions and a mechanism for assigning blame for theft. CoinShuffle [106] creates a decentralized CoinJoin-like model using layered encryption. CoinShuffle++ [107] makes CoinShuffle efficient by using Random DiceMix-style shuffling. TumbleBit [108] allows users to send payments without being connected by using an untrusted third party and cryptographic puzzles, and utilizing a cut-and-choose procedure to allow the sender to verify the payment without sharing payment information with the recipient. Zero-knowledge proof techniques based on commitment schemes (e.g., Bulletproofs [28]) can be used to hide the transaction amount. Systems that use ring signatures or zero-knowledge proofs can hide both the sender and the receiver while allowing the public to verify legitimacy. For example, Zerocoin [26] introduces an accumulation-based mint and spend, while Zerocash [27] introduces a shielded transaction model using compact zk-SNARKs. Groth16 [109] creates efficient proofs for general circuits, while ZEXE [110] can operate on encrypted data. Privacy-preserving solutions for smart contracts include Hawk [79], which compiles private smart contracts with an off-chain manager. Also, zkLedger [82] provides private auditable ledgers for enterprises, and Zether [111] enables confidential payments for account-based blockchains. Privacy-preserving solutions can also be enabled using two distinct phases (i.e., stem and fluff) for message diffusion. Dandelion [80], [112] reduce information about a message's origin by separating diffusion into stem and fluff.

Several surveys provide a network-centric view of these privacy mechanisms. Conti et al. [113], [114] examine security and privacy issues in Bitcoin, including anonymization techniques and countermeasures. Zhang et al. [115] provide an overview of security and privacy mechanisms in blockchain platforms, from consensus mechanisms and data-structure security to higher-level protocols. Bernal et al. [116] systematize privacy-preserving solutions at the protocol and system levels, and Satybaldy and Nowostawski summarize privacy-preserving blockchain systems and deployment limitations [117]. Yin

et al. [118] provide a complementary survey focusing on performance trade-offs in permissioned and permissionless environments. Differential privacy (DP) introduces calibrated noise, enabling auditors to produce assurance about data accuracy [29], [30], [119]. DP-SGD [31] enables practical DP through per-example training. Reconstruction bounds [120] regulate what can be inferred from responses, while secure multi-party computation and trusted execution enable analytics and audits as in SecureML [121] and Ekiden [122]. Authenticated oracles, such as Town Crier [69], connect on-chain policy to off-chain data via attestation.

For regulation, cryptographic privacy can provide precise evidence of violations when proofs of policy compliance are included with transactions. Long-lived telemetry systems can rely on differential privacy because it is provably composable and attenuates privacy loss [29], [30]. Network-layer privacy includes origin metadata, but if relay paths and batching are not aligned with supervision, observability may be reduced. Ordering effects and MEV [75], [76] indicate that transaction ordering can be used to assess privacy or fairness. Protocol choices such as randomizing relay paths, randomizing batch order, or committing before transactions become public affect leakage and observability [88], [41], [9]. zkFabLedger [83] integrates zero-knowledge methods with Hyperledger Fabric to provide an audited ledger with privacy protection and regulatory compliance.

Overall, privacy mechanisms enable regulated information release, but they can under- or over-represent regulated signals via ambiguous cross-layer evidence-generation frameworks.

D. Findings

The literature we review highlights three characteristics of effective cryptocurrency regulation: (i) a short, reliable alert-to-enforcement pathway to facilitate pre-finality ex-ante interventions, (ii) clearly defined cross-layer evidence-edge models that incorporate auditable trust and membership assumptions in multi-chain environments, and (iii) privacy mechanisms that ensure verifiable but constrained information release (e.g., zk proofs for minimal identity release and differential privacy for long-term usage telemetry and reputation). Given this, we design a dual-layer architecture with a permissioned auditing consortium blockchain, cross-layer anchoring to a public transactional layer, and privacy-preserving reputation release.

IV. Case Study: AuditDC

This section instantiates the above concepts via a case study of AuditDC. It is designed atop a multi-layer network in which a public transaction layer and a consortium auditing layer are coupled by cross-layer anchoring edges [123]. In AuditDC, auditing contracts act as network intervention operators that throttle, freeze, or isolate risky flows before settlement, while privacy mechanisms operate as local perturbations that preserve utility. The goal is an auditable cryptocurrency transaction scheme that supports transaction-data auditing and reputation recording while preserving the privacy of reputation values [124]. Table II lists symbols and their descriptions.

TABLE II
Symbols and Descriptions

Symbols	Descriptions
Tx_{ij}	Transaction between user i and j.
Sig_i	Signature under user i's private key.
$PubKey^k_{UTXO}$	Address of the k-th unspent transaction output.
$PubKey_{new}$	Newly generated transaction address.
Sn_{out}	Serial number of outgoing transaction $Tx_{ij,out}$.
Sn_{in}	Serial number of change transaction $Tx_{ij,in}$.
$T(t)$	Total number of transactions of the user.
$H(t)$	Number of normal transactions of the user.
$Gcred$	Global reputation value.
$Lcred$	Local reputation value.
Δf	Sensitivity in differential privacy.
ε	Privacy budget in the Laplace mechanism.
$\hat{M}_r$	Noise-added reputation value model.

A. System Overview

Definition 1. Transaction (Tx_{ij}): A transaction generated on the cryptocurrency transaction blockchain. Let Tx_{ij} denote a transaction between users i and j, where Sig_i is the signature under the outgoing user's private key, Sn is the transaction serial number, $PubKey^k_{UTXO}$ is the address of the k-th Unspent Transaction Output (UTXO), $PubKey_{new}$ is the newly generated transaction address, and $Value$ is the change amount. A transaction Tx_{ij} can be expressed as

$$Tx_{ij} = \left\{ \begin{array}{l} Sn, (PubKey^k_{UTXO}, Sig_i), \\ (PubKey_{new}, Value) \end{array} \right\}. \quad (1)$$

Definition 2> Outgoing transaction ($Tx_{ij,out}$): A transaction fragment stored on the public blockchain for validation. It includes the transaction serial number Sn_{out}, a payment address $PubKey^k_{UTXO}$, and an unlocking proof (e.g., Sig_i):

$$Tx_{ij,out} = \left\{ Sn_{out}, \left(PubKey^k_{UTXO}, Sig_i \right) \right\}. \quad (2)$$

Definition 3. Change transaction ($Tx_{ij,in}$): A transaction fragment generated together with the outgoing transaction. It includes the transaction serial number Sn_{in}, the new public key address $PubKey_{new}$, and the change amount $Value$:

$$Tx_{ij,in} = \{ Sn_{in}, (PubKey_{new}, Value) \}. \quad (3)$$

Definition 4. Transaction verification. This process verifies the unlocking signature and UTXO for each $Tx_{ij,out}$: For UTXO-based cryptocurrencies, each UTXO is recursively traced back to earlier UTXOs and ultimately to an initial mining reward. For account-based systems (*e.g.*, Ethereum), the account balance must also be checked.

Definition 5. Transaction audit: This process constructs an audit model from authoritative regulatory rules or machine-learning techniques and deploys an audit smart contract on the consortium blockchain. Once the contract is validated by consensus, it evaluates transaction type, amount, and user reputation, applies differentiated audit methods, and outputs whether a transaction is normal or abnormal.

AuditDC Architecture: Fig. 2 shows AuditDC, an automatic and decentralized cryptocurrency regulation scheme. Its backbone is a dual-layer architecture consisting of a cryptocurrency

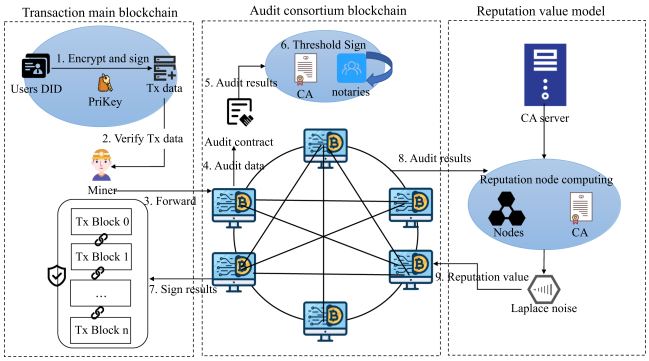

Fig. 2. The architecture overview of AuditDC.

transaction blockchain and an audit consortium sidechain [125]. The transaction blockchain establishes trusted transaction data among users via distributed consensus and cryptographic mechanisms. The audit consortium sidechain executes audit contracts and records audit outcomes, reputation values, and transaction events, enabling traceability and follow-up audits [126]. More specifically, the audit consortium sidechain is a permissioned blockchain operated by a set of known entities (e.g., regulated institutions). To satisfy certificate issuance requirements, a regulator-operated Certificate Authority (CA) service is introduced as a logically centralized component. It maintains user identity registrations, binds verified identities to transaction public keys or addresses, and provides these bindings to the consortium auditing layer for compliance checks. Cross-layer audit evidence is authenticated via threshold signatures produced by notaries, implemented as a subset of permissioned consortium nodes.

Algorithm 1: Audit Smart Contract.

Input : Tx_type, $user\_credit$, audit judge model $\mathcal{M}$
Output: Audit result $\mathcal{R}_a$

if $user\_credit \geq$ threshold **then**
  if $Tx\_type =$ Cross_border **then**
    $\mathcal{R}_a \leftarrow \mathcal{M}_{\text{Trusted\_cross\_border\_Audit}}()$;
  else
    $\mathcal{R}_a \leftarrow \mathcal{M}_{\text{Trusted\_domestic\_Audit}}()$;
  end
else
  if $Tx\_type =$ Cross_border **then**
    $\mathcal{R}_a \leftarrow \mathcal{M}_{\text{Untrusted\_cross\_border\_Audit}}()$;
  else
    $\mathcal{R}_a \leftarrow \mathcal{M}_{\text{Untrusted\_domestic\_Audit}}()$;
  end
end
return transaction audit result $\mathcal{R}_a$;

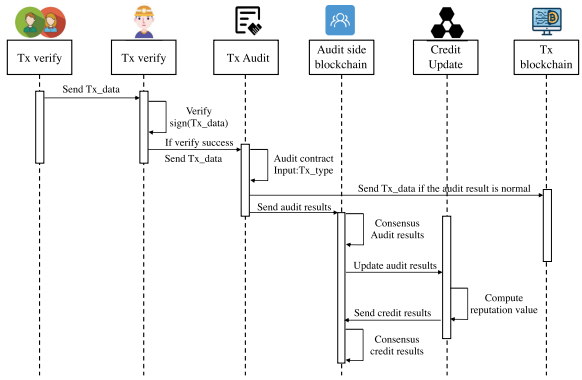

Fig. 3. The cryptocurrency transaction and audit workflow.

B. Audit Workflow and Contract Design

To avoid collusion risks, most cryptocurrencies adopt public blockchains as their base platforms to enhance credibility and security through open participation and monetary incentives. In AuditDC, users involved in cryptocurrency transactions must complete real-name registration to enable regulation and traceability. AuditDC classifies transactions into different types and treats each transaction request as a pair of outgoing and change fragments $Tx_{ij,out}$ and $Tx_{ij,in}$. The audit goal is to build an audit judge model $\mathcal{M}$ that provides auditing functions for different transaction types. The audit judge model $\mathcal{M}$ can be programmed from simple regulatory rules and trained using standard machine-learning algorithms such as deep learning, support vector machines, random trees, random forests, and Gradient Boosting Decision Trees [127]. The audit smart contract is sketched in Algorithm 1.

AuditDC integrates an audit consortium sidechain with the cryptocurrency transaction blockchain. Transaction information is periodically transmitted to the consortium sidechain via the audit contract. After identity verification by an authentication node, the information is broadcast to other consortium nodes, which reach consensus and store reputation values and audit decisions. Nodes obtain decryption keys for transaction information via voting, enabling regulation and traceability. The overall transaction and audit workflow is shown in Fig. 3.

1) *User real-name authentication:* The user submits identity information Msg_{user} to the CA service on the audit consortium sidechain. Upon successful authentication, the server signs the user's transaction data with its private key to generate a user certificate $Cer_{user} = sig(id_{msg}, PK_{userId})$, and returns the response.
2) *Cryptocurrency transaction:* The payer retrieves the payee's public key address $PubKey_{new_i}$ and change address $PubKey_{new_j}$, and sends the UTXO with an unlocking proof (signature) to the verifier.
3) *Transaction verification (Tx Verify):* On receiving transaction information, the verifier checks balance, assigns a transaction number, and signs the transaction. The verifier is a leader node on the transaction blockchain, selected among worker nodes by the consensus protocol.
4) *Transaction audit (Tx Audit):* After verification, transaction information is sent to the audit server and handled by the audit contract. The contract classifies transaction type (domestic vs. cross-border), assesses the user's credit status, and applies different audit rigor accordingly. The audit result ("*normal*" or "*abnormal*") is forwarded to the next stage.

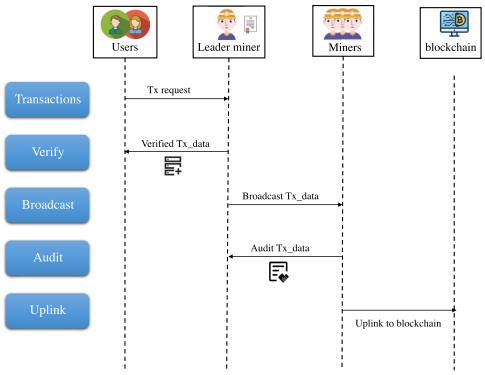

Fig. 4. The overall consensus workflow.

5) *Transaction block generation:* If the audit result is normal, the verifier broadcasts the transaction, which is stored in the transaction blockchain through consensus.
6) *User reputation update:* In parallel, audit results are fed to the reputation evaluation model, signed, and propagated on the audit consortium sidechain, where the user's reputation value is updated by consensus.

C. Proof of Contribution Consensus Algorithms

Supervising cryptocurrency transactions via an audit consortium sidechain offers two main advantages. First, user and audit information is only broadcast to consortium nodes rather than to all public blockchain nodes, mitigating linkability and preserving transaction anonymity between users. Second, audit contracts can automatically analyze transaction parameters, detect anomalies, and require consensus verification among consortium nodes. Existing consensus algorithms [128], [129] are ill-suited for audit smart contracts, since they either consume excessive resources or exacerbate the *Matthew effect*.

To mitigate resource waste, AuditDC employs a Proof of Contribution (PoC) consensus algorithm tailored to transaction audits. The overall consensus workflow is shown in Fig. 4. PoC uses audit contributions of consortium nodes to establish consensus, improving efficiency in computing and communication across the network. Upon receiving a transaction request, a leader node is elected among participating nodes based on contribution values, with higher-contribution nodes prioritized within a time window. The elected leader coordinates consensus and disseminates audit information. Worker-node contributions are determined by online availability and completed audits. This design reflects two considerations: (1) the stability and security of a blockchain are positively correlated with the number of active nodes, and (2) greater online participation fosters wider adoption and higher reliability.

In AuditDC, the online-time contribution of a node is computed as its online duration multiplied by a coefficient. The overall contribution is evaluated by combining online-time contribution with the number of completed audits:

$$C = \alpha \cdot T_{online} + (1-\alpha)\, C_{audit}, \tag{4}$$

where $\alpha \in [0,1]$ is a system parameter, T_{online} is the node's online time, and C_{audit} is the number of audit tasks completed by the node. Thus, larger α prioritizes stability and responsiveness (more online participation), while smaller α prioritizes audit workload contribution.

D. Reputation Value Model

The system evaluates the credibility of user transactions from historical records. Accumulated reputation reduces the probability of abnormal behavior and supports differentiated regulation: users with high credit scores can be subject to lower-intensity policies, while low-score users are regulated more tightly. Policies are implemented exclusively via smart contracts.

In the reputation evaluation module, a CA service is included in the consortium setting to satisfy Know Your Customer (KYC) demands. The CA establishes mappings between verified identities and cryptocurrency public key addresses, and issues certificates when required. A reputation-calculation contract dynamically updates each user's score from post-transaction information and stores it across consortium sidechain nodes. The CA service is operated by the regulator to meet trust requirements, while reputation computation and storage are replicated across permissioned consortium sidechain nodes. This separation avoids centralizing the audit decision log or the reputation ledger, while still enabling compliant identity binding and accountable evidence anchoring.

After each transaction, a user's record is labeled as a "*normal*" or "*abnormal*" transaction by the audit contract. The local reputation value $Lcred$ is then updated as

$$Lcred = \frac{H(t)}{T(t)}, \tag{5}$$

where $T(t)$ is the total number of transactions, and $H(t)$ is the number of normal transactions.

Once a transaction completes, the global reputation value is updated by automatically triggering the credit-update contract. After execution, the result is signed and uploaded to the consortium. The global reputation model at the i-th update $Gcred_i$ is

$$Gcred_i = Sig\left(\frac{Lcred}{Gcred_{i-1}}\right), \tag{6}$$

The sequence $\{Gcred_i\}_{\forall i}$ forms a credibility blockchain for user reputation, and after each transaction the audit contract updates $Gcred_i$ and synchronizes it with the consortium blockchain. Based on these models, the overall reputation value model M_r is defined as:

$$M_r = \sum\nolimits_i Gcred_i. \tag{7}$$

E. Privacy Protection Via DP and Obfuscation

AuditDC combines transaction obfuscation with Differential Privacy (DP) to hide transaction amounts and reputation values [130], as shown in Fig. 5.

Transaction amount hidden: Transaction obfuscation is a standard privacy method in Bitcoin-like systems. A complete

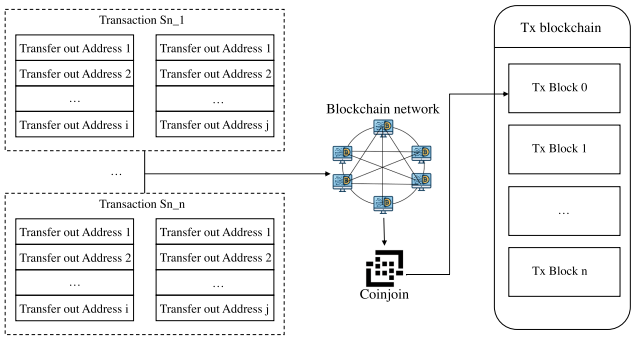

Fig. 5. The process of transaction obfuscation and traceability.

transaction may contain multiple outgoing and incoming fragments. After verifying integrity, incoming and outgoing sub-transactions are numbered. The k-th complete transaction Sn_k includes all transfer-in and transfer-out sub-transactions:

$$Sn_k \in \{Sn_{out_0}, Sn_{out_1}, \ldots, Sn_{out_i}, Sn_{in_0}, Sn_{in_1}, \ldots, Sn_{in_j}\}. \quad (8)$$

when packaging a transaction, sub-transaction numbers are attached to fragments and encrypted before being recorded on-chain. Sub-transactions in a block are then shuffled and broadcast in batches to public blockchain nodes as ciphertext fragments, which provides obfuscation while preserving verifiability. The encryption key is stored at regulatory nodes using a key-sharing mechanism, and traceability is initiated via a voting mechanism.

Reputation value hidden: To protect reputation privacy on the consortium sidechain, AuditDC integrates a DP mechanism into reputation hiding. Let datasets D and D' have the same schema and differ in at most one record (neighboring datasets). A randomized algorithm M satisfies ε-differential privacy if for any set of outputs SM and neighboring D, D',

$$Pr[M(D) \in SM] \leq \exp(\varepsilon) \cdot Pr[M(D') \in SM], \quad (9)$$

where ε is the privacy budget. In this paper, the Laplace mechanism is applied to user reputation values [131]:

$$\hat{M}_r = M_r + Laplace\left(\frac{\Delta f}{\varepsilon}\right), \quad (10)$$

where $Laplace(\Delta f/\varepsilon)$ has scale $b = \Delta f/\varepsilon$ and probability density

$$P(x) = \frac{\exp\left(-\frac{|x|}{b}\right)}{2b}, \quad (11)$$

and Δf is the sensitivity:

$$\Delta f = \max \|M_r(D) - M_r(D')\|_1. \quad (12)$$

Intuitively, Δf captures the maximum change in the released reputation statistic when one transaction in the input dataset is modified (between neighboring D and D'). Larger Δf or smaller ε implies a larger noise scale (stronger privacy but lower utility), while larger ε implies smaller noise (weaker privacy but higher accuracy). The noise-added reputation value model $\hat{M}_r$ is then broadcast as a consortium blockchain transaction to audit nodes for consensus. This repeats until consensus is reached or the time limit expires [132]. Algorithm 2 summarizes the DP-based reputation-hiding procedure.

Algorithm 2: Reputation Value Hiding Based on DP.

Input : $H(t)$, $T(t)$, sensitivity Δf, privacy budget ε
Output: Reputation value V_r

$time \leftarrow$ Current Time;
Iteration times $i \leftarrow 0$;
while Current Time $-time <$ time commitment **do**
  if $i = 0$ **then**
    $Lcred \leftarrow \frac{H(t)}{T(t)}$;
    $Gcred_0 \leftarrow \mathrm{Sig}_{\mathrm{audit\_node}}(Lcred)$;
  else
    $Gcred_i \leftarrow \mathrm{Sig}\left(\frac{Lcred}{Gcred_{i-1}}\right)$;
  end
  $V_r \leftarrow \frac{1}{i+1}\sum_{k=0}^{i} Gcred_k + \mathrm{Laplace}\left(\frac{\Delta f}{\varepsilon}\right)$;
  $i \leftarrow i + 1$; Sleep(t);
  $time \leftarrow$ Current Time;
end
return *reputation value*V_r;

V. Scheme Analysis and Evaluation

This section presents the security analysis and an illustrative evaluation of AuditDC.

A. Security Analysis

AuditDC incorporates regulatory compliance into Bitcoin's trustless architecture to reduce integrity and trust risks associated with decentralized cryptocurrency oversight while adhering to blockchain security risk classifications.

Decentralized trust: Audit smart contracts are executed by the consortium and have replicated audit logs that provide tamper-evident logging without reliance on a sole operator. Practical BFT-style solutions (e.g., Tendermint [42] and HoneyBadgerBFT [133]) demonstrate how replicated agreement can survive network latency and message reordering, while cross-ledger SoK studies [134] illustrate how committee selection and threshold signing shape mutually agreed trust assumptions among inter-chain infrastructures. The audit consortium blockchain operates with known entities under a standard permissioned-consensus trust model.

Transaction data security: AuditDC uses a consortium audit sidechain to maintain a standing order of audit processes, providing evidence that audits were ordered based on system performance. The block containing proposed audit decisions and corresponding evidence is finalized using PoC. Block proposal and confirmation are limited to authenticated members of the auditor consortium. Once confirmed, audit results are replicated across auditors, providing integrity evidence against adversarial manipulation and demonstrating adherence to best-practice verification mechanisms [135].

Privacy protection: The protection of confidential information is achieved through: (i) the obfuscation of transactions and encrypting all audit m-ready metadata to prevent those who perform trades on the settlement layer from discovering their identity, and (ii) the use of differential privacy when publishing reputation releases. Each update satisfies ε-differential privacy. Under the basic composition theorem, publishing k updates yields an overall privacy guarantee of at most $k\varepsilon$-differential privacy. Thus, AuditDC is designed to consider reputation publication as an ongoing stream of data, with an explicit global privacy budget that applies to all releases. This is intended to ensure that cumulative leakage remains manageable while maximizing utility for auditing.

B. Evaluation Setup

Experimental environment: We design and deploy an audit smart contract on a Hyperledger Fabric network to support an auditable supervisory mechanism. The contract detects anomalous or non-compliant transactions by evaluating transaction-derived parameters, aiming to balance regulatory assurance and operational efficiency. Specifically, we use an Ethereum-based network as the transaction substrate and realize the UTXO-style transaction abstraction, while a Hyperledger Fabric network serves as the consortium audit sidechain. The audit smart contract is implemented in Go and deployed on Fabric. Transactions from the transaction layer are relayed to Fabric via a cross-chain relay service, which invokes the audit contract to record audit decisions. To link the two ledgers, the relay periodically anchors a timestamped commitment on the consortium audit sidechain, enabling consistent cross-chain referencing and subsequent verification.

Workloads: The workload consists of a 7:3 ratio of domestic to cross-border payments. Transaction sizes follow an exponential distribution with an average of 300 bytes of metadata per audit call. The audit model $\mathcal{M}$ executes 12 rule checks by default. We vary the transaction arrival rate, number of audit rules, and privacy budget ε to test performance under diverse conditions. For metrics, we report (i) throughput as the number of transactions that complete the audit-and-commit pipeline per second, and (ii) latency as the time from a transaction being submitted to its audit decision being finalized on the consortium audit blockchain and anchored for verification.

Baselines: We compare AuditDC against three schemes: (1) **zkCross** [136] is a cross-chain privacy-preserving auditing framework based on zero-knowledge proofs and atomic swap verification; (2) **BASIA** [102] is a blockchain-assisted searchable integrity auditing mechanism with large-scale data arbitration; (3) **zkFabLedger** [83] is a privacy-preserving, regulation-compliant auditing system with zk-SNARK verification.

Overall performance: As shown in Figs. 6 and 7, AuditDC gets an average throughput improvement of 47.3% over the baselines and reduces latency by 30.2% on average. The superior performance of AuditDC primarily arises from its dual-layer audit pipeline and lightweight cross-layer evidence verification. Unlike zkCross and zkFabLedger, which require generating and verifying computationally intensive zero-knowledge proofs for each audit task, AuditDC leverages lightweight threshold signatures and batched audit execution to significantly reduce computational overhead. Furthermore, compared to BASIA, which relies on arbitration and searchable indexing mechanisms that constrain scalability under high workloads, AuditDC achieves higher concurrency and lower end-to-end delay by decoupling transaction validation from audit computation and optimizing cross-chain synchronization without compromising audit accuracy or security.

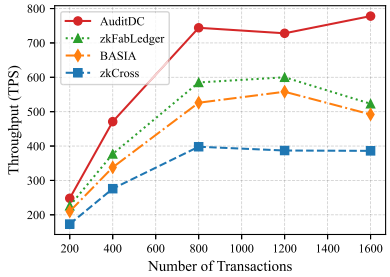

Fig. 6. Throughput with varying transaction volumes against baselines.

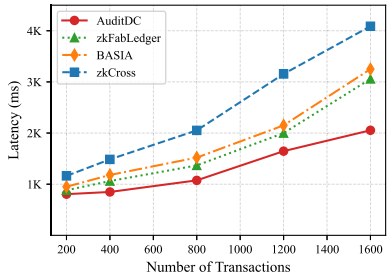

Fig. 7. Latency with varying transaction volumes against baselines.

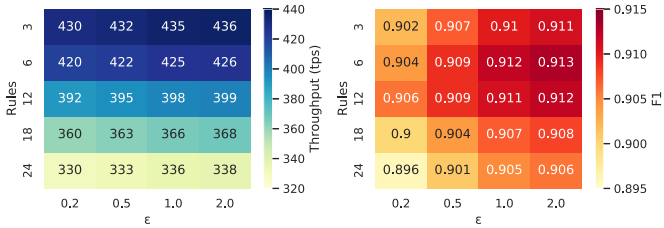

(a) Throughput. (b) F1 Score.

Fig. 8. Rule-privacy interaction with varying audit rules and budgets.

Rules-privacy interaction: Fig. 8 presents the trade-off between the number of audit rules $R \in \{3, 6, 12, 18, 24\}$ and privacy budget $\epsilon \in \{0.2, 0.5, 1.0, 2.0\}$. For each configuration, we generate 1,200 transactions and evaluate both throughput and detection F1-score. We treat the audit outcome as a binary detection task: a transaction is labeled positive if it violates at least one rule in $\mathcal{M}$, and negative otherwise. As expected, increasing R raises computational overhead, while larger ϵ slightly improves F1 by reducing noise. A clear Pareto frontier emerges: (R=12, ϵ=1.0) provides a near-optimal trade-off (~400 tps, F1≈0.911), whereas (R=18, ϵ=0.2) marginally lowers F1 but significantly decreases throughput.

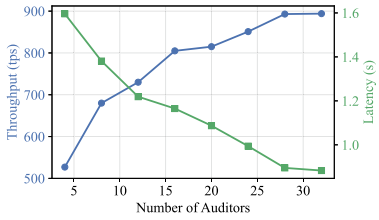

Fig. 9. Throughput and latency with different numbers of auditors.

Scalability with audit nodes: As shown in Fig. 9, we scale the number of auditors from 4 to 32 to evaluate the system's scalability. This scalability advantage stems from AuditDC's contribution-driven consensus mechanism and parallel verification pipeline. By dynamically distributing computationally intensive cryptographic checks (e.g., signature aggregation and rule verification) across a larger set of auditors, the system effectively amortizes the coordination cost. The latency decreases correspondingly as the parallel processing capability expands, with a performance plateau observed near 32 nodes, indicating the point where network synchronization costs begin to balance out the benefits of parallel computation.

VI. Discussion

Evidence and finality: An important challenge is optimizing evidence anchoring to reduce verification costs and increase security. Key tasks include defining the granularity of anchoring, avoiding replay attacks on anchors, and timestamped anchors so that anchoring latency is balanced against reorganization risk, which differs across finality regimes. Future research should develop cross-layer budgeting models that optimize anchor cost, data freshness, and rollback security under settlement-side constraints.

Decision and privacy: The latency-sufficient, calibrated, and audit-ready detection pipelines are required for valid ex-ante intervention. Open problems remain in sustaining calibration to maintain decision-path verifiability under data drift [137]. There are also issues in generating machine-readable evidence. Privacy must be achieved through persistent accounting, budget allocation for repeated releases, and binding guarantees attached to verifiable audit trails to guard against iterative disclosure attacks.

Governance: The governance of membership and resilience will be critical for permissioned auditing layers. Proposed solutions could include establishing verifiable audit logs for policy revisions, outlining admission criteria, and linking audit actions to operator attestation. Settlement governance needs to formalize procedures for upgrades and emergencies. The goal is to remove ambiguities related to single points of failure.

VII. Conclusion

This tutorial offers examples using a network-centric view of cryptocurrency regulation in a multi-layered, interdependent network. We model the transaction layer as connected to the consortium auditing layer via an explicit anchor. We consolidate prior work on governance, detection, and privacy as coordinated network mechanisms rather than stand-alone analyses. From this perspective, we use smart contracts as intervention primitives, and implement privacy as constrained information flow with evidence paths. As a case study, AuditDC demonstrates this method by enabling ex-ante checks on an auditing blockchain, anchoring consortium authority back to the transaction layer, while maintaining differential privacy.

References

[1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," Whitepaper, 2008. [Online]. Available: https://bitcoin.org/bitcoin.pdf
[2] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum Project Yellow Paper*, vol. 151, no. 2014, pp. 1–32, 2014.
[3] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten, "SoK: Research perspectives and challenges for bitcoin and cryptocurrencies," in *Proc. IEEE Symp. Secur. Privacy*, 2015, pp. 104–121.
[4] N. Amarasinghe, X. Boyen, and M. McKague, "A survey of anonymity of cryptocurrencies," in *Proc. Australas. Comput. Sci. Week Multiconference*, 2019, pp. 1–10.
[5] S. Meiklejohn et al., "A fistful of bitcoins: Characterizing payments among men with no names," in *Proc. Conf. Internet Meas. Conf.*, 2013, pp. 127–140.
[6] D. Ron and A. Shamir, "Quantitative analysis of the full bitcoin transaction graph," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2013, pp. 6–24.
[7] E. Androulaki et al., "Hyperledger fabric: A distributed operating system for permissioned blockchains," in *Proc. 13th EuroSys Conf.*, 2018, pp. 1–15.
[8] M. Castro et al., "Practical byzantine fault tolerance," in *Proc. OsDI*, 1999, vol. 99, no. 1999, pp. 173–186.
[9] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, "HotStuff: BFT consensus with linearity and responsiveness," in *Proc. ACM Symp. Princ. Distrib. Comput.*, 2019, pp. 347–356.
[10] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, "On the security and performance of proof of work blockchains," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2016, pp. 3–16.
[11] A. Back et al., "Enabling blockchain innovations with pegged sidechains," vol. 72, pp. 201–224, 2014. [Online]. Available: https://www.semanticscholar.org/paper/Enabling-Blockchain-Innovations-with-Pegged-Back-Corallo/1b23cd2050d5000c05e1da3c9997b308ad5b7903
[12] M. Herlihy, "Atomic cross-chain swaps," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2018, pp. 245–254.
[13] J. Poon and T. Dryja, "The bitcoin lightning network: Scalable off-chain instant payments," 2016. [Online]. Available: https://bit.ly/2Hppzos
[14] A. Zamyatin, D. Harz, J. Lind, P. Panayiotou, A. Gervais, and W. Knottenbelt, "Xclaim: Trustless, interoperable, cryptocurrency-backed assets," in *Proc. IEEE Symp. Secur. Privacy*, 2019, pp. 193–210.
[15] M. De Domenico et al., "Mathematical formulation of multilayer networks," *Phys. Rev. X*, vol. 3, no. 4, 2013, Art. no. 041022.
[16] R. Albert, H. Jeong, and A.-L. Barabási, "Error and attack tolerance of complex networks," *Nature*, vol. 406, no. 6794, pp. 378–382, 2000.
[17] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, "Catastrophic cascade of failures in interdependent networks," *Nature*, vol. 464, no. 7291, pp. 1025–1028, 2010.
[18] J. Gao, S. V. Buldyrev, H. E. Stanley, and S. Havlin, "Networks formed from interdependent networks," *Nature Phys.*, vol. 8, no. 1, pp. 40–48, 2012.
[19] Y.-Y. Liu et al., "Controllability of complex networks," *Nature*, vol. 473, no. 7346, pp. 167–173, 2011.
[20] Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási, "Observability of complex systems," in *Proc. Nat. Acad. Sci.*, vol. 110, no. 7, pp. 2460–2465, 2013.
[21] A. E. Motter and Y.-C. Lai, "Cascade-based attacks on complex networks," *Phys. Rev. E*, vol. 66, no. 6, 2002, Art. no. 065102.
[22] X. Yue et al., "HyFaaS: Accelerating serverless workflows by unleashing hybrid resource elasticity," *IEEE Trans. Parallel Distrib. Syst.*, vol. 37, no. 1, pp. 272–286, Jan. 2026.
[23] Q. Ding, X. Yue, Q. Zhang, Z. Xiong, J. Chang, and H. Zheng, "Bc ^{2} FL: Double-layer blockchain-driven federated learning framework for agricultural IoT," *IEEE Internet Things J.*, vol. 12, no. 4, pp. 4362–4374, Feb. 2025.

[24] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling byzantine agreements for cryptocurrencies," in *Proc ACM 26th Symp. Operating Syst. Princ.*, 2017, pp. 51–68.

[25] A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-of-stake blockchain protocol," in *Proc. Annu. Int. Cryptol. Conf.*, 2017, pp. 357–388.

[26] I. Miers, C. Garman, M. Green, and A. D. Rubin, "Zerocoin: Anonymous distributed e-cash from bitcoin," in *Proc. IEEE Symp. Secur. Privacy*, 2013, pp. 397–411.

[27] E. B. Sasson et al., "Zerocash: Decentralized anonymous payments from bitcoin," in *Proc IEEE Symp. Secur. Privacy*, 2014, pp. 459–474.

[28] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell, "Bulletproofs: Short proofs for confidential transactions and more," in *Proc. IEEE Symp. Secur. Privacy*, 2018, pp. 315–334.

[29] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. IEEE Trans. Cloud Comput.*, 2006, pp. 265–284.

[30] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations Trends Theor. Comput. Sci.*, vol. 9, no. 3–4, pp. 211–407, 2014.

[31] M. Abadi et al., "Deep learning with differential privacy," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2016, pp. 308–318.

[32] Q. Feng, D. He, S. Zeadally, M. K. Khan, and N. Kumar, "A survey on privacy protection in blockchain system," *J. Netw. Comput. Appl.*, vol. 126, pp. 45–58, 2019.

[33] M. U. Hassan, M. H. Rehmani, and J. Chen, "Differential privacy in blockchain technology: A futuristic approach," *J. Parallel Distrib. Comput.*, vol. 145, pp. 50–74, 2020.

[34] A. Qashlan, P. Nanda, and M. Mohanty, "Differential privacy model for blockchain based smart home architecture," *Future Gener. Comput. Syst.*, vol. 150, pp. 49–63, 2024.

[35] X. Yue, S. Yang, L. Zhu, S. Trajanovski, and X. Fu, "Demeter: Fine-grained function orchestration for GEO-distributed serverless analytics," in *Proc. IEEE Conf. Comput. Commun.*, 2024, pp. 2498–2507.

[36] A. Augusto, R. Belchior, M. Correia, A. Vasconcelos, L. Zhang, and T. Hardjono, "SoK: Security and privacy of blockchain interoperability," in *Proc. IEEE Symp. Secur. Privacy*, 2024, pp. 3840–3865.

[37] R. Pass and E. Shi, "Thunderella: Blockchains with optimistic instant confirmation," in *Proc. Annu. Int. Conf. Theory Appl. Cryptogr. Techn.*, 2018, pp. 3–33.

[38] B. David, P. Gaži, A. Kiayias, and A. Russell, "Ouroboros Praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain," in *Proc. Annu. Int. Conf. Theory Appl. Cryptogr. Techn.*, 2018, pp. 66–98.

[39] D. Ongaro and J. Ousterhout, "In search of an understandable consensus algorithm," in *Proc. IEEE USENIX Annu. Tech. Conf.*, 2014, pp. 305–319.

[40] M. Kivelä, A. Arenas, M. Barthélemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, "Multilayer networks," *J. Complex Netw.*, vol. 2, no. 3, pp. 203–271, 2014.

[41] I. Eyal, A. E. Gencer, E. G. Sirer, and R. Van Renesse, "Bitcoin-NG: A scalable blockchain protocol," in *Proc. USENIX Symp. Netw. Syst. Des. Implementation*, 2016, pp. 45–59.

[42] E. Buchman, "Tendermint: Byzantine fault tolerance in the age of blockchains," Master's thesis, Univ. Guelph, Guelph, ON, Canada, 2016.

[43] A. Singh, K. Click, R. M. Parizi, Q. Zhang, A. Dehghantanha, and K.-K. R. Choo, "Sidechain technologies in blockchain networks: An examination and state-of-the-art review," *J. Netw. Comput. Appl.*, vol. 149, 2020, Art. no. 102471.

[44] P. Han, Z. Yan, W. Ding, S. Fei, and Z. Wan, "A survey on cross-chain technologies," *Distrib. Ledger Technol., Res. Pract.*, vol. 2, no. 2, pp. 1–30, 2023.

[45] H. Yuan, S. Fei, and Z. Yan, "Technologies of blockchain interoperability: A survey," *Digit. Commun. Netw.*, vol. 11, no. 1, pp. 210–224, 2025.

[46] C. Decker and R. Wattenhofer, "A fast and scalable payment network with bitcoin duplex micropayment channels," in *Proc. Symp. Self-Stabilizing Syst.*, 2015, pp. 3–18.

[47] A. Miller, I. Bentov, S. Bakshi, R. Kumaresan, and P. McCorry, "Sprites and state channels: Payment networks that go faster than lightning," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2019, pp. 508–526.

[48] L. Yang et al., "An access control model based on blockchain master-sidechain collaboration," *Cluster Comput.*, vol. 27, no. 1, pp. 477–497, 2024.

[49] L. Harn, "Group-oriented (t, n) threshold digital signature scheme and digital multisignature," *IEE Proc.-Comput. Digit. Techn.*, vol. 141, no. 5, pp. 307–313, 1994.

[50] B. Bünz, L. Kiffer, L. Luu, and M. Zamani, "FlyClient: Super-light clients for cryptocurrencies," in *Proc. IEEE Symp. Secur. Privacy*, 2020, pp. 928–946.

[51] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford, "OmniLedger: A secure, scale-out, decentralized ledger via sharding," in *Proc. IEEE Symp. Secur. Privacy*, 2018, pp. 583–598.

[52] L. Luu et al., "A secure sharding protocol for open blockchains," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2016, pp. 17–30.

[53] M. Zamani, M. Movahedi, and M. Raykova, "RapidChain: Scaling blockchain via full sharding," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2018, pp. 931–948.

[54] V. Bagaria, S. Kannan, D. Tse, G. Fanti, and P. Viswanath, "Prism: Deconstructing the blockchain to approach physical limits," in *Proc. IEEE ACM Conf. Comput. Commun. Secur.*, 2019, pp. 585–602.

[55] A. Kiayias, A. Miller, and D. Zindros, "Non-interactive proofs of proof-of-work," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2020, pp. 505–522.

[56] A. Kiayias, N. Lamprou, and A.-P. Stouka, "Proofs of proofs of work with sublinear complexity," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2016, pp. 61–78.

[57] A. Kate, G. M. Zaverucha, and I. Goldberg, "Constant-size commitments to polynomials and their applications," in *Proc. Int. Conf. Theory Appl. Cryptol. Inf. Secur.*, 2010, pp. 177–194.

[58] Y. Lindell, "Fast secure two-party ECDSA signing," *J. Cryptol.*, vol. 34, no. 4, 2021, Art. no. 44.

[59] R. Gennaro, S. Goldfeder, and A. Narayanan, "Threshold-optimal dsa/ecdsa signatures and an application to bitcoin wallet security," in *Proc. Int. Conf. Appl. Cryptogr. Netw. Secur.*, 2016, pp. 156–174.

[60] G. Golan Gueta et al., "SBFT: A scalable and decentralized trust infrastructure," in *Proc. 49th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, 2019, pp. 568–580.

[61] C. Stathakopoulou, T. David, and M. Vukolic, "Mir-BFT: High-throughput BFT for blockchains," 2019, *arXiv:1906.05552*.

[62] D. Boneh, B. Lynn, and H. Shacham, "Short signatures from the weil pairing," *J. Cryptol.*, vol. 17, no. 4, pp. 297–319, 2004.

[63] X. Yue, Q. Ding, J. Zhu, and Y. Ding, "TimeLink: Enabling dynamic runtime prediction for Flink iterative jobs," *J. Supercomputing*, vol. 80, no. 11, pp. 16546–16573, 2024.

[64] S. Dziembowski et al., "Perun: Virtual payment hubs over cryptocurrencies," in *Proc. IEEE Symp. Secur. Privacy*, 2019, pp. 106–123.

[65] E. K. Kogias, P. Jovanovic, N. Gailly, I. Khoffi, L. Gasser, and B. Ford, "Enhancing bitcoin security and performance with strong consistency via collective signing," in *Proc. USENIX Secur.*, 2016, pp. 279–296.

[66] H. Sun, H. Mao, X. Bai, Z. Chen, K. Hu, and W. Yu, "Multi-blockchain model for central bank digital currency," in *Proc. 18th Int. Conf. Parallel Distrib. Comput., Appl. Technol*, 2017, pp. 360–367.

[67] L. Luu, D.-H. Chu, H. Olickel, P. Saxena, and A. Hobor, "Making smart contracts smarter," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2016, pp. 254–269.

[68] P. Tsankov, A. Dan, D. Drachsler-Cohen, A. Gervais, F. Buenzli, and M. Vechev, "Securify: Practical security analysis of smart contracts," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2018, pp. 67–82.

[69] F. Zhang, E. Cecchetti, K. Croman, A. Juels, and E. Shi, "Town crier: An authenticated data feed for smart contracts," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2016, pp. 270–282.

[70] L. Akoglu, H. Tong, and D. Koutra, "Graph-based anomaly detection and description: A survey," *Data Mining Knowl. Discov.*, vol. 29, no. 3, pp. 626–688, 2015.

[71] M. Weber et al., "Anti-money laundering in bitcoin: Experimenting with graph convolutional networks for financial forensics," 2019, *arXiv:1908.02591*.

[72] P. Veličković et al., "Deep graph infomax," in *Proc. Int. Conf. Learn. Representations*, 2019, pp. 1–17.

[73] E. Manzoor, S. M. Milajerdi, and L. Akoglu, "Fast memory-efficient anomaly detection in streaming heterogeneous graphs," in *Proc. 22th ACM SIGKDD Int. Conf. Knowl. Discov Data Mining*, 2016, pp. 1035–1044.

[74] P. Daian et al., "Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges," 2019, *arXiv:1904.05234*.

[75] P. Daian et al., "Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability," in *Proc. IEEE Symp. Secur. Privacy*, 2020, pp. 910–927.

[76] K. Qin, L. Zhou, P. Gamito, P. Jovanovic, and A. Gervais, "An empirical study of DEFI liquidations: Incentives, risks, and instabilities," in *Proc. 21st ACM Int. Meas. Conf.*, 2021, pp. 336–350.

[77] M. Bartoletti, B. Pes, and S. Serusi, "Data mining for detecting bitcoin ponzi schemes," in *Proc. Crypto Valley Conf. Blockchain Technol.*, 2018, pp. 75–84.

[78] J. Bonneau, A. Narayanan, A. Miller, J. Clark, J. A. Kroll, and E. W. Felten, "Mixcoin: Anonymity for bitcoin with accountable mixes," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2014, pp. 486–504.

[79] A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou, "Hawk: The blockchain model of cryptography and privacy-preserving smart contracts," in *Proc. IEEE Symp. Secur. Privacy*, 2016, pp. 839–858.

[80] S. Bojja Venkatakrishnan, G. Fanti, and P. Viswanath, "Dandelion: Redesigning the bitcoin network for anonymity," in *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 1, no. 1, pp. 1–34, 2017.

[81] Ú. Erlingsson, V. Pihur, and A. Korolova, "RAPPOR: Randomized aggregatable privacy-preserving ordinal response," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2014, pp. 1054–1067.

[82] N. Narula, W. Vasquez, and M. Virza, "zkLedger: Privacy-Preserving auditing for distributed ledgers," in *Proc. USENIX Symp. Netw. Syst. Des. Implementation*, 2018, pp. 65–80.

[83] X. Yang, J. Hou, L. Xu, and L. Zhu, "zkFabLedger: Enabling privacy preserving and regulatory compliance in hyperledger fabric," *IEEE Trans. Netw. Service Manag.*, vol. 22, no. 2, pp. 2243–2263, Apr. 2025.

[84] C. Decker and R. Wattenhofer, "Information propagation in the bitcoin network," in *Proc. IEEE P2P Proc.*, 2013, pp. 1–10.

[85] A. Biryukov, D. Khovratovich, and I. Pustogarov, "Deanonymisation of clients in bitcoin P2P network," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2014, pp. 15–29.

[86] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, "Eclipse attacks on {Bitcoin's }{ peer-to-peer} network," in *Proc. USENIX Secur.*, 2015, pp. 129–144.

[87] M. Apostolaki, A. Zohar, and L. Vanbever, "Hijacking bitcoin: Routing attacks on cryptocurrencies," in *Proc. IEEE Symp. Secur. Privacy*, 2017, pp. 375–392.

[88] G. Naumenko, G. Maxwell, P. Wuille, A. Fedorova, and I. Beschastnikh, "Erlay: Efficient transaction relay for bitcoin," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2019, pp. 817–831.

[89] J. Garay, A. Kiayias, and N. Leonardos, "The bitcoin backbone protocol: Analysis and applications," *J. ACM*, vol. 71, no. 4, pp. 1–49, 2024.

[90] Y. Sompolinsky and A. Zohar, "Secure high-rate transaction processing in bitcoin," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2015, pp. 507–527.

[91] A. E. Gencer et al., "Decentralization in bitcoin and ethereum networks," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2018, pp. 439–457.

[92] V. Chandola et al., "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, no. 3, pp. 1–58, 2009.

[93] S. Ranshous et al., "Anomaly detection in dynamic networks: A survey," *Wiley Interdiscipl. Rev., Comput. Statist.*, vol. 7, no. 3, pp. 223–247, 2015.

[94] L. Akoglu, M. McGlohon, and C. Faloutsos, "oddball: Spotting anomalies in weighted graphs," in *Proc. Pacific-Asia Conf. Knowl. Discov. Data Mining*, 2010, pp. 410–421.

[95] M. U. Hassan, M. H. Rehmani, and J. Chen, "Anomaly detection in blockchain networks: A comprehensive survey," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 1, pp. 289–318, Firstquarter 2022.

[96] M. Bartoletti, S. Carta, T. Cimoli, and R. Saia, "Dissecting Ponzi schemes on ethereum: Identification, analysis, and impact," *Future Gener. Comput. Syst.*, vol. 102, pp. 259–277, 2020.

[97] W. Chen et al., "Detecting Ponzi schemes on Ethereum: Towards healthier blockchain technology," in *Proc. World Xwide Web Conf.*, 2018, pp. 1409–1418.

[98] A. Sharma, P. K. Singh, E. Podoplelova, V. Gavrilenko, A. Tselykh, and A. Bozhenyuk, "Graph neural network-based anomaly detection in blockchain network," in *Proc. Int. Conf. Comput., Commun., Cyber-Secur.*, 2022, pp. 909–925.

[99] S. Chen, Y. Liu, Q. Zhang, Z. Shao, and Z. Wang, "Multi-distance spatial-temporal graph neural network for anomaly detection in blockchain transactions," *Adv. Intell. Syst.*, vol. 7, no. 8, 2025, Art. no. 2400898.

[100] D. Eswaran and C. Faloutsos, "SedanSpot: Detecting anomalies in edge streams," in *Proc. IEEE Int. Conf. Data Mining*, 2018, pp. 953–958.

[101] T. Xie et al., "zkBridge: Trustless cross-chain bridges made practical," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2022, pp. 3003–3017.

[102] Y. Miao, K. Gai, Y.-a. Tan, L. Zhu, and W. Meng, "Blockchain-assisted searchable integrity auditing for large-scale similarity data with arbitration," *IEEE Trans. Dependable Secure Comput.*, vol. 22, no. 6, pp. 6012–6027, Nov./Dec. 2025.

[103] X. Yue, S. Yang, L. Zhu, S. Trajanovski, F. Li, and X. Fu, "Exploiting wide-area resource elasticity with fine-grained orchestration for serverless analytics," *IEEE/ACM Trans. Netw.*, vol. 33, no. 1, pp. 398–413, Feb. 2025.

[104] Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec, "GNNExplainer: Generating explanations for graph neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, vol. 32, pp. 9244–9255.

[105] D. Luo et al., "Parameterized explainer for graph neural network," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, vol. 33, pp. 1–15.

[106] T. Ruffing, P. Moreno-Sanchez, and A. Kate, "CoinShuffle: Practical decentralized coin mixing for bitcoin," in *Proc. Eur. Symp. Res. Comput. Secur.*, 2014, pp. 345–364.

[107] T. Ruffing et al., "P2P mixing and unlinkable bitcoin transactions," *Cryptol. ePrint Arch.*, 2016.

[108] E. Heilman, L. AlShenibr, F. Baldimtsi, A. Scafuro, and S. Goldberg, "TumbleBit: An untrusted bitcoin-compatible anonymous payment hub," in *Proc. NDSS*, 2017, pp. 1–15.

[109] J. Groth, "On the size of pairing-based non-interactive arguments," in *Proc. Annu. Int. Conf. Theory Appl. Cryptographic Techn.*, 2016, pp. 305–326.

[110] S. Bowe, A. Chiesa, M. Green, I. Miers, P. Mishra, and H. Wu, "ZEXE: Enabling decentralized private computation," in *Proc. IEEE Symp. Secur. Privacy*, 2020, pp. 947–964.

[111] B. Bünz, S. Agrawal, M. Zamani, and D. Boneh, "Zether: Towards privacy in a smart contract world," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2020, pp. 423–443.

[112] G. Fanti et al., "Dandelion lightweight cryptocurrency networking with formal anonymity guarantees," in *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 2, no. 2, pp. 1–35, 2018.

[113] M. Conti, E. S. Kumar, C. Lal, and S. Ruj, "A survey on security and privacy issues of bitcoin," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 4, pp. 3416–3452, Fourthquarter 2018.

[114] S. Quamara and A. K. Singh, "A systematic survey on security concerns in cryptocurrencies: State-of-the-art and perspectives," *Comput. Secur.*, vol. 113, 2022, Art. no. 102548.

[115] R. Zhang, R. Xue, and L. Liu, "Security and privacy on blockchain," *ACM Comput. Surv.*, vol. 52, no. 3, pp. 1–34, 2019.

[116] J. B. Bernabe, J. L. Canovas, J. L. Hernandez-Ramos, R. T. Moreno, and A. Skarmeta, "Privacy-preserving solutions for blockchain: Review and challenges," *IEEE Access*, vol. 7, pp. 164908–164 940, 2019.

[117] A. Satybaldy and M. Nowostawski, "Review of techniques for privacy-preserving blockchain systems," in *Proc. ACM Int. Symp. Blockchain Secure Crit. Infrastructure*, 2020, pp. 1–9.

[118] R. Yin, Z. Yan, X. Liang, H. Xie, and Z. Wan, "A survey on privacy preservation techniques for blockchain interoperability," *J. Syst. Architecture*, vol. 140, 2023, Art. no. 102892.

[119] H. Corrigan-Gibbs et al., "Prio: Private, robust, and scalable computation of aggregate statistics," in *Proc. USENIX Symp. Netw. Syst. Des. Implementation*, 2017, pp. 259–282.

[120] I. Dinur and K. Nissim, "Revealing information while preserving privacy," in *Proc. 22nd ACM SIGMOD-SIGACT-SIGART Symp. Princ. Database Syst.*, 2003, pp. 202–210.

[121] P. Mohassel and Y. Zhang, "SecureML: A system for scalable privacy-preserving machine learning," in *Proc. IEEE Symp. Secur. Privacy*, 2017, pp. 19–38.

[122] R. Cheng et al., "Ekiden: A platform for confidentiality-preserving, trustworthy, and performant smart contracts," in *Proc. IEEE Eur. Sympo. Secur. Privacy*, 2019, pp. 185–200.

[123] T. Neudecker and H. Hartenstein, "Network layer aspects of permissionless blockchains," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 838–857, Firstquarter 2019.

[124] F. Liu, "A statistical overview on data privacy," *Notre Dame JL Ethics Pub. Pol'y*, vol. 34, 2020, Art. no. 477.

[125] S. Bano et al., "SoK: Consensus in the age of blockchains," in *Proc. ACM Conf. Adv. Financial Technol.*, 2019, pp. 183–198.

[126] H. Kalodner et al., "BlockSci: Design and applications of a blockchain analysis platform," in *Proc. USENIX Secur.*, 2020, pp. 2721–2738.

[127] N. Atzei, M. Bartoletti, and T. Cimoli, "A survey of attacks on ethereum smart contracts (SoK)," in *Proc. Princ. Secur. Trust*, 2017, vol. 10204, pp. 164–186.

[128] G. Zhang, F. Pan, S. Tijanic, and H. Jacobsen Arno, "PrestigeBFT: Revolutionizing view changes in BFT consensus algorithms with reputation mechanisms," in *Proc. IEEE 40th Int. Conf. Data Eng.*, 2024, pp. 1930–1943.

[129] X. Wu, Z. Wang, X. Li, and L. Chen, "DBPBFT: A hierarchical PBFT consensus algorithm with dual blockchain for IoT," *Future Gener. Comput. Syst.*, vol. 162, 2025, Art. no. 107429.

[130] I. Mironov, "Rényi differential privacy," in *Proc. IEEE Comput. Secur. Foundations Symp.*, 2017, pp. 263–275.

[131] R. Bassily, A. Smith, and A. Thakurta, "Private empirical risk minimization: Efficient algorithms and tight error bounds," in *Proc. IEEE 55th Annu. Symp. Foundations Comput. Sci.*, 2014, pp. 464–473.
[132] Y.-X. Wang, B. Balle, and S. Kasiviswanathan, "Subsampled Rényi differential privacy and analytical moments accountant," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, 2019, pp. 1226–1235.
[133] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song, "The honey badger of BFT protocols," in *Proc. ACM Conf. Comput. Commun. Secur.*, 2016, pp. 31–42.
[134] A. Zamyatin et al., "SoK: Communication across distributed ledgers," in *Int. Conf. Financial Cryptogr. Data Secur.*, 2021, pp. 3–36.
[135] X. Li, P. Jiang, T. Chen, X. Luo, and Q. Wen, "A survey on the security of blockchain systems," *Future Gener. Comput. Syst.*, vol. 107, pp. 841–853, 2020.
[136] Y. Guo et al., "zkCross: A novel architecture for cross-chain privacy-preserving auditing," in *Proc. USENIX Secur.*, 2024, pp. 6219–6235.
[137] M. Cai et al., "Popularity-aware alignment and contrast for mitigating popularity bias," in *Proc. 30th ACM SIGKDD Int. Conf. Knowl. Discov Data Mining*, 2024, pp. 187–198.

Qingyang Ding received the Ph.D. degree from the Central University of Finance and Economics, Beijing, China, in 2020. He is currently an Associate Professor with the School of Management, Beijing Union University, Beijing. His research interests include federated learning, blockchain, and digital transformation.

Qinnan Zhang (Member, IEEE) received the Ph.D. degree from the Central University of Finance and Economics, Beijing, China, in 2023. She is currently a Postdoctoral Fellow with the Institute of Artificial Intelligence, Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing, Beihang University, Beijing. Her research interests include blockchain, federated learning, incentive mechanism, game theory, and edge intelligence.

Xiaofei Yue (Graduate Student Member, IEEE) received the M.E. degree from Northeastern University, Shenyang, China, in 2022. He is currently working toward the Ph.D. degree with the School of Computer Science and Technology, Beijing Institute of Technology, Beijing, China. His research interests include distributed systems, cloud/serverless computing, and data analytics.

Jianbin Chen received the Ph.D. degree from the Beijing Institute of Technology, Beijing, China, in 2005. He is currently a Professor with the School of Management, Beijing Union University, Beijing. His research interests include digital asset management and knowledge governance.

Ziming Zhao is currently a ZJU 100 Young Professor with the School of Software Technology, Zhejiang University, Hangzhou, China. He has authored or coauthored more than 50 papers in international journals and conference proceedings, including IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, MobiCom, SIGCOMM, DAC, WWW, INFOCOM, KDD, CCS, RTSS, IJCAI, IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, IEEE TRANSACTIONS ON MOBILE COMPUTING, IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, IEEE TRANSACTIONS ON NETWORKING, IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY, IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING, COSE, DATE, ESE, and AAAI. His research interests include machine learning, traffic identification, network security, and mobile computing.

Zhaoxuan Li is currently an Associate Researcher with the State Key Laboratory of Cyberspace Security Defense, Institute of Information Engineering (IIE), Chinese Academy of Sciences (CAS), Beijing, China. He has authored or coauthored more than 50 papers in international journals and conference proceedings, including IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY, IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING, IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, IEEE TRANSACTIONS ON MOBILE COMPUTING, IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, SIGCOMM, CCS, INFOCOM, and WWW. His research interests include blockchain security and privacy-preserving.

Xiaochen Hao received the bachelor's degree from the University of International Business and Economics, Beijing, China, in 2020. She is currently working toward the master's degree majoring in library and information science with Beijing Union University, Beijing. Her research interests include digital servitization and digital transformation.

Xinyu Zhang received the bachelor's degree in management from the Institute of Science and Technology, Hebei University of Science and Technology, Shijiazhuang, China, in 2024. She is currently working toward the master's degree majoring in library and information science with Beijing Union University, Beijing, China. Her research interests include data analysis and digital transformation.

Nan Wang received the bachelor's degree from Weifang University, Weifang, China, in 2022. She is currently working toward the master's degree majoring in business administration with the School of Management, Beijing Union University, Beijing, China. Her research focuses on digital asset management.