Scalability Meets Regulation: UTXO-Based Sharding and Zero-Knowledge Proofs for Regulated Digital Currencies

Si Yuan Jin[1,2], Yong Xia[2*], Bo Tong Xu[3]

[1]School of Business and Management, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong, China.

[2*]HSBC Laboratory, HSBC, Guangzhou, China.

[3]Department of Electronic Business, South China University of Technology, Guangzhou, China.

*Corresponding author. E-mail: yong.xia@hsbc.com;

Abstract

FinTech and RegTech integration is essential for developing scalable and compliant Regulated Digital Currency (RDC) systems. We propose an Unspent Transaction Output (UTXO)-based sharding method for scalability challenges in RDCs by minimizing cross-shard transactions. Results show UTXO-based sharding delivers linear throughput increase and consistently low latency. However, frequent user onboarding across shards under UTXO-based sharding complicates Know-Your-Customer (KYC) processes, highlighting regulatory inefficiencies. We find integrating zero-knowledge proofs can seamlessly streamline customer onboarding and overcome regulatory burdens. By doing so, we can deploy high-performance and compliant RDC systems.

Keywords: Regulatory Technology (RegTech), Financial Technology (FinTech), Scalability, Regulation

1 Introduction

Blockchain technologies have the potential to transform finance by enabling asset ownership digitalization and reconstructing the web's infrastructure [1]. Yet, regulatory complexities have constrained their broader adoptions [2]. Integrating RegTech tools can overcome these obstacles and create responsive compliance systems [3]. While FinTech continues evolving, critical knowledge gaps remain regarding how to ensure FinTech innovations and regulations. The current study examines how to ensure technology innovation and regulation in the emerging RDCs.

Most RDCs such as Central Bank Digital Currencies (CBDCs) commonly adopt the blockchain technology to record transfers of token ownership [1]. However, these systems face scalability challenges stemming from the potentially billions of parallel transaction requests from users [4]. To address the scalability limitations, sharding methods that divide a blockchain into multiple smaller sub-blockchains have been explored [5]. However, most sharding research has focused on the context of unregulated digital currency without considering regulatory requirements. We propose a novel solution that addresses both the scalability and regulation challenges for practical RDC systems.

Two main sharding approaches have emerged for blockchain systems: account-based and UTXO-based sharding [6]. In account-based sharding, accounts are partitioned into shards [7], which

can increase throughput [7] but induce extra cross-shard transactions that have high transaction latency due to the extra communication costs across shards, similar to the banking systems [8]. In UTXO-based sharding, UTXOs are partitioned into shards to minimize cross-shard transactions, thereby promising low latency [9]. However, for KYC regulation, UTXO-based sharding may face challenges in efficiently onboarding new users as RDC payees may have never transacted within certain shards.

To address the scalability and regulatory challenges of RDC systems, we find that zero-knowledge proofs fit well with the UTXO-based sharding approach that reconciles scalability with the regulatory requirements of RDCs. We first empirically verified UTXO-based sharding delivered linearly scalable transaction throughput. We then incorporated a zero-knowledge proof, which are a cryptographic protocol that allows for the verification of knowledge without revealing the underlying information [10]. The integration can mitigate KYC redundancies and simplify customer onboarding across shards.

Empirical tests demonstrate the advantages of our integrated UTXO-based sharding and zero-knowledge proofs approach over account-based alternatives. For scalability, UTXO-based sharding achieved linearly increasing throughput with the number of shards. In a simulated RDC system with 1 million wallets, UTXO-based sharding attained approximately 2000 Transactions Per Second (TPS), which shows a 2x throughput improvement over account-based sharding. For customer onboarding, our models reduces cross-shard identity verification redundancy as users can privately prove their identity across shards without repeatedly customer onboarding. In summary, by synergistically combining UTXO-based sharding for scalability and zero-knowledge proofs for efficient regulation, our methodology reconciles the performance and compliance requirements for practical RDC system deployment.

This research contributes to the FinTech literature by addressing scalability and compliance barriers for RDC [11]. While prior work recognizes these challenges [12], we present specific techniques integrating zero-knowledge proofs with UTXO-based sharding. Our solution enables privacy-preserving identity verification across shards while optimizing throughput. This novel application of cutting-edge information technologies to reconcile efficiency and regulation expands the boundaries of FinTech research.

Our paper also advances RegTech research. RegTech is an emerging research area that uses information technology and digital innovation to provide regulation more efficiently and effectively [13–15]. Although previous work primarily explores using Artificial Intelligence (AI) to streamline regulatory procedures [16–18] and ensure fairness [19] in other contexts, our research applies the blockchain to fortify regulation in RDCs. To our knowledge, this research represents one of the first RegTech applications in an RDC context, opening up new research directions at the intersection of FinTech and RegTech.

This research also contributes to the sharding literature. Sharding is a horizontal scaling technique, following the idea of divide and conquer [20]. With account-based sharding, cross-shard transactions exceed 99.98% when the number of shards reaches 16 [7], resulting from payers and payees rarely sharing shards [6]. While many efforts have been devoted to designing more efficient communication and processing methods, including the two-phase commit approach [6, 21, 22], transaction split-based approach [7], and relay transaction-based approach [23], our paper works to minimize the number of cross-shard transactions by leveraging the setting of RDCs. While some recent works have explored UTXO-based sharding methodologies for RDCs [24], our study differs by also considering regulatory efficiency as an additional design requirement.

The practical implications of our research are profound. By improving transaction performance and curtailing regulatory redundancies, we suggest a more efficient mechanism for implementing RDCs. This potentially influences practitioners, policymakers, and a wide array of stakeholders in the FinTech and RegTech industries.

2 Background

2.1 Regulated Digital Currency

RDCs are issued and regulated by authorized institutions and backed by stable-price assets like fiat money or traded commodities [25, 26]. For instance, CBDC is backed by fiat money and

issued by central banks [27]. Conversely, unregulated digital currency like Bitcoin use decentralized issuance via mining [28].

RDC systems can be either UTXO-based or account-based designs [29]. The distinction between both comes from the way to calculate the individual balance [30]. UTXO-based design calculates individual balances via UTXO accumulation, while the account-based method maintains an account ledger where each account has a real-time balance. UTXO, introduced by Bitcoin [28], refers to the unspent transaction outputs that users can transact with others. Most RDCs adopt UTXO-based designs, similar to digital cash [31].

Tokens can be fungible or non-fungible [32]. Fungible tokens, like cryptocurrencies, offer a standard medium of exchange, while non-fungible tokens uniquely represent assets. This study focuses on fungible tokens due to their enhanced functionality from value divisibility.

Tokens can be created and managed without using the blockchain [30]. However, the blockchain provides a decentralized, transparent ledger for recording token ownership and transfers, which prevents double-spending and makes the token transactions immutable [33]. So even though tokens are technically independent of the blockchain, the blockchain is commonly used as the underlying infrastructure when implementing tokens to benefit from its security properties.

2.2 Blockchain and Sharding

RDCs use the blockchain to eliminate trust issues and reduce information friction [34]. The blockchain can have different design options (Table 1). Most RDCs adopt permissioned blockchains for better regulation control. The permissioned blockchain can be public or private, depending on whether the transactions are openly visible [33]. If the blockchain is private, only authorized participants can view and submit transactions, thus the network is more efficient with limited traffic. If the blockchain is public, the general public can also view and submit transactions with the potential consequence of heavy traffic flow but a more innovative environment.

Improving blockchain performance can involve reducing communication and computation overhead, adding resources to a single node (vertical scaling), or adding more nodes to a network (horizontal scaling) [35]. Sharding, a horizontal scaling technique, attempts to minimize performance issues by dividing tasks [20].

Our work explores two types of sharding methods and their implications on scalability and regulation efficiency, as shown in Table 2. An account-based sharding method can efficiently meet regulatory requirements but it may face lots of cross-shard transactions, while a UTXO-based sharding method can minimize cross-shard transactions. However, the UTXO-based sharding system faces the challenges of frequent KYC procedures for onboarding new users. We used zero-knowledge proofs to tackle the regulation inefficiency problem.

2.3 Regulation Technology

RegTech leverages modern technologies to make regulatory compliance more efficient and effective [14], handling challenges introduced by large data throughput and size [36]. A report by Thomson Reuters highlights consistent year-over-year increase in the number of regulatory alerts, underlining the growing complexity of compliance [1].

One pivotal aspect of regulatory compliance is the KYC process, which is designed to verify the identities of clients [4]. The KYC is vital for preventing financial crimes including money laundering and the financing of terrorism [3].

The advance of FinTech presents significant challenges to KYC for retail banks due to its need for considerable technology, staff training, and monitoring investments [15, 37, 38]. Emerging technologies like distributed ledger technology and smart contracts propose solutions to these challenges [39, 40]. Despite its broad coverage, research exploring RDC implementations in RegTech remains limited. This research introduces a regulatory model for a sharded two-tier RDC, leveraging zero-knowledge interactive proofs to speed onboarding new users in UTXO-based sharding systems [41–43].

Zero-knowledge proof uses a limited amount of information from a prover to a verifier [44] and has been used to verify identify information [10]. Zero-knowledge proofs have been used in digital currency protocols to efficiently verify the identity

[1] Regulatory Intelligence Feeds, Thomson Reuters (2020)

Table 1: Comparison of different blockchain designs for RDCs.

	Public	**Private**
Permissionless	Anyone can view and submit transactions; Every node can be validator;	-
Permissioned	Anyone can view and submit transactions; Authorized nodes can be validators;	Authorized nodes can view and submit transactions; Authorized nodes can be validators;

Table 2: Comparison of UTXO-based and account-based sharded blockchain.

	Account-based sharding	**UTXO-based sharding (Our contribution)**
Definition	Accounts are partitioned into shards;	UTXOs are partitioned into shards;
Transaction	More cross-shard transactions (high latency)	Less cross-shard transactions (low latency)
Customer Onboarding	No new customer onboarding; cross-shard KYCs	Frequent new customer onboarding (enhanced by zero-knowledge proof)
UTXO Searching	Account-based search - low search cost	UTXO-based search - high search cost (optimized by wallet ID with UTXOs*)

* We use "Wallet ID" to present the public address of UTXO owners (see details in Appendix A).

information of digital currency users [45]. While prior studies mainly focus on using zero-knowledge proofs to address privacy concerns [45–47], our study leverages zero-knowledge proofs to address the inefficiency of the frequent user onboarding from UTXO-based sharding design.

3 Methodology

3.1 Sharded Blockchain

The RDC system adopts a two-tier model [48] (Fig. 1), consisting of a single private and permissioned tier-1 network, led by a tier-1 leader, and multiple public and permissioned tier-2 networks, each led by a tier-2 leader [49]. Any customer can view and submit transactions to tier-2 networks, but do not operate nodes since they lack fixed servers or IP addresses. Tier-2 networks can issue RDCs that are backed by reserves held within the tier-1 network. Tier-2 servers preprocess transaction requests from customers and submit these to their corresponding tier-2 network leader.

3.2 Transaction

Fig. 2 shows a transaction within a tier-2 network, where public user A pays RDC UTXO(1) to public user B, generating new RDCs for both public users within the same tier-2 network. At time n, user

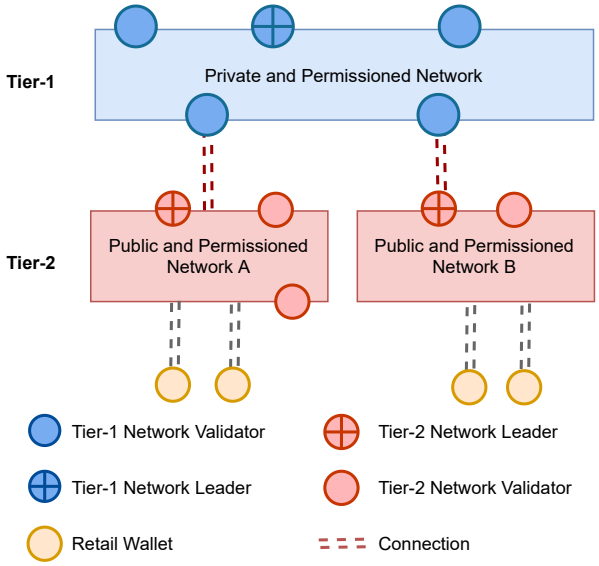

Fig. 1: An RDC two-tier hybrid network.

A has 10$ (example amount) RDC token labeled UTXO(1). To pay user B with 6$, the system first destroys UTXO(1). Then the network mints a new 6$ RDC token UTXO(6) and assigns its ownership to user B. It also mints a 4$ RDC change token UTXO(5) and returns it to user A. After the transaction at time n+1, user B now has the new 6$ RDC token UTXO(6). User A has the 4$

RDC change token UTXO(5). The transaction is recorded on the ledger of the tier-2 network.

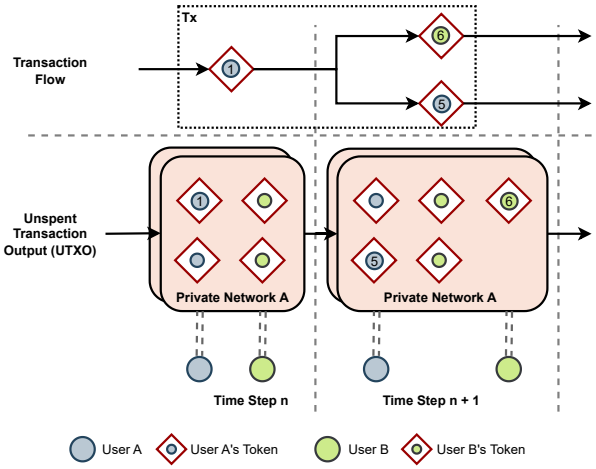

Fig. 2: A transaction within a single tier-2 network: user A pays RDC UTXO(1), to user B. The transaction generates a new RDC UTXO(6) for user B and a change RDC UTXO(5) for User A. Both new RDCs remain in tier-2 network A.

Our UTXO-based sharding method can execute transactions in parallel for better scalability and minimize the number of cross-shard transactions for low transaction latency. Fig. 3 shows that sharding could serve as a solution by distributing transaction requests across different tier-2 networks. UTXO-based sharding partition tokens into shards. Public customers can connect to different tier-2 networks based on their token information and spend their RDCs by sending requests to the corresponding network. In comparison, account-based sharding can also improve system capacity, but additional shards might increase the number of cross-shard transactions [50], which can in turn impact transaction latency.

Algorithm 1 describes the transaction process in the account-based sharding method. The inputs include the payer wallet ID (P), payee wallet ID (Q), transaction amount (A), account-to-shard hashmap (S_{account}), UTXO (U), and tier-2 network leader (N). The sharding hashmap S_{account} identifies the sharded networks N_P and N_Q for P and Q, respectively. N_P then identifies all available RDCs for P and verifies the sufficiency of funds against amount A. If the payer and payee are on the same sharded network ($N_P == N_Q$), N_P can update its ledger directly. However, if $N_P \neq N_Q$, N_P must coordinate with N_Q to update both ledgers, which requires interoperability techniques, like hashed time lock smart contract [51] for synchronization.

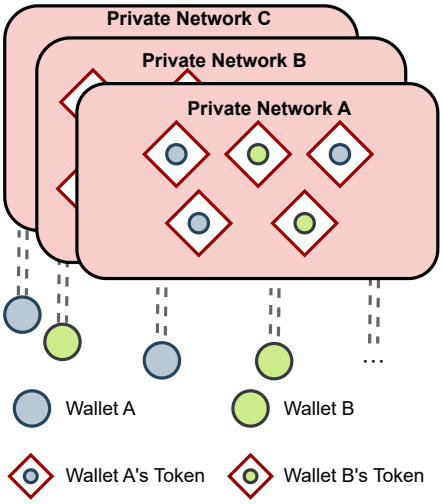

Fig. 3: Sharded blockchain: token distribution.

Algorithm 1 Account-based sharding transaction

Require: payer wallet ID (P), payee wallet ID (Q), transaction amount (A), account-based sharding hashmap (S_{account}), UTXO (U), network leader (N).
Ensure: transaction result
$N_P \leftarrow S_{\text{account}}[P]$
$N_Q \leftarrow S_{\text{account}}[Q]$
$U_P \leftarrow$ Searching P's available RDCs in N_P
if $A >$ Total amount of U_P **then**
Return transaction failure due to insufficient funds
end if
if $N_P == N_Q$ **then**
N_P updates its ledger by creating two new UTXO tokens: one transfers amount A to payee Q and one returns change to payer P
else
N_P and N_Q coordinate to update ledgers across networks
end if
Return transaction success

Account-based sharding partitions accounts across shards. As the number of shards (n)

increases, the likelihood of cross-shard transactions also rises, approaching $\frac{n-1}{n}$. With 16 shards, cross-shard transfers already exceed 99.98% [7].

In RDC systems with many tier-2 networks, account-based sharding will frequently encounter inter-shard transaction requests. Each cross-shard transfer requires coordination between the payer and payee's networks to synchronize ledger updates. This communication overhead can substantially increase transaction latency as the shard count scales up.

Therefore, while additional shards improve throughput via parallelization, the explosion of cross-shard transactions under account-based approaches ultimately hampers performance gains due to the high latency. This scaling limitation motivates exploring alternative sharding methodologies tailored for RDC system properties.

Algorithm 2 UTXO-based sharding transaction

Require: payer wallet ID (P), payee wallet ID (Q), transaction amount (A), UTXO-based sharding hashmap (S_{UTXO}), UTXO (U), network leader (N).
Ensure: transaction result
$U_P \leftarrow$ Searching P's available RDCs in N_P
$N_P \leftarrow S_{\text{UTXO}}[U_P]$
if $A >$ Total amount of U_P **then**
Return transaction failure due to insufficient funds
end if
N_P updates its ledger by creating two new UTXO tokens: one transfers amount A to payee Q and one returns change to payer P
Return transaction success

Algorithm 2 outlines the UTXO-based sharding transaction process. The inputs include the payer wallet ID (P), payee wallet ID (Q), transaction amount (A), UTXO-to-shard hashmap (S_{UTXO}), UTXOs (U) and network leader (N). Each sharded network manages a distinct set of RDC tokens. To spend an RDC, the payer P must send the request to the network N_P, determined by the UTXO-based sharding mapping (S_{UTXO}) using the RDC ID U_P rather than the wallet ID (or user ID). Since RDCs are allocated to shards based on their ID, the network can directly modify the RDC record, avoiding cross-shard transactions. When the payee later utilizes the received RDC, they likewise route based on the RDC ID to the same network, which allows different networks to process transactions concurrently without additional cross-shard communications.

While UTXO-based sharding minimizes cross-shard traffic and enhances performance, it complicates KYC regulations if the payee is not already a customer within the payer's shard. To address this, we propose integrating zero-knowledge proofs for identity verification, which allows fast onboarding of new payees across shards without repetitive manual processes.

3.3 Streamlined Customer Onboarding

The onboarding processes often face redundancy due to a lack of a unified data-sharing protocol. We propose using a zero-knowledge proof within our sharded two-tier regulated network to minimize regulatory redundancy [52].

The tier-1 network secures the entire network, and asks the tier-2 network leader to collect the necessary KYC information when a public customer opens a wallet. In practice, tier-1 network leaders act as regulators of the networks, possessing strong incentives to oversee regulation. Tier-2 networks also would like to cooperate with the tier-1 network to maintain a good reputation [53]. The KYC data can be encrypted and stored via the zero-knowledge proof on the tier-1 network, which allows other tier-2 networks to onboard the same customers without real-time manual verification by using a zero-knowledge proof within the tier-1 network.

When a new user registers with a wallet app and completes KYC onboarding on one shard network, a zero-knowledge proof associated with their verified identity is created. This proof can then be used when the user needs to transact on other shard networks where they have not yet been onboarded. Instead of repeating lengthy KYC processes, the user simply presents the zero-knowledge proof, allowing the new shard network to privately verify their identity. No personal information is revealed to the verifying network. In this manner, zero-knowledge proofs mitigate redundant KYC procedures and simplify cross-shard onboarding. The user enjoys faster activation on new networks while preserving their privacy. Regulators also benefit from rigorous unified KYC

standards applied cohesively across all shards via the shared zero-knowledge proofs.

Fig. 4 outlines the step-by-step process for a customer to acquire RDC through the three-layer issuance system:

1. The customer initiates a request via the wallet app to obtain RDC using conventional funds like cash.
2. The app checks if the customer already has a registered wallet ID. If yes, it sends the RDC issuance request to the customer's tier-2 network server. If no ID exists, the app collects the required KYC information and forwards it to the tier-2 server.
3. For an existing ID, the tier-2 server mints the RDC on its ledger and confirms success to the customer. If no wallet ID exists, the tier-2 server requests the tier-1 network to verify the customer's KYC information using a zero-knowledge proof. Upon successful verification, the tier-2 server registers a new wallet ID, mints, and deposits the RDC, and confirms success. If verification fails, the customer has to re-provide KYC details and repeat the process.
4. In the KYC process, the tier-2 server registers the wallet ID and an associated KYC zero-knowledge proof, sharing it with both the customer and the tier-1 network for future verification needs.

Fig. 5 illustrates the transaction process, which combines UTXO-based sharding with KYC zero-knowledge proofs to meet both performance and regulatory requirements. As Algorithm 2 shows, the payee may not be onboarded to the payer's network (A). To verify the payee's identity, the tier-1 network provides a zero-knowledge proof that tier-2 networks can verify.

After each transaction, network A shares zero-knowledge proofs with the tier-1 network for the further onboarding process across other tier-2 networks. The detailed description is as follows:

- The payer initiates a transaction with the payee.
- The payee confirms the transaction, and provides their wallet ID and associated KYC zero-knowledge proof to the payer, assuming the payee has previously completed the KYC process within the tier-2 network.
- The payer submits the information to its sharded network A. Network A verifies if the payee has completed KYC within its network. If not, it uses the KYC zero-knowledge proof to validate the payee's onboarding information with the tier-1 network.
- Upon confirming the KYC information, the tier-2 network updates its ledger with the transaction. Then, the transaction is successful. Otherwise, network A requests the payee's KYC information.

In summary,our methodology employs UTXO-based sharding to minimize cross-shard transactions and enhance scalability. Integrating zero-knowledge proofs for KYC enables quick, private identity verification to onboard new users across shards, overcoming regulatory inefficiencies. This novel combination of UTXO-based sharding to enhance performance and zero-knowledge proofs to streamline regulation provides a holistic solution tailored for scalable and compliant RDC system deployment.

4 Experiment Design

Our model requires a UTXO-based blockchain system, and we used Corda [54] to construct an RDC infrastructure. Corda, a permissioned blockchain platform, maintains a UTXO data structure, which aligns with UTXO-based sharding requirements. Other blockchain platforms, such as Besu, operate as an account-based system and lack a UTXO data structure design.

Transactions require signatures from both the payee and the payer. Customers submit transaction requests to network servers that process transactions with merchant nodes. Our experiments simulate requests by randomly generating them with real-world merchant time intervals to closely emulate application conditions.

We implemented our method to evaluate scalability and latency, as shown in Fig. 6. The experimental network includes multiple tier-2 network leaders who are responsible for transaction validation and tier-2 network validators that ensure leader validity. We used Jmeter[2] to simulate public customers, initiating transaction requests to the tier-2 network via a Remote Procedure Call (RPC) port. Each tier-2 network has multiple nodes (P) and each node P has the corresponding smart contracts (S) and ledger (L).

[2]https://jmeter.apache.org/

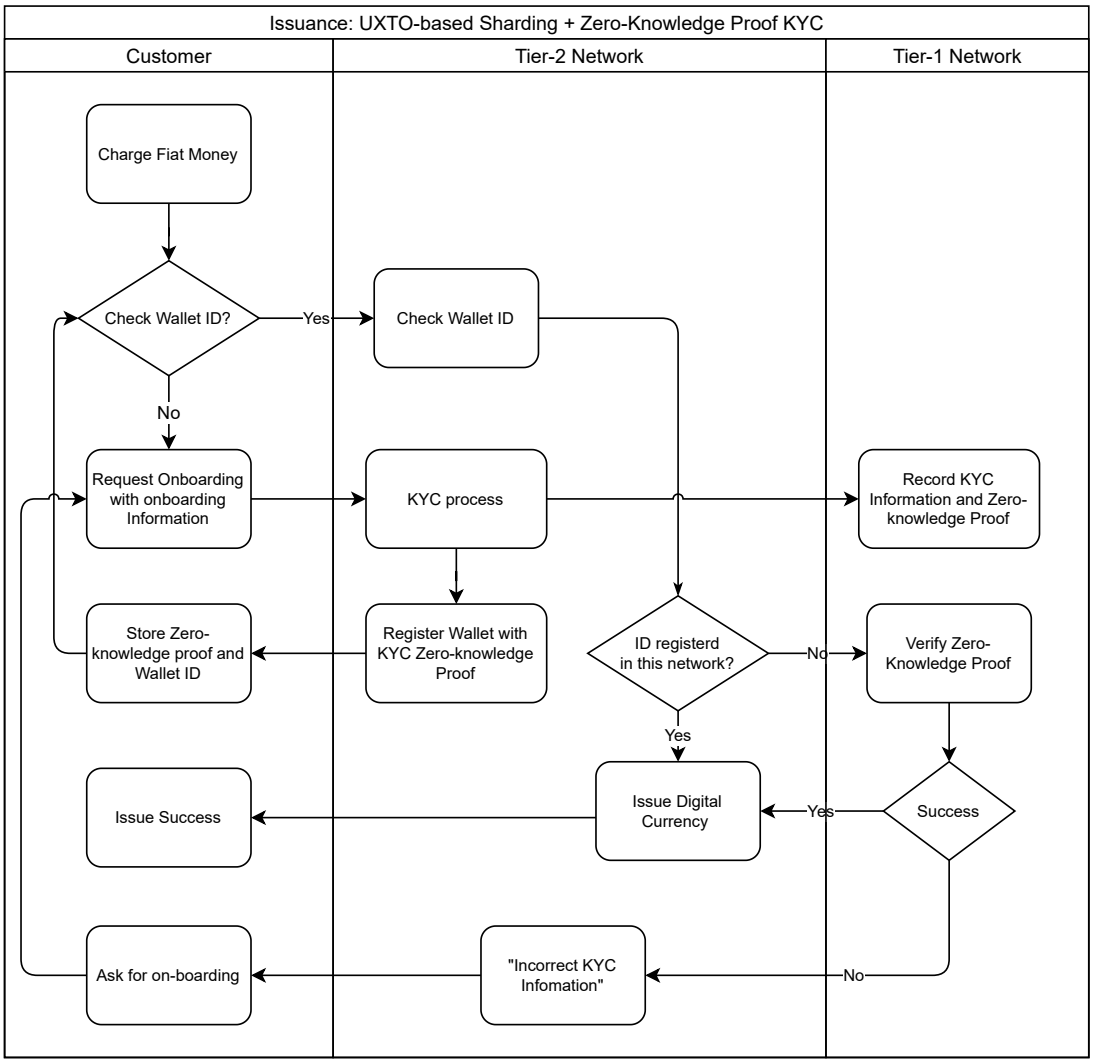

Fig. 4: Issuance: UTXO-based sharding + zero-knowledge proof KYC.

We simulated two types of merchants to cover various payment scenarios: HHP merchants and HMP merchants, both of which can own their nodes to process transactions or own wallets to connect with network servers.

- Human-Human Payment (HHP) merchants: Both payer and payee are humans and they use human-human interaction to conduct the transaction.
- Human-Machine Payment (HMP) merchants: The payer is a human while the payee could be a machine, such as a vending machine.

Each experiment followed these steps:

- The tier-2 network leader distributed 1,000 RDCs to every wallet in its network to initialize balances.
- We evenly distributed customer tokens across the tier-2 network servers to balance the load for later UTXO-based sharding.
- A script randomly selected customer wallets to initiate transactions with merchants.
- As the simulation was finished, we captured performance data including latency and throughput from the servers.

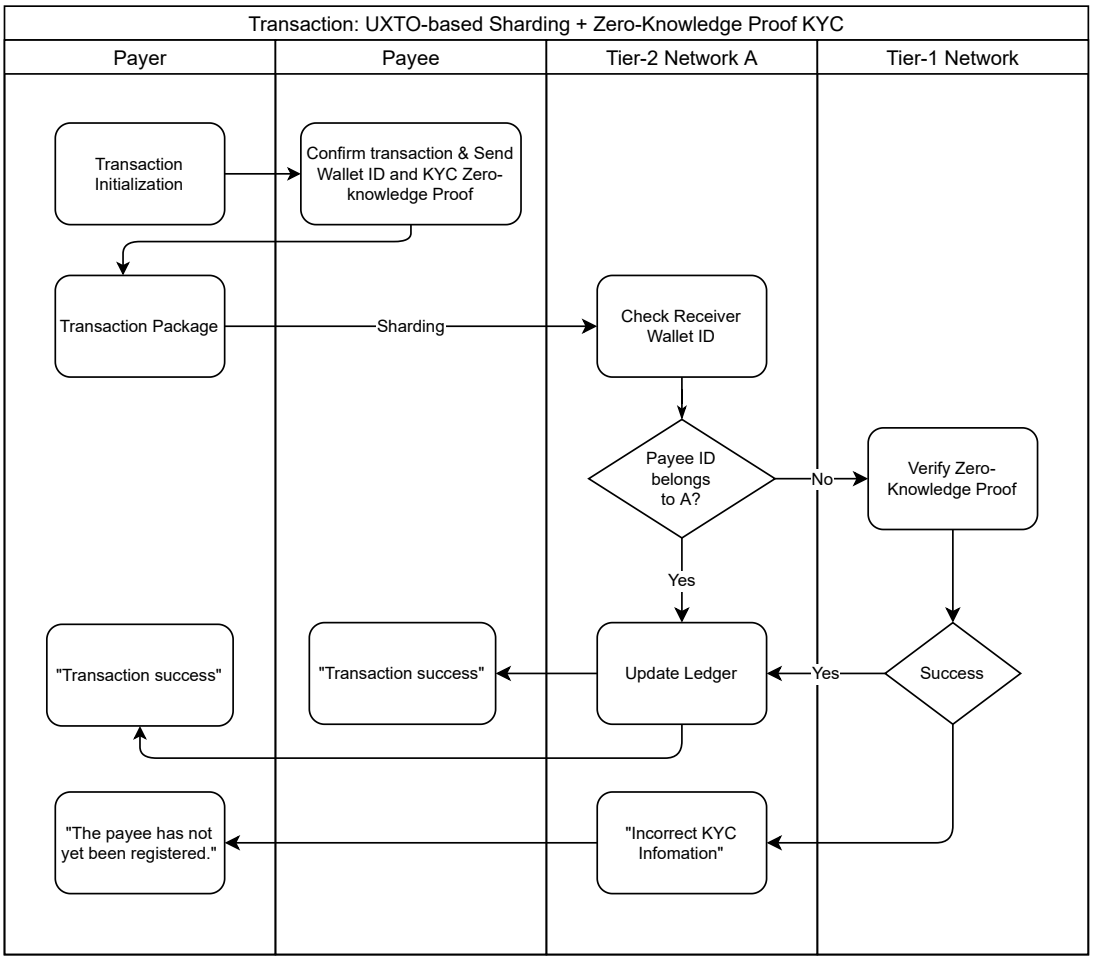

Fig. 5: Transaction: UTXO-based sharding + zero-knowledge proof KYC.

Experiments were conducted on a cloud infrastructure (Fig. 7) with 50 Amazon EC2 servers[3] to simulate different networks. Each EC2 server was connected with an Relational Database Service database[4] to record transaction ledger data. We used a single Admin EC2 instance to distribute the network configuration settings and execute commands to run the simulations across the 50 servers. We leveraged Spring Boot to implement node APIs for request handling. AWS CLI scripts deployed and managed the test network.

[3]https://aws.amazon.com/ec2/instance-types/t2/
[4]https://aws.amazon.com/rds/

5 Results

Account-based systems can efficiently find balances using unique wallet IDs (account addresses). However, in UTXO-based systems, searching by token IDs is inefficient due to the huge number of UTXOs. Instead, we used wallet IDs (addresses) as basic identity units (see clarification in Appendix A). We create a table mapping token IDs to wallet IDs. When searching tokens, we first retrieved the wallet ID, then used token IDs to do the second-round search.

As shown in Fig. 8, our wallet-based search achieves near-constant latency versus linear growth for UTXO-based searching. By enabling efficient token retrieval from networks, our improvement enables the practical deployment of UTXO-based systems.

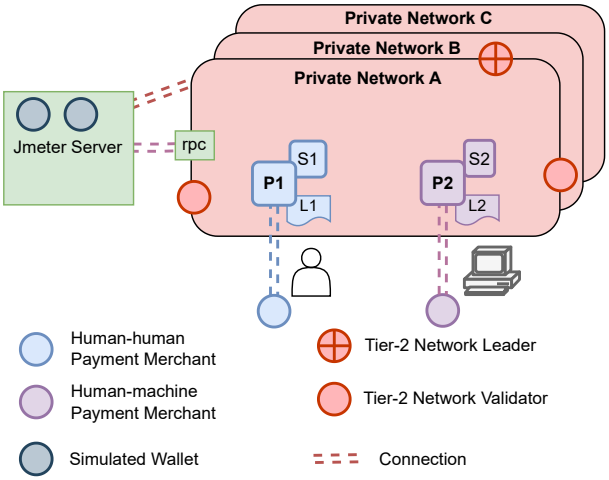

Fig. 6: Experiment architecture.

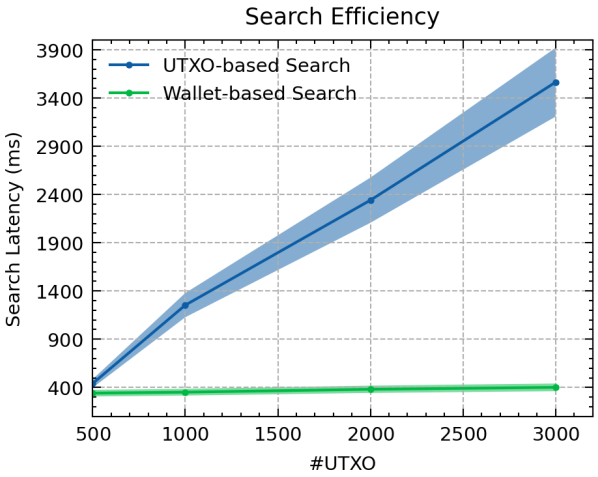

Fig. 8: Search cost comparison.

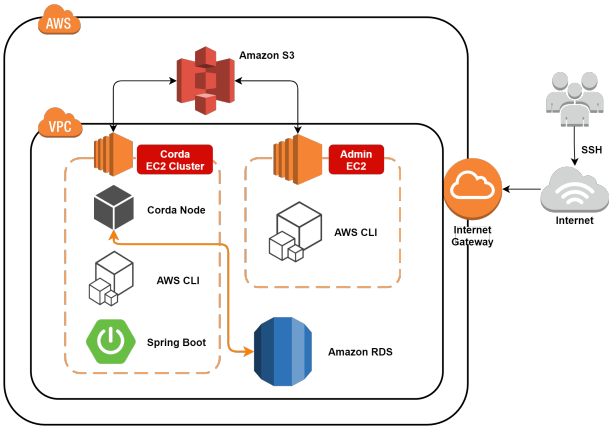

Fig. 7: Cloud deployment.

To evaluate the transaction processing capacity of each Corda node, we utilized the control variable method (Table 3). Since open-source Corda only supports single-threaded nodes, we used AWS t2.medium virtual machines (2 vCPUs, 4GB RAM) to simulate node performance [5]. Our results in Table 4 show that each Corda network server sustained approximately 4.5 TPS. Tier-2 validator nodes achieved 8.4 TPS. In comparison, enterprise Corda typically reaches over 200 TPS, equivalent to 8.4 TPS in our open-source system.

Field experiments were conducted to determine transaction intervals for HHPs and HMPs and found: 20 seconds for HHPs and 10 seconds for HMPs. Our simulation incorporated these real-world timings where we initiate transactions with time intervals according to the merchant types.

After establishing the node performance baseline, we evaluated overall network transaction scalability and latency. To achieve a target throughput of 100 TPS, our network configuration comprised 24 network servers, 12 tier-2 network leaders, HHP/HMP merchants, and approximately one million wallets.

As we increased the network/shard count, our UTXO-based approach demonstrated superior linear scalability versus account-based sharding. As shown in Fig. 9, with just 12 shards, we achieved 100 TPS with nearly constant latency, equivalent to 2,000 TPS with Corda Enterprise nodes. Fig. 10 depicts the latency distribution, which remained low despite the increase in load. Notably, the account-based sharding method experienced relatively high latency due to cross-shard transactions.

6 Discussion

Overall, sharding improves the performance of RDC systems in three primary ways:

1. High throughput: Sharding enables concurrent transaction handling across servers, boosting throughput. This parallelization allows the system to scale linearly with additional shards.
2. Low latency: Distributing requests across shards maintains low latency despite heavy

[5] https://aws.amazon.com/ec2/instance-types/t2/

Table 3: Experiment deployment.

No.	Node deployment			
1	1 Network server	4 HHP merchant	2 HMP merchant	1 Tier-2 network leader
2	2 Network server	4 HHP merchant	2 HMP merchant	1 Tier-2 network leader
3	4 Network server	4 HHP merchant	2 HMP merchant	1 Tier-2 network leader
4	2 Network server	4 HHP merchant	2 HMP merchant	2 Tier-2 network leader
5	4 Network server	4 HHP merchant	2 HMP merchant	2 Tier-2 network leader

Table 4: Experiment performance measurement.

No.	TPS	Average latency (ms)	90th percentile latency (ms)
1	4.7	2,026	2,485
2	8.1	2,452	2,743
3	8.4	3,482	3,804
4	9.1	1,985	3,514
5	19.4	2,910	3,476

loads. By evening out transaction volumes, shards prevent congestion and delays.

3. Low cost: Horizontal scaling is generally cheaper than vertical scaling with hardware limits. Adding low-cost shards is more economical for meeting demand surges [55].

UTXO-based sharding facilitates RDC parallel circulation. The experiment validates our proposed method's linearly increasing scalability and low latency in diverse scenarios, supporting robust and efficient transaction processing. Critically, we demonstrate near-constant latency up to 100 TPS with 12 shards. This reliability despite increasing loads showcases the effectiveness of the approach.

Our experiments utilized a static UTXO allocation to shards. In practice, transaction volumes fluctuate over time, creating variable loads across shards. An intelligent dynamic load balancer could monitor traffic in real-time and automatically redistribute UTXOs to evenly distribute the workload. This would prevent congestion on overloaded shards to further improve latency.

While users may initially have their UTXOs distributed across shards, we provide consolidation functions that can merge split tokens into unified holdings within a single shard. This consolidation helps optimize performance by reducing cross-shard transactions.

Another scenario involves users with multiple addresses wanting to jointly spend funds across their various holdings. This requires multi-party signatures to atomically process such transactions, even when all parties belong to one legal entity. However, this complex multi-party case is beyond the scope of our current model. A simpler alternative allows users to first transfer their tokens into a single address under their control. They can then readily spend the consolidated funds from one unified wallet, avoiding multi-party coordination.

We do not include customer onboarding during performance testing, as it remains relatively independent from transaction throughput and latency tests. The onboarding process for a new customer typically consumes different times depending on the type of the RDC wallet. China's e-CNY requires an onsite KYC process for one type of wallet [56]. Thus, its impact on transaction performance testing is hard to measure. Our zero-knowledge proof method can streamline the process by avoiding onsite KYC processes.

7 Conclusion

Our research proposes a UTXO-based sharding method for enhancing the scalability and streamlining the KYC process of RDCs. The application of our model can extend not only to RDC systems but also to other regulated digital asset infrastructures, including digital bonds. Our research has shed light on the possibilities of implementing RegTech within the FinTech landscape. By aligning regulatory compliance with scalability requirements, we believe our research opens doors

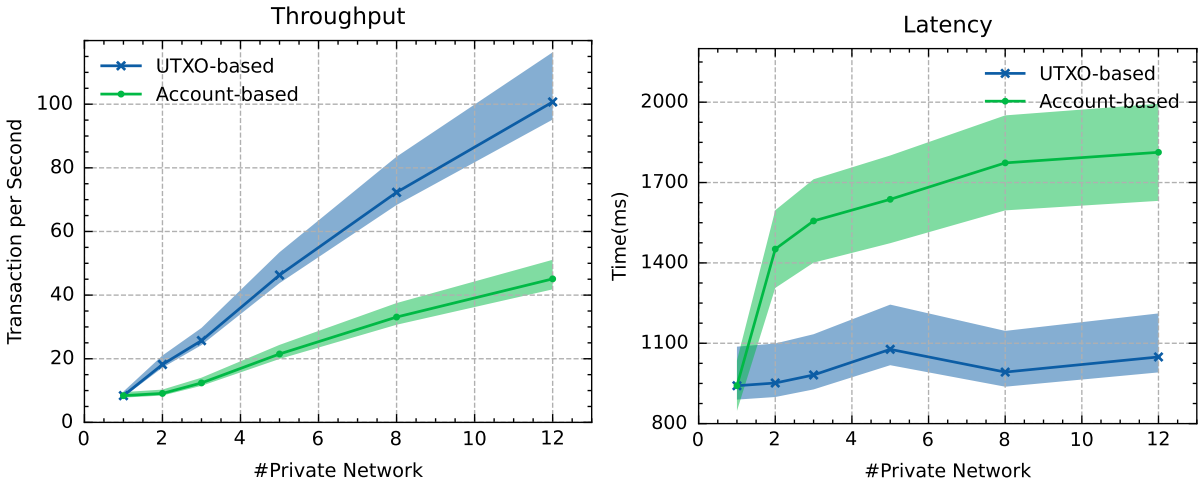

(a) Throughput.

(b) Latency.

Fig. 9: Performance comparison.

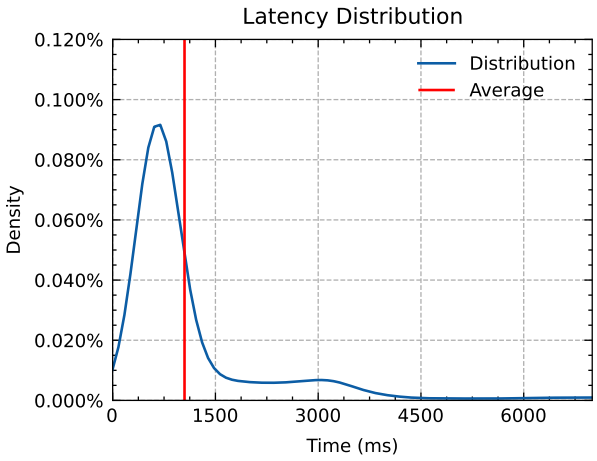

Fig. 10: Latency density with 12 tier-2 networks: the transaction latency is not significantly impacted and has an average value of 1.05 seconds.

for more efficient and secure RDC ecosystems and future RegTech and FinTech applications.

For central banks and monetary authorities, our system offers a path towards implementing digital fiat currencies at a nationwide scale. The ability to securely onboard millions of users while maintaining transaction latency is pivotal for the viability of central bank digital currencies. Policymakers can also benefit from streamlined compliance procedures that reduce redundancies without sacrificing rigor.

For incumbent financial institutions and new FinTech entrants, our techniques unlock opportunities to build innovative services and products leveraging regulated digital tokens. The improved scalability empowers banks to deliver next-generation payment solutions, while simplified identity verification facilitates seamless customer onboarding.

Consumers stand to gain convenient access to regulated digital money that provides the security of sovereign backing. Merchants and businesses will also benefit from faster settlement times and reduced transaction costs. Beyond payments, tokenized assets can expand to encompass bonds, securities, and smart contracts.

While the results have shown the effectiveness of our model in improving scalability and regulation efficiency, further studies are needed to explore other implications. As seen in ongoing RDC projects, blockchain faces trade-offs between scalability, privacy, and resilience [5,27,57]. Thus, future investigations can delve deeper into other facets of the UTXO-based sharding method, such as security and privacy. These aspects are pivotal in evaluating the model's potential for broader deployment in the digital finance field.

Thus far we have focused on simple transaction smart contracts. However, RDCs may also need to

enable more complex smart contract functionality for programmatic transactions. Sharded smart contract systems introduce additional complexity around cross-shard communications. Our model could be expanded to shard smart contract execution while keeping data local to optimize efficiency. Future work can study the detailed design of such systems for practical infrastructure building.

A Terminologies

RDC systems employ a variety of technical jargon that can be misinterpreted. This section aims to clarify these terms.

Address vs wallet

In both UTXO and account-based models, unique addresses represent user identities. In UTXO-based systems, addresses (public keys) are associated with specific UTXOs. In account-based systems, addresses (public keys) are linked to accounts. In this paper, "wallet ID" refers to a user's public address. In Algorithm 1, "wallet ID" denotes an account's public address. In Algorithm 2, "wallet ID" represents a UTXO owner's public address.

The original Bitcoin whitepaper by Satoshi Nakamoto [28] did not introduce the term "wallet" but referred to individuals as "owners of public keys". In contrast, the Ethereum whitepaper, as detailed by Buterin (2014) [58], acknowledges the concept of "wallet" in a UTXO-based context where "wallet contains UTXO", and distinguishes Ethereum's account-based approach where identities are tied to accounts with unique 20-byte public addresses. To avoid confusion, "wallet ID" in this paper refers to the public address of a user.

UTXO-based vs account-based

RDC systems can use either a UTXO-based model or an account-based model for asset verification and balance calculation [4]. These terms refer to the underlying authentication methods, not the assets themselves. While RDC is still emerging, there is a tendency among some academics to conflate "UTXO-based" systems with "token-based" platforms. However, token-based platforms primarily utilize digital tokens to establish and transfer asset ownership, as outlined by [1]. The key difference is that the token-based systems emphasize token usage, while UTXO and account-based models focus on the way of balance validation.

Acknowledgments. The author wishes to thank the anonymous reviewers whose thoughtful critiques and suggestions greatly improved the quality of this paper. S. Jin also wishes to thank his two undergraduate supervisors for their invaluable guidance on his undergraduate thesis, which formed the basis for this paper. Their mentorship has been instrumental in the author's academic growth. S. Jin also expresses his profound gratitude to his partner and family for their unwavering support, patience, and belief in him throughout his studies and research.

Declarations

- Funding: No funding was received to assist with the preparation of this manuscript.
- The authors have no relevant financial or non-financial interests to disclose.
- Ethics approval: Not applicable.
- Consent to participate: Not applicable.
- Consent for publication: Not applicable.
- Availability of data and materials: The data produced during this study have been disclosed in the tables and figures. Detailed data can be made available from the authors upon reasonable request.
- Code availability: The code produced during the current study is not publicly available due to confidentiality but is available from the authors on reasonable request.
- Authors' contributions: All authors contributed to the study conception and design. S. Jin performed material preparation, algorithm design, data collection, and analysis under the supervision of Y. Xia. B. Xu provides the AWS computing resources for the experiment. S. Jin wrote the first draft of the manuscript. All authors read and approved the final manuscript.

References

[1] Jin, S.Y., Li, Z.T., Huang, H.A., Tam, K.Y.

Token-Based Platforms and Green Dilemma: Examining the Role of Community Perceptions and Web Page Environmental Disclosures. *Available at SSRN 4569995*, 2023.

[2] Henderson, M.T., Raskin, M. A regulatory classification of digital assets: toward an operational Howey test for cryptocurrencies, ICOs, and other digital assets. *Colum. Bus. L. Rev.*, page 443, 2019.

[3] Hong Kong Monetary Authority and KPMG. Transforming risk management and compliance: Harnessing the power of regtech. Technical report, H.K.M.A whitepaper, 2020.

[4] Allen, S., Capkun, S., Eyal, I., Fanti, G., Ford, B.A , Grimmelmann, J., Juels, A., Kostiainen, K., Meiklejohn, S., Miller, A., and others. Design choices for central bank digital currency: Policy and technical considerations. Technical report, NBER, 2020.

[5] Jin, S.Y., Xia, Y. CEV Framework: A Central Bank Digital Currency Evaluation and Verification Framework With a Focus on Consensus Algorithms and Operating Architectures. *IEEE Access*, 10:63698–63714, 2022.

[6] Liu, Y., Liu, J., Salles, M.A.V., Zhang, Z., Li, T., Hu, B., Henglein, F., Lu, R. Building blocks of sharding blockchain systems: Concepts, approaches, and open problems. *Comput. Sci. Rev.*, 46, 2022.

[7] Zamani, M., Movahedi, M., Raykova, M. RapidChain: Scaling blockchain via full sharding. In *ACM Conference on Computer and Communications Security (CCS)*, pages 931–948, 2018.

[8] Huang, H., Peng, X., Zhan, J., Zhang, S., Lin, Y., Zheng, Z., Guo, S. BrokerChain: A cross-shard blockchain protocol for account/balance-based state sharding. In *IEEE INFOCOM*, pages 1968–1977, 2022.

[9] Jin, S.Y., Xu, B.T., Intallura, P., Xia, Y. A UTXO-based Sharding Method for Stablecoin. In *International Conference on Blockchain Computing and Applications (BCCA)*, pages 195–199, 2022.

[10] Fiege, U., Fiat, A., Shamir, A. Zero knowledge proofs of identity. In *ACM Symposium on Theory of Computing (STOC)*, pages 210–217, 1987.

[11] Hendershott, T., Zhang, X., Zhao, J.L., Zheng, Z. FinTech as a game changer: Overview of research frontiers. *Inf. Syst. Res.*, 32(1):1–17, 2021.

[12] Samudrala, R.S., Yerchuru, S.K. Central bank digital currency: risks, challenges and design considerations for India. *CSI Trans. ICT*, 9(4):245–249, 2021.

[13] Williams, J.W. Regulatory technologies, risky subjects, and financial boundaries: Governing 'fraud' in the financial markets. *Account. Org. Soc.*, 38(6-7):544–558, 2013.

[14] Anagnostopoulos, I. Fintech and regtech: Impact on regulators and banks. *J. Econ. Bus.*, 100:7–25, 2018.

[15] Butler, T., O'Brien, L. Understanding RegTech for digital regulatory compliance. *Disrupt. Finance*, pages 85–102, 2019.

[16] Aziz, S., Dowling, M. Machine learning and AI for risk management. *Disrupt. Finance*, pages 33–50, 2019.

[17] Lee, J. Access to finance for artificial intelligence regulation in the financial services industry. *Eur. Bus. Org. Law Rev.*, 21:731–757, 2020.

[18] Kurum, E. RegTech solutions and AML compliance: what future for financial crime? *J. Financial Crime*, 30(3):776–794, 2023.

[19] Siering, M. Explainability and fairness of RegTech for regulatory enforcement: Automated monitoring of consumer complaints. *Decis. Support Syst.*, 158:113782, 2022.

[20] Zhang, Q., Wang, S., Zhang, D., Wang, J., Sun, J. FortunChain: EC-VRF-based scalable blockchain system for realizing state sharding. *IEEE Trans. Netw. Serv. Manag.*, 2023.

[21] Kokoris-Kogias, E., Jovanovic, P., Gasser, L., Gailly, N., Syta, E., Ford, B. OmniLedger: A secure, scale-out, decentralized ledger via sharding. In *IEEE Symposium Security and Privacy*, pages 583–598, 2018.

[22] Liu, Y., Liu, J., Li, D., Yu, H., Wu, Q. FleetChain: A secure scalable and responsive blockchain achieving optimal sharding. In *International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP)*, pages 409–425, 2020.

[23] Wang, J., Wang, H. Monoxide: Scale out Blockchains with Asynchronous Consensus Zones. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 95–112, 2019.

[24] Lovejoy, J., Virza, M., Fields, C., Karwaski, K., Brownworth, A., Narula, N. Hamilton: A High-Performance Transaction Processor for Central Bank Digital Currencies. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 901–915, 2023.

[25] Arner, D.W., Auer, R., Frost, J. Stablecoins: risks, potential and regulation. *BIS working paper*, 2020.

[26] Hardjono, T., Lipton, A., Pentland, A. Toward an interoperability architecture for blockchain autonomous systems. *IEEE Trans. Eng. Manag.*, 67(4):1298–1309, 2019.

[27] Auer, R., Böhme, R. Central bank digital currency: the quest for minimally invasive technology. Technical report, BIS, 2021.

[28] Nakamoto, S. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, page 21260, 2008.

[29] Boar, C., Wehrli, A. Ready, steady, go? — Results of the third BIS survey on central bank digital currency. *BIS papers*, 2021.

[30] Lee, A., Malone, B., Wong, P. Tokens and accounts in the context of digital currencies. *FEDS Notes*, pages 12–23, 2020.

[31] Boar, C., Holden, H., Wadsworth, A. Impending arrival — a sequel to the survey on central bank digital currency. *BIS paper*, 2020.

[32] Belchior, R., Vasconcelos, A., Guerreiro, S., Correia, M. A survey on blockchain interoperability: Past, present, and future trends. *ACM Comput. Surv.*, 54(8):1–41, 2021.

[33] Ziolkowski, R., Miscione, G., Schwabe, G. Decision problems in blockchain governance: Old wine in new bottles or walking in someone else's shoes? *J. Manag. Inf. Syst.*, 37(2):316–348, 2020.

[34] Auer, R., Böhme, R. The technology of retail central bank digital currency. *BIS Q. Rev.*, 2020.

[35] Yu, G., Wang, X., Yu, K., Ni, W., Zhang, J.A., Liu, R.P. Survey: Sharding in blockchains. *IEEE Access*, 8:14155–14181, 2020.

[36] Arner, D.W., Barberis, J.N., Buckley, R.P. The emergence of RegTech 2.0: From know your customer to know your data. 2016.

[37] Gill, M., Taylor, G. Preventing money laundering or obstructing business? Financial companies' perspectives on 'know your customer' procedures. *Brit. J. Criminol.*, 44(4):582–594, 2004.

[38] Teichmann, F., B., S., Sergi, B.S. RegTech–Potential benefits and challenges for businesses. *Technol. Soc.*, 72:102150, 2023.

[39] Parra Moyano, J., Ross, O. KYC optimization using distributed ledger technology. *Bus. Inf. Syst. Eng.*, 59:411–423, 2017.

[40] De Filippi., P., Hassan, S. Blockchain technology as a regulatory technology: From code is law to law is code. *arXiv preprint arXiv:1801.02507*, 2018.

[41] Ben-Or, M., Tiwari, P. A deterministic algorithm for sparse multivariate polynomial interpolation. In *ACM Symposium on Theory of Computing (STOC)*, pages 301–309, 1988.

[42] Rackoff, C., Simon, D.R. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In *Annu. Int. Cryptol. Conf.*, pages 433–444, 1991.

[43] Goldreich, O., Micali, S., Wigderson, A. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3):690–728, 1991.

[44] Chor, B., Goldwasser, S., Micali, S., Awerbuch, B. Verifiable secret sharing and achieving simultaneity in the presence of faults. In *26th Annual Symposium on Foundations of Computer Science (SFCS)*, pages 383–395. IEEE, 1985.

[45] Pocher, N., Veneris, A. Privacy and Transparency in CBDCs: A Regulation-by-Design AML/CFT Scheme. *IEEE Trans. Netw. Serv. Manag.*, 19(2):1776–1788, 2021.

[46] Li, W., Guo, H., Nejad, M., Shen, C. Privacy-preserving traffic management: A blockchain and zero-knowledge proof inspired approach. *IEEE Access*, 8:181733–181743, 2020.

[47] Sun, X., Yu, F.R., Zhang, P., Sun, Z., Xie, W., Peng, X. A survey on zero-knowledge proof in blockchain. *IEEE Network*, 35(4):198–205, 2021.

[48] Bindseil, U. Tiered CBDC and the financial system. *ECB Working Paper Series 2351*, 2020.

[49] Calle, G., Eidan, D. Central Bank Digital Currency: an innovation in payments. *R3 White Paper*, April 2020.

[50] Kr'ol, M., Ascigil, O., Rene, S., Sonnino, A., Al-Bassam, M., Rivi'ere, E. Shard scheduler: object placement and migration in sharded account-based blockchains. In *ACM Conference Advances in Financial Technologies*, pages 43–56, 2021.

[51] Robleh Ali. Cellular structure for a digital fiat currency. In *Digital Currency Economics And Policy*, pages 89–102. World Scientific, 2021.

[52] BIS Annual Economic Report. III. Blueprint for the future monetary system: improving the old, enabling the new. Technical report, BIS, 2023.

[53] Hansen, L.L. Corporate financial crime: social diagnosis and treatment. *J. Financ. Crime*, 16(1):28–40, 2009.

[54] Hearn, M., Brown, R.G. Corda: A distributed ledger. *Corda Technical White Paper*, 2016, 2016.

[55] Yang, J., Liu, C., Shang, Y., Cheng, B., Mao, Z., Liu, C., Niu, L., Chen, J. A cost-aware auto-scaling approach using the workload prediction in service clouds. *Inf. Syst. Frontiers*, 16:7–18, 2014.

[56] Xu, J. Developments and implications of central bank digital currency: The case of China e-CNY. *Asian Econ. Policy Rev.*, 17(2):235–250, 2022.

[57] Sethaput, V., Innet, S. Blockchain application for central bank digital currencies (CBDC). *Cluster Comput.*, pages 1–15, 2023.

[58] Buterin, V. A next-generation smart contract and decentralized application platform. *White Paper*, 3(37):2–1, 2014.