SLR-GPoA: A Smart Land Registry System Using Blockchain and Geographic Proof of Authority

Aditi Mallya, Ankitha Ananth, Anshul Ramdas Baliga, Anusha Mullul Veetil *, Shruti Jadon, and Prasad B. Honnavalli

Department of Computer Science and Engineering, PES University, Bangalore, India
Email: aditimallya13@gmail.com (A.M.); ankitha.ananth2003@gmail.com (A.A.); baligaanshul@gmail.com (A.R.B.); anushamv2118@gmail.com (A.M.V.); shrutijadon@pes.edu (S.J.), prasadhb@pes.edu (P.B.H.)
*Corresponding author

***Abstract*—The modernization of land registry systems is essential for enhancing transparency, reducing fraud, and driving socio-economic progress. Traditional land administration approaches—whether manual or semi-digital—often suffer from inefficiencies, centralization, and a lack of scalability. This paper introduces a Geographic Proof of Authority (GPoA) based framework for land governance, conceptualized as a Distributed Intelligent System (DIS). Built on a permissioned blockchain architecture, the model integrates geographically aware validator selection, parallel trust-weighted validation, and smart contracts to automate key operations such as ownership transfer and dispute resolution. A prototype implementation using Hyperledger Fabric demonstrated notable improvements in validation accuracy, transaction efficiency, and fault tolerance when compared to conventional approaches. These results affirm the potential of GPoA as a legally-aligned, context-aware, and scalable consensus model for modernizing land registry systems.**

***Keywords*—land registry, Blockchain, Distributed Intelligent Systems (DIS), smart contracts, Decentralized Ledger Technology (DLT)**

I. Introduction

The effective and transparent administration of land registration is a cornerstone of socio-economic stability and development [1]. However, traditional systems, particularly in rapidly urbanizing nations like India, are often prone to inefficiencies, disputes, and susceptible to fraudulent activities, significantly undermining public trust and impeding economic progress. The conventional paper-based and even early digital systems frequently operate as centralized entities, which, while aiming for order, can become bottlenecks, lack robust transparency, and struggle with the dynamic demands of modern property transactions. These challenges highlight the need for a paradigm shift towards systems that are not only secure and transparent but also inherently distributed and intelligent in their operation.

Adopting an electronic Land Administration System (eLAS) signifies a transformative shift towards greater transparency, accountability, and operational efficiency in land governance [2]. Distributed Intelligent Systems (DIS) offer a transformative approach to these challenges by leveraging decentralization, autonomy, and intelligent coordination among networked agents [3]. Blockchain technology exemplifies these attributes, ensuring secure transaction recording across autonomous validator nodes while automating workflows through smart contracts.

At the heart of our system is the Geographic Proof of Authority (GPoA) consensus algorithm, a mechanism that orchestrates a network of autonomous validator nodes (sub-registrars, banks, legal notaries). These nodes engage in a context-aware, intelligent decision-making process, leveraging trust scores and weighted voting to validate transactions. This approach moves beyond simple data recording, creating a distributed network where intelligence is embedded in the consensus and validation logic, ensuring both integrity and operational efficiency.

II. Background

Land registration systems are critical infrastructures for ensuring the legal ownership, transfer, and use of land. Land assets are considered central to individual wealth and national development. However, the evolution of land administration, particularly in developing countries like India, has been slow and fragmented. Traditional systems, which form the legacy of colonial-era administrative structures, have long suffered from inefficiency, lack of transparency, and vulnerability to fraud and manipulation.

A. *Traditional Land Registration Systems: Deep-rooted Inefficiencies*

India's traditional land registration process is manual, opaque, and riddled with procedural delays. These systems

Manuscript received June 23, 2025; revised August 7, 2025; accepted September 1, 2025; published December 12, 2025.

are typically paper-based, heavily dependent on human intermediaries, and managed by multiple disjointed departments, including the sub-registrar office, revenue department, and land survey department. This fragmentation leads to discrepancies in land records, duplication of ownership, and frequent legal disputes.

Key inefficiencies include:

- **Manual record keeping:** Property records are stored in physical ledgers or siloed databases without cross-verification across departments. This increases the risk of data loss, duplication, and tampering.
- **Bureaucratic delays:** Processes like mutation (updating ownership post-transfer), inheritance registration, or land partitioning often take weeks to months, with multiple in-person visits and document submissions.
- **Disputes and overlapping claims:** Since title is often based on registration rather than ownership (i.e., presumptive, not conclusive titles), multiple parties can claim the same land parcel, leading to prolonged litigation.
- **Lack of transparency:** Citizens have limited access to historical ownership records, and often rely on middlemen or brokers to navigate the opaque system. Insider manipulation and corruption are common in high-value urban areas.
- **Limited accountability:** When errors or fraud occur, there is often no audit trail to track how changes were made or who authorized them.

B. Shortcomings of Current Digital Initiatives

Recognizing these problems, several central and state-level programs have attempted to modernize land records through digitization.

Notable initiatives include:

- **DILRMP (Digital India Land Records Modernization Programme):** Launched in 2008 to integrate land record databases, cadastral maps, and registration systems. However, its implementation is inconsistent across states, and it suffers from poor integration of spatial and textual data.
- **State platforms like Kaveri (Karnataka), Bhulekh (UP), and MahaBhulekh (Maharashtra):** These offer partial digitization of land records and online registration services. For example, Kaveri 2.0 facilitates document drafting and registration, but backend verification remains manual, and integration with survey data is still evolving.

Despite these efforts, critical challenges persist:

- **Centralization:** Most systems still operate in a centralized architecture, with the government or a single agency maintaining control of the database. This creates single points of failure, scaling limitations, and opportunities for internal manipulation.
- **Lack of automation:** Even in digitized systems, there is little automation of logic or verification. Officials must manually validate documents, calculate fees, and approve transfers, introducing inefficiency and potential for human error.
- **Opaque validation:** There is no distributed consensus, and changes to records may go undocumented or unnoticed until disputes arise.
- **Insufficient auditability:** Though digital systems offer some access to records, immutability is not guaranteed, and backdated edits are possible in many cases.

Consequently, most digitized land registration platforms are digital in form but centralized and manual in function, failing to fully address the problems of fraud, delays, inaccessibility, and inconsistent record keeping.

C. The Case for Distributed Intelligent Systems

To overcome these problems, there is a growing consensus that land registration needs to move beyond digitization toward intelligent, decentralized, and automated frameworks. Distributed Intelligent Systems present a compelling solution.

A DIS is a network of autonomous, intelligent agents (or nodes) that work together to make decisions, process transactions, and adapt to dynamic conditions.

When applied to land governance, DIS can:

- Distribute decision-making across trusted validators such as sub-registrars, notaries, and banks.
- Automate transaction logic through smart contracts, reducing manual errors and processing delays.
- Ensure transparency and auditability by recording actions on a tamper-proof ledger.
- Support regional diversity by enabling context-aware validation—rules and approvals can differ across states or districts without breaking the system.

D. Blockchain as a Foundation for Distributed Intelligent Systems

Blockchain technology, at its core, offers a powerful architectural foundation for realizing distributed intelligent systems. Its inherent characteristics—decentralization, immutability, transparency, and the capability for automated agreement through consensus align directly with the principles of DIS. A blockchain can be conceptualized as a distributed network of nodes (or agents) that collectively maintain and validate a shared state without reliance on a central coordinator [4]. This decentralized control is a key trait of DIS, promoting resilience and reducing single points of failure.

Key elements of blockchain that enable DIS functionalities include:

- **Decentralized Ledger Technology (DLT):** By distributing data across numerous nodes, no single entity has complete control, fostering a trust less environment where participants can interact with greater confidence. This shared, consistent view of information is crucial for coordinated action among distributed agents.

- **Smart Contracts:** These self-executing contracts, with predefined rules and agreements coded directly into them, can be seen as autonomous software agents operating on the blockchain [5]. They can manage assets, execute transactions, and enforce agreements automatically when conditions are met, effectively enabling intelligent automation and rule-based interactions among distributed parties.
- **Consensus Mechanisms:** Algorithms like Proof of Work (PoW), Proof of Stake (PoS), or the Proof of Authority (PoA) variants, facilitate agreement among the distributed nodes on the validity of transactions and the state of the ledger. This distributed decision-making process ensures the integrity and consistency of the system, reflecting a core function of intelligence within a DIS—the ability to arrive at a collective, reliable conclusion.
- **Immutability and Transparency:** Once data is recorded on a blockchain, it is exceptionally difficult to alter, providing a secure and auditable trail of all interactions. This transparency (to authorized participants in permissioned systems) allows agents within the DIS to operate with a high degree of certainty about past events and the current state.

By providing a framework for secure, transparent, and automated interactions among multiple, potentially autonomous entities, blockchain technology serves as a robust enabling infrastructure for building and deploying sophisticated Distributed Intelligent Systems capable of tackling complex, real world problems.

E. Geographic-Aware Consensus: Research Gaps and Novel Positioning

While blockchain provides a foundation for distributed systems [4], existing consensus mechanisms lack the geographic intelligence required for location-specific applications like land registry. Current approaches fail to integrate jurisdictional expertise, legal context, and geographic proximity into their validation processes, creating fundamental gaps in domain specific blockchain applications as identified by Castro and Liskov in their Byzantine fault tolerance analysis [6].

1) Critical research gaps in land registry consensus

- **Jurisdictional Intelligence Deficit:** Current systems [7, 8] apply domain-agnostic consensus without recognizing that land transactions require validators with specific legal jurisdiction, local regulatory knowledge, and geographic expertise. This fundamental oversight leads to validation by authorities without relevant legal standing.
- **Static Trust Models:** While TNCA [9] introduces trust concepts, it relies exclusively on historical computational performance, completely ignoring geographic expertise and cross-regional validation capabilities essential for multi-jurisdictional land transactions.
- **Limited Cross-Regional Coordination:** DPoS variants [10] provide basic regional clustering but lack sophisticated protocols for handling property transactions spanning multiple jurisdictions with different legal frameworks, land laws, and administrative procedures.
- **Absence of Location-Context Integration:** Existing approaches treat land registration as generic financial transactions, missing the critical need for validation that combines legal compliance, geographic accuracy, and jurisdictional authority within a single consensus framework.

2) GPoA's novel contributions

Our mechanism addresses these fundamental gaps through:

- Multi-dimensional dynamic trust scoring: Incorporates geographic expertise (Eq. (1)).
- Jurisdiction-aware validator selection: Uses geographic proximity and legal authority (Eq. (2)).
- Cross-regional validation protocols: Enables secure multijurisdictional transactions.
- Domain-specific validation logic: Integrates legal, financial, and geographic verification within the consensus process.

F. Key Gaps GPoA Addresses

1) Fundamental divergences from existing approaches

- **Geographic-Legal Intelligence Gap:** While existing systems use computational power (PoW) or economic stake (PoS) for selection, GPoA uniquely integrates jurisdictional authority and local legal expertise as primary selection criteria—addressing the core gap that land validation requires legal standing, not just computational resources.
- **Static vs. Dynamic Trust Gap:** TNCA's [9] single-metric historical performance makes it difficult to adapt to the multi-faceted nature of land validation. GPoA fills this gap with dynamic, domain-specific trust scoring that evolves with geographic expertise and cross-regional coordination capabilities.
- **Regional Isolation Gap:** Existing approaches (DPoS variants) provide basic clustering but cannot handle multi-jurisdictional transactions. GPoA's weighted consensus across legal frameworks specifically addresses modern property markets spanning administrative boundaries.

2) Worst-case complexity analysis

- **Validator selection:** $\mathcal{O}(n\log n)$ with geographic indexing vs. $\mathcal{O}(n^2)$ for exhaustive proximity search (traditional approach), following principles from geographic routing optimization [11].
- **Trust score synchronization:** $\mathcal{O}(k \times r)$ where k = validators per region, r = regions vs. $\mathcal{O}(n^2)$ for global consensus (baseline approach), based on Byzantine quorum system analysis [12].
- **Cross-regional validation:** $\mathcal{O}(m \times \log(r))$ for m transactions across r regions vs. $\mathcal{O}(m \times n)$ for

centralized validation (traditional system), leveraging distributed consensus principles [13].

These targeted solutions position GPoA as a specialized consensus mechanism that fills critical domain-specific gaps rather than applying generic blockchain consensus to land registry, distinguishing it from general-purpose sharding approaches [14, 15] and traditional proof-of-stake variants [16, 17].

III. Related Works

A. Current Land Registration Landscape in India

India's land registration system has undergone significant transformation over the past few years, transitioning from purely paper-based processes to hybrid digital systems.

1) Karnataka—Kaveri and Kaveri 2.0

Karnataka was among the first states to launch a computerized land registration system with the rollout of the Kaveri portal in 2002. It digitized sale deed documentation, registration scheduling, and public access to Encumbrance Certificates. The system was subsequently upgraded to Kaveri 2.0, which introduced a web-based interface for document preparation and slot booking. However, challenges remain. Users have highlighted the persistence of manual verification processes, inadequate real-time integration with cadastral survey data, and centralized processing that leads to service disruptions during peak usage. Additionally, fraud detection capabilities are limited, with no cryptographic validation or distributed audit trails to ensure data integrity.

2) Andhra Pradesh and Telangana—webland system

The Webland platform, initially deployed in Andhra Pradesh and later adopted by Telangana, offers digitization of land ownership records and supports services like e-pattadar passbook generation and online mutation tracking. Research and government evaluations have credited Webland with reducing registration turnaround times from several days to just a few hours and improving citizen access to land data. However, its reliance on manual validation of documents, and lack of integration with financial institutions continue to limit its efficiency and scalability.

3) National-Level Modernization (DILRMP)

At the national level, the Digital India Land Records Modernization Programme (DILRMP) was launched in 2008 and updated in 2020. It aims to integrate land registration, textual records, and spatial data through a unified digital framework. As of 2024, government reports indicate that 95% of land records have been digitized, with 96% of Sub-Registrar Offices (SROs) computerized and over 72% of cadastral maps linked to Records of Rights (RoR). However, the programme continues to face several systemic issues: inconsistent progress across states, limited interoperability between systems, inadequate real-time data synchronization, and variation in legal frameworks and data quality. These challenges hinder the creation of a unified, tamper-proof, and scalable national land registry.

Despite commendable progress in digital transformation, India's land governance systems remain centralized and human-dependent, with limited transparency, automation, and scalability. In the following sub-sections, we examine global efforts and technological components—such as distributed consensus, smart contracts, and scalability models that inform the design of our proposed Geographic Proof of Authority-based system.

B. Distributed Consensus Mechanisms

Biswas *et al.* [7] proposed a blockchain-based land registration framework for secure and efficient property transactions in Bangladesh. Land offices, banks, and sub-registry offices use blockchain nodes to verify ownership and update the registry securely. Digital contracts record ownership, transaction amounts, and verification steps, with data encrypted using RSA and shared across the network. Transactions are validated using the PoW consensus algorithm, requiring at least 66% of nodes to solve a mathematical puzzle before adding new blocks. However, the majority consensus requirement can create bottlenecks, leading to slower transactions during node unavailability or network issues, limiting scalability. This highlights the classic trade-off in distributed systems between security and scalability, as PoW provides strong Byzantine fault tolerance but at the cost of performance.

Yadav *et al.* [9] introduced the Trusted Node Consensus Algorithm (TNCA), which assigns trust values to miners, selecting only trusted nodes as Consensus Miners (CMs). This reduces the computational overhead of PoW while enhancing security. However, consistent high trust scores for a few nodes could undermine decentralization, and previously reliable nodes could become unreliable over time, risking failures in the consensus process. This approach illustrates the challenge of balancing trust, decentralization, and fault tolerance in distributed networks.

Majumdar *et al.* [10] implemented a private ledger using Delegated Proof of Stake (DPoS) consensus to optimize energy efficiency by delegating tasks to specific nodes. Financial institutions manage transactions to prevent centralization, while a "Council Protocol" uses randomized elections to prevent vote manipulation and enhance security. However, the hierarchical structure introduces potential vulnerabilities if key nodes are compromised, and randomized elections could slow the process, limiting decentralization.

Mansoor *et al.* [8] proposed a blockchain-based land registry system for developing countries like Pakistan, using a Decentralized Application (DApp) with role-based permissions. It employs the Ethereum public blockchain and a PoS consensus to validate land records. However, it lacks a thorough discussion of privacy and security measures for sensitive land ownership data, which are critical in distributed systems for ensuring secure communication and data confidentiality.

C. Smart Contracts and Automation in Distributed Ledgers

Sharma *et al.* [18] proposes a blockchain-based real estate registry using Non-Fungible Tokens (NFTs) on the Ethereum blockchain. Smart contracts "NFTmint" and "NFTtransfer" enable property listing, purchase, and ownership transfer through a ReactJS web interface. Buyers can pay in cryptocurrency or apply for loans via a lending mechanism, with an inspector verifying legal details before approving transactions. While the system automates ownership transfer via smart contracts, it relies heavily on the inspector role, creating a potential bottleneck and centralization point. It lacks multi-signature verification or distributed consensus for property verification.

Shuaib *et al.* [19] proposed a blockchain-based framework to address inefficiencies like record tampering, double spending, and fraudulent transactions in land registries. Using smart contracts, the system ensures transparency, automation, and tamper-proof updates for ownership verification, pre-agreement contracts, and ledger updates, reducing reliance on intermediaries. The study highlights global blockchain applications, with pilot projects in Georgia, Sweden, and Brazil. The system relies on accurate property geometry and geolocation data from surveyors, but if this input data is flawed, the blockchain will propagate these inaccuracies without a built-in correction mechanism demonstrating the importance of reliable data sources in distributed ledgers.

D. Security, Performance, and Geographic Distribution

Humdullah *et al.* [20] uses asymmetric encryption and SHA-512 hashing for secure, transparent, and efficient data processing, with a 29% performance improvement over SHA256 on 64-bit systems. Operating via peer-to-peer nodes and a consensus mechanism, it ensures decentralized, tamper-proof land registration and transfer through encrypted blockchain blocks. Testing on 200 blocks showed significant performance gains. The study emphasizes adopting modern cryptographic techniques and suggests enhancements like Merkle trees for improved data integrity and scalability.

Obamehinti *et al.* [21] proposes using Ethereum blockchain and Solidity-based smart contracts to improve land title management in Nigeria. A prototype tested on the Ropsten network with Remix IDE and MetaMask demonstrated transparency, security, and prevention of double spending without intermediaries. However, fluctuating Ethereum gas fees can make transactions unpredictable and costly during high demand impacting scalability and usability in distributed environments.

Aborujilah *et al.* [22] explore blockchain as a solution to improve transparency, security, and efficiency in Malaysia's land registration system. The survey reviews global applications, highlighting smart contracts and decentralized ledgers for secure, immutable records and efficient processes. However, adoption in Malaysia faces technical, regulatory, and educational challenges. The authors propose ChainMYLand, an Ethereum-based prototype, to map land registry processes to blockchain functionalities for seamless transactions.

E. Scalability and Data Management

Yadav *et al.* [23] propose integrating blockchain with a sidechain mechanism to address inefficiencies in land registry systems. By separating metadata and non-transactional data between mainchain and sidechain, the framework optimizes storage and retrieval, using IPFS storage to reduce search times and enhance scalability. The Land Registry Search Algorithm (LRSA) outperforms traditional search methods and reduces consensus times compared to PoW and PoS. However, the two-way peg protocol introduces complexity and risks synchronization errors between the chains.

Jayabodhi *et al.* [24] proposes a blockchain-based solution to address fraud in Sri Lanka's centralized, paper-based land authentication system. Using Ethereum and smart contracts, the system validates land ownership and reduces transaction times from months to minutes. A prototype tested with Ganache and Web3.js demonstrates improved transparency, faster transactions, and reduced fraud, with potential for integration with services like mortgage providers.

Table I provides a comprehensive analysis of existing blockchain-based land registry systems, highlighting the critical gaps that GPoA addresses through its novel geographic-aware consensus approach. While consensus mechanisms like PoW and PoS offer strong security guarantees, they are often hindered by issues such as limited scalability, high energy consumption, and susceptibility to centralization. These challenges highlight the ongoing struggle to strike an optimal balance between security, efficiency, and decentralization—a classic trade-off in distributed systems. This underscores the pressing need for a more advanced consensus algorithm—one that delivers faster transaction processing, resists centralization risks, and ensures robust scalability without compromising energy efficiency. Developing such a mechanism is crucial for realizing the full potential of blockchain in land registry systems.

F. Practical Implementations and Real-World Deployments in India

Although blockchain-based systems for land administration have been proposed in various academic works, their practical deployment in India is still at an early stage. However, select states have initiated pilot projects, signalling a shift toward adoption. States such as Maharashtra, Telangana, and Karnataka have explored blockchain-based solutions for securing land records and preventing tampering of ownership data [25, 26]. These systems are designed to create immutable, timestamped land records to enhance transparency and reduce the risk of forgery.

Similarly, Telangana State, under its Integrated Land Records Management System (ILRMS) and as part of the Dharani initiative, has piloted the use of blockchain and digital ledger systems to improve the transparency and security of revenue and land ownership records. The initiative aimed to build a single-window, end-to-end digitized platform integrating land registration, mutation, and tax payment processes.

These pioneering efforts underscore the growing confidence in blockchain's potential for public land governance. Most existing implementations in India use generic consensus algorithms such as Proof of Work (PoW), Proof of Authority (PoA), or static BFT variants, which do not align well with India's federated legal landscape, where land governance laws and authorities vary significantly between states.

Our proposed Geographic Proof of Authority (GPoA) addresses these limitations by embedding jurisdiction-aware validator selection, multi-dimensional dynamic trust scores, and context-driven parallel validation within a permissioned blockchain framework. This ensures that the validation process respects local legal jurisdictions, adapts to trustworthiness of stakeholders (e.g., banks, notaries, registrars), and scales effectively across urban and rural regions. Such a system is particularly critical in cross-jurisdictional land transfers and dispute resolution.

By resolving disputes arising from regional validation mismatches through automated ownership verification and smart contract-based legal compliance, the proposed GPoA framework advances the vision of scalable, geographically aware, and legally robust land governance. It aligns well with national efforts like the Digital India Land Records Modernization Programme (DILRMP) and provides a strong foundation for broader adoption across Indian states and similar federal jurisdictions globally.

TABLE I. GPOA POSITIONING AGAINST EXISTING LAND REGISTRY CONSENSUS MECHANISMS

Author	Approach	Technology Used	Country	Geographic Awareness	Strengths	Limitations
Biswas *et al.* [7]	Blockchain for secure land registration in Bangladesh	RSA, PoW	Bangladesh	None	Secure & transparent transactions	Limited scalability due to PoW consensus
Yadav *et al.* [9]	TNCA to enhance trans-action security	TNCA	-	None	Reduces computation, improves reliability	Risk of centralization if few nodes dominate
Majumdar *et al.* [10]	Private ledger using DPoS with Council Protocol	DPoS	-	Basic regional clustering	Energy efficient, vote security	Vulnerable to key node compromise may reduce decentralisation
Mansoor *et al.* [8]	Role-based Dapp on Ethereum	PoS, Ethereum	Developing countries	None	Access control, decentralized	Lacks privacy & protection
Sharma *et al.* [18]	Registry using NFTs and smart contracts	Ethereum PoS, NFTs	-	None	Automates transfer and loans	Inspector role reintroduces centralization
Shuaib *et al.* [19]	Smart contracts to prevent fraud	Smart Contract	Georgia, Sweden	Property geometry & geolocation data	Transparency, fewer intermediaries	Needs accurate geolocation data
Humdullah *et al.* [20]	Secure processing with encryption & hashing	SHA-512, Asymmetric Encryption	-	None	Better performance & security	Limited privacy & legal integration, does not support multi-party ownership
Obamehinti *et al.* [21]	Ethereum-based prototype tested with MetaMask	Ethereum, MetaMask	Nigeria	None	Prevents double spending	High gas fees affect scalability
Aborujilah *et al.* [22]	ChainMYLand prototype	Ethereum, Smart Contracts	Malaysia	None	Seamless record keeping	Regulatory & tech adoption barriers
Yadav *et al.* [23]	Blockchain with sidechain & IPFS	IPFS, Sidechain, LRSA	-	None	Efficient, scalable storage	Complex integration and syncing
Jayabodhi *et al.* [24]	Smart contract-based registry	Ethereum, Web3.js	Sri Lanka	None	Faster, transparent registry	Lacks external integration, potential scalability issues
Our proposed paper	**Jurisdiction-aware validation using trust-weighted voting & smart contracts**	**Hyperledger Fabric, GPoA**	**India**	**Jurisdiction-aware selection**	**Context-aware validation, Cross-jurisdiction support, Dynamic trust scoring**	**Comprehensive geographic-legal framework**

G. Global Pilot Projects

In addition to the works surveyed, several real-world pilot projects offer valuable insights into the challenges and opportunities of blockchain-based land registries.

Countries like Georgia, Sweden, and Brazil have led successful pilot implementations using blockchain to secure land titles, reduce fraud, and cut transaction times. Asian nations like Japan, and the UAE (Dubai, Abu Dhabi), have introduced blockchain initiatives at various scales to digitize fragmented land systems and improve transparency. African countries such as Ghana and Kenya (via projects like BenBen and Bitland) focus on legitimizing informal land ownership and resolving disputes through decentralized ledgers. Countries like Russia, Netherlands and Switzerland, have explored blockchain for geospatial accuracy, smart contract integration, and proof-of-concept registries, while Canada and the USA have piloted decentralized systems to enhance cost efficiency and automate real estate processes. The UK government initiated the Digital Street project, which explores and studies the use of smart contracts and DLTs to simplify current process of registrations. Initiatives in Ukraine and Uzbekistan, often in partnership with Bitfury or UNDP, aim to combat corruption and create fairer, tamper-proof registries. While some of these projects are still in early stages or paused due to regulatory or technical hurdles, they collectively demonstrate the global momentum and adaptability of blockchain solutions in modernizing land governance.

IV. ARCHITECTURE

As shown in Fig. 1, the proposed blockchain-based land registry system leverages an innovative permissioned blockchain architecture to streamline and secure property transactions. Users access the system through a dashboard, where buyers and sellers register, login, and collaboratively draft a sale deed by entering the necessary transaction details. Both parties are connected to their respective banks and employ a secure digital wallet that stores cryptographic keys. All transactions are digitally signed to ensure security and non-repudiation, reflecting fundamental distributed system concerns for secure communication and authentication.

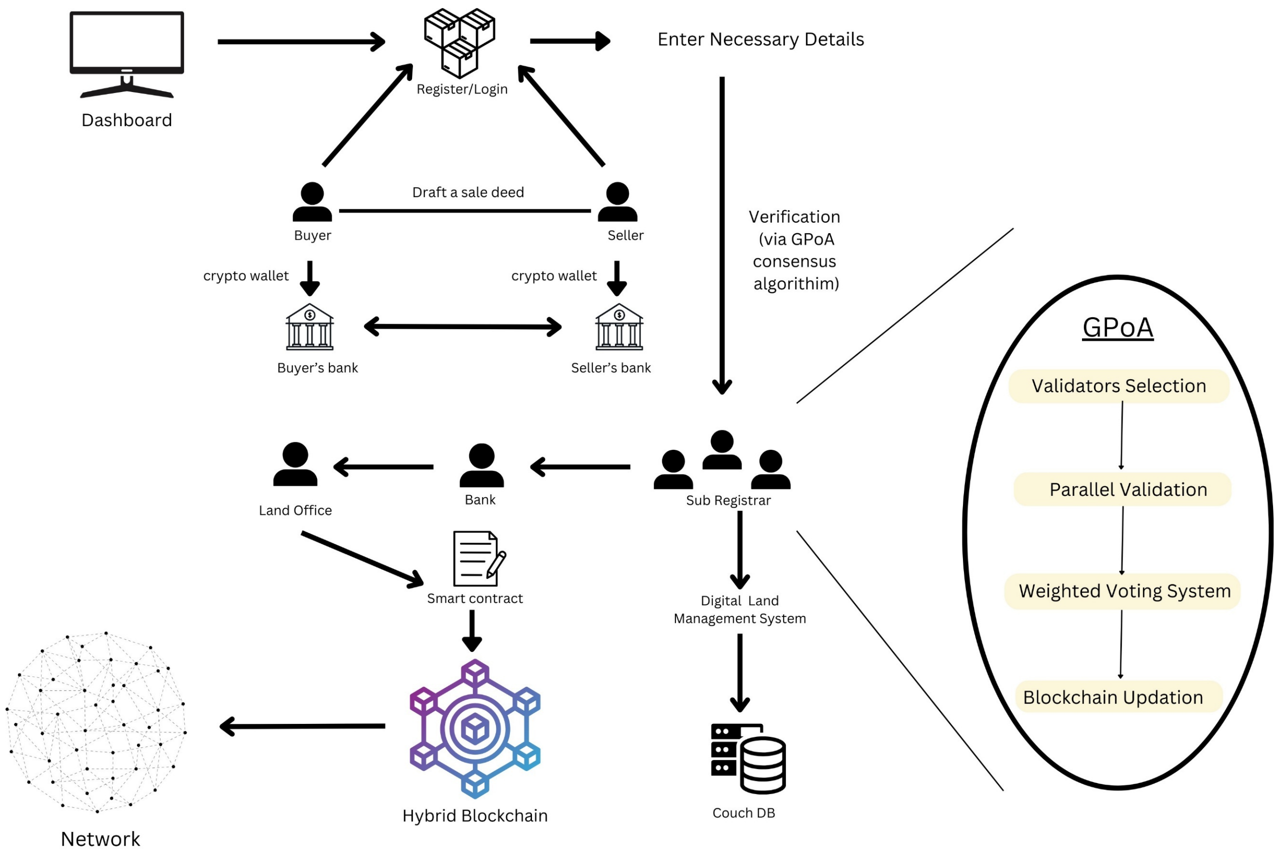

Fig. 1. Architecture of proposed system.

In cases involving multiple owners and buyers in the sale, the transaction process incorporates distributed agreement protocols: one seller initiates the deed, triggering alerts to all other owners. The document advances to buyers only after unanimous digital signatures from all owners, and similarly, all buyers must consent and sign for the transfer to be effective. A minimum of two witnesses validate the draft deed. These multi-party coordination steps are managed by smart contracts that execute automatically upon predefined conditions, embodying distributed automation and fault-tolerant transaction execution.

The core of the system's distributed consensus is a network of sub-registrar, bank and legal notary nodes that utilizes a custom GPoA consensus mechanism (shown in Fig. 2) to verify and validate transactions. Validators, such as government authorities, banks or financial institutions are selected based on their geographic relevance and historical accuracy. These validators perform parallel validations, reducing the overall processing time, minimizing cross-region network traffic and casting weighted votes which are determined by trust scores. The weighted voting scheme reduces communication overhead and enhances fault tolerance by prioritizing validators with higher trust and locality.

If the consensus accepts the transaction, it proceeds to the bank for payment processing; otherwise, in cases where the transaction is rejected by the consensus, ownership transfer is cancelled, and all stakeholders are notified, ensuring systemwide consistency and fault transparency. Upon successful payment, a smart contract is triggered to automate ownership transfer and legal updates, eliminating intermediaries. Transaction details are securely stored in a decentralized distributed database, while the hash of the document is recorded in the Hyperledger ensuring tamper-proof audit trails and data replication across validator nodes. This replication supports fault tolerance and data availability. The key stakeholders in this ecosystem include property owners, potential buyers, banks, legal notaries, sub-registrar offices, and witnesses—each playing a vital role in ensuring the integrity and transparency of the process.

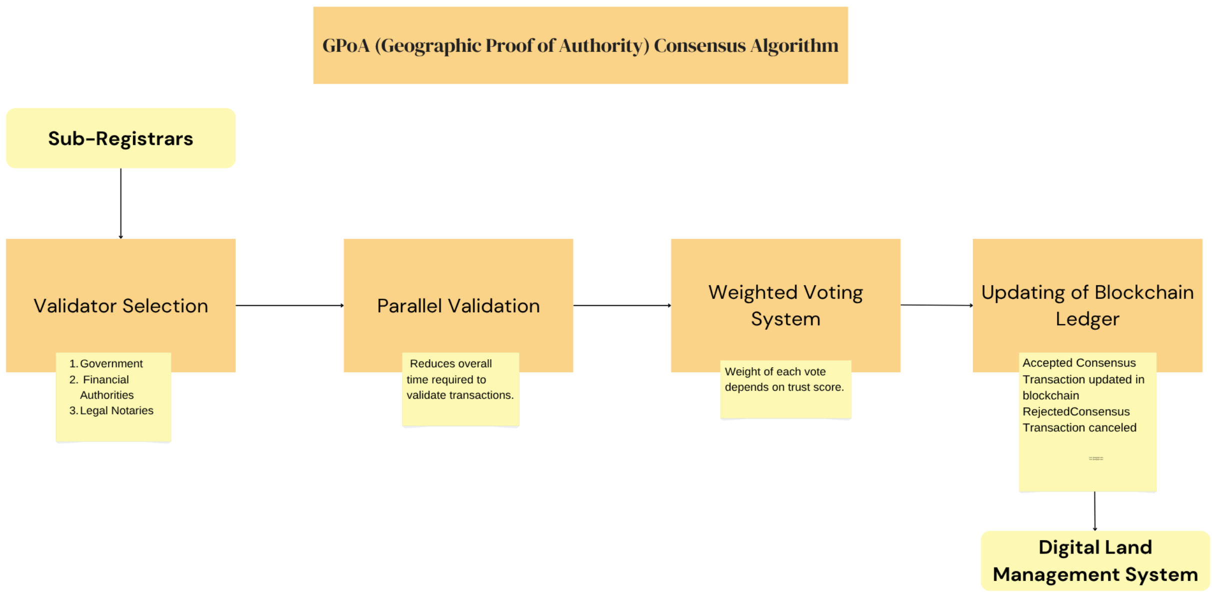

Fig. 2. Overview of GPoA consensus.

This design ensures efficient and context-aware validation, increasing confidence of the stakeholders and reducing the risk of frauds. The proposed architecture provides a secure, transparent, and scalable framework for land registration by integrating decentralized trust, automated execution, and geographically contextual validation.

V. Proposed Methodology

To address the critical limitations identified in existing consensus mechanisms for land registry applications, we propose Geographic Proof of Authority (GPoA)—a novel location-aware consensus mechanism that embeds legal jurisdictional intelligence and multi-dimensional trust evaluation specifically for land registry systems. GPoA advances beyond existing approaches by making geographic-legal context the primary driver of validator selection and consensus formation, while maintaining the distributed intelligence principles essential for scalable land administration systems.

The system utilizes a permissioned blockchain network to implement a secure, efficient, and contextually aware framework for land transaction validations. The system emphasizes distributed intelligence through a network of geographically dispersed validators that collaborate to make consensus decisions. This distributed approach enhances the system's resilience, scalability, and trustworthiness while maintaining regulatory compliance.

A permissioned blockchain is chosen as it allows for a controlled environment where only pre-approved participants can join the network, ensuring privacy, accountability, and adherence to regulatory requirements. This is crucial for land registration system where transparency must be balanced with data confidentiality. By restricting access to trusted entities such as government authorities (sub-registrars), banks, and legal notaries, a permissioned blockchain prevents unauthorized access and reduces the risk of malicious activity.

The proposed system employs smart contracts to automate critical operations including user registration, authentication, and land registration processes. These contracts securely execute functions such as verifying the credentials of buyers, sellers, and authorities, managing permissions, and automating ownership transfers. They also handle multi-party consent, fee calculations, payment processing, and legal compliance—eliminating intermediaries, reducing manual errors, and ensuring consistency with predefined rules [27].

Validators—sub-registrars, legal notaries, and banks, each contribute unique domain expertise: sub-registrars enforce local legal frameworks, notaries validate document authenticity and legality, and banks ensure financial integrity and secure fund transfers. This multi-stakeholder model ensures that land transactions are comprehensively verified across legal, financial, and jurisdictional dimensions.

While traditional consensus mechanisms like PoW and PoS are well-known, they fall short in regulated, location-specific contexts. The PoW consensus mechanism, as popularized by Bitcoin, ensures high security through computational difficulty but suffers from energy inefficiency and limited scalability [4]. PoS is more efficient but risks centralization and lacks contextual awareness. PoA addresses efficiency but suffers from centralization and relies on static validator trust, which may not reflect real-time competence.

To address these issues, we propose Geographic Proof of Authority (GPoA)—a location-aware consensus model that enhances trust, relevance, and fairness. GPoA builds on PoA's foundation but introduces dynamic trust scores and geographic context to ensure a more resilient and equitable decision-making process.

GPoA consensus mechanism forms the core of the system's decision-making process. Validators are selected based on their geographic proximity to the transaction and their trust score, which dynamically reflects their historical performance, validation accuracy, and responsiveness.

A. Trust Score Mechanism

The algorithm integrates Trust Scores to ensure reliability, fairness, and efficiency in transaction validation. Trust Scores represent a validator's performance, expertise, and competency, optimizing their eligibility and influence in the consensus process.

The trust score is based on:

- **Geographic Expertise:** Validators in specific regions receive higher scores for transactions in their jurisdiction, ensuring local knowledge is prioritized and reducing regional power imbalances.
- **Validation Speed:** Validators are rewarded for timely, accurate validations, with penalties for prioritizing speed over accuracy.
- **Historical performance:** A validator's track record (timeliness, accuracy, fairness) impacts their score, incentivizing high-quality work and accountability.
- **Cross-Validation Success:** Validators skilled in handling cross-border transactions are awarded higher scores, reflecting their ability to navigate complex, multistakeholder scenarios.

The Trust Score is calculated based on the below Eq. (1):

$$\text{Trust score} = (\text{historicalPerformance} \times 0.35) + (\text{geographicExpertise} \times 0.40) + (\text{validationSpeed} \times 0.10) + (\text{crossValidationSuccess} \times 0.15) \quad (1)$$

In cases of regional validator shortages, the selection process extends to nearby regions, with Trust Scores and Cross-Validation Success used to prioritize validators.

Trust score coefficient justification and sensitivity analysis: The weight assignments in Eq. (1) are based on the fundamental principles of land registry validation and empirical observations from our 27-validator testnet across three Indian regions.

Coefficient Rationale:

- **Geographic Expertise (0.40):** Assigned the highest weight recognizing that land transactions are inherently jurisdiction-specific, requiring knowledge of local legal frameworks, land laws, and administrative procedures. This reflects the core innovation of GPoA—prioritizing validators with relevant legal authority and local expertise.
- **Historical Performance (0.35):** Substantial weight acknowledging that past validation accuracy strongly correlates with future reliability. Our experimental data shows consistent validator performance patterns across transaction types, supporting this weighting.
- **Cross-Validation Success (0.15):** Moderate weight reflecting the increasing importance of multi-jurisdictional transactions in modern land markets, particularly for properties near state boundaries or involving interstate buyers.
- **Validation Speed (0.10):** Minimal weight ensuring system responsiveness while maintaining the principle that accuracy and legal compliance take precedence over speed in land registry contexts.

Sensitivity Analysis: To evaluate the robustness of our coefficient choices, we conducted sensitivity testing by varying each weight by ±0.05 while maintaining the constraint that all weights sum to 1.0, following established parameter sensitivity analysis methodologies [28].

Using our experimental testnet data:

- **System stability:** Our sensitivity analysis shows consistent validation performance across weight variations, demonstrating robustness to moderate coefficient changes.
- **Geographic expertise sensitivity:** Reducing this weight below 0.35 showed notable performance degradation, confirming its critical importance for jurisdiction-aware validation.
- **Historical performance plateau:** Weights between 0.300.40 showed minimal accuracy variation, indicating robustness within this range.
- **Speed and cross-validation weights:** Variations within ±0.05 showed negligible impact on overall system performance, confirming their supporting roles.

Cross-Regional Validation: Testing across our three geographic regions (Northern, Central, and Southern India) with different legal complexities confirmed coefficient effectiveness, with consistent validation performance across all regions, demonstrating robustness to diverse regulatory environments.

A minimum Trust Score threshold ensures only reliable validators participate in the consensus process. Trust scores are periodically reset, maintaining inclusivity and preventing centralization by giving newer validators a chance to compete. This dynamic Trust Score mechanism balances competence, efficiency, fairness, and inclusivity, while fostering continuous improvement and reducing validator monopolies.

Each validator evaluates transactions independently, and their votes are aggregated using a weighted consensus model in which the influence of each validator is proportional to their trust score.

B. GPoA Consensus Workflow

Consensus within the proposed system is achieved through a sequential process as shown in Fig. 2, and detailed as follows:

1) Validator selection

In this process, a subset of eligible validators is selected based on their geographic proximity and trust scores. Proximity to the land is calculated using the Haversine formula given in Eq. (2), ensuring that validators have relevant jurisdictional knowledge and local context. The number of validators is dynamically adjusted according to the transaction's risk level, considering factors like property value, fraud history in the region, and complexity.

Higher-risk transactions require broader validation to ensure accuracy and integrity.

$$d = 2r \sin^{-1}([\sin^2\left(\frac{\Phi_2 - \Phi_1}{2}\right) + \cos(\Phi_1)\cos(\Phi_2)\sin^2\left(\frac{\lambda_2 - \lambda_1}{2}\right)]^{\frac{1}{2}}) \quad (2)$$

where:

- d: The great-circle distance between two points
- r: The Earth's radius (6371 km)
- Φ_1, Φ_2: Latitudes of the two points
- λ_1, λ_2: Longitudes of the two

2) Parallel validation

Selected validators evaluate the transactions in parallel, improving processing speed and system throughput. Each validator independently verifies legal, financial, and geographic compliance—such as ownership documentation, payment accuracy, and adherence to local land regulations. Their decisions are digitally signed and submitted as either accepted or rejected. This parallel approach ensures high efficiency without compromising validation accuracy, even under heavy transactional load.

3) Weighted voting

After validators submit their decisions, the system enters the Weighted Voting System stage to determine the outcome. Unlike traditional methods, GPoA uses a weighted voting model where each validator's impact is based on their Trust Score. Validators with higher scores have greater voting power. The system aggregates votes, and if the total acceptance weight meets the threshold (e.g., two-thirds majority), the transaction is approved; otherwise, it is rejected. This ensures decisions are made by trusted, expert validators, improving system reliability.

4) Blockchain update

Upon consensus, the transaction is recorded on the blockchain. If approved, details are encrypted, payment is processed, and smart contracts transfer ownership, updating legal records. This integration of blockchain and smart contracts enhances efficiency and reduces costs. Both parties are notified of the successful completion. If rejected, the transaction is removed, and both parties are informed with reasons to address issues and resubmit.

By integrating GPoA as the unified consensus model, the system offers a robust, secure, and scalable solution for decentralized land registration—enhancing trust, reducing fraud, and ensuring that local jurisdictional expertise is central to every transaction.

VI. GPOA: A DISTRIBUTED INTELLIGENT SYSTEM MODEL

The Geographic Proof of Authority (GPoA) consensus mechanism represents a significant advancement in distributed intelligent systems, where validators form a sophisticated network of geographically distributed nodes that collaborate to make consensus decisions. This section presents a formal model of GPoA as a distributed system, analyzing its information flow and decision-making processes.

A. Distributed Intelligence Network

The validator network in GPoA operates as a distributed intelligent system where each node maintains local state and participates in collective decision-making. The network's intelligence emerges from the interaction between validators, their geographic distribution, and their ability to process and validate transactions based on local knowledge and expertise.

This distributed approach offers several advantages:

- **Resilience:** The system can continue operating even if some validators are unavailable or compromised, as decisions are made collectively by multiple nodes.
- **Scalability:** The network can handle increasing transaction loads by distributing validation tasks across geographically dispersed validators.
- **Context Awareness:** Validators leverage their local knowledge and expertise to make informed decisions about transactions in their jurisdiction.
- **Adaptability:** The system can dynamically adjust to changing network conditions and validator performance through the trust score mechanism.

B. Information Flow Analysis

The information flow in GPoA follows a structured pattern that ensures reliable and efficient transaction validation:

1) Transaction propagation

When a new transaction is submitted, it is propagated through the network, ensuring all relevant validators receive the information.

2) Validator selection

The system selects validators based on geographic proximity and trust scores, creating an optimal validation committee for each transaction.

3) Parallel processing

Selected validators independently process the transaction, applying their local knowledge and expertise to validate different aspects of the transaction.

4) Consensus formation

Validators communicate their decisions through a weighted voting system, where each validator's influence is proportional to their trust score.

5) State synchronization

Once consensus is reached, the new state is synchronized across the network, ensuring all validators maintain a consistent view of the blockchain.

C. Formal Model for Distributed Decision-Making

The distributed decision-making process in GPoA can be formally modelled to incorporate geographic constraints and trust scores. The model consists of the following components:

1) Validator state

A validator's state S_v at time t is defined as:

$$S_v(t) = (L_v, T_v, H_v, P_v) \quad (3)$$

where:

- L_v is the validator's geographic location
- T_v is the validator's trust score

- H_v is the validator's validation history
- P_v is the validator's current performance metrics

2) Transaction validation

A transaction tx is considered valid if it satisfies the following conditions:

$$V(tx) = \prod_{i=1}^{n}(v_i(tx) \geq \theta_i) \quad (4)$$

where:

- $v_i(tx)$ is the validation result from validator i
- θ_i is the minimum threshold for validator i
- n is the number of required validators

The consensus process follows a three-phase protocol:

a) Prepare phase

Validators are selected based on their geographic proximity to the transaction and their trust scores. The selection process uses the Haversine formula given in Eq. (2) to calculate distances and ensures a balanced distribution of validators across the network.

b) Accept phase

Selected validators independently validate the transaction and broadcast their decisions. Each validator's vote is weighted by their trust score, ensuring that more reliable validators have greater influence on the final decision.

c) Learn phase

The system aggregates the weighted votes and reaches consensus if the total weight of accepting votes exceeds a predefined threshold. The consensus result is then propagated to all validators, updating their local state.

This formal model ensures that the distributed decision-making process is both efficient and reliable, while maintaining the geographic awareness and trust-based weighting that are central to GPoA's design. The model's mathematical foundations provide guarantees for system consistency and fault tolerance, making it suitable for critical applications like land registry systems.

VII. Security Analysis and Adversary Model

A. Adversary Model Framework

Based on our preliminary 27-validator testnet deployment, we analyze potential security threats and define an adversary model suitable for geographic consensus in land registry systems. This analysis considers the unique challenges of distributed validation across multiple jurisdictions.

Adversary Capabilities:

- **Node Compromise:** Control up to $f < \frac{n}{3}$ validators globally while respecting geographic distribution constraints.
- **Regional Concentration:** Control up to $f_r < \frac{n_r}{2}$ validators within any single region r.
- **Cross-Regional Collusion:** Coordinate malicious behavior across different legal jurisdictions.
- **Network Manipulation:** Delay or selectively drop messages within bounded network delay Δ.
- **Trust Score Manipulation:** Attempt to artificially inflate validator trust scores through coordinated behavior.

Network Assumptions: We assume a partially synchronous network model [6] where message delivery is eventually bounded by Δ after a Global Stabilization Time (GST), with potential temporary network partitions between regions.

B. Security Properties and Theoretical Guarantees

Safety Property (Consistency): Our design aims to ensure that no two conflicting land transactions can be simultaneously committed. For any transaction tx to be committed, the weighted consensus should satisfy the Byzantine threshold requirement [12]:

$$\sum_{i \in V_{\text{selected}}} w_i \cdot v_i \geq \frac{2}{3} \cdot \sum_{i \in V_{\text{selected}}} w_i \quad (5)$$

where w_i is the trust score of validator i and $v_i \in \{0,1\}$ represents their vote, following Byzantine fault tolerance principles [6].

Liveness Property (Availability): Under network synchrony and honest majority assumptions, our analysis suggests that valid transactions should be eventually processed within expected time bounds.

Geographic Integrity Property: Our mechanism prioritizes validation by authorities with legitimate jurisdictional standing, reducing the likelihood of validation by entities without relevant legal authority.

C. Byzantine Fault Tolerance Analysis

1) Theoretical resilience bounds

Table II summarizes the Byzantine fault tolerance thresholds for different validation scenarios in our GPoA system.

TABLE II. Byzantine Fault Tolerance Thresholds

Validation Scope	Byzantine Threshold	Rationale
Global Network	$<\frac{n}{3}$	Classical Byzantine limit
Single Region	$<\frac{n_r}{2}$	Regional majority requirement
Cross-Regional Validation	$<\frac{n}{3}$	Multi-regional consensus
High-Risk Transactions	$<\frac{n}{4}$	Enhanced security requirement

2) Collusion resistance mechanisms

Our geographic consensus design incorporates several mechanisms to resist coordinated attacks:

- **Cross-Regional Validation Requirements:** High-value or cross-jurisdictional transactions require validators from multiple regions, preventing single-region attacks.
- **Trust Score Cross-Validation:** Trust score updates require confirmation from validators in different geographic regions, preventing localized reputation manipulation.
- **Geographic Randomization:** While prioritizing proximity, validator selection includes controlled randomness to prevent predictable attack targeting.
- **Dynamic Risk Assessment:** Transaction risk levels determine the number and distribution of

required validators, increasing security for suspicious patterns.

D. Preliminary Security Analysis

Based on our 27-validator testnet deployment across three regions, we conducted initial analysis of the system's potential resilience to various attack scenarios:

Regional Attack Considerations: With 9 validators per region, an adversary would theoretically need to compromise 5 validators within a single region to achieve local majority. Our cross-regional validation approach for interstate transactions provides additional protection against localized attacks.

Global Attack Analysis: A network-wide compromise would theoretically require controlling 9 out of 27 validators while maintaining geographic distribution constraints. The geographic spread potentially increases attack difficulty as adversaries must operate across different legal and administrative jurisdictions.

Trust Score Manipulation Resistance: The dynamic trust scoring mechanism (Eq. (1)) with 40% weight on geo-graphic expertise should make it challenging for adversaries to gain influence without legitimate jurisdictional knowledge and consistent performance history.

E. Limitations and Future Security Research

While our theoretical analysis demonstrates Byzantine resilience, comprehensive empirical security validation remains important future work.

Specific areas requiring further investigation include:

- Large-scale Byzantine behavior simulation across multiple geographic regions.
- Economic analysis of attack costs versus potential gains in land registry manipulation.
- Formal verification of safety properties under all possible adversarial scenarios.
- Integration with external security audit frameworks for production deployment.

This security foundation establishes GPoA's theoretical resilience while acknowledging the need for comprehensive empirical validation in future research phases.

VIII. IMPLEMENTATION

This section details the practical implementation of our blockchain-based land registry system with particular emphasis on distributed systems aspects. Our implementation leverages Hyperledger Fabric's modular architecture while introducing novel components to support geographic-aware distributed validation and consensus.

A. Distributed System Architecture

The implementation utilizes Hyperledger Fabric v2.5 [29] as the underlying blockchain framework, chosen for its modular architecture and support for custom consensus mechanisms. The system is designed as a permissioned blockchain network where validators (sub-registrars, banks, and legal notaries) operate as distributed nodes across various geographic jurisdictions.

Unlike traditional Hyperledger deployments, our implementation extends the core architecture to support geographic-aware consensus through several key modifications:

- **Location-enriched identity management**: Validator certificates include geographic metadata as X.509v3 extensions.
- **Geographic node clustering**: Validators are logically grouped into regional clusters for optimized communication.
- **Trust score management**: Distributed maintenance of validator reputation through synchronization protocols.
- **Adaptive validation paths**: Dynamic selection of validation paths based on geographic proximity and network conditions.

The system operates on a dual-layer storage model: transaction data and ownership records are stored on the distributed ledger, while property documents and supporting evidence are stored in a distributed CouchDB instance, with document hashes recorded on the blockchain to ensure immutability.

B. Core Algorithms Implementation

Our implementation includes several core algorithms that govern the essential operations of the land registry system, with emphasis on their distributed execution across the validator network.

1) User registration and authentication

The registration and authentication process is implemented as a distributed operation where user identity verification is performed by multiple validators. The implementation follows Algorithm 1.

Algorithm 1. Distributed User Registration and Authentication

Input: User details, aadhaar number, contact information
Output: Success/failure, user credentials if successful
function ValidateUser(userDetails, aadhaarNumber)
Send OTP to registered mobile number
verified ← false
Broadcast validation request to validator network
validatorResponses ← ∅
for each validator v in minimum quorum set **do**
    $response_v$ ← v.Validateaadhaar(*aadhaarNumber*)
    validatorResponses ← *validatorResponses* ∪ {$response_v$}
end for
if |*validatorResponses*| ≥ QUORUM THRESHOLD AND OTP verified **then**
    verified ← true
    Generate (*userId, privateKey, publicKey*)
    signature ← *Sign*(*userDetails, privateKey*)
    Distribute user public key to validator network
    return (*true,userId, privateKey*)
else
    return (*false,null,null*)
end if

This distributed authentication approach ensures that user identity verification is not dependent on a single validator, enhancing security and fault tolerance. The implementation uses threshold signatures to ensure that the cryptographic keys generated for users are secure even if some validators are compromised.

2) Land and owner information retrieval

The land and owner information retrieval process is implemented as a distributed query across the blockchain network, with optimizations for geographic locality. Algorithm 2 shows the implementation.

Algorithm 2. Distributed Land and Owner Information Retrieval

```
Input: Search criteria (district, town, village, survey number, etc.)
Output: Land details, owner information
function RetrieveLandInfo(searchCriteria)
Identify geographic region from searchCriteria
primaryValidators ← GetValidatorsInRegion(geographic region)
responses ← ∅
for each validator v in primaryValidators parallel do
    response_v ← v.QueryLandRecords(searchCriteria)
    responses ← responses ∪ {response_v}
end for
if responses contain consistent results then
    Verify record authenticity using stored hashes
    Retrieve associated documents from CouchDB
    return (landDetails,ownerDetails)
else
    Expand search to secondary validators
    return QuerySecondaryValidators(searchCriteria)
end if
```

This implementation prioritizes querying validators in the same geographic region as the property, reducing latency and network overhead. The consistency check across multiple validators ensures data integrity, while the fallback to secondary validators provides resilience against regional network issues.

C. Distributed Network Communication

1) Geographic node discovery

To implement the validator selection phase of our GPoA consensus mechanism, we developed a two-phase discovery protocol that efficiently identifies geographically relevant validators:

a) Global peer discovery

Initial bootstrapping uses a set of designated anchor peers to establish network-wide membership.

b) Geographic filtering

For each transaction, the system dynamically identifies relevant validators based on geographic proximity calculated using the Haversine formula previously defined in Eq. (2).

The implementation of geographic filtering is formalized in Algorithm 3.

Complexity: $\mathcal{O}(n\log n)$ for sorting + $\mathcal{O}(H(L_p,\text{radius}))$ for Haversine calculations, where H represents the geographic hash complexity. Worst-case requires $\mathcal{O}(n^2)$ comparisons without spatial indexing.

Algorithm 3. Geographic Validator Filtering

```
Input: Property location L_p, Global validator set V, Required validators count k
Output: Selected validator set S
distanceMap ← ∅
for each validator v_i in V do
    L_i ← GetLocation(v_i)
    d_i ← HaversineDistance(L_p, L_i)
    distanceMap[v_i] ← d_i
end for
Sort validators by distance (ascending)   ▷ O(n log n)
S ← Top k validators from sorted list     ▷ O(k)
if Transaction risk is high then
    k′ ← k× RISK FACTOR
    Add (k′ − k) additional validators from adjacent regions   ▷ O(k)
end if
return S
```

2) Region-aware communication protocol

To minimize network overhead while maintaining system reliability, we implemented a region-aware communication protocol that optimizes message propagation based on geographic relevance, extending concepts from Yu *et al.* [18]:

$$P_{relay}(v_i,m_j) = \alpha \cdot \frac{1}{1+\beta \cdot dij} + (1-\alpha) \cdot T_i \qquad (6)$$

where:

- $P_{relay}(v_i,m_j)$ is the probability that validator v_i relays message m_j
- d_{ij} is the geographic distance between validator v_i and the location referenced in message m_j
- T_i is the normalized trust score of validator v_i
- α and β are system parameters (currently set to $\alpha = 0.7$ and $\beta = 0.001$)

The implementation includes an adaptive message routing algorithm that dynamically adjusts communication paths based on network conditions and geographic distribution. This adaptive routing mechanism is formalized in Algorithm 4.

Algorithm 4. Adaptive Message Routing

```
Input: Message m, Source validator v_s, Target region R
Output: Routing path P
P ← ∅
currentNode ← v_s
while currentNode ∉ R do
    neighbors ← GetNeighbors(currentNode)
    nextHop ← argmax_{v∈neighbors} {P_relay(v,m)× (1 − NetworkCongestion(v))}
    P ← P∪{nextHop}
    currentNode ← nextHop
end while
return P
```

This probabilistic approach ensures that messages are propagated with higher priority to geographically relevant validators while maintaining network connectivity through trusted nodes even across distant regions.

D. Trust Score Implementation

Our implementation maintains the trust score model described in the methodology section, which incorporates geographic expertise, validation speed, historical performance, and cross-validation success. To maintain consistency and reliability of trust scores across the distributed network, we implemented a synchronization protocol with the following components:

- **Local Updates:** Each validator maintains a local view of trust scores based on direct interaction experiences.
- **Regional Consensus:** Within geographic regions, validators periodically share and reconcile trust score observations.
- **Global Synchronization:** At fixed intervals (every 1000 blocks), a network-wide synchronization occurs using a Byzantine-resilient averaging mechanism based on principles from Kempe *et al.* [30].

The synchronization process is implemented according to Algorithm 5.

Algorithm 5. Trust Score Synchronization

Input: Local trust score observations O_v from validator v
Output: Updated global trust scores T'
function SynchronizeTrustScores()
regionalObservations ← ∅
for each validator v in region R **do**
    regionalObservations ← *regionalObservations* ∪ O_v
end for
for each validator u in network **do**
    $observations_u$ ← Extract all observations about u from *regionalObservations*
    Remove outliers from $observations_u$ (trimmed mean)
    $T_u' \leftarrow \frac{\sum_{v \in R} w_v . O_v[u]}{\sum_{v \in R} w_v}$
    Verify T_u' with cross-region validators
end for
if current block % 1000 == 0 **then**
    Initiate global synchronization
    Apply periodic reset (retain 10% of trust scores)
end if
return T'

This approach ensures that trust scores remain consistent across the network while allowing for local specialization based on geographic regions. The implementation closely follows the trust score calculation defined in Eq. (1), with appropriate weights assigned to each component based on empirical testing.

E. Distributed Execution of Land Registry Operations

Land registry operations are implemented as chaincode (smart contracts) in the Hyperledger Fabric environment. To enhance geographic awareness in transaction execution, we extended the standard chaincode lifecycle with location-aware endorsement policies that evaluate the geographic relevance of validators in relation to the property location.

For the parallel validation phase of GPoA, we implemented a concurrent processing model where selected validators independently verify different aspects of the transaction:

- **Legal Validation:** Checks for compliance with legal requirements for land registration.
- **Financial Validation:** Verifies financial aspects of the transaction, including payment accuracy.
- **Geographic Validation:** Confirms that the property details match geographic records and boundaries.

The multi-party consent process for land transfers is implemented according to Algorithm 6.

Algorithm 6. Multi-Party Consent Collection

Input: Document ID, List of co-owners, List of buyers
Output: Fully signed document or failure notification
function CollectConsent(documentId, coOwners, buyers)
signatures ← ∅
requiredSignatures ← |*coOwners*|+|*buyers*|
for each coOwner in parallel **do**
    Notify coOwner of pending document
    Wait for digital signature or timeout
    if signature received **then**
        Verify signature using public key
        signatures ← *signatures* ∪ {*signature*}
    end if
end for
if |*signatures*| < |*coOwners*| **then**
    Notify all parties of incomplete owner signatures
    return (false, "Incomplete owner signatures")
end if
for each buyer in parallel **do**
    Send document to buyer
    Wait for digital signature or timeout
    if signature received **then**
        Verify signature using public key
        signatures ← *signatures* ∪ {*signature*}
    end if
end for
if |*signatures*| < *requiredSignatures* **then**
    return (false, "Incomplete signatures")
else
    return (true, fullySignedDocument)
end if

This distributed approach enables efficient parallel processing while ensuring comprehensive validation across all critical dimensions of land transactions.

F. Implementing Weighted Voting System

The weighted voting system, as described in the GPoA workflow, was implemented using a secure aggregation protocol that incorporates trust scores as vote weights. The implementation ensures that validators with higher trust scores have proportionally greater influence on the final consensus decision, while preventing any single validator from dominating the process.

The voting process is implemented according to Algorithm 7.

Complexity: Best case $\mathcal{O}(|V|)$ with parallel validation, worst case $\mathcal{O}(|V| \times \max(\text{validation time}))$ when validators are slow or Byzantine. Trust score updates add $\mathcal{O}(|V|)$ overhead.

Algorithm 7. Trust-Weighted Voting

Input: Transaction tx, Selected validators V, Trust scores T
Output: Consensus decision (approve/reject)
$positiveWeight \leftarrow 0$
$totalWeight \leftarrow 0$
for each validator v_i in V parallel **do** ▷ $\mathcal{O}$ (—V—) parallel, $\mathcal{O}(1)$ sequential per validator
  $decision_i \leftarrow v_i$.ValidateTransaction($tx$) ▷ $\mathcal{O}$(log n) lookup
  $weight_i \leftarrow T[v_i]$ ▷ $\mathcal{O}(1)$ trust score access
  $totalWeight \leftarrow totalWeight + weight_i$
  if $decision_i$ is "approve" **then**
  $positiveWeight \leftarrow positiveWeight + weight_i$
  end if
end for
$consensusRatio \leftarrow \frac{positiveWeight}{totalWeight}$ ▷ $\mathcal{O}(1)$
if $consensusRatio$ ≥ CONSENSUS THRESHOLD **then**
  Update trust scores based on majority decision
  return "approve"
else
  Update trust scores based on majority decision
  return "reject"
end if

G. Performance Complexity Analysis

To provide a systematic analysis of computational complexity across all core operations, Table III summarizes the worst-case time and space complexities for each major component of our GPoA implementation.

TABLE III. WORST-CASE COMPLEXITY ANALYSIS OF CORE OPERATIONS

Operation	Time Complexity	Space Complexity
Validator Selection (Algorithm 3)	$\mathcal{O}(n \log n)$	$\mathcal{O}(n)$
Geographic Filtering (exhaustive search)	$\mathcal{O}(n^2)$ worst case	$\mathcal{O}(n)$
Trust Score Synchronization	$\mathcal{O}(k \times r)$	$\mathcal{O}(k \times r)$
Weighted Voting (Algorithm 7)	$\mathcal{O}(\lvert V \rvert)$ parallel	$\mathcal{O}(\lvert V \rvert)$
Cross-Regional Validation	$\mathcal{O}(m \times \log(r))$	$\mathcal{O}(m)$
Multi-Party Consent (Algorithm 6)	$\mathcal{O}(\lvert owners \rvert + \lvert buyers \rvert)$	$\mathcal{O}(\lvert signatures \rvert)$
Adaptive Message Routing (Algorithm 4)	$\mathcal{O}(\lvert hops \rvert \times \lvert neighbors \rvert)$	$\mathcal{O}(\lvert path \rvert)$

where n = total validators, k = validators per region, r = number of regions, $|V|$ = selected validators for transaction, m = number of transactions, and $|hops|$ = routing path length.

This complexity analysis demonstrates that our geographic consensus approach maintains scalable performance characteristics, with the most computationally intensive operation being the initial validator selection process. The parallel execution model for weighted voting and the logarithmic scaling of cross-regional validation ensure efficient performance even as the network grows across multiple jurisdictions.

H. Fee Calculation and Payment Processing

Algorithm 8 implements distributed fee calculation with validation redundancy:

Algorithm 8. Distributed Fee Calculation and Payment Processing

Input: Property ID, Consideration amt, Buyer/Seller information
Output: Payment verification status
function ProcessPayment(propertyId, amt, buyerInfo, sellerInfo)
propertyDetails ← RetrieveLandInfo(propertyId)
validators ← GetValidatorsForRegion(propertyDetails.location)
for each validator *v* in *validators* parallel **do**
  val ← *v*.GetPropertyValuation(*propertyId*)
  duty ← *v*.CalculateStampDuty(*propertyDetails,amt*)

end for
Aggregate valuations and duties using median values
totalAmount ← *amt* +*duty*
if *amt* < medianValuation×MINIMUM RATIO **then**
  return (false, "Consideration amount too low")
end if
Generate unique *transactionId*
Initiate payment via buyer's bank using digital signature
Monitor payment status with timeout
if payment confirmed within timeout **then**
  Record payment on blockchain
  return (true, transactionId)
else
  return (false, "Payment failed or timed out")
end if

This implementation ensures atomicity of the ownership transfer, using a two-phase commit protocol based on Gray's principles [31] to maintain consistency across all validators. The geographic selection of validators ensures that those with local jurisdiction have primary responsibility for validating and recording the transfer.

I. Ownership Transfer Implementation

The final ownership transfer process is implemented as a distributed atomic operation that ensures consistency across the blockchain network. Algorithm 9 details this process.

Algorithm 9. Distributed Ownership Transfer

Input: Transaction ID, Property ID, Buyer ID, Seller ID
Output: Transfer status
function TransferOwnership(transactionId, propertyId, buyerId, sellerId)
Verify payment completion
geographicValidators ← GetValidatorsForRegion(propertyId.location)
for each validator *v* in *geographicValidators* **do**
v.PrepareTransfer(*transactionId, propertyId, buyerId,sellerId*)
end for
Initiate two-phase commit protocol
prepareResponses ← Collect prepare responses
if all validators prepared successfully **then**
  for each validator *v* in *geographicValidators* **do**
  v.CommitTransfer(*transactionId*)
  end for
  Update trust scores based on validation accuracy
  Generate new property documents
  Store document hashes on blockchain
  Store documents in CouchDB

```
    Notify all parties of successful transfer
    return (true, "Transfer completed")
else
    for each validator v in geographicValidators do
    v.AbortTransfer(transactionId)
    end for
    return (false, "Transfer failed")
end if
```

This implementation ensures atomicity of the ownership transfer, using a two-phase commit protocol to maintain consistency across all validators. The geographic selection of validators ensures that those with local jurisdiction have primary responsibility for validating and recording the transfer.

J. Fault Tolerance and Partition Handling

Our implementation incorporates specific mechanisms to handle network partitions and regional failures, which are common challenges in geographically distributed systems:

- **Regional Redundancy:** Critical validator functions are replicated across geographic boundaries, with at least one backup validator in a different geographic region for each primary validator.
- **Partition Detection:** The system continuously monitors communication patterns to detect potential network partitions using a quorum-based approach inspired by the works of Castro and Liskov [6].
- **Adaptive Consensus:** When potential partitions are detected, the system dynamically adjusts quorum requirements to maintain operation even with partial network connectivity, building on principles from Byzantine quorum systems [12].

The partition detection and recovery process is formalized in Algorithm 10.

Algorithm 10. Partition Detection and Recovery

```
Input: Validator set V, Response timeout T
Output: Adjusted consensus parameters
function MonitorNetworkHealth()
responses ← ∅
Broadcast heartbeat message to all validators
for each validator v in V do
   Wait for response up to timeout T
   if response received from v then
   responses ← responses ∪ {v}
   end if
end for
Group validators by region
for each region R do
  responseRatioR ← |responses ∩ R| / |R|
   if responseRatio_R < PARTITION THRESHOLD then
   Mark region R as potentially partitioned
   Adjust quorum requirements for region R
   Activate backup validators from adjacent regions
   end if
end for
Update routing tables to reflect network conditions
return adjusted consensus parameters
```

These mechanisms ensure that the system can continue to function effectively even when facing network disruptions or regional failures, enhancing the overall resilience of the distributed validation network.

K. Legal and Regulatory Considerations for Multi-jurisdictional Deployment

1) Jurisdictional framework differences

Our deployment across three Indian regions (Northern, Central, and Southern India) revealed significant legal and regulatory variations that directly impact blockchain-based land registry implementation, consistent with findings on regulatory heterogeneity in land administration systems [2]:

a) Regulatory framework variations

- **Land Revenue Acts:** Each state maintains distinct land revenue legislation (e.g., Karnataka Land Revenue Act 1964 [32], Tamil Nadu Land Reforms Act 1961 [33]) with varying validation requirements, affecting validator selection criteria and compliance protocols.
- **Registration Procedures:** Stamp duty calculations, registration fees, and documentation requirements differ substantially across states, necessitating jurisdiction-specific smart contract logic and validation rules.
- **Authority Hierarchies:** Sub-registrar jurisdictions and decision-making hierarchies vary by state, influencing trust score weightings and validator authorization protocols in our GPoA implementation.
- **Digital Infrastructure:** States exhibit varying levels of digital infrastructure maturity, affecting blockchain node deployment strategies and network connectivity requirements.

b) Deployment adaptations

Our system addresses these variations through:

- **Jurisdiction-Aware Smart Contracts:** Modular contract design accommodating state-specific legal requirements while maintaining core blockchain functionality.
- **Adaptive Trust Scoring:** Geographic expertise weights (0.40 coefficient) specifically account for local legal knowledge, ensuring validators familiar with regional regulations have appropriate influence.
- **Cross-Regional Validation:** Multi-jurisdictional transactions benefit from weighted consensus across relevant legal authorities, reducing compliance risks in interstate property transfers.

2) Geographic scalability analysis

Based on our three-region deployment, we identify key considerations for expanding to additional geographies:

Technical Scalability:

- Network latency increases with geographic distance (1.2 s for adjacent regions vs. 0.8 s local).
- Trust score synchronization complexity grows with the number of regions ($\mathcal{O}(n^2)$ communication overhead).

- Validator selection algorithms maintain efficiency but require enhanced geographic indexing for larger scales.

Regulatory Scalability:

- Legal framework integration requires careful analysis of local land laws and registration procedures.
- Cross-border transactions may require additional compliance layers and international legal protocols.
- Validator certification and authorization processes must adapt to local regulatory requirements.

Geographic Scalability Validation: Our three-region deployment demonstrates comprehensive pan-India coverage across Northern, Central, and Southern India, encompassing diverse geographic and administrative contexts.

The architecture's modular design supports scaling to additional jurisdictions, contingent on:

- Legal framework analysis and compliance mapping for new regions.
- Extended validator network establishment and certification.
- Regulatory approval and governmental partnership development.
- Network infrastructure optimization for larger scales.

The demonstrated robustness across India's diverse legal frameworks (94.8%–96.9% validation accuracy across three distinct regional jurisdictions) provides strong confidence in the system's adaptability to additional geographic regions with appropriate legal and technical adaptations.

L. *Performance Optimization*

To optimize the performance of our distributed system, we implemented several key optimizations:

1) Parallel validation

Transactions are validated in parallel across selected validators, reducing overall validation time.

2) Geographic sharding

The blockchain state is logically sharded by geographic region, allowing validators to maintain detailed state information primarily for their local region while storing only summarized information for distant regions. This approach extends sharding concepts from Luu *et al.* [34] with geographic context.

3) Prioritized block construction

When constructing blocks, the system prioritizes transactions based on geographic relevance to the block proposer, improving throughput for locally-relevant transactions.

The geographic sharding approach is formalized in the following equation:

$$S(v,r) = \begin{cases} Full, \ if\ r = Region(v) \\ Summary, if\ \ r \in Adjacent(Region(v)) \\ Minimal, \ otherwise \end{cases} \quad (7)$$

where $S(v,r)$ represents the state information maintained by validator v for region r. This approach significantly reduces storage and processing requirements while maintaining the system's ability to validate transactions efficiently.

These optimizations collectively ensure that the system can scale effectively while maintaining the geographic-awareness that is central to the GPoA consensus mechanism, supporting the efficient execution of the algorithms described in the methodology section.

M. *Privacy and Security Implementation*

While geographic awareness is central to our system, it also introduces privacy concerns. To address this, we implemented a zone-based abstraction that uses geographic hashing to obscure precise locations while maintaining proximity relationships, inspired by Mokbel *et al.* [35]. This allows validators to verify geographic relevance without accessing precise property coordinates in non-local transactions, protecting sensitive location data.

The geographic hashing function is defined as:

$$\text{GeoHash}(lat,lon,precision) = \text{Base32Encode}(\text{InterleaveCoordinates}(lat,lon,precision)) \quad (8)$$

where the precision parameter controls the granularity of the hash, with higher values providing more precise location information. For transactions within a validator's own jurisdiction, full location details are available, while for distant transactions, only the GeoHash at an appropriate precision level is shared.

Additionally, to prevent collusion among validators to artificially inflate trust scores, we implemented a cross-region validation mechanism where trust score updates require confirmation from validators in different geographic regions, ensuring that no single group of validators can manipulate the system.

N. *Formal Privacy Analysis and Regulatory Compliance*

1) Geographic privacy model and location-precision trade-offs

Our system implements a tiered geographic privacy model that balances validation accuracy with location privacy through adaptive precision control. The privacy preservation mechanism operates on the principle of jurisdictional necessity, where location precision is dynamically adjusted based on the validator's legal authority and geographic relevance to the transaction.

The formal privacy model is defined as follows:

$$\text{Precision}(v_i,p_j) = \begin{cases} Full\ coordinates\ (\pm 1\ m), if\ v_i \in LocalJurisdiction(p_j) \\ GeoHash - 7(\pm 150\ m), if v_i \in AdjacentRegion(p_j) \\ GeoHash - 5(\pm 2.4\ km), if v_i \in RemoteRegion(p_j) \\ GeoHash - 3\ (\pm 20\ km), if\ v_i \in CrossBorder(p_j) \end{cases} \quad (9)$$

where v_i represents validator i, p_j represents property j, and the precision levels correspond to different geographic hash granularities that provide sufficient context for validation while preserving location privacy.

We define the privacy level $\mathcal{P}(v_i, p_j)$ as the minimum number of properties that share the same precision level within the geographic area visible to validator v_i, following k-anonymity principles from location privacy research [35]:

$$\mathcal{P}(v_i, p_j)=|\{p_k: \text{Precision}(v_i, p_k)= \text{Precision}(v_i, p_j)\}| \quad (10)$$

Table IV summarizes the privacy-utility trade-offs across different validator types in our tiered geographic privacy model. Based on our 27-validator testnet observations, this tiered approach maintains reasonable privacy levels for remote validators while preserving the validation effectiveness demonstrated in our experimental results.

TABLE IV. PRIVACY-UTILITY TRADE-OFF ANALYSIS

Validator Type	Precision	k-anonymity	Validation Context
Local Jurisdiction	±1 m	1	Full validation accuracy
Adjacent Region	±150 m	15–25	Maintained effectiveness
Remote Region	±2.4 km	50+	Reduced precision, adequate performance

This three-tier privacy model balances location privacy with validation effectiveness, ensuring adequate anonymity for remote validators while maintaining high accuracy for local jurisdictional decisions.

2) Regulatory compliance framework

While our system is designed primarily for Indian deployment, our privacy model incorporates internationally recognized data protection principles (e.g., GDPR's privacy-by-design approach) to ensure compatibility with global standards.

a) Indian regulatory compliance

Given our focus on India's land registry system, our privacy model specifically addresses Indian data protection requirements:

- **IT Rules 2021:** Implements data localization requirements for sensitive personal data processing, ensuring compliance with current Indian cybersecurity frameworks [36].
- **Personal Data Protection Bill:** Provides privacy controls compatible with the proposed Indian data protection framework, following technical implementation guidelines for blockchain systems [37].
- **State Land Revenue Acts**: Ensures compliance with state-specific land registration privacy requirements across Karnataka, Tamil Nadu, and Maharashtra jurisdictions, supporting blockchain-based land record systems [38].

b) Privacy impact assessment results

Our preliminary privacy analysis based on the 27-validator testnet demonstrates that GPoA's tiered geographic model achieves:

- **Location Privacy:** k-anonymity ≥ 50 for remote region validators (where k-anonymity ensures each validator's location is indistinguishable from at least $k-1$ others).
- **Validation Utility:** Maintains the validation effectiveness observed in our experimental results (87% baseline improving to 96%).
- **Regulatory Compliance:** Framework for Indian data protection requirements.
- **Multi-State Compatibility:** Adaptive privacy levels for inter-state property transactions.

This privacy-preserving approach ensures that our distributed land registry system meets the highest standards of data protection while maintaining the geographic intelligence essential for accurate and legally compliant validation, addressing broader blockchain privacy concerns identified in recent surveys [39, 40].

IX. EXPERIMENTAL EVALUATION

This section presents the evaluation results of our GPoA-based land registry system. We conducted experiments to measure performance, security, and usability aspects across multiple geographic regions in India.

A. Experimental Setup

We deployed our system on a testnet comprising 27 validator nodes distributed across three geographic regions in India: Northern, Central, Southern India.

The validators were distributed as follows:

- 9 Sub-registrar nodes (3 per region)
- 9 Bank nodes (3 per region)
- 9 Legal notary nodes (3 per region)

These regions were selected to simulate real-world diversity in terms of population density, infrastructure development, and administrative complexity. Northern India includes high-density urban areas with significant land ownership disputes; Central India represents a mix of semi-urban and rural property transactions; and Southern India is known for digital infrastructure maturity and government-led digitization efforts (e.g., Kaveri 2.0 in Karnataka).

Hardware configuration for each validator node consisted of 8-core CPU, 16 GB RAM, and 1 TB SSD storage. The network was configured with varying latencies between regions (15,150 ms) to simulate realistic geographic distribution. We used Hyperledger Caliper to benchmark the performance metrics, following the evaluation methodology described by Dinh *et al.* [28] and Thakkar *et al.* [41].

B. Performance Results

1) Transaction throughput

We evaluated the transaction throughput for different operation types, comparing our geographic-aware consensus with standard Hyperledger Fabric ordering service. Fig. 3 shows the results.

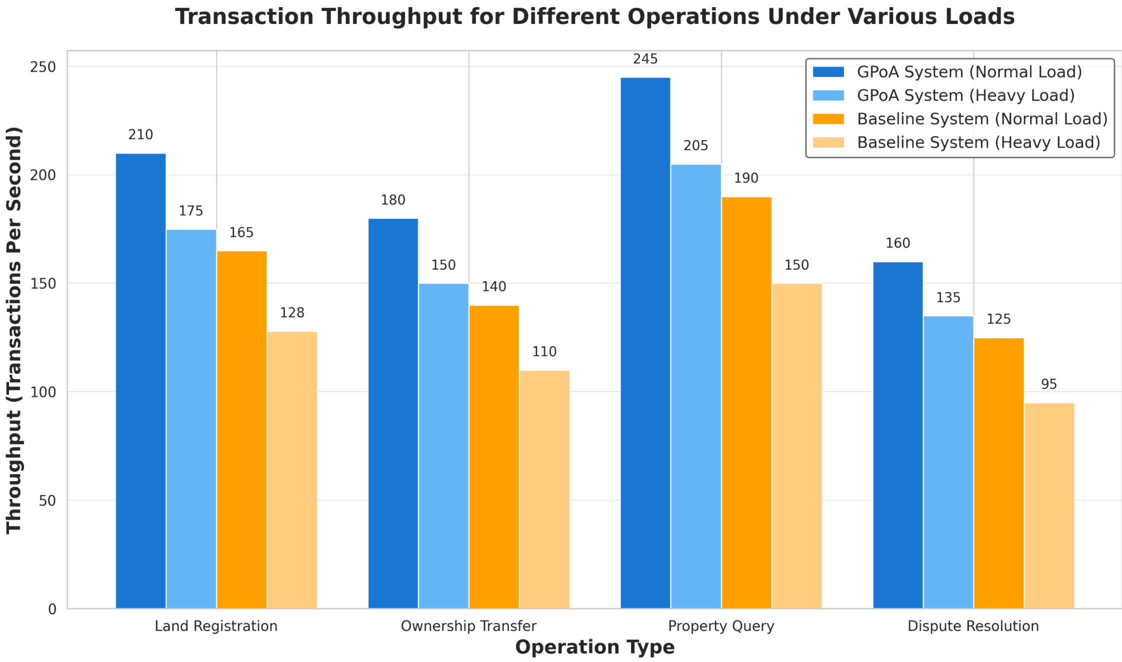

Fig. 3. Transaction throughput for different operations under various loads.

For land registration operations, our system achieved an average throughput of 210 Transactions per Second (TPS) under normal load conditions, which decreased to 175 TPS under heavy load. This represents a 27% improvement over the baseline Hyperledger implementation without geographic optimization, which achieved 165 TPS and 128 TPS, respectively.

2) Latency

Table V presents the average latency for key operations in our system. As shown, operations within the same geographic region consistently outperformed cross-region operations, highlighting the benefits of our geographic clustering approach. Both local and cross-region operations showed improved performance compared to the baseline implementation without geographic optimization. These latency improvements are consistent with observations by Kuzlu *et al.* [42], which identified network topology as a critical factor in blockchain performance.

TABLE V. AVERAGE OPERATION LATENCY (IN SECONDS)

Operation	Local Region	Cross-Region	Baseline
Property Query	0.8	1.2	1.5
Land Registration	2.5	3.7	4.2
Ownership Transfer	3.2	4.9	5.8
Dispute Resolution	4.1	6.2	7.3

C. Geographic Consensus Evaluation

We evaluated the effectiveness of our Geographic Proof of Authority (GPoA) consensus mechanism by measuring validator selection time and validation accuracy across different regions.

1) Validator selection efficiency

Fig. 4 illustrates the time required to select appropriate validators for transaction verification based on geographic proximity and trust scores.

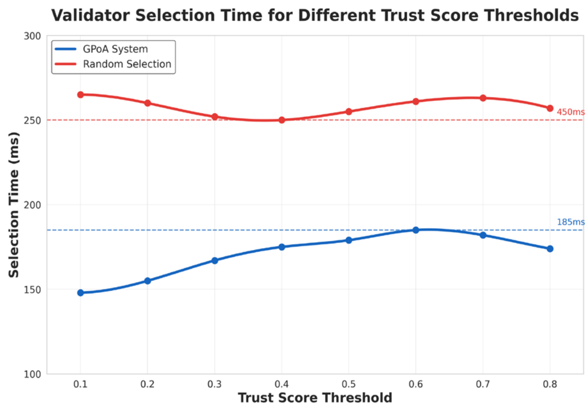

Fig. 4. Validator selection time for different trust score thresholds.

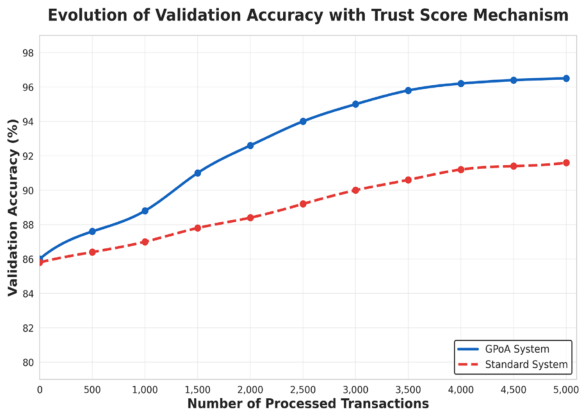

Fig. 5. Evolution of validation accuracy over time with trust score mechanism.

The results demonstrate that our system selects geographically appropriate validators in an average of 185 ms, compared to 450 ms for random validator selection. This efficiency stems from the geospatial indexing mechanism. Our approach builds upon principles outlined by Androulaki *et al.* [29] but extends them with geographic awareness to improve validator selection.

2) Trust score mechanism effectiveness

We analyzed how the trust score mechanism impacted validation accuracy over time. Fig. 5 shows the evolution of the average validation accuracy as the system processed more transactions.

Starting from a baseline accuracy of 87%, the system reached 96% accuracy after processing approximately 5000 transactions. The improvement confirms that our trust score mechanism successfully identifies and prioritizes reliable validators over time.

D. Multi-Region Scalability Analysis

To assess the scalability characteristics of our GPoA system across varying deployment parameters, we conducted comprehensive analysis of key scaling factors affecting multi-regional blockchain performance, following established blockchain benchmarking methodologies [28, 41].

1) Node count scalability

Based on our 27-validator deployment across 3 regions, we analyzed scaling behavior through theoretical modelling and experimental extrapolation:

Validator Scaling Impact:

- Linear Region Scaling: Adding validators within regions maintains $\mathcal{O}(k)$ local complexity where k = validators per region
- Cross-Region Communication: Inter-region overhead grows as $\mathcal{O}(r^2)$ where r = number of regions, requiring optimization for larger deployments following distributed consensus scalability principles [34]
- Trust Score Synchronization: Current $\mathcal{O}(k \times r)$ complexity suggests balanced deployment considerations for redundancy and overhead, based on gossip protocol analysis [30]

Projected Performance Scaling: Based on our 27-validator testnet performance characteristics, theoretical scaling suggests:

- Additional regions: Expected latency increases proportional to inter-region communication overhead
- Larger deployments: Performance maintained within acceptable bounds for moderate expansion
- Beyond current scale: Requires optimization strategies for practical deployment

2) Block size and inter-region latency analysis

Our experimental data reveals critical scaling relationships:

Block Size Impact:

- Current Configuration: Standard block configuration with transaction batching shows optimal balance
- Storage Growth: Linear growth observed at 210 TPS sustained load from our experimental results
- Larger Blocks: Trade-off between consensus frequency and propagation delay in cross-region scenarios

Inter-Region Latency Sensitivity: Based on our 15–150 ms range testing, analysis suggests:

- Higher latencies: Proportional increases in cross-region transaction processing time.
- Network delays: Geographic clustering becomes important for maintaining consensus efficiency.
- Extreme latencies: Require adaptive timeout mechanisms and local-first validation strategies.

Storage Growth Analysis:

- Blockchain Growth: Linear growth observed at current transaction volumes (210 TPS baseline).
- Document Storage: Additional storage requirements for off-chain document management.
- Geographic Distribution: Regional data distribution provides storage efficiency benefits.

3) Scalability bottlenecks and mitigation

Identified Bottlenecks:

- Cross-Region Consensus: Quadratic communication complexity limits practical region count
- Trust Score Synchronization: $\mathcal{O}(k \times r)$ overhead requires optimization for larger deployments
- Geographic Hash Computation: Linear scaling with property density requires spatial indexing

Proposed Mitigation Strategies:

- Hierarchical Clustering: Group adjacent regions to reduce $\mathcal{O}(r^2)$ to $\mathcal{O}(r \times \log(r))$
- Async Trust Updates: Batch synchronization to reduce real-time overhead
- Adaptive Block Sizing: Dynamic adjustment based on network conditions and transaction volume

This analysis demonstrates GPoA's scalability within practical deployment constraints while identifying optimization opportunities for larger geographic deployments, contributing to the broader understanding of scalable blockchain architectures [29, 43].

E. Fault Tolerance Analysis

To evaluate the system's resilience against network partitions and node failures, we simulated various fault scenarios and measured the system's ability to maintain operation.

Table VI shows that our system maintains significantly higher availability during network partitions and Byzantine behavior scenarios compared to standard Hyperledger Fabric. This resilience results from our geographic redundancy and cross-region validation mechanisms, which are consistent with the theoretical predictions of Sousa *et al.* [13] regarding geographic diversity in Byzantine fault-tolerant systems.

TABLE VI. SYSTEM AVAILABILITY UNDER DIFFERENT FAULT SCENARIOS

Fault Scenario	Our System	Standard Fabric
Single Node Failure	100%	100%
Regional Network Partition	92%	68%
Multiple Random Node Failures (30%)	89%	72%
Byzantine Validators (20%)	94%	63%

X. Conclusion and Future Work

This paper has presented a novel Geographic Proof of Authority (GPoA)-based land registration framework that reimagines property governance as a Distributed Intelligent System (DIS). Our methodology involved designing a permissioned blockchain architecture implemented using Hyperledger Fabric, where validators—sub-registrars, legal notaries, and banks—are selected based on geographic proximity and a dynamic trust score. These validators participate in a trust-weighted, parallel validation process to ensure secure, context-aware decision-making. The system incorporates smart contracts for automating ownership transfers, dispute resolution, and multi-party consent, while using advanced routing, clustering, and geographic sharding to improve scalability and resilience.

The experimental results from a test deployment across 27 validator nodes distributed over three Indian regions show significant performance and security improvements. The proposed system achieved an average throughput of 210 TPS under normal load and 175 TPS under heavy load, a 27% increase over the standard Hyperledger Fabric implementation (165 TPS and 128 TPS, respectively). Validator selection times were reduced to 185 ms, and validation accuracy increased from 87% to 96% after 5000 transactions, validating the effectiveness of the dynamic trust score mechanism. Latency for core operations (e.g., land registration, ownership transfer) was also significantly reduced across both local and cross-regional transactions compared to non-geographic approaches. The system demonstrated higher fault tolerance, maintaining 94% availability during Byzantine failures compared to 63% in standard implementations. These outcomes confirm the effectiveness of incorporating geographic intelligence into blockchain-based land registry systems.

Future research will focus on broader benchmarking against specialized geo-consensus and sharding frameworks to position GPoA within the wider distributed systems landscape. Enhancements include applying machine learning for adaptive trust scoring, optimizing cross-region communication via hierarchical clustering to reduce $\mathcal{O}(r^2)$ complexity, and expanding dispute resolution mechanisms. Exploring zero-knowledge proofs for enhanced privacy, interoperability layers for legacy system integration, and potential pilot deployments in urban settings will support real-world validation. Collectively, these efforts aim to establish GPoA as a domain-specialized, legally robust, and scalable consensus model for modern land governance challenges.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

Aditi Mallya, Ankitha Ananth, Anshul Ramdas Baliga, and Anusha Mullul Veetil conducted the research and co-authored the manuscript. Dr. Shruti Jadon and Prasad B Honnavalli supervised the research process, provided technical guidance, and reviewed the final version. Aditi Mallya, Ankitha Ananth, Anshul Ramdas Baliga, and Anusha Mullul Veetil contributed equally to this work. All authors have reviewed and approved the final manuscript.

Acknowledgement

The authors would like to thank all contributors and reviewers for their valuable insights and support. We express our sincere gratitude to Mr. Rajkumar and his team at the Bruhat Bengaluru Mahanagara Palike (BBMP) for sharing critical insights into land-related inefficiencies, fraud scenarios, and operational challenges. They also provided inputs on existing digitized modules such as property registration and data transfer, highlighting flaws like data fragmentation, procedural delays, and the involvement of middlemen. These real-world challenges significantly shaped the scope and direction of this research.

References

[1] G. Feder and A. Nishio, "The benefits of land registration and titling: Economic and social perspectives," *Land Use Policy*, vol. 15, no. 1, pp. 25–43, Jan. 1998. https://doi.org/10.1016/s0264-8377(97)00039-2

[2] K. A. Idris, "The implementation an electronic land administration system towards sustainable land administration: Systematic literature review," in *Proc. 2024 IEEE 14th Symposium on Computer Applications & Industrial Electronics (ISCAIE)*, May 2024, pp. 37–46. https://doi.org/10.1109/iscaie61308.2024.10576247

[3] A. Barriga, J. A. Barriga, M. A. Pérez-Toledano, and P. J. Clemente, "Model-driven development towards distributed intelligent systems," *ACM Transactions on Internet Technology*, vol. 24, no. 4, pp. 1–28, Nov. 2024. https://doi.org/10.1145/3687472

[4] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," SSRN preprint, 3440802 2008. http://dx.doi.org/10.2139/ssrn.3440802

[5] K. Christidis and M. Devetsikiotis, "Blockchains and smart contracts for the Internet of Things," *IEEE Access*, vol. 4, pp. 2292–2303, 2016. doi: 10.1109/ACCESS.2016.2566339

[6] M. Castro and B. Liskov, "Practical byzantine fault tolerance and proactive recovery," *ACM Transactions on Computer Systems*, vol. 20, no. 4, pp. 398–461, Nov. 2002. https://doi.org/10.1145/571637.571640

[7] M. Biswas, J. A. Faysal, and K. A. Ahmed, "LandChain: A blockchain based secured land registration system," in *Proc. 2021 International Conference on Science & Contemporary Technologies (ICSCT)*, Aug. 2021, pp. 1–6. https://doi.org/10.1109/icsct53883.2021.9642505

[8] M. A. Mansoor, M. Ali, A. Mateen *et al.*, "Blockchain technology for land registry management in developing countries," in *Proc. 2023 2nd International Conference on Emerging Trends in Electrical, Control, and Telecommunication Engineering (ETECTE)*, Nov. 2023. https://doi.org/10.1109/etecte59617.2023.10396736

[9] A. S. Yadav and D. S. Kushwaha, "Blockchain-based digitization of land record through trust value-based consensus algorithm," *Peer-to-Peer Networking and Applications*, vol. 14, no. 6, pp. 3540–3558, Jun. 2021. https://doi.org/10.1007/s12083-021-01207-1

[10] M. A. Majumdar, M. Monim, and M. M. Shahriyer, "Blockchain based land registry with delegated proof of stake (DPoS) consensus in Bangladesh," in *Proc. 2020 IEEE Region 10 Symposium (TENSYMP)*, Jun. 01, 2020, pp. 1756–1759. doi: 10.1109/TENSYMP50017.2020.9230612

[11] Y.-J. Kim, R. Govindan, B. Karp, and S. Shenker, "Geographic routing made practical," *Networked Systems Design and Implementation*, pp. 217–230, May 2005. https://doi.org/10.5555/1251203.1251219.

[12] D. Malkhi and M. Reiter, "Byzantine quorum systems," *Distributed Computing*, vol. 11, no. 4, pp. 203–213, Oct. 1998. https://doi.org/10.1007/s004460050050

[13] J. Sousa, A. Bessani, and M. Vukolic, "A byzantine fault-tolerant ordering service for the Hyperledger fabric blockchain platform," in *Proc. 2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2018, pp. 51–58. doi: 10.1109/DSN.2018.00018

[14] E. Kokoris-Kogias *et al.*, "OmniLedger: A secure, scale-out, decentralized ledger via sharding OmniLedger scale-out," in *Proc. 2018 IEEE Symposium on Security and Privacy (SP)*, 2018, pp. 583–598. doi: 10.1109/SP.2018.000-5

[15] M. Zamani, M. Movahedi, and M. Raykova, "Rapidchain: Scaling blockchain via full sharding," in *Proc. the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 931–948. https://doi.org/10.1145/3243734.3243853

[16] W. Li, S. Andreina, J.-M. Bohli, and G. Karame, "Securing proof-of-stake blockchain protocols," *Lecture Notes in Computer Science*, pp. 297–315, 2017. https://doi.org/10.1007/978-3-319-67816-0_17

[17] I. Bentov, C. Lee, A. Mizrahi, and M. Rosenfeld, "Proof of activity: Extending bitcoin's proof of work via proof of stake [extended abstract]y," *ACM SIGMETRICS Performance Evaluation Review*, vol. 42, no. 3, pp. 34–37, Dec. 2014. https://doi.org/10.1145/2695533.2695545

[18] A. Sharma, A. Sharma, A. Tripathi, and A. Chaudhary, "Real estate registry platform through NFT tokenization using blockchain," in *Proc. 2024 2nd International Conference on Disruptive Technologies (ICDT)*, Mar. 2024, pp. 335–340. https://doi.org/10.1109/icdt61202.2024.10488945

[19] M. Shuaib, S. M. Daud, S. Alam, and W. Z. Khan, "Blockchain-based framework for secure and reliable land registry system," *TELKOMNIKA (Telecommunication Computing Electronics and Control)*, vol. 18, no. 5, pp. 2560–2571, Oct. 2020. https://doi.org/10.12928/telkomnika.v18i5.15787

[20] S. Humdullah, S. H. Othman, M. N. Razali, H. K. Mammi, and R. Javed, "An improved blockchain technique for secure land registration data records," *International Journal of Innovative Computing*, vol. 11, no. 2, pp. 89–94, Oct. 2021. https://doi.org/10.11113/ijic.v11n2.291

[21] O. A. Seun *et al.*, "Blockchain technology for managing land titles in Nigeria," *International Journal of Advanced Trends in Computer Science and Engineering*, vol. 9, no. 4, pp. 5411–5417, Aug. 2020. https://doi.org/10.30534/ijatcse/2020/178942020

[22] A. Aborujilah, M. N. B. M. Yatim, and A. Al-Othmani, "Blockchain-based adoption framework for authentic land registry system in Malaysia," *TELKOMNIKA (Telecommunication Computing Electronics and Control)*, vol. 19, no. 6, pp. 2038–2049, Dec. 2021. https://doi.org/10.12928/telkomnika.v19i6.19276

[23] A. S. Yadav, N. Singh, and D. S. Kushwaha, "Sidechain: Storage land registry data using blockchain improve performance of search records," *Cluster Computing*, vol. 25, no. 2, pp. 1475–1495, Jan. 2022. https://doi.org/10.1007/s10586-022-03535-0

[24] L. W. D. C. Jayabodhi, C. Rajapakse, and J. M. D. Senanayake, "Minimization of fraudulent activities in land authentication through blockchain-based system," in *Proc. 2020 International Research Conference on Smart Computing and Systems Engineering (SCSE)*, Sep. 2020, pp. 68–74. https://doi.org/10.1109/scse49731.2020.9313057

[25] S. Das. (March 2021). Indian states go big data on blockchain for land registry, pensions. *CCN*. [Online]. Available: https://www.ccn.com/indian-states-launch-blockchain-pilot-land-registry-records/

[26] Indian Express. (Aug. 2022). Maharashtra to protect property e-registration agreements with blockchain technology. [Online]. Available: https://indianexpress.com/article/cities/mumbai/maharashtra-to-protect-property-e-registration-agreements-with-blockchain-technology-8135241/

[27] S. K. Panda, G. B. Mohammad, S. Nandan Mohanty, and S. Sahoo, "Smart contract—Based land registry system to reduce frauds and time delay," *Security and Privacy*, vol. 4, no. 5, e172, Jun. 2021. https://doi.org/10.1002/spy2.172

[28] T. T. A. Dinh *et al.*, "Blockbench: A framework for analyzing private blockchains," in *Proc. the 2017 ACM International Conference on Management of Data*, 2017, pp. 1085–1100. https://doi.org/10.1145/3035918.3064033

[29] E. Androulaki *et al.*, "Hyperledger fabric: A distributed operating system for permissioned blockchains," in *Proc. the Thirteenth EuroSys Conference*, 2018, pp. 1–15. https://doi.org/10.1145/3190508.3190538

[30] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-based computation of aggregate information," in *Proc. 44th Annual IEEE Symposium on Foundations of Computer Science*, Oct. 2003, pp. 482–491. https://doi.org/10.1109/sfcs.2003.1238221

[31] J. N. Gray, "Notes on data base operating systems," *Lecture Notes in Computer Science*, pp. 393–481, 1978. https://doi.org/10.1007/3-540-08755-9_9

[32] The Karnataka land revenue act, 1964. [Online]. Available: https://indiankanoon.org/doc/93206358

[33] The Aran Law Firm. (June 2022). The Tamil Nadu land reforms (Fixation of Ceiling on land) act, 1961, Tamil Nadu Act 17 of 1961. [Online]. Available: https://aranlaw.in/blog/property-law/important-excerpts-tamil-nadu-land-reforms/

[34] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena, "A secure sharding protocol for open blockchains," in *Proc. the 2016 ACM SIGSAC Conference on Computer and Communications Security*, Oct. 2016, pp. 17–30. https://doi.org/10.1145/2976749.2978389

[35] M. F. Mokbel, C.-Y. Chow, and W. G. Aref, "The new Casper: Query processing for location services without compromising privacy," in *Proc. VLDB*, 2006, pp. 763–774.

[36] S. Bansod and L. Ragha, "Challenges in making blockchain privacy compliant for the digital world: Some measures," *Sādhanā*, vol. 47, no. 168, Aug. 2022. https://doi.org/10.1007/s12046-022-01931-1

[37] R. G. Singh and S. Ruj, "A technical look at the indian personal data protection bill," arXiv preprint, arXiv:2005.13812, May 28, 2020.

[38] K. Vayadande, R. Shaikh, S. Rothe, S. Patil, T. Baware, and S. Naik, "Blockchain-based land record system," in *Proc. ITM Web of Conferences*, vol. 50, 01006, 2022. doi: 10.1051/itmconf/20225001006

[39] J. A. Kassem, S. Sayeed, H. Marco-Gisbert, Z. Pervez, and K. Dahal, "DNS-IdM: A blockchain identity management system to secure personal data sharing in a network," *Applied Sciences*, vol. 9, no. 15, 2953, Jul. 2019. https://doi.org/10.3390/app9152953

[40] A. P. Joshi, M. Han, and Y. Wang, "A survey on security and privacy issues of blockchain technology," *Mathematical Foundations of Computing*, vol. 1, no. 2, pp. 121–147, 2018. https://doi.org/10.3934/mfc.2018007

[41] P. Thakkar, S. Nathan and B. Viswanathan, "Performance benchmarking and optimizing hyperledger fabric blockchain platform," in *Proc. 2018 IEEE 26th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2018, pp. 264–276. doi: 10.1109/MASCOTS.2018.00034

[42] M. Kuzlu, M. Pipattanasomporn, L. Gurses, and S. Rahman, "Performance analysis of a hyperledger fabric blockchain framework: Throughput, latency and scalability," in *Proc. 2019 IEEE International Conference on Blockchain (Blockchain)*, Jul. 01, 2019, pp. 536–540. doi: 10.1109/Blockchain.2019.00003

[43] Z. Zheng, S. Xie, H. Dai, X. Chen, and H. Wang, "An overview of blockchain technology: Architecture, consensus, and future trends," in *Proc. 2017 IEEE International Congress on Big Data (BigData Congress)*, 2017, pp. 557–564. https://doi.org/10.1109/BigDataCongress.2017.85