

LoChain: A Decentralized and Privacy-Preserving Blockchain Protocol for Mobility Data Management

Merouane Mohamed Smaine Bouderbala¹, Didem Demirag¹, Sébastien Gambs¹

¹ Université du Québec à Montréal (UQAM), Département d'informatique, Montréal, Québec, Canada
bouderbala.merouane@courrier.uqam.ca, gambs.sebastien@uqam.ca, demirag.didem@uqam.ca

Keywords: Blockchain, Location Privacy, Hyperledger Fabric, Civic Data Governance, Identity Rotation, Decentralized Autonomous Organization, Tor, Geoaddress Abstraction.

Abstract

Mobility data has become a strategic asset in urban planning, crisis management and smart city operations. However, centralized systems for mobility tracking raise severe privacy concerns as they have the ability to directly link individuals to their movements. To address these issues, we propose LoChain, a decentralized protocol that enables the privacy-preserving collection and processing of mobility data based on blockchain technology. More precisely, LoChain replaces precise coordinates with standardized geoaddresses, associate user movements to disposable identities, communication them via Tor routing and stores the resulting data across a decentralized network built on Hyperledger Fabric. The system also employs a novel geopool and multi-channel architecture to simulate sharding, enabling localized data ingestion, inter-district communication and global statistical aggregation without compromising individual privacy. Localized position obfuscation and pseudo-random identity purging are used to further prevent re-identification. A proof-of-concept prototype, including an Android app, blockchain backend and visualization layer was developed and evaluated using synthetic data from 10,000 virtual users. The experiments results obtained from the simulation highlight the LoChain's ability to preserve user privacy while maintaining analytical utility. Finally, we also introduce an incentive model as well as a decentralized governance structure to ensure long-term scalability, regulatory compliance and participatory control.

1. Introduction

The rapid growth of smartphones and mobile devices has dramatically increased the collection and use of mobility data. Mobility data are essential for a wide range of applications beyond simple navigation. For instance, urban planning relies on mobility data for infrastructure optimization, emergency services leverage it for real-time response coordination, while businesses use it to tailor services to individual behaviors. Yet, the benefits of mobility data are increasingly overshadowed by substantial concerns regarding privacy, data security and ethical usage.

Traditional services for managing mobility data, such as the ones handled by Google and Apple, store extensive amounts of sensitive personal data on centralized servers. While these services offer convenience and detailed insights into population mobility and behavior, a centralized architecture means that any security breach or malicious use can expose users' personal movements, preferences and habits. For example, incidents such as Google's unauthorized location data collection despite users disabling location history (Lohr, 2018) or Apple's continued data collection through their "Significant Locations" feature highlight the existing tensions with transparency as well as user privacy (O'Flaherty, 2020).

Moreover, numerous studies have demonstrated that even so-called anonymous mobility data can be re-identified by correlating it with external datasets, thus compromising user privacy. Notably, researchers have successfully re-identified individuals from seemingly anonymous datasets by analyzing a few distinct positions (De Montjoye et al., 2013), raising serious doubts about the effectiveness of traditional anonymization practices.

To address these limitations, we propose LoChain, a decentralized system specifically designed to protect user privacy while

maintaining the usability and analytical utility of mobility data. More precisely, LoChain adopts an architecture built upon Hyperledger Fabric, enabling secure data sharing and robust privacy controls that are not possible in centralized models. In addition, by implementing privacy-preserving techniques such as disposable identities, Tor-based anonymization and noise injection, LoChain mitigates the risk of user identification and unauthorized data exploitation.

The outline of the paper is as follows. First, we review briefly in Section 2 the relevant related work on mobility data management systems. Afterwards in Section 3, we introduce our proposed system LoChain, first presenting its fundamental principles and the building blocks upon which it is built before describing in more details how it manages mobility data in a decentralized and privacy-preserving manner. We describe respectively the prototype implementation of LoChain and its experimental evaluations in Sections 4 and 5. Then, we propose an incentive model in Section 6 before discussing the possible decentralized governance structure in Section 7. Finally, we conclude by discussing possible future works in Section 8.

2. Related Work

Mobility data can be collected through a variety of means, including but not limited to navigation apps. It can also be extracted from photo metadata, inferred from IP addresses or estimated through triangulation of cellular and Wi-Fi network signals. This diversity of collection methods raises important questions regarding data accuracy and individual privacy.

For instance, Google Maps is a widely adopted mapping and navigation solution, thanks in part to its native integration into Android, the dominant smartphone operating system, which

is used by over a billion unique users each month. Google Maps offers a variety of services, including real-time navigation, traffic information, street and satellite views and business listings. However, Google's centralized model for managing location data has been criticized for its lack of transparency. In 2018, Google disclosed a security breach in Google+ that exposed the personal data of 500,000 users (Lohr, 2018). Additionally, the Australian Competition and Consumer Commission (ACCC) sued Google for misleading consumers about the collection of their location data (McKinnell, 2021).

Apple Maps, integrated into iOS devices, is a popular alternative to Google Maps. Apple emphasizes user privacy, but some features, such as the "Significant Locations" service, have raised concerns. This service automatically records frequently visited locations to offer personalized services, such as navigation suggestions to important places like home or work. However, in 2020, a study revealed that Apple continued to collect some location data even when Location Services were disabled (O'Flaherty, 2020). While these practices are intended to improve the user experience, they have reinforced criticism regarding the lack of transparency in mobility data management.

OpenStreetMap (OSM) is a collaborative project that aims to become the equivalent of Wikipedia for geolocation data (Haklay and Weber, 2008). To this end, it offers access to map-editing software, allowing users to add roads, paths, buildings and other points of interest using satellite imagery, personal GPS readings, or other data sources, as well as enrich maps with tags. OSM has proven extremely valuable during humanitarian crises, such as the 2010 Haiti earthquake, where it created the most detailed digital map of Haiti at the time in just two days, becoming an indispensable tool for all NGOs involved in the humanitarian aid effort. OSM's free nature and collaborative nature represent both its main strengths and limitations. While this approach allows for remarkable responsiveness and usefulness in emergency situations, the lack of substantial financial and hardware support is a challenge for the long-term sustainability of OSM.

FOAM (FOAM, 2018) is a decentralized protocol built on Ethereum (Buterin, 2014), utilizing specialized hardware called anchors to provide accurate and decentralized location-based services. FOAM employs economic incentives, called FOAM tokens to encourage secure and decentralized location validation, supported by Byzantine fault-tolerant synchronization to secure time and location proofs. Despite its innovative approach, FOAM relies on specific hardware, introducing complexity and costs that limit widespread adoption.

Other approaches based on crowdsensing aim at harnessing the power of collective participation to enable large-scale data collection through mobile devices. However, this approach also raises challenges, particularly in terms of confidentiality. Thus, the collection and processing of personal data pose significant security issues, requiring measures to protect sensitive information while still providing useful data. For example, in OpenStreetMap, one of the security measures implemented is user pseudonymization, in which contributions are associated with pseudonyms rather than real identities, thus reducing the risk of personal data disclosure. Furthermore, Transport Layer Security (TLS) encryption is used to secure communications between users and servers, thus preventing data interception during transmission.

Additionally, Xiong and co-authors have showed that differential privacy is an approach that can be applied to limit the risk

of user re-identification (Xiong et al., 2014). This technique introduces random noise into the collected data to ensure that no individual information can be isolated while still enabling reliable aggregated analyses. This method is particularly useful for massive datasets where re-identification attacks are possible. Privacy-preserving solutions like CrowdBLPS (Zou et al., 2019) combine blockchain with advanced cryptographic techniques, such as homomorphic encryption and zero-knowledge proofs (Goldreich et al., 1991), to preserve user privacy during participatory sensing. However, these approaches often require explicit user coordination, trusted setup or heavy computation, limiting their scalability and real-time applicability.

LoChain adopts a different philosophy by avoiding custom hardware and heavyweight cryptography in favor of standard mobile GPS, spatial abstraction via geoadresses and disposable identities combined with network-layer anonymization through Tor. Its architecture enables the decentralized ingestion through spatially partitioned Fabric channels and provides statistical utility without exposing individual paths. Hyperledger Fabric has been selected for its permissioned consensus, channel isolation and modularity, which align well with a decentralized data governance and privacy-preserving architecture.

3. Proposed System: LoChain

To address the critical privacy issues and centralization risks identified earlier, we propose LoChain, a blockchain-based mobility management system that emphasizes user privacy and data decentralization. Instead of attempting to replace existing navigation tools, LoChain acts as a privacy-enhancing layer compatible with widely adopted GPS services, maximizing privacy while minimizing disruption.

Fundamental principles. At its core, LoChain is built around five architectural tenets:

- *Privacy at the point of capture.* Mobility data is anonymized on-device before transmission, eliminating reliance on the backend for privacy enforcement.
- *No raw coordinate sharing.* The underlying application implementing LoChain never exposes or transmits actual GPS coordinates, rather the positions are encoded to refer to abstracted geoadresses.
- *Decentralized ingestion and aggregation.* The data is submitted to district-specific fabric channels and later aggregated through a global ledger.
- *Delayed disclosure.* The aggregated data becomes publicly accessible only after a 24-hour buffer to prevent real-time exploitation and privacy breaches that could be caused by timing information.
- *Composable governance.* The governance and revenue logic are modular, supporting stakeholder-defined access, upgrades and incentives via DAO mechanisms.

3.1 Building Blocks

Blockchain. A blockchain is a decentralized, distributed, and often public digital ledger (Nakamoto, 2008). This technology records transactions in linked and secured blocks using cryptography, ensuring data integrity and immutability. Each

block contains a cryptographic hash of the previous block, a timestamp, and the transaction data, thus forming a continuous and secure chain. For example, in a financial transaction, if Alice sends bitcoins to Bob, this transaction will be recorded in a block containing the transaction information, the precise timestamp of the event, and a cryptographic hash linking this block to all previous blocks. This hash ensures that any modification made to one block would affect all subsequent blocks, making manipulation extremely difficult and detectable. This structure effectively prevents any retroactive modification without altering all subsequent blocks, thus ensuring the transparency and reliability of the recorded data.

Each network participant has a complete copy of the blockchain and can independently verify and audit transactions. For example, in the case of a transaction between Alice and Bob, any participant can view and validate this transaction without needing to trust a central authority. This reduces the costs associated with auditing and verification, as built-in cryptographic mechanisms ensure data accuracy. Data in the blockchain is managed autonomously using a peer-to-peer network and a distributed timestamping server. Transactions are authenticated by a distributed consensus mechanism (Castro and Liskov, 1999), in which network participants (or “nodes”) verify and validate each transaction. This process is fueled by personal incentives, such as rewards for miners in the case of Proof of Work or gains for validators in the case of Proof of Stake. These mechanisms encourage participants to act honestly for the network.

Hyperledger Fabric. Hyperledger Fabric (Hyperledger Fabric — hyperledger.org, n.d.) represents a nuanced and flexible approach, designed specifically for permissioned environments, that distinguishes itself from public blockchains like Ethereum or Bitcoin due to its modularity, customizable consensus protocols and enterprise-oriented architecture. Unlike public blockchains, Hyperledger Fabric relies on robust digital identities managed by Certificate Authorities to authenticate participants. Each participating organization has a digital certificate issued by a CA, ensuring that only authorized actors can access and transact on the network. This approach provides granular control over permissions and improves overall system security by minimizing the risk of unauthorized access.

Additionally, Hyperledger Fabric is built around channels, a concept for partitioning data within the network. More specifically, a channel is a private subnetwork that only certain participants can access, and transactions within a channel are visible only to its members, providing increased privacy without sacrificing overall transparency. In our project, this functionality is leveraged to create ingestion, edge and global channels, enabling localized and secure data management while ensuring global consistency. Another key aspect of Hyperledger Fabric is its modular architecture. This modularity is expressed in the ability to customize consensus protocols, smart contract languages (such as Go, Java, or Node.js), and the databases used to store the ledger state. For example, Hyperledger Fabric supports LevelDB and CouchDB as database options for storing state, providing flexibility based on project requirements.

Geoadresses. Mobility data are highly sensitive due to their high precision, which can lead to re-identification attacks. To mitigate this, LoChain relies on *geoadresses*, a standardized set of geographic references corresponding to fixed points (typically road intersections or known geographic landmarks). Geoadresses abstract exact positions into uniform points, ensuring

consistent and anonymized representation of mobility data, which facilitates analysis while preserving individual privacy.

Disposable identities. Disposable identities correspond to regularly renewed digital identifiers, which are inspired by privacy practices used in blockchain-based cryptocurrencies (e.g., Bitcoin’s (Nakamoto, 2008) usage of disposable addresses). Temporary identifiers are generated periodically and associated with users’ geolocation transactions. After a defined period or a set number of transactions, these disposable identities are discarded and replaced, effectively preventing long-term user profiling.

The primary goal of implementing disposable identities is to improve user privacy. Indeed, in traditional location-based services, a user’s movements are often linked to a unique and persistent identifier (such as an account or device ID). This link allows users to be tracked and profiled over time, which can potentially lead to privacy issues. Disposable identities seek to prevent this by ensuring that any attempt to build a user’s movement profile is limited to the lifetime of a unique and temporary identity, significantly reducing the risk of the collected data being used for invasive tracking or profiling. The different stages of the disposable identity process are as follows:

1. *Identity generation.* Each time the service is used, users receive a unique temporary identity, similar to a public address in blockchain terms. This identity does not reveal the user’s permanent personal identifiers.
2. *Transaction representation.* The user’s movements are represented by transactions linked to this disposable identity, using a list of geoadresses to map the geographical space traveled. A transaction is generated each time a new nearest address is identified during the trip.
3. *Identity disposal and renewal.* After a certain number of transactions or a specific period, the current temporary identity is abandoned. A new identity is generated for future activities, thus preventing the correlation of data collected over different periods with a single user.

Location proofs. Rather than sharing raw coordinates, each transaction references a pseudonymous geoadress and includes a digital signature proving the device’s knowledge of the corresponding private key. This allows third parties to verify location claims without direct access to mobility data. LoChain integrates a proof-of-location feature, developed to authenticate each mobility data point without compromising user anonymity. The process is based on the following steps:

1. When a user enters an area covered by our system, a disposable identity is generated for the user. This identity is used to record movements as anonymous transactions.
2. Each transaction is associated with key information, including the disposable identity used to sign the transaction, the timestamp indicating the exact time of the transaction, the geographic coordinates in the form of a geoadress, ensuring uniform granularity and sufficient abstraction to preserve privacy as well as a unique transaction identifier generated to prevent duplication or tampering.
3. A cryptographic signature is applied to each transaction using the private key associated with the disposable identity.

The system uses asymmetric cryptography to secure transactions. Specifically, a signature based on the Elliptic Curve Digital Signature Algorithm (ECDSA) is used. This choice is motivated by its efficiency, its compatibility with decentralized systems such as blockchains, and its robustness against known cryptographic attacks (Johnson et al., 2001). Each transaction is signed with a private key, while the corresponding public key, linked to the disposable identity, is used to validate the signature.

Anonymous communication network. To achieve anonymity of the communications at the network level, LoChain integrates the *Tor network* (Dingledine et al., 2004) through the Orbot Android application (The Tor Project, 2025). Transactions from the client application are routed through Tor, thus dissociating disposable identities from the user’s actual IP address, which contributes to reinforce the user privacy.

3.2 Decentralized Data Management

Unlike centralized systems (*e.g.*, Google Maps), LoChain leverages the blockchain technology—specifically Hyperledger Fabric—to decentralize the mobility data storage and processing. This decentralized model significantly reduces vulnerabilities related to single points of failure and centralized data breaches. To achieve optimal scalability and local performance, LoChain employs a carefully designed channel architecture integrated with geo-pools (detailed below).

Geo-pools. The *geo-pool* concept is central in LoChain’s architecture, enabling robust scalability, localized data management and the protection of privacy simultaneously. Each geo-pool corresponds to a specific geographic zone or district and is managed by dedicated blockchain nodes and ordering nodes belonging exclusively to that local organization (ideally multiple nodes per district for fault tolerance and data safety). More precisely, each geo-pool leverages Hyperledger Fabric’s unique channel architecture, which was chosen for its modularity, performance under permissioned setups and flexible access policies. This structure effectively simulates *sharding* by utilizing Fabric’s channel system capability—each channel maintaining its independent ledger—and exploiting inherent geographic constraints to restrict node communication and data propagation. Unlike traditional blockchain setups in which global connectivity among nodes is standard, LoChain explicitly restricts node communication geographically through dedicated channel assignments and access control policies, reducing congestion and improving fault isolation.

Data ingestion channels are exclusive to each geo-pool (district), involving only nodes managed by the respective local organization. They collect raw user transactions representing movement within their district. This enforces localized control, reduces data leakage risk and simplifies jurisdictional governance. The *neighboring district channels* also called (peripheral channels) are created among geographically adjacent districts to enable secure data coordination between neighbors. This setup facilitates data exchanges for enhanced user mobility insights across adjacent districts, ensuring smooth transitions and regional data consistency. Finally, all geo-pool nodes participate to the *global channel*, sharing exclusively aggregated and anonymized data without references to any temporary identity. By aggregating and anonymizing locally before transmission, the system reduces global network load, improving scalability and ensuring user privacy at a global scale, while enabling macro-scale analytics without exposing raw movement traces.

The architecture scales organically as densely populated areas naturally accrue more nodes and geo-pools, increasing processing granularity and improving local responsiveness without central coordination. This aligns well with the decentralized DAO governance model discussed later.




Figure 1. Spatial clustering of intersections into geo-pools over the Montreal region. Each color represents a simulated district for the LoChain prototype.

Each district maintains its own dedicated set of *orderers* responsible for handling multiple types of transactions. In particular in local data ingestion channels, orderers exclusively process transactions related to raw mobility data within their respective district, which ensures local control and prevents raw movement data from leaking to other districts. For peripheral (*i.e.*, neighboring district) channels, orderers from each pair of geographically adjacent districts jointly handle transactions in their shared peripheral channel. This joint responsibility ensures efficient and manageable inter-district communication. Finally, all districts have orderers collectively participating in the global channel, processing aggregated and anonymized data summaries.

This strategic use of orderers structurally restricts interactions to geographically adjacent districts at most, making distant inter-district communication impossible by design and simplifying the network architecture compared to traditional blockchain systems. This geographic confinement inherently creates a sharding-like structure, segmenting the network into manageable subnetworks. Unlike traditional sharding methods that require complex cross-shard protocols and validations, LoChain circumvents such complexity due to its inherently localized mobility data. Thus, by limiting cross-region interactions to neighboring districts, LoChain achieves significant scalability benefits, avoiding typical blockchain bottlenecks associated with distant cross-shard transactions. The global channel receives only pre-aggregated anonymized statistics, further minimizing congestion and ensuring efficient network operation.

LoChain replaces precise GPS coordinates with geoaddresses, which are standardized references to road intersections. This abstraction ensures that all users are mapped to a shared set of public locations, avoiding the disclosure of individual-specific information while facilitating meaningful data aggregation. Despite the privacy protections provided by disposable identities

and geoaddresses, some privacy risks remain due to the possibility of identifying repetitive patterns in user movements. For example, consistent starting and ending points in trajectories can be used to infer sensitive details such as users' home or workplace locations.

To mitigate this vulnerability, LoChain additionally injects noise to obfuscate such recurring patterns. In particular, the LoChain app infers locally users' frequent destinations (e.g., home or workplace) and generates multiple alternative geoaddresses within a predefined radius around these sensitive points. When users complete their journeys at these critical locations, the system randomly selects alternative neighboring geoaddresses, which introduces ambiguity and prevents the precise identification of the user's points of interests.

Another complementary technique involves intentionally generating random movement patterns for certain temporary identities. This technique is especially relevant for users with predictable or stationary behavior (e.g., those working remotely from home). By periodically generating pseudo-random journeys unrelated to typical user behavior, LoChain further disrupts pattern recognition attempts. To prevent the distortion of the outputs generated, purged identities and substituted geoaddresses are flagged internally. This allows aggregate analytics (e.g., heatmaps or mobility metrics) to retain validity while still preserving user privacy.

These methods were partially validated through simulations using synthetic populations and OpenStreetMap-based geography. Full implementation of these techniques, particularly the generation of realistic randomized movement patterns, will leverage external services such as OpenRouteService to enhance realism and maintain usability alongside privacy guarantees.

4. Prototype Implementation

To validate the feasibility of LoChain, a functional prototype was developed with three interconnected components: (1) an Android client application for secure data capture, (2) a permissioned blockchain network using Hyperledger Fabric for storage and validation and (3) a Node.js/Angular-based visualization platform for analytics and heatmap rendering.

More precisely, the Android client application captures, anonymizes and securely transmits geolocation data. More precisely, by employing disposable identities, location proofs authenticated through ECDSA (Johnson et al., 2001) and routing all communications through the Tor network via Orbot, the application anonymizes user-generated mobility transactions locally before their submission. This means that raw GPS coordinates or persistent identifiers are never transmitted, thus ensuring privacy from the point of capture. The decentralized backend infrastructure leverages Hyperledger Fabric's modular architecture, organized into specialized channels aligned with geo-pools, which effectively simulates sharding and enables efficient localized data management.

Finally, the visualization component, built with Node.js (backend) and Angular (frontend), transforms aggregated and anonymized blockchain data into intuitive visualizations (e.g., temporary identity activity, heatmaps or statistical summaries). The modular architecture allows the deployment of a Node.js backend instance with each district's blockchain node, offering significant optimization opportunities detailed below. Deploying Node.js

backends at each organizational node provides powerful optimization capabilities while enabling the following technical optimization.

- *Local pre-processing and aggregation.* When posting aggregated data onto the global channel, each backend instance aggregates locally the organization's raw mobility data, summarizing movement statistics before blockchain submission. This reduces unnecessary blockchain transaction loads, optimizes orderer performance and enables local error correction.
- *Decentralized caching for visualization.* The Node.js backends maintain local caches of frequently requested visualizations. This reduces query latency, providing near-instant visualizations while minimizing blockchain queries and network load.
- *Asynchronous batch submissions.* Transactions are buffered and submitted asynchronously in batches (e.g., every 10 minutes or after a fixed number of transactions). This batching technique reduces transaction overhead while enhancing blockchain scalability.
- *Distributed validation and data quality control.* Thanks to the standard data structure leveraging the geoaddresses, local backend instances perform validation checks (e.g., anomaly detection or data consistency) prior to blockchain submission. This reduces errors in the blockchain ledger, while preserving data quality and minimizing storage inefficiencies.
- *Peer-to-peer (P2P) communication.* Node.js backends communicate directly with geographically adjacent nodes, synchronizing recent updates locally to reduce redundant blockchain interactions and global orderer congestion.

Together, these optimizations highlight the practical benefits and strategic importance of LoChain's Node.js/Angular visualization architecture.

5. Experimental Evaluation

Evaluation setting. To realistically evaluate LoChain, we rely on data based on real-world movement patterns. Initially, the mobility data of a small set of real users was collected at regular intervals to establish a realistic reference dataset. This dataset served as a baseline to simulate 10,000 virtual users, each performing 1,000 transactions, distributed randomly across multiple geographic districts. More precisely, transactions were generated by assigning disposable identities to simulated users and mapping their movements to intersections extracted from OpenStreetMap road networks. Each transaction included disposable identity references, geoaddress coordinates and precise timestamps, thus creating realistic and scalable transaction volumes suitable for a preliminary evaluation.

To assess the robustness of privacy mechanisms, we employed a dual indexing strategy. First, thanks to the encoding of movements through transactions and the standardization of the number of authorized transactions for each identity, transactions were evenly distributed among disposable identities, creating

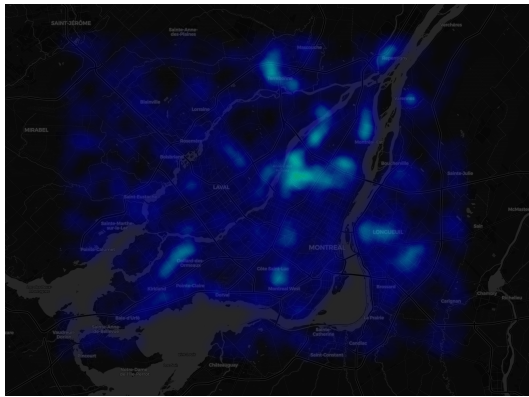


Figure 2. Aggregated heatmap of simulated user movements generated by the LoChain visualization platform.

uniform data volumes across identities. This uniformity effectively mitigates statistical re-identification attacks by removing identifiable usage patterns associated with individual identities. Conversely, indexing transactions by geoaddresses produced naturally varied data distributions, reflecting realistic patterns based on intersection popularity. This method provided the foundation for generating meaningful heatmaps and mobility statistics, confirming LoChain’s analytical utility.

The simulation utilized road intersections as starting and ending points, implicitly modeling vehicle-based or road-centric movements. While effective for initial evaluation, this assumption limits the representation of other mobility forms (*e.g.*, walking, cycling or public transportation). Future evaluation will incorporate pedestrian, cycling and public transit models to reflect broader movement patterns..

Summary of evaluation results. The evaluation of the LoChain prototype in the simulated environment yielded the following insights. First, disposable identities, geoaddresses abstraction, Tor integration, noise injection and uniform indexing successfully reduced simulated re-identification risks, validating LoChain’s core privacy mechanisms. Second, geo-pool and channel architectures effectively simulated sharding, confirming LoChain’s potential to scale beyond traditional Hyperledger Fabric node-count limits while maintaining local data management efficiency. Finally, the visualization interface effectively translates anonymized blockchain data into actionable insights (*e.g.*, heatmaps or statistical summaries), demonstrating the compatibility of privacy measures with practical analytical utility.

Overall, this assessment confirms LoChain’s technical feasibility and demonstrates its potential as a robust privacy-preserving geolocation management system ready for further validation in real-world scenarios.

6. Incentive Model

A fundamental challenge for decentralized and privacy-preserving technologies, such as LoChain, lies in establishing a sustainable economic model that encourages widespread adoption and ongoing support. For instance, traditional privacy-focused systems like the Tor network rely heavily on volunteers, often limiting their growth, performance and sustainability. To overcome these limitations, LoChain proposes an incentive model specifically aligned with its architectural strengths and current privacy standards.

Indeed currently, location-based services typically follow a centralized model, in which a user mobility data is continuously collected and stored directly by companies. This approach entails significant privacy risks and substantial responsibilities around data management and compliance with regulations such as GDPR (Union, 2016) or Law 25 (du Québec, 2021). To achieve this compliance, businesses must invest heavily in secure data storage, protection measures, and compliance frameworks, incurring considerable operational overhead and liability risks.

Instead of this traditional centralized storage and direct management of user mobility data, LoChain introduces an innovative model leveraging temporary digital identities and blockchain-based location verification. This model allows businesses to securely verify real-time user location without needing to directly store, manage or control sensitive location data themselves.

Illustrative example. To illustrate how this works, consider a scenario involving a ride-hailing application like Uber operating on the LoChain network:

1. *User generates pointer.* A user who wishes to book a ride securely generates an encrypted “pointer” referencing their latest mobility transaction stored on the blockchain. This pointer is cryptographically signed using the user’s temporary identity private key, ensuring authenticity.
2. *Business requests verification.* The company receives this encrypted pointer and submits it to the blockchain network through a verification request signed with the company’s private key.
3. *Blockchain verification.* The blockchain network authenticates the company’s request using public keys and verifies the referenced geolocation transaction. The network then securely confirms the user’s current position directly to the company in real-time, without revealing raw data beyond the requested verification.

Benefits of incentive model. This approach introduces several compelling advantages for all stakeholders. First, user mobility data remains encrypted and private with public visibility delayed by at least 24 hours. This prevents real-time surveillance and mitigates risks associated with data misuse or user profiling. Second, companies no longer bear the responsibility or expense associated with extensive mobility data storage and management. This greatly simplifies regulatory compliance with privacy laws such as GDPR and reduces operational costs.

Third, after a 24-hour privacy buffer, aggregated and anonymized mobility data becomes publicly available. This data serves broader public interests, supporting urban planning, transportation optimization, infrastructure development and academic research. Finally, each location verification request from a business incurs a small transaction fee payable to the blockchain network. These fees create a sustainable economic incentive for stakeholders operating nodes and orderers, ensuring long-term financial viability without relying on purely volunteer-driven support or centralized entities.

The successful adoption of this incentive model depends critically on balancing the privacy interests of users, operational needs of businesses and the public benefits of aggregated location data. In particular, measures should be established to prevent potential misuse, such as businesses excessively querying

user locations or using verification requests to indirectly track users. Compliance with regulatory frameworks, including clear and explicit user consent management, remains paramount to maintain trust and ensure ongoing operational legitimacy.

To summarize by combining strong privacy preservation, reduced business liability and a sustainable economic incentive structure, we believe that LoChain provides a practical and attractive approach for deploying decentralized geolocation services at scale.

7. Decentralized Governance

The current market for mobility data is characterized by opacity and a significant lack of public awareness. Users are typically unaware of how their personal location data is collected, sold and utilized. Existing privacy regulations (such as GDPR) are frequently bypassed or inadequately enforced, leading to systemic privacy abuses and widespread mistrust. Moreover, the mobility data market is dominated by centralized entities whose economic incentives often conflict directly with user privacy. The inherent centralized control over sensitive data creates significant risks related to misuse, security breaches and irresponsible data management.

LoChain's architecture inherently addresses many privacy concerns at the source by anonymizing and decentralizing mobility data management. However, beyond mere technical decentralization, LoChain's unique structure also offers an opportunity to establish a new and transparent industry standard that fundamentally reshapes economic incentives and governance. Instead of relying on a single entity or group controlling the network and the associated data standards, LoChain can be operated as a consortium-based network structured as a *Decentralized Autonomous Organization* (DAO). Such a structure allows for cooperative, democratic governance involving multiple stakeholders, including businesses, local authorities, NGOs, research institutions and even users—operated nodes.

Proposed DAO structure. Under the DAO framework, the LoChain network would function as follows:

- *Distributed control and ownership.* Multiple stakeholders become DAO members, each controlling and operating nodes within the LoChain network. Geo-pools serve as jurisdictional units for node assignment and revenue distribution, aligning technical structure with governance logic. Decision-making power, economic incentives and network governance are distributed, avoiding any single point of control.
- *Transparent economic incentives.* Businesses accessing location verification services pay transaction fees (as described earlier). These fees fund the operation, maintenance and expansion of the DAO-controlled network. DAO members democratically decide on fee structures, reinvestment priorities and network improvements.
- *Fair and democratic decision-making.* Key governance decisions, including data-handling standards, privacy protocols and incentive structures, are made collectively by DAO members through a transparent voting mechanism. Stakeholders participate according to clearly defined roles and responsibilities, ensuring accountability and fairness.
- *User participation and protection.* Users can become direct participants or stakeholders in the DAO, ensuring their interests (especially regarding privacy and transparency) are directly represented in governance decisions. User rights and protections are codified within the DAO governance documents, reinforcing trust and transparency.

Implementing LoChain as a DAO-managed network would deliver several substantial benefits. In particular, by collectively defining transparent privacy-preserving standards for mobility data collection and verification, the DAO provides a universal reference point for businesses, offering them a regulatory-friendly alternative to data hoarding by removing direct control over personal data and replacing it with verifiable proofs on a collectively managed network. In addition, democratic governance ensures that decisions are transparent and accountable. Users and businesses can see clearly how and why data is accessed, shared or monetized, increasing public trust and adoption.

A clear incentive model—in which members receive proportional rewards for operating nodes, contributing resources or enhancing network security and utility—ensures long-term network sustainability. Incentives directly align with stakeholder interests toward maintaining a robust and trusted network. By anonymizing data at the source and establishing a collectively managed standard for data handling, the DAO model drastically reduces risks associated with violating data privacy laws. Thus, businesses benefit from reduced operational overhead related to regulatory compliance and data security.

Potential challenges. While promising, the DAO approach must carefully address certain challenges to remain fair and viable. First, determining fair initial distribution of voting power or tokens among stakeholders is essential. Economic incentives must be structured to avoid centralization of control or disproportionate influence by powerful entities. Second, transparent rules for decision-making, voting, conflict resolution and stakeholder participation must be precisely defined to ensure efficient and effective governance. Finally, DAOs and tokenized governance may face regulatory complexities. Careful legal structuring and clear compliance guidelines will be essential for widespread adoption.

Integrating DAO-based governance into LoChain's incentive model presents a groundbreaking opportunity to fundamentally reshape how mobility data is managed, monetized and regulated. By democratizing control, establishing transparent economic incentives and creating cooperative industry standards, LoChain can drive a more sustainable and equitable data economy.

8. Conclusion and Future Work

This paper introduced LoChain, a decentralized blockchain-based solution specifically designed to preserve user privacy in mobility data management. By leveraging Hyperledger Fabric's multi-channel architecture to achieve geographic data confinement, LoChain provides robust scalability and privacy protections simultaneously. Through innovative use of disposable identities, standardized geoaddresses, strategic noise injection, Tor network integration and geographically-constrained channel structures, LoChain addresses critical privacy vulnerabilities inherent in current centralized mobility data management.

A prototype developed as a proof-of-concept demonstrated the technical feasibility of LoChain's privacy-preserving mechanisms within a controlled and simulated environment. More precisely, the experimental evaluation has validated that LoChain maintains individual anonymity while delivering analytically valuable aggregated insights through intuitive visualizations.

LoChain's architecture uniquely emulates blockchain sharding by geographically restricting node communication inherently segmenting data processing into manageable subnetworks without introducing complex cross-shard protocols. By simplifying the challenges typically associated with blockchain scalability, we believe that LoChain represents a compelling and practical solution for privacy-conscious use of mobility data across diverse use cases such as urban analytics, transportation optimization, and infrastructure management. To summarize, LoChain sets the foundation for a scalable, privacy-preserving, and efficient geolocation data management framework that aligns with modern privacy expectations and decentralized system architectures.

Future work. Future avenues of research will include:

- *Real-world deployment and validation.* In the future, we aim at conducting extensive testing and validation in realistic, geographically distributed settings to confirm system robustness and privacy protections under genuine operational conditions. Practical trials will help identify unforeseen challenges, enabling iterative refinement based on empirical observations.
- *Enhanced simulation models.* We will integrate more diverse and realistic mobility models, covering various types of movement beyond vehicle-based road networks, such as pedestrian pathways, cycling routes and public transportation systems. Expanding the simulation scenarios will significantly enhance the accuracy and realism of LoChain's evaluation data.
- *Optimized noise injective.* We will further refine the noise injection methodologies, which will benefit from the enhanced simulation models, particularly the pseudo-random movement pattern generation, fully leveraging external services (such as OpenRouteService) to ensure optimal privacy protection without compromising data utility. Future enhancements will also consider Sybil resistance strategies as well to prevent the abuse of disposable identities by malicious entities.
- *Development of economic and governance models.* Finally, we will investigate DAO frameworks, leveraging geopools as natural organizational units to foster community-driven governance, decision-making and economic incentives. The integration of token-based or incentive-compatible economic models could increase user participation, strengthen platform adoption and ensure sustainable operations.

Acknowledgments

Sébastien Gambs acknowledges the support of NSERC as well as the Canada Research Chair program.

References

Buterin, V., 2014. Ethereum: A next-generation smart contract and decentralized application platform. *Ethereum White Paper*, 3(37), 2–1. <https://ethereum.org/en/whitepaper/>.

Castro, M., Liskov, B., 1999. Practical byzantine fault tolerance. *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI)*, USENIX Association, 173–186.

De Montjoye, Y.-A., Hidalgo, C. A., Verleysen, M., Blondel, V. D., 2013. Unique in the Crowd: The Privacy Bounds of Human Mobility. *Scientific Reports*, 3, 1376.

Dingledine, R., Mathewson, N., Syverson, P., 2004. Tor: The second-generation onion router. *Proceedings of the 13th USENIX Security Symposium*, USENIX Association, San Diego, CA, USA, 303–320.

du Québec, G., 2021. Loi modernisant des dispositions législatives en matière de protection des renseignements personnels (loi 25). Adoptée le 22 septembre 2021, entrée en vigueur progressive jusqu'en 2024.

FOAM, 2018. Foam whitepaper.

Goldreich, O., Micali, S., Wigderson, A., 1991. Proofs that Yield Nothing But Their Validity ... Zero-Knowledge Proof Systems. *Journal of the ACM*, 38(3), 691–729. <https://dl.acm.org/doi/10.1145/116825.116852>.

Haklay, M., Weber, P., 2008. OpenStreetMap: User-Generated Street Maps. *IEEE Pervasive Computing*, 7(4), 12–18.

Hyperledger Fabric — hyperledger.org, n.d. <https://www.hyperledger.org/projects/fabric>. [Accessed 17-05-2024].

Johnson, D., Menezes, A., Vanstone, S., 2001. *The Elliptic Curve Digital Signature Algorithm (ECDSA)*. Springer, Berlin, Heidelberg.

Lohr, S., 2018. Google+ data breach exposed data of up to 500,000 users. <https://www.nytimes.com/2018/10/08/technology/google-plus-security-disclosure.html>. [Accessed 17-05-2024].

McKinnell, J., 2021. Google faces massive fines after world-first data breach ruling.

Nakamoto, S., 2008. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*. <https://bitcoin.org/bitcoin.pdf>.

O'Flaherty, K., 2020. Apple iphone location controversy: What you need to know. <https://www.forbes.com/sites/kateoflahertyuk/2020/01/24/apple-iphone-location-controversy-what-you-need-to-know/>. [Accessed 17-05-2024].

The Tor Project, 2025. Orbot: Tor for android. Online. Accessed: 2025-01-02.

Union, E., 2016. Regulation (EU) 2016/679 ... (General Data Protection Regulation). *Official Journal of the European Union*. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

Xiong, P., Zhu, T., Pan, L., Niu, W., Li, G., 2014. Privacy preserving in location data release: A differential privacy approach. D.-N. Pham, S.-B. Park (eds), *PRICAI 2014: Trends in Artificial Intelligence*, Springer International Publishing, Cham, 183–195.

Zou, S., Xi, J., Wang, H., Xu, G., 2019. CrowdBLPS: A blockchain-based location-privacy-preserving mobile crowdsensing system. *IEEE Transactions on Industrial Informatics*, 16(6), 4206–4218.