



Journal of Global IT Dynamics

Volume 1, Issue 1 (2026)



Designing a Scalable Blockchain-Based Framework for Secure Smart City Infrastructure Management

Muhammad Sameer *

Department of Science, Sultan Qaboos University, Oman

Article Info

Abstract

Smart cities require the efficient and secure integration of key infrastructure domains such as water, energy, transportation, smart lighting and waste management deploying a myriad of data-generating IoT sensors and devices. Current management systems feature single points of failure, lack of auditability, insufficient privacy protection and lack of scalability as IoT nodes increase in density. In this paper, we present SmartChain, a three-tier multilayered blockchain-based architecture that incorporates a permissioned distributed ledger, an AI-based anomaly detection module (ADM) and a dual-layered privacy preservation approach that combines Zero-Knowledge Proofs (ZKP) and Ciphertext-Policy Attribute-Based Encryption (CP-ABE). SmartChain is tested on a large dataset - 2000 timestamped transactions involving Smart Cities' five infrastructure types across several zones of a city. The results show a mean throughput of 3,421 transactions per second (TPS), a mean transaction latency of 3,847 milliseconds and a mean privacy score of 82.4 out of 100. The machine learning based anomaly module produces an F1-Score of over 97% and an AUROC score of 0.991 with Random Forest as the classifier. Benchmarking against Hyperledger Fabric 2.5, Ethereum 2.0 and the IOTA Tangle demonstrate the scalability, security, privacy and efficiency of SmartChain. This research renders SmartChain a practical production-level platform for management of new smart city infrastructure.

Keywords

Blockchain; Smart City; IoT Security; Anomaly Detection; Zero-Knowledge Proof; Attribute-Based Encryption; Scalability; Privacy-Preserving Computing; Hyperledger; Machine Learning

1. Introduction

The notion of smart cities has progressed from an ambit claim to tangible engineering problem

of global scale. Cities now instrument their infrastructure with billions of IoT sensors and

actuators generating telemetry data, constantly flowing over a variety of networks. In turn, this creates data streams that are used to optimise traffic signal cycles in real time, predict faults and leaks in water supply systems, balance loads in power grids, automatically route garbage trucks, and dynamically manage street lighting systems [1]. By 2030, the global number of IoT devices to be deployed in infrastructure is expected to surpass 30 billion, generating exabytes of data per year [2].

However, security of current smart city platforms is woefully deficient. Deployments generally take a centralised cloud-based approach in which a single management server or a cluster of servers receives data in a variety of infrastructure domains. This structure creates a valuable target for security attacks: by breaching the central database, an adversary can access the operational parameters of the water, power and transportation systems of the entire city. In addition, once data is committed to a centralised data silo, auditability is poor, tamper evidence is weak and forensic reconstruction is expensive [3]. Privacy is also a consideration: data originating from utilities can, when combined across domains and over time, expose critical details about daily routines and behaviours of individual citizens [4].

Distributed ledgers, and especially blockchain, present a new architecture. Blockchain systems eliminate the trusted intermediary by distributing validation among peers and using a cryptographically secure chain of blocks of transactions that cannot be altered without the network's consent [5]. Permissioned blockchain networks like Hyperledger Fabric also enable the enforcement of fine-grained access control policies at the protocol level, making them apt for public utilities to operate [6].

But there are challenges to using blockchain in smart cities. Large-scale telemetry of infrastructure may not be processed through traditional blockchain consensus algorithms which are designed for financial transactions rather than high rates of sensor data [7]. Elaborations for privacy (such as zero-knowledge proofs) are expensive, and must be carefully managed to prevent deferring. Anomaly detection systems to detect cyber-physical attacks that pose threats to infrastructure must run in near real-time and are robust to the skewed class distribution in datasets where attack events are rare [8].

The SmartChain Framework is a new scalable blockchain framework that uses the power of distributed ledger technology (DLT) to solve the problems associated with managing smart cities. More specifically, This paper contributes to smart cities in three ways:

- (1) A 3-Tier Blockchain Model - that allows parts of the framework to grow separately from each other (i.e., some Systems of Record will need lots more storage than some Systems of Verification) which creates a solution that is scalable horizontally.
- (2) An Anomaly Detection Module that uses AI methods (Isolation Forest and Random Forest Classification) for detecting potential abnormal activity; trained on a data set of 2000 records covering 5 different types of infrastructure, resulting in high performance with an F1-Score of 97.4% and an AUROC of .991.
- (3) A Security Framework Using Dual Privacy Preservation Models that use ZKP to verify all transactions without needing to disclose any data (ZKP), and CP-ABE to provide the ability to restrict access at a more granular level (e.g., streets versus buildings)
- (4) Comprehensive empirical evaluation including throughput and latency profiling, spatial zone-level analysis, temporal anomaly pattern characterisation, principal component analysis of the feature space, and competitive benchmarking against three state-of-the-art blockchain platforms.

The structure of the remainder of this paper is as follows. In Section 2 a review of the literature related to blockchain applications in smart cities, and Internet of Things (IoT)

security and privacy-preserving distributed systems is presented. In Section 3 the SmartChain architecture and constituent modules are described. In Section 4 the dataset, experimental configuration, and evaluation methodology utilised in this study are outlined. In Section 5 the experimental results are presented and analysed. In Section 6 a comparative benchmarking discussion is provided. In Section 7 limitations and future work are discussed. In Section 8 the paper will be concluded.

2. Related Work

2.1 Blockchain for IoT and Smart City Applications

For almost 10 years, researchers have conducted research regarding IoT security using blockchain technology. For example, Dorri et al. [9] created a light-weight blockchain architecture that was specific to IoT with the goal of proving that permissioned ledgers could provide the throughput required for real-time device telemetry along with providing tamper-evident audit logs. However, this architecture did not include the ability to use privacy preserving query mechanisms, nor was there an inclusion of anomaly detection capabilities. In addition, Reyna et al. [10] have provided an overall view of the integration challenges that blockchain has experienced when integrated into IoT; as the challenges of energy consumption, scalability, and standardisation were identified as barriers in the integration process. Furthermore, our classification of blockchain-IoT integration patterns provided by Reyna et al. will serve as a basis/reference for our own SmartChain Framework, which addresses many of the challenges presented above.

Biswas and Muthukkumarasamy [11] explored the potential for blockchain technology as a means of securing heterogeneous communications networks in the context of smart city applications. They studied the access control problem and proposed a policy-based key management scheme; however, their study lacked any empirical evaluation on a large scale and did not account for the multi-domain infrastructure present in modern smart cities. Sharma et al. [12] continued this line of work by proposing an integrated approach to cybersecurity for smart cities, using a combination of software-defined networking (SDN) and blockchain, and reporting positive latency results based on simulation evaluations. However, their evaluation was limited to simulation techniques without taking into consideration AI-driven detection of security threats.

2.2 Privacy-Preserving Mechanisms in Blockchain Systems

Researchers have devoted substantial effort to studying how distributed ledgers can maintain privacy and protect data. Kosba et al. introduced the Hawk system for smart contract blockchain technology that allows for private transaction processing via zero-knowledge proofs, demonstrating the feasibility of executing private smart contracts without the use of a trusted intermediary. Ben-Sasson et al. showed that zk-SNARKS (zero-knowledge succinct non-interactive arguments of knowledge) are practically usable for verifying computations without disclosing input values. This is precisely what happened with the ZKP layer of the SmartChain. Bethencourt et al. described Ciphertext-Policy Attribute-Based Encryption (CP-ABE) and proved its utility in managing access control to data in distributed systems where data owners want to establish policies that limit the ability to decrypt encrypted data based on a set of attributes instead of on a person's identity. The proposed system uses CP-ABE to create role aware access control for smart cities.

Recently, Zhang and others [16] proposed a method to share private data in a secure, anonymous way by using both blockchain and homomorphic encryption. They say that they were able to achieve strong privacy protections through the use of differential privacy methods; however, this incurred

a very high cost in terms of computing power. Liu and others [17] looked into how much transparency should be given to citizens when using blockchain in smart cities and recommend implementing a system of selective disclosure based on commitment schemes. Our work (SmartChain) seeks to provide an efficient operation by balancing the strength of the privacy mechanism against the amount of runtime under both frameworks (ZKP and CP-ABE).

2.3 Machine Learning for Anomaly Detection in Cyber-Physical Systems

Various machine learning paradigms have been used for anomaly detection in industrial and cyber-physical systems. Isolation Forest algorithm by Liu et al. [18] uses feature space recursive partitioning techniques, and has been shown to have improved performance over traditional anomaly detection methods for high-dimensional datasets with low levels of anomaly contamination, specifically as found within transaction logs of an infrastructure. Random Forests by Breiman [19] are also known for their outstanding generalisation performance for tabular classification tasks; in the ICS security area, Umer et al. [20] used Random Forest classifiers to obtain recall rates of >96% for SCADA logs. Additionally, Junejo & Goh [21] successfully applied ML techniques to detect cyber-attacks in water treatment facilities.

Long Short-Term Memory (LSTM) networks (a type of Recurrent Neural Network) have been used for temporal anomaly detection in streams of IoT data. Malhotra et al. [22] indicated that LSTM Autoencoders could detect variations from normal operation sequences as an indication that the equipment is likely to fail or that there was an injection of synthetic adversarial input, using only a subset of the sequence has been shown to be effective. The smart chain's ADM architecture utilizes some of the temporal reasoning principles of LSTM-based systems; however, the empirical data and results presented in this paper utilized Isolation Forest and Random Forest Classifiers. The advantages of these classifiers include generating results that are more interpretable than LSTM-based systems, as well as significantly shorter inference latencies. These properties are essential for monitoring real-time infrastructures.

2.4 Comparative Blockchain Platforms for Infrastructure Applications

Many blockchain technology implementations have been designed or updated to support smart city and infrastructure type applications. Hyperledger Fabric [6] is from the Linux Foundation and has a modular, permissioned architecture, complete with a variety of pluggable methods for consensus and endorsement policies, making it an extremely common enterprise-grade blockchain technology. Ethereum 2.0 [23] has significantly improved energy efficiency compared with its predecessor by transitioning to a Proof of Stake consensus method; Ethereum 2.0 maintains full Turing-complete (fully featured) capabilities to create smart contracts. The IOTA Tangle [24] is based upon a Directed Acyclic Graph (DAG) model for distributed ledgers; therefore, there are no mining fees and high throughput of transactions enables massive numbers of micro-transactions to take place, making the IOTA Tangle one of the best candidates for deployment at the IoT scale. In Section 6, we will perform a comparative analysis of the SmartChain against each of the three presented blockchain platform technologies using the metrics of; throughput, latency, energy consumption, and privacy.

3. SmartChain: Proposed Framework Architecture

3.1 Architectural Overview

SmartChain has been architected with a three-tier distributed system that will support the separation of concerns for the data ingestion, ledger management, and consensus layer, and the service and analytics layer. This model allows each layer to be scaled independently horizontally, as well as to provide modularity for replacement or upgrades to individual components without causing disruption to the end-to-end pipeline. SmartChain has defined these three tiers as follows: (i) the IoT Data Ingestion Layer, (ii) the Blockchain Core Layer, and (iii) the Service and Analytics Layer.

The IoT Data Ingestion Layer consists of lightweight edge agents (or deployed edge devices) at each infrastructure node (such as water sensor nodes, smart meter nodes, traffic controller nodes, luminosity sensor nodes, and waste bin fill-level sensor nodes) that perform the functions of data sampling, performing local preprocessing (outlier filtering and normalising units), and submitting authenticated transactions to the blockchain network. Edge agents generate a cryptographic hash for their payload and sign the transaction using its device certificate, thereby providing assurance of end-to-end integrity of the data from the sensor to the ledger.

The core blockchain layer implements a permissioned hyperledger type of distributed ledger and uses a Byzantine Fault Tolerant (BFT) consensus protocol; multiple channels can be created based upon the type of infrastructure that allows for multiple transactions to be executed in parallel, eliminating the need for cross-domain transactions, and allowing the nodes that are part of the consensus agreement to be ignorant of the data being processed if they are missing appropriate decryption attribute based credentials.

The services and analytics layer exposes a RESTful API to city operations centre and authorised third party applications. The AI based anomaly detection module runs at this level consuming transaction streams to apply trained anomaly detection model outputs and generate security alerts, there is also a requirement for zero-knowledge proof verification to audit aggregate statistics (e.g. total water consumption; total grid load) without accessing the underlying transaction records.

3.2 AI-Driven Anomaly Detection Module (ADM)

SmartChain's ADM is an important aspect of SmartChain's security. It will detect cyber-physical attacks (spoofing, replay attack, false data injection) and physical anomalies (infrastructure failure, operational deviation). The ADM works in two phases. The first phase is an unsupervised phase that uses Isolation Forest to detect anomalies in real-time streaming data and uses a pre-trained Random Forest classifier in the second phase to classify and categorize events into potential threats.

The Isolation Forest model is appropriate for use in the first phase of the ADM, as it does not require any labelled training data; has a small memory footprint; produces anomaly scores in $O(\log n)$ time for each transaction, and is feasible for real-time use. The contamination parameter for the Isolation Forest model is set to the empirically measured anomaly rate from the historical data. The Random Forest classifier in the second phase of the ADM is trained with 300 estimators, and has a maximum depth of 12, on labelled historical training data. The 10 features produced through feature engineering that are input into the Random Forest classifier include: TPS

throughput; Millisecond latency; Number of nodes; Privacy score; Gas fee in Wei; Size of data in kilobytes; Infrastructure-specific measurable metric value; Infrastructure type encoded; Zone encoded; Hour of the day.

Class imbalance is an important factor to consider when designing a system. In operational infrastructure datasets, there will be fewer transactions that actually represent false positives than transactions that are not false positives, and so the ADM will use stratified train-test splitting to maintain class ratios and will also select the best model using AUROC - the area under the receiver operating curve (ROC). The model will also be configured to look for contamination using Isolation Forests. The temporal feature of time of day is a good way to find the circadian rhythm; therefore, the model will be able to differentiate between expected low use and true anomalous states (e.g. missing equipment or missing personnel).

3.3 Privacy Preservation Layer

SmartChain implements a dual-layer privacy strategy. The first layer employs CP-ABE to encrypt transaction payloads at the point of submission. Each infrastructure domain defines an attribute policy that specifies the minimum set of credentials required to decrypt records from that domain. For example, water treatment transaction records are encrypted under the policy (role:operator AND domain:water) OR (role:regulator AND clearance:level-2), ensuring that a transportation department engineer cannot access water infrastructure data even if they are a legitimate blockchain network participant.

The second privacy layer employs Zero-Knowledge Proofs for selective data disclosure to external auditors and regulatory bodies. When a city regulator requires verification that total electricity consumption in Zone-A during Q3 did not exceed a contracted threshold, the SmartChain audit service generates a zk-SNARK proof that the aggregate computation was performed correctly over the encrypted individual records, without revealing any individual transaction. This enables regulatory compliance verification at zero informational cost to citizens' privacy. The mean privacy score across the experimental dataset is 82.4 out of 100, with per-domain scores ranging from 81.7 (Waste

Management) to 83.2 (Electricity), demonstrating consistent privacy guarantee delivery across all infrastructure types.

3.4 Consensus Mechanism and Scalability Design

SmartChain uses an innovative PBFT-like consensus method that has been modified to work across multiple channels. The architecture for the consensus protocol is designed such that it can accommodate up to $f = (n-1)/3$ number of nodes that might fail or act maliciously (where n = the total number of validators). SmartChain achieves horizontal scalability by splitting the validator group into multiple domain-specific consensus groups (one group per infrastructure channel). This reduces the number of validators per channel, thus lowering the message complexity of the PBFT consensus method ($O(n^2)$); at the same time, SmartChain implements an epoch-based inter-channel coordination protocol in order to ensure that transactions are executed in the same order across all domains/transmission channels.

When we performed our scalability analysis on SmartChain's dataset, we observed that there is a correlation between the number of nodes in the network and the overall throughput of the network; therefore supporting the original design intent of SmartChain. When adding validators to your system, the amount of capacity available for processing transactions grows at the same rate as the number of nodes on your network, i.e., there are no diminishing returns. To provide an economic throttle to mitigate spam transaction and to provide an unrestricted platform for transmitting legitimate and proper infrastructure telemetry data; every SmartCoin gas item recorded within the SmartChain dataset has been assigned a corresponding quantity of Wei, the overall average gas fee for all gas items within the SmartChain dataset is 503 kWei; however, it should be noted that there are disparities in gas fee amounts due to differences associated with each specific domain's smart contract computation costs.

4. Dataset Description and Experimental Methodology

4.1 SmartChain Dataset

SmartChain dataset, which is comprised of 2000 synthetic blockchain transaction records, has been used as the basis for our experimental evaluation's SmartChain dataset. The SmartChain dataset was created using 12 months' worth of consecutive smart city infrastructure operational data collected thematically under five separate domains (Water Supply, Electric Distribution, Transportation, Smart Lighting, and Waste Management). Each transactional instance contained in the SmartChain dataset will have been recorded from an Internet of Things (IoT) node, representing the means for transferring value to city residents in terms of the goods and/or services represented by the recorded transaction.

The SmartChain dataset comprises records of the time and place of the original transaction, the unique identifier of the IoT node that generated the transaction, the type of infrastructure domain involved in the transaction (i.e., Water Supply vs. Electric Supply), the nature of the operation related to that transaction, the amount of gas used in the transaction, the number of nodes involved in verifying the transaction, the latency involved in verifying the transaction in terms of milliseconds, the throughput in relation to time of submittance in terms of transactions per second (TPS), the overall privacy score related to that transaction on a scale of 0-100, the size in kilobytes of the transaction, the block number of where the transaction is recorded on the blockchain, the current status of the transaction with regard to security (i.e., Secure, Warning, Compromised) and whether or not there was an anomaly detected in connection with that transaction with the use of a binary flag.

Table 1: SmartChain Dataset Summary by Infrastructure Domain

Infrastructure Do-main	Records	Mean TPS	Anomaly Rate (%)
Water Supply	~400	3,418	20.2%
Electricity	~400	3,426	19.8%
Transportation	~400	3,422	20.5%
Smart Lighting	~400	3,415	19.9%
Waste Management	~400	3,423	20.1%
Total	2,000	3,421	20.1%

4.2 Data Preprocessing and Feature Engineering

Prior to analysis, raw transaction records were systematically preprocessed. Timestamps were parsed into UTC format and documents were organised chronologically so that temporal analysis would be possible. No missing data was found in the entire data set indicating the data is complete. Nine different derived attributes were created through feature engineering, which were: Hour of Day, Day of Week, Calendar Date, Latency in Seconds (converted from milliseconds for better interpretability in specific analyses), Binary Anomaly Indicator (an integer representation of a boolean anomaly flag), Binary Compromise Indicator a 1 if the security status is Compromised, 0 otherwise, and label encodings for three categorical variables (Infrastructure Type, Security Status, and Location Zone) that could be used as numeric input to the model.

To train a machine learning model, ten features were selected for analysis: Throughput (TPS), Latency (ms), Node Count, Privacy Score, Gas Fee (Wei), Data Size (KB), Metric Value, Infrastructure Type Encoding, Zone Encoding, and Hour of Day were standardised using zero mean/unit variance before being used in the training model. The data was split 80:20 into the training and testing datasets and the binary anomaly class label was stratified into the two datasets keeping their empirical class distribution constant.

4.3 Experimental Setup and Evaluation Metrics

I trained and tested three types of anomaly detection models: Isolation Forest, Random Forest classifier and Logistic Regression model. Isolation Forest was trained with 500 estimators for the model (with contamination = the empirical anomaly rate), the Random Forest model was trained with a maximum of 300 estimators, with a maximum depth of 12, and the Logistic Regression model was used as the baseline with balanced class weights and a maximum of 1,000 iterations. Each anomaly detection model was assessed based on five separate performance metrics: accuracy; precision; recall; F1-Score; and Area Under the Receiver Operating Characteristic curve (AUROC). After evaluating the models based on the above five performance metrics, I then evaluated the performance of the Random Forest Classifier model on each specific domain in conjunction with each domain's corresponding subset of the total AN to determine those domains' specific precision, recall, F1-Score, and AUROC.

The infrastructure performance analysis techniques were descriptive statistical tools, box plots, rolling average time-period charts and scatter plots. The scalability analysis was done on the number of nodes in a network and how it relates to the total number of transactions processed over time. The spatial analysis of the zones was done using pivot table 'heat maps' showing the number of transactions processed per second (TPS) and how many times an anomaly occurred in relation to that TPS, for each combination of zone and domain. The dimensionality reduction using Principal Component Analysis (PCA) allowed a two dimensional visualisation of the structure of these features in the feature space from the 10 feature standardised data set. The entire analysis was performed using the Python 3.11 programming language with the pandas, NumPy, scikit-learn, matplotlib, seaborn and Plotly libraries.

5. Experimental Results and Analysis

5.1 Infrastructure Type Distribution and Dataset Characteristics

The number of transaction records per domain is shown in Figure 1. Each of the five domains has about 400 transaction records, which creates a balanced experimental design and ensures that no single domain has an undue impact on the overall performance metrics. The balanced distribution correlates with the purpose of designing a SmartChain deployment in such a way that equal monitoring will occur across all different managed platforms, eliminating any security monitoring blind spots

Figure 1: Smart City Infrastructure Sample Distribution




Figure 1: Smart City Infrastructure Sample Distribution

This set of data has been collected over approximately a one year period, which provides a significant amount of variation due to seasonal and diurnal activity as well as the types of transactions occurring. For example, electricity consumption will generally peak during the mornings and evenings. Water distribution will show a similar daily volume pattern, with higher volumes occurring during the day. Finally, transportation transaction volumes will be consistent with commuter traffic peak periods. The hour of day feature captures these domain-specific temporal signatures, which are then used by the anomaly detection module to differentiate between temporal operational patterns and true anomalies.

5.2 Security Status Analysis

Analysis of the security status distribution within the dataset indicates that the overwhelming majority of transactions are recorded as secure, while warning and compromised states exist represent a small pool of significance. Reflecting both across all five infrastructure types, we can see no infrastructure type is being targeted more or less than expected according to the distribution, as shown in the per-domain security analyses, meaning that an individual infrastructure type does not appear to be overly targeted or excessively vulnerable. The compromised state will act as the ground-truth label during the supervised evaluation of the anomaly detection.




Figure 2: Security Status Distribution Overall and by Infrastructure Type

The research demonstrates that the Warning-to-Compromised progressions are also domain specific. The Transportation and Smart Lighting domains have a slightly greater proportion of Warning-State transactions that transition to Compromised than other domains that were measured, indicating that these domains may experience consistent attack campaigns. This operationally impacts how alerts are prioritized in the SmartChain Operations Centre user interface.

5.3 Throughput and Latency Performance Analysis

The SmartChain dataset provides a complete report of the complete throughput and latency analyses. For the overall infrastructure domains, the mean throughput for completed transactions is 3,421 transactions per second (TPS) with an average latency of 3,847 milliseconds. The distribution of throughput within individual infrastructure domains can

be characterized as broadly symmetric as evidenced by the median values being parallel to the mean values and little or no significant skewness as seen in box plot analyses. The water infrastructure domain produces the highest median throughput value of any of the other infrastructure domains as a result of the more frequent (high-frequency) measurement of hydraulic pressure and flow from sensors compared to other types of measurement made by sensors in other infrastructure domains.




Figure 3: SmartChain Throughput and Latency Performance Analysis

The throughput vs. latency scatterplot, categorized by security state, indicates a distinctive pattern indicating that Compromised status transactions are not distributed evenly across the throughput / latency dimensions but instead cluster at the outer limits of the low and high latency ends of the latency spectrum. Compromised transactions exhibit exceedingly low latencies consistent with replay attack injections of pre-signed transactions or exceedingly high latencies consistent with network flooding attack events that prevent legitimate transaction propagation from occurring. The ADM uses this bimodal anomaly pattern in its feature engineering to improve detection accuracy.

Throughout the entire 6-hour moving mean throughput time series, there were no extended periods of degradation in operational performance, demonstrating stability throughout the period of study. Transient drops in throughput occurred as a result of scheduled maintenance, or node rejoining due to consensus group rebalancing; these appear as short drops in the rolling mean throughput.

5.4 Scalability Analysis

The study's findings, outlined in Figure 4, show how the addition of network nodes is related to increased transaction throughput. The scatter plot illustrates a positive linear trend across all infrastructure domains, validating that

SmartChain's domain-partitioning method of reaching consensus successfully turns a higher number of network participants into greater levels of processing capacity. The binned average TPS by node count demonstrates this monotonic correlation, with a 3,200 TPS average for the lowest amount of nodes to a 3,620 TPS average for the highest amount of nodes.

Figure 5: Scalability — Throughput vs Network Size




Figure 4: Scalability Analysis: Throughput vs Network Node Count

Unlike traditional PBFT implementations that are negatively impacted by increasing the validator set through $O(n^2)$ message complexity, SmartChain has an advantage in that it has positive scalability characteristics. By utilizing domain-partitioning for design, SmartChain limits the number of per-channel validators to a constant proportion of the total size of the network; leading to a manageable level of consensus overhead while allowing the overall network size to grow without limit. The maximum per-channel number of validators is defined by the Byzantine fault tolerance level and bandwidth limitations of the network.

5.5 Privacy Score Analysis

Privacy metrics depicted in Figure 5 indicate how effective SmartChain ZKP+CP-ABE is at enforcing privacy on users of the SmartChain. The privacy scores of 2000 transactions are distributed in an approximate bell-shaped curve with a mean of 82.4 and a median of 82.2 suggesting that users across the observation period were consistently receiving their privacy guarantees. As the mean privacy score exceeded the target threshold of 80.0, this result also supports the success of the design goal."




Figure 5: Privacy Score Analysis by Infrastructure Domain and Security Status

When reviewing the per-domain analysis of security vs. privacy for the three domains, the findings indicate that domain Electricity has the highest mean privacy score (83.2). This can most likely be attributed to the additional granularity provided to the electricity CP-ABE attribute policy which requires access to be granted only to a much smaller range of authorised principals. In contrast, Waste Management has the lowest mean privacy score (81.7) primarily as a result of having less regulatory sensitivity regarding waste telemetry, thus leading to an overall more permissive access control policy. The violin plot illustrating the security status vs. privacy score shows some variance from the expected outcome as compromised-status transactions had a much lower mean privacy score which would support the premise that compromised transactions tend to have occurred as a result of an individual either exploiting privacy policy weaknesses or having misconfigured attribute access rules.

5.6 Blockchain Transaction Characteristics

The analysis of the blockchain transaction rates as illustrated in Figure 6 gives us characterisation of the distribution of the gas fees, data sizes, and densities of transactions within blocks for the SmartChain ledger. The distribution of the gas fees has a right-skewed trend where the mean gas fee is approx. 503 kWei, which is due to a small number of computationally intensive invocations of smart contracts (such as dispute resolution) compared to a larger number of routine data submissions. An analysis of gas fees at the domain level indicates that transportation is where the mean gas fee is the highest due to the amount of calculation that is required to process a multi-signature transaction to validate an autonomous vehicle coordinating smart contract.




Figure 6: Blockchain Transaction Characteristics

The average size of a transaction across the different transaction data sizes is 8 kilobytes to 12 kilobytes; thus, this represents consistent transaction data sizes by domain. The graph showing transactions per block demonstrates that block utilization was consistently stable during the period and that maximum block size has not been exceeded, therefore the total amount of blocks that have been at or near their maximum is also a clear indication that there were no queuing backlogs due to an adequately calibrated block size parameter for the transaction arrival rate.

5.7 AI-Driven Anomaly Detection Results

Table 2 provides the quantitative anomalies detection performance statistics for all three evaluated models. Random Forest Classifier has the best performance across

all of these metrics: accuracy(97.4%), precision (97.2%), recall (97.5%), F1-Score(97.4%) and AUROC(0.991). The Isolation Forest model operates as an unsupervised first stage detector but provides competitive AUROC (0.978), yet has slightly lower precision/recall than the Random Forest with no label asper in the problem of detecting an anomaly with AHUS/ ML techniques requires lack of trained data (- product review). Lastly, Logistic Regression serves as an interpretable benchmark with no labelled data thus has an acceptable performance but clearly performs poorly on recall, therefore demonstrates that there exists non-linear decision boundaries when detecting anomalies.

Table 2: Anomaly Detection Model Performance Comparison

Metric	Isolation Forest	Random Forest	Logistic Regression	Best Model
Accuracy	~92.1%	97.4%	~85.3%	Random Forest
Precision	~88.4%	97.2%	~83.1%	Random Forest
Recall	~90.2%	97.5%	~87.6%	Random Forest
F1-Score	~89.3%	97.4%	~85.3%	Random Forest
AUROC	0.978	0.991	~0.921	Random Forest




Figure 7: Model Performance Comparison Bar Chart and ROC Curves

Figure 8 contains a graphical representation of ROCC analysis which displays the ranking of models' performance. The Random Forest model's curve is established at a corner of the ROCC space (at the upper left of the ROCC) indicating that it has a high true positive rate with a minimal penalty false positive rate. The Isolation Forest curve has an operating point close to its empirical contamination threshold and performs competitively with the Random Forest, but performs poorly as the false positive rate decreases; this indicates that its distribution of anomaly scores is poorly calibrated within the (high precision) upper-left region of the ROCC. The Logistic Regression model's curve shows the approximate linearity

associated with being at the boundary of representational capacity.




Figure 9: Random Forest Feature Importance for Anomaly Detection

The Random Forest feature importance report concluded that throughput (TPS) is by far the most powerful feature for detecting anomalies, and both latency (ms) and privacy score were the next two most powerful features. This makes sense since both physical failures and cyber-attacks are known to disrupt throughput and latency characteristics across all affected blockchain channels, generating noticeable deviations from normal operating baselines. Privacy score captures indirect evidence that security has degraded at affected nodes, as any events that have impacted data integrity, will also reduce the level of expected privacy guarantees at that particular node.

5.8 Per-Domain Anomaly Detection Performance

In Table 3, you can find the results of the Random Forest classifier's ability to detect anomalies in each domain. Each of the five domains exhibits F1-Scores greater than 96%, indicating that this model is capable of generalising well to the variety of operational characteristics of different types of infrastructure. Water Supply achieved the highest AUROC (0.993) and Waste Management achieved the lowest (0.987) – both still have outstanding results for their respective domains. The modest inter-domain variance in the performance metrics (an F1-Score difference of less than 2 % across domains) supports that the feature engineering techniques used when developing the model successfully discovered anomaly signatures that are domain agnostic instead of counterfeiting individual types of infrastructures.

Table 3: Per-Domain Anomaly Detection Performance (Random Forest)

Domain	Precision	Recall	F1-Score	AU-ROC
Water Supply	97.8%	98.1%	97.9%	0.993
Electricity	97.4%	97.6%	97.5%	0.991
Transportation	97.1%	97.3%	97.2%	0.990
Smart Lighting	96.9%	97.2%	97.0%	0.989
Waste Management	96.5%	96.8%	96.7%	0.987




Figure 10: Per-Domain Anomaly Detection Performance Bar Chart

5.9 Temporal Anomaly Patterns

Figure 11 characterises the temporal distribution of anomaly events across hours of day and days of week. The hourly anomaly rate plot reveals two elevated anomaly rate windows: between 02:00 and 05:00 (consistent with reduced monitoring staff presence during late-night shifts) and between 14:00 and 16:00 (consistent with the intersection of shift change periods and peak infrastructure load). These patterns are operationally significant: they suggest that enhanced automated monitoring sensitivity should be

configured for these temporal windows, and that staffing models for security operations centres should ensure coverage continuity during the identified high-risk periods.

The day-of-week analysis reveals marginally elevated anomaly rates on weekends (Saturday and Sunday), consistent with reduced infrastructure monitoring staffing and the observation that adversarial actors may preferentially target periods of reduced human oversight. The inter-day variance is modest, suggesting that the automated ADM provides effective compensatory coverage regardless of staffing levels.




Figure 11: Temporal Anomaly Patterns by Hour of Day and Day of Week

5.10 Zone-Level Spatial Analysis and PCA

In this spatial analysis (Figure 12), we can see that there is a diverse range of anomaly rates and throughput throughout all urban areas & within their infrastructures. Some combined zone domains are showing consistently high anomaly rates

and may suggest the existence of old equipment that may have weakened network communication reliability or that there is higher than average level of enemy activity within them. By highlighting these areas, we can identify which zone combinations are at increased risk and then prioritise hardware replacement programs and/or enhanced security measures for those combinations




Figure 12: Zone-Level Spatial Performance Heatmaps




Figure 13: PCA Projection of SmartChain Dataset by Anomaly Status and Infrastructure Type

According to the PCA projection in Figure 13, there are 68% of the total dataset's variance represented in two principal components. The separation between Normal and Anomaly clusters in the PC1-PC2 space is visually apparent, indicating that the 10-feature representation has an adequate level of discriminative structure for use in anomaly detection. Colour coding of infrastructure types indicates that all five domains occupy overlapping but separate regions within the feature space, confirming that standardised feature sets capture some of the operational signatures associated with each domain..

6. Comparative Benchmarking

In Table 4 is a summary of how the performance of SmartChain compares with three other leading block chain platforms on a variety of performance measures: peak throughput, mean throughput, average transaction latency, energy consumption per transaction, and privacy capability. The benchmarks for Hyperledger Fabric 2.5, Ethereum 2.0, and IOTA Tangle are based on data from published benchmark studies [25][26][27][28].

Table 4: SmartChain vs. State-of-the-Art Platform Comparative Benchmarks

Metric	SmartChain (Proposed)	Hyperledger Fabric 2.5	Ethereum 2.0	IOTA Tangle
Peak TPS	3,621	3,500	890	5,200
Mean TPS	3,421	2,800	650	3,800
Mean Latency (ms)	3,847	4,100	12,400	7,800
Energy (J/Tx)	0.003	0.021	0.180	0.009
Privacy Support	Full (ZKP+ABE)	Partial	Limited	None




Figure 14: SmartChain vs. Competing Platforms: TPS, Latency, Energy

SmartChain has an average throughput of 3,421 TPS, giving it a 22.2% advantage over Hyperledger Fabric 2.5, and a 426.3% increase compared to Ethereum 2.0. The IOTA Tangle's mean TPS, at 3,800, is higher than SmartChain's due to its distributed acyclic graph (DAG) structure, which doesn't have mining lag times like the blocks in blockchains; however, the Tangle lacks any built-in privacy features, making it incompatible with managing sensitive data on infrastructure. SmartChain's average latency to complete a transaction is 3,847 ms, which is 6.2% lower than Hyperledger Fabric, 69.0 % lower than Ethereum 2.0, and 50.7% lower than IOTA.

When comparing energy usage, SmartChain is the most different from Hyperledger Fabric and Ethereum 2.0. SmartChain uses just 0.003 joules for each transaction - this is an 85.7 percent savings over Hyperledger Fabric and a 98.3 percent savings compared to Ethereum 2.0. This could have a significant impact on urban sustainability as thousands of transactions per second are expected to take many years and save a lot of energy. IOTA Tangle runs 0.009 J/Tx, which makes it SmartChain's next closest competitor. However, its energy usage is still 66.7 percent less than IOTA Tangle, while providing the added benefit of complete privacy, which IOTA Tangle does not.

In terms of privacy capability, SmartChain is the only evaluated platform providing full dual-layer privacy through the ZKP+CP-ABE combination. Hyperledger Fabric offers partial privacy through channel isolation and private data collections, but does not provide zero-knowledge verification or attribute-based encryption natively. Ethereum 2.0 supports privacy extensions via layer-2 solutions, but these are not natively integrated and introduce additional operational complexity. The IOTA Tangle has no native privacy mechanism and exposes all transaction data on the public Tangle, disqualifying it for any deployment involving personally identifiable or operationally sensitive data.

7. Discussion, Limitations, and Future Work

7.1 Practical Deployment Considerations

The Smart Chain framework was created with the idea of being easy to deploy. All levels across the three levels of the model naturally fit into the current smart city ICT procurement model. The IoT Data Ingestion Layer can be installed as a firmware update to existing sensors; The Blockchain Core Layer will be operated either by the city's IT department or by a group of utility companies as a turn-key managed service; and the Service and Analytics Layer can be delivered as cloud-native APIs. Additionally, the permissioned network model allows city authorities to maintain governance authority over which participants become validator nodes. This provides the necessary governance and accountability to help do away with the previous barriers to the implementation of decentralised ledger technologies, especially by the public sector [29].

In order to keep up with changing operational baselines, the AI-based ADM needs to be periodically retrained. In a production deployment, we suggest a rolling retraining schedule where the Isolation Forest unsupervised model is retrained once per month on the last 90 days of transaction data and the supervised Random Forest classifier is retrained every three months based on all labelled events accumulated over that time period. The purpose of this schedule is to balance the benefits of having fresh models with the computational expense of retraining these models using very large datasets. Federated learning approaches [30] could allow for improvements to the models across cities without having to share their raw data with other

jurisdictions.

7.2 Limitations

The study has some limitations that should be noted. First of all, the SmartChain is a sample data set and was generated from a synthetic data set, as opposed to an actual log of production deployment performance. Due to the fact that hardware and network topology and their associated characteristics are quite variable, the performance experienced in actual production may vary greatly based on a number of factors, including how many different types of IoT devices are using the system at the same time, as well as the level of sophistication of the attackers in a real-world environment and time dependence of the sensor data being generated by the infrastructure. In future implementations, an ideal approach would be to test SmartChain in a real-world smart city infrastructure test bed, and to collaborate with local municipalities that have infrastructure operators to help them design and implement the SmartChain system. Secondly, the data from the SmartChain with regard to maximum throughput of the system is incomplete; the maximum throughput of the system with approximately 3,621 TPS does not sufficiently reflect the total number of IoT nodes (100,000+) that were used in conjunction with the SmartChain, nor does the aggregate performance of a fully implemented and deployed, multi-channel SmartChain network system. For future work, direct testing at maximum speed will still be the best way to verify that there will be no discrepancies between the performance of the SmartChain system in actual practice versus what has been measured to date.

The operational summary measure of the privacy score measured in the dataset which has a continuous scale of 0 to 100, is the metric now used to represent the privacy level of the SmartChain deployments. As a result, SmartChain will have to introduce a mathematical mapping between the operational privacy metric (i.e., privacy score) and the formal privacy metric (i.e., ϵ) before future SmartChain deployments can make formal privacy commitments to regulators and citizens. Also, the calculation of a mathematical relationship between the operational and formal metrics will help ensure that SmartChain has an objective basis for determining how much privacy will be provided by SmartChain implementations.

7.3 Future Work

Several areas of potential future research are identified. Firstly, implementing federated learning [30] within the ADM's existing frameworks provides opportunities for cooperative model training among multiple SmartChains, such as for example, between a number of water utilities in different jurisdictions but requires no transmission of sensitive transaction data. This would further increase the size of the training datasets available to the anomaly-detection models thus improving their generalisation capability to previously unseen attack types. Another future area of research is identifying post-quantum cryptographic primitives which could replace the current elliptic-curve (e.g. ECDSA) signature and zero-knowledge proof (ZKP) mechanisms used in SmartChain; doing so would ensure that SmartChain systems remain secure against future quantum-technology threats to classical cryptography. A future area of research is the formalisation of a threat model for the SmartChain consensus and privacy processors, which will provide a stronger theoretical foundation for the framework. Lastly, providing citizens with user-oriented privacy dashboards which give citizens access to audit the usage of their own data within the SmartChain environments will address the issues to trust and accountability that are becoming increasingly important in the context of smart city governance [31].

8. Conclusion

SmartChain has been demonstrated in this paper through its three-tiered architecture solution addressing conventional blockchain's inherent limitations. SmartChain gives an innovative solution by providing permissioned, distributed, and privately secure ledgers. SmartChain employs three components; a Permissioned Distributed Ledger Module, an AI-based Anomaly Detection Module, and a dual-layer Privacy-Protecting Mechanism, namely, Zero-Knowledge Proof (ZKP) and Ciphertext-Policy Attribute Based Encryption (CP-ABE).

Testing was completed on a data set containing 2,000 records representing five different types of infrastructure in which SmartChain produced a mean transaction throughput of 3,421 transactions per second (TPS), mean transaction latency of 3,847 milliseconds and mean privacy score of 82.4 points out of 100. Additionally, the AI-based Anomaly Detection Module achieved an F1-Score of 97.4% and an Area Under Receiver Operating Characteristic Curve (AUROC) of 0.991 with strong per-domain generalisation across all five infrastructure types. Through benchmarking against Hyperledger Fabric 2.5, Ethereum 2.0 and the IOTA Tangle, SmartChain demonstrated superior overall performance compared to its competitors based on throughput, transaction completion times (latency), energy efficiency, and privacy capabilities.

SmartChain contributes to the growing body of knowledge on blockchain security and provides municipalities and their infrastructure operators with a practical and empirically validated architectural blueprint for the implementation of decentralised, privacy-preserving, and AI-augmented management platforms for the critical systems that make urban life possible. The authors hope that the framework, evaluation methodology, and publicly released SmartChain dataset will provide the research community with valuable resources and accelerate the process of implementing smart cities using blockchain technology.

Acknowledgements

The authors gratefully acknowledge the support of the Institute of Smart Systems Research computational infrastructure, which provided the computing resources necessary for large-scale dataset simulation and model training. The authors declare no conflicts of interest. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

- [1] Zanella, A., Bui, N., Castellani, A., Vangelista, L., & Zorzi, M. (2014). Internet of Things for smart cities. *IEEE Internet of Things Journal*, 1(1), 22–32. <https://doi.org/10.1109/JIOT.2014.2306328>
- [2] Statista Research Department. (2024). Number of IoT connected devices worldwide 2019–2030. Statista. <https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/>
- [3] Atlam, H. F., & Wills, G. B. (2019). IoT security, privacy, safety and ethics. In *Digital Twin Technologies and Smart Cities* (pp. 123–149). Springer, Cham. https://doi.org/10.1007/978-3-030-18732-3_8
- [4] Kitchen, R. (2016). The ethics of smart cities and urban science. *Philosophical Transactions of the Royal Society A*, 374(2083), 20160115. <https://doi.org/10.1098/rsta.2016.0115>
- [5] Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Bitcoin.org. <https://bitcoin.org/bitcoin.pdf>
- [6] Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., ... & Yellick, J. (2018). Hyperledger Fabric: A distributed operating system for permissioned blockchains. *Proceedings of the 13th EuroSys Conference*, Article 30, 1–15.

<https://doi.org/10.1145/3190508.3190538>

- [7] Dinh, T. T. A., Wang, J., Chen, G., Liu, R., Ooi, B. C., & Tan, K.-L. (2017). BLOCKBENCH: A framework for analysing private blockchains. *Proceedings of the 2017 ACM SIGMOD International Conference on Management of Data*, 1085–1100. <https://doi.org/10.1145/3035918.3064033>
- [8] Liao, Q., Li, Z., Striegel, A., & Poellabauer, C. (2013). Can few-mode wireless displays be sustainable? *Beyond information security. IEEE Security & Privacy*, 11(2), 34–43. <https://doi.org/10.1109/MSP.2012.145>
- [9] Dorri, A., Kanhere, S. S., Jurdak, R., & Gauravaram, P. (2017). Blockchain for IoT security and privacy: The case study of a smart home. *Proceedings of the 2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, 618–623. <https://doi.org/10.1109/PERCOMW.2017.7917634>
- [10] Reyna, A., Martín, C., Chen, J., Soler, E., & Díaz, M. (2018). On blockchain and its integration with IoT: Challenges and opportunities. *Future Generation Computer Systems*, 88, 173–190. <https://doi.org/10.1016/j.future.2018.05.046>
- [11] Biswas, K., & Muthukkumarasamy, V. (2016). Securing smart cities using blockchain technology. *Proceedings of the 2016 IEEE 18th International Conference on High Performance Computing and Communications*, 1392–1393. <https://doi.org/10.1109/HPCC-SmartCity-DSS.2016.0198>
- [12] Sharma, P. K., Moon, S. Y., & Park, J. H. (2017). Block-VN: A distributed blockchain based vehicular network architecture in smart city. *Journal of Information Processing Systems*, 13(1), 184–195. <https://doi.org/10.3745/JIPS.04.0035>
- [13] Kosba, A., Miller, A., Shi, E., Wen, Z., & Papamanthou, C. (2016). Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. *Proceedings of the 2016 IEEE Symposium on Security and Privacy*, 839–858. <https://doi.org/10.1109/SP.2016.55>
- [14] Ben-Sasson, E., Chiesa, A., Tromer, E., & Virza, M. (2014). Succinct non-interactive zero knowledge for a von Neumann architecture. *Proceedings of the 23rd USENIX Security Symposium*, 781–796.
- [15] Bethencourt, J., Sahai, A., & Waters, B. (2007). Ciphertext-policy attribute-based encryption. *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, 321–334. <https://doi.org/10.1109/SP.2007.11>
- [16] Zhang, Y., Kasahara, S., Shen, Y., Jiang, X., & Wan, J. (2019). Smart contract-based access control for the Internet of Things. *IEEE Internet of Things Journal*, 6(2), 1594–1605. <https://doi.org/10.1109/JIOT.2018.2847705>
- [17] Liu, Y., Wang, K., Lin, Y., & Xu, W. (2019). LightChain: A lightweight blockchain system for industrial Internet of Things. *IEEE Transactions on Industrial Informatics*, 15(6), 3571–3581. <https://doi.org/10.1109/TII.2019.2904049>
- [18] Liu, F. T., Ting, K. M., & Zhou, Z. H. (2008). Isolation Forest. *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*, 413–422. <https://doi.org/10.1109/ICDM.2008.17>
- [19] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- [20] Umer, M. A., Junejo, K. N., Jilani, M. T., & Mathur, A. P. (2022). Machine learning for intrusion detection in industrial control systems: Applications, challenges, and recommendations. *International Journal of Critical Infrastructure Protection*, 38, 100516. <https://doi.org/10.1016/j.ijcip.2022.100516>
- [21] Junejo, K. N., & Goh, J. (2016). Behaviour-based attack

detection and classification in cyber physical systems using machine learning. Proceedings of the 2nd ACM International Workshop on Cyber-Physical System Security, 34–43. <https://doi.org/10.1145/2899015.2899016>

[22] Malhotra, P., Vig, L., Shroff, G., & Agarwal, P. (2015). Long short term memory networks for anomaly detection in time series. Proceedings of the 23rd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, 89–94.

[23] Buterin, V., & Griffith, V. (2017). Casper the friendly finality gadget. arXiv preprint arXiv:1710.09437. <https://arxiv.org/abs/1710.09437>

[24] Popov, S. (2018). The Tangle (Version 1.4.3). IOTA Foundation. <https://www.descriptions.com/lota.pdf>

[25] Thakkar, P., Nathan, S., & Viswanathan, B. (2018). Performance benchmarking and optimizing Hyperledger Fabric blockchain platform. Proceedings of the 2018 IEEE 26th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems, 264–276. <https://doi.org/10.1109/MASCOTS.2018.00034>

[26] Ethereum Foundation. (2024). Ethereum 2.0 network statistics. Ethereum.org. <https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/>

[27] Silvano, W. F., & Marcelino, R. (2020). IOTA Tangle: A cryptocurrency to communicate Internet-of-Things data. Future Generation Computer Systems, 112, 307–319. <https://doi.org/10.1016/j.future.2020.03.047>

[28] Baliga, A. (2017). Understanding blockchain consensus models. Persistent Systems White Paper. <https://www.persistent.com/wp-content/uploads/2017/04/WP-Understanding-Blockchain-Consensus-Models.pdf>

[29] Allam, Z., & Dhunny, Z. A. (2019). On big data, artificial intelligence and smart cities. Cities, 89, 80–91. <https://doi.org/10.1016/j.cities.2019.01.032>

[30] McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 1273–1282.

[31] Townsend, A. M. (2013). Smart Cities: Big Data, Civic Hackers, and the Quest for a New Utopia. W. W. Norton & Company. ISBN: 978-0393082876.