International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(22s) 2026 pp. 1776-1790
https://doi.org/10.70917/ijcisim-2026-5777
Article

A Privacy-Preserving Markov Reputation Framework with Local Identity Binding for OIOT Networks

Abraham Tetteh[1], Archival J. Sebial[1]

[1]Department of CISM, University of San Carlos Cebu City, Philippines. 21104407@usc.edu.ph, ajsebial@usc.edu.ph

Abstract: Focusing on store-carry-forward communication and trust-aware routing when Internet of things is sporadically connected, opportunistic Internet of Things (OIOT) networks can enable identity-reputation correlation, behavioural profiling as well as reputation fraud. This paper introduces an AES-GCM-encrypted reputation storage/private key, HMAC-SHA256 integrity checks, and PID-associated authenticated data/privacy-preserving Markov reputation model that involves node-local pseudonymous identity binding, no global reputation exchange, and HMAC-SHA256 integrity checks. One of the considered insider forwarding threats is blackhole behaviour. In the Opportunistic Network Environment simulator, the framework was tested on 30 paired runs of the four scenarios: Baseline Epidemic, Blackhole, Markov, and Privacy-Markov. Privacy-Markov achieved 0.28 delivery probability, 2.17 overhead ratio, 6,297s latency, and 1.267 hops, with low variability observed across the paired runs. No identity-exposure, reputation-sharing, integrity-verification-failure, or AAD-authentication-failure events were recorded under the evaluated conditions; however, cryptographic execution time and energy consumption were not modelled.
Keywords: Opportunistic Internet of Things; Markov Reputation; Identity Privacy; Pseudonymous Identity Binding; AES-GCM; Blackhole Attack; Trust-Aware Routing

1. Introduction

In Opportunistic Internet of Things (OIOT) networks, mobile and resource-constrained devices exchange data through short-range encounters [1]. Intermittent connectivity, network partitions, and changing topologies distinguish OIOT from infrastructure-based IoT and support applications such as disaster response, vehicular sensing, smart cities, and human-centric IoT [2]. Trust and reputation mechanisms can improve resilience by deriving forwarding decisions from observed cooperation [3], but persistent identity-based reputation can expose users to identity disclosure, reputation manipulation, replay, and long-term behavioural or mobility profiling [4]. Existing approaches frequently employ global reputation exchange, aggregation, or explicit identity publication [5], increasing communication cost and exposure to false reporting. Computationally intensive or deterministic trust updates may also produce unstable decisions under intermittent contacts and temporary failures [6]. Although probabilistic Markov models provide a suitable basis for trust-state evolution, privacy-preserving identity management has rarely been integrated throughout the reputation lifecycle. A lightweight local framework that avoids global reputation exchange is therefore required.

The study aimed to design a privacy-sensitive reputation system using local pseudonymous identity binding, encrypted and integrity-protected reputation storage, and probabilistic Markov trust evolution. Its contributions include HMAC-derived PIDs without global reputation exchange, AES-128-GCM/HMAC-SHA256 protection with PID-bound AAD, and integration with trust-constrained forwarding. Implemented in THE ONE, four scenarios were evaluated through 30 paired runs. Compared with Blackhole routing, delivery increased by 34.2%, while overhead and hop count decreased by 95.9% and 60.4%, respectively, with a 33.6% latency increase.

2. Related Works



Copyright: © 2026 by the authors

Trust-based routing, attack mitigation, privacy protection, and Markov trust models are critically reviewed in this section, and the technical gap addressed by the proposed framework is derived.

2.1 Trust-Based Routing and Insider-Attack

Mitigation approaches have addressed major OIOT security concerns, including intermittent connectivity, dynamic topology, surveillance, and malicious participation [7]. Reputation-based methods such as BiRep detect non-cooperation through forwarding evidence [8], while AI-enabled social routing [9], greedy-node detection [10], credit incentives [11], and contract-theory mechanisms [12] improve cooperation but may require persistent identifiers, reporting, or additional control [13].

2.2 Privacy-Aware and Lightweight OIOT Routing

Privacy-aware opportunistic routing remains comparatively underexplored. ML-Fresh applied machine learning to forwarding selection [14], while social and encounter information reduced routing costs [15], but neither provided PID-indexed, encrypted reputation. Privacy-preserving security mechanisms have been explored for Cloud-IoT environments [16]. However, these approaches focus on access control and data protection rather than integrating trust-aware routing and reputation management. Other studies have focused on decentralized intrusion detection [17], energy-efficient social routing [18], and context-aware trust prediction for relay selection [22].

2.3 Markov-Based Trust Models and Critical Research Gap

Probabilistic trust modelling is valuable under uncertain OIOT contacts. Hidden Markov models have been used to represent behavioural transitions [19] and noisy attack detection [20], while lightweight trust routing [21], community trust, blackhole reporting and deep trust management have addressed relay selection and malicious behaviour. However, these approaches generally associate trust states with known entities and do not provide protected local pseudonym-reputation binding. In the proposed framework, each peer is initialized with reputation 0.5 in the SUSPICIOUS state, updated probabilistically from forwarding evidence, classified as GOOD, SUSPICIOUS, or MALICIOUS using explicit thresholds, and locally bound to a PID before authenticated encrypted storage.

Broader OIOT security research has also examined context-aware trust routing [22]. Although these approaches improve delivery, classification, contextual assessment, or integrity, they do not jointly provide protected local reputation storage and routing control. Resilient network security mechanisms have also been investigated to address diverse attack surfaces [24]. However, these approaches differ from the proposed lightweight framework, which focuses on privacy-preserving reputation and trust-aware routing. Blockchain-based mechanisms further introduce synchronization, replicated storage, consensus, and connectivity requirements that are difficult to sustain during intermittent contacts [23]. IOTA-based trust protection similarly differs from the proposed lightweight symmetric-cryptography model [25].

3. Materials and Methods

A privacy-preserving Markov routing methodology is specified in this section. Local PIDs, encrypted reputation records, integrity verification, and trust-based forwarding are integrated under the assumption that behavioural evidence is already available.

3.1 Design Goals and Assumptions

i) *Design Goals:* Five goals were specified: identities were to be protected during trust evaluation (G1); manipulation and rebinding were to be limited (G2); global reputation exchange was to be avoided (G3); Markov evolution was to be preserved (G4); and constant-time local cryptographic processing was to be maintained, with hardware costs left for future validation (G5).

ii) *System Assumptions:* Behavioural evidence was assumed to be supplied by an existing misbehaviour detector. Non-colluding nodes, private local reputation states, and the absence of a central authority were assumed.

3.2 Network and Trust Model

A decentralised OIOT network with intermittent short-range contacts and store-carry-forward routing was assumed. Forwarding evidence was used by a Markov reputation process to assign GOOD, SUSPICIOUS, or MALICIOUS states. Identity-indexed plaintext storage and reputation exchange were excluded.

3.3 Privacy-Preserving Identity Protection Framework

i) *Identity and Reputation Threats:* Threats were defined as identity tracing, identity-reputation correlation, stored-record tampering, replay, and cross-identity record substitution.

ii) *Local Pseudonymous Identity Binding (PID):* Each observing node $i$generates a local pseudonymous identifier for every encountered peer j. The PID is derived using HMAC-SHA256 with a node-local secret key and is never transmitted or shared with other nodes. The derivation implemented in the framework is expressed as

$$PID_{i,j} = Trunc_{12hex}\left[\left(Addr_j \parallel " \mid " \parallel S_{PID}\right)\right], \qquad (1)$$

where $PID_{i,j}$ denotes the pseudonymous identifier assigned by node $i$to peer j, $K_i^H$is the HMAC key derived from the local secret of node i, $Addr_j$is the simulator address of peer j, $S_{PID}$is the configured PID salt, and $Trunc_{12hex}$retains the first 12 hexadecimal characters of the HMAC output.

Because each node possesses a distinct local secret, different observing nodes generate different PIDs for the same peer. The PID is used as the exclusive index for local reputation records, thereby preventing the direct storage of reputation values under persistent node addresses. This mechanism protects the local identity–reputation association; it does not provide complete network-layer anonymity because the ONE simulator continues to use node addresses internally for communication and mobility operations.

At router initialisation, each node obtained a distinct 32-byte node-local master secret. The implementation supports three mutually exclusive secret-generation modes. When the fixed Secret parameter is configured, is computed as SHA-256 over the UTF-8 encoding of the configured secret concatenated with the label |node= and the node address. When the secretSeed parameter is configured, is computed as SHA-256 over the UTF-8 encoding of SIMULATION_SECRET|run=, the configured secret seed, |node=, and the node address. This deterministic mode provides reproducible but node-distinct simulation secrets. When neither parameter is configured, a fresh 32-byte secret is generated using Java SecureRandom. The resulting master secret remains unchanged for the lifetime of the router instance within a simulation run. Across-run reuse depends on the configured mode and value: identical fixedSecret or secretSeed values reproduce the same node-specific secrets when node addresses remain unchanged, whereas the SecureRandom mode generates new secrets.

iii) *Secure Local Reputation Storage:* Each local PID-indexed record contains contact and forwarding-evidence counts, reputation, trust state, timestamp, and HMAC. Before storage, the serialised record is encrypted with AES-128-GCM as defined in Eq. (2).

$$C_{i,j} = AES - GCM_{K_i^E}\left(V_{i,j}AAD = PID_{i,j}\right), \qquad (2)$$

where $C_{i,j}$is the sealed reputation record, $K_i^E$is the 128-bit AES key of node i, $R_{i,j}$is the serialized reputation record, and $IV_{i,j}$is a randomly generated 96-bit initialization vector. A 128-bit GCM authentication tag is used. The PID is supplied as Additional Authenticated Data, which binds the ciphertext to the corresponding local pseudonym.

Plaintext reputation tables are prohibited. Records are retrieved only through local PID-indexed decryption after contact or forwarding events.

iv) *Integrity and PID-Binding Protection:* Before encryption, HMAC-SHA256 is computed over the canonical record:

$$H_{i,j} = \left(PID_{i,j} \parallel c_{i,j} \parallel f_{i,j} \parallel r_{i,j} \parallel s_{i,j} \parallel t_{i,j}\right), \qquad (3)$$

where $c_{i,j}$is the contact count, $f_{i,j}$is the forwarding-evidence count, $r_{i,j}$is the reputation value, $s_{i,j}$is the trust state, and $t_{i,j}$is the record timestamp.

The framework therefore applies two integrity controls. First, HMAC-SHA256 detects modification of the reputation-record fields after successful decryption. Second, AES-GCM binds the ciphertext to $PID_{i,j}$through Additional Authenticated Data. A record encrypted for one PID cannot be successfully decrypted using a different PID because GCM authentication fails. Authentication or integrity failures cause the affected record to be rejected and reinitialized to its neutral state.

HMAC and AES-GCM operations, integrity failures, and AAD-authentication failures are recorded. Zero failures are interpreted only for the evaluated conditions because active tampering and record swapping were not injected.

3.4 Privacy-Preserving Markov Reputation Workflow

The local sequence is illustrated in Figure 1. After a peer is encountered, its PID-indexed record is decrypted with AES-GCM and the PID as AAD, and the HMAC is verified. Failed records are rejected and reinitialised; valid records are updated, resealed, stored locally, and used for forwarding. Thus, identity binding, integrity verification, reputation evolution, and routing selection are invoked in one deterministic module sequence.

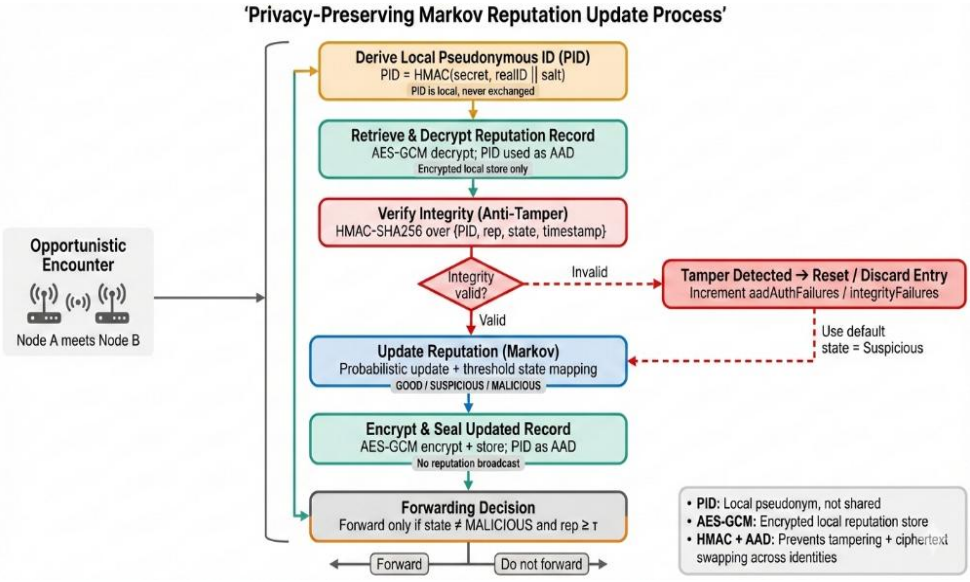


Fig. 1. Privacy-Preserving Markov Reputation Update Process (Identity-Protection Flow).

The implemented workflow is summarised in Figure 1. PID derivation, authenticated record retrieval, Markov updating, resealing, and the forwarding decision are performed locally; no PID or reputation record is transmitted.

3.4.1 Markov Reputation Evolution

For peer j, node $i$maintains a contact count $c_{i,j}$and a positive forwarding-evidence count $f_{i,j}$. A Laplace-smoothed reputation estimate is calculated as

$$\hat{r}_{i,j} = \{0.5, c_{i,j} = 0, \frac{f_{i,j}+1}{c_{i,j}+2}, c_{i,j} > 0. \quad (4)$$

The initial reputation is 0.5, corresponding to the SUSPICIOUS state. Following a positive forwarding event, the reputation is updated probabilistically as

$$r'_{i,j} = \{min(\hat{r}_{i,j} + \delta^{+}), u < p_{improve}, \hat{r}_{i,j}, u \geq p_{improve}, \quad (5)$$

where $u \sim U(1)$, $p_{improve} = 0.85$, and $\delta^{+} = 0.08$.

Following a contact event without positive forwarding evidence, the update is

$$r'_{i,j} = \{max(\hat{r}_{i,j} - \delta^{-}), u < p_{worsen}, \hat{r}_{i,j}, u \geq p_{worsen}, \quad (6)$$

where $p_{worsen} = 0.55$and $\delta^{-} = 0.06$.

The updated reputation is mapped to a discrete trust state:

$$S_{i,j} = \{GOOD, r_{i,j} \geq 0.50, MALICIOUS, r_{i,j} \leq 0.30, SUSPICIOUS, 0.30 < r_{i,j} < 0.50. \quad (7)$$

A forwarding connection is selected only when the encountered peer is not classified as MALICIOUS and its reputation is at least the configured GOOD threshold of 0.50. Reputation updates occur whenever a connection becomes active and whenever a message is successfully received from a peer. The stochastic generator used for reputation-transition decisions was linked to the documented run-specific Markov seed. Movement seeds ranged from 1 to 30, while the corresponding Markov-transition seeds ranged from 1001 to 1030. A deterministic node-specific seed derivation was used to maintain reproducibility while providing distinct random streams across nodes.

3.4.2 Algorithm 1: Privacy-Preserving Markov Reputation Update

Input: observing node i, encountered peer j, connection or forwarding observation e
Output: updated encrypted reputation record and forwarding decision

Step-1: Derive the node-local pseudonym $PID_{i,j}$using Eq. (1).

Step-2: Search the local encrypted store for the record indexed by $PID_{i,j}$.

Step-3: If no record exists, initialise $c_{i,j} = 0$, $f_{i,j} = 0$, $r_{i,j} = 0.5$, and $s_{i,j} = SUSPICIOUS$.

Step-4: Otherwise, decrypt the stored record using AES-128-GCM with $PID_{i,j}$as AAD.

Step-5: Verify the GCM authentication tag and the HMAC value defined in Eq. (3).

Step-6: If verification fails, reject the record and restore a neutral reputation entry.

Step-7: If a new connection is established, increment $c_{i,j}$.

Step-8: If a message is successfully received from peer j, increment $f_{i,j}$.

Step-9: Recalculate the smoothed reputation $r_{i,j}$using Eq. (4).

Step-10: Apply the probabilistic reward or penalty using Eq. (5) or Eq. (6).

Step-11: Map $r_{i,j}$to $GOOD$, $SUSPICIOUS$, or $MALICIOUS$using Eq. (7).

Step-12: Compute a new HMAC over the updated reputation record.

Step-13: Encrypt and store the updated record locally using AES-128-GCM.

Step-14: Forward messages through peer $j$only when $s_{i,j} \neq MALICIOUS$and $r_{i,j} \geq \theta_G$, where θG = 0.50.

3.4.3 Computational and Storage Complexity

PID derivation requires one HMAC-SHA256 operation and therefore has $O(1)$computational complexity for each encountered peer. Loading or storing a reputation entry requires one AES-GCM decryption or encryption and

one HMAC computation. Because each reputation record has a fixed number of fields, each local update also has $O(1)$time complexity.

For a node maintaining reputation records for $n_i$encountered peers, the encrypted local store requires $O(n_i)$memory. The framework introduces no reputation-broadcast messages because all reputation records and PIDs remain local. Its additional cost therefore consists primarily of local cryptographic processing and encrypted-record storage rather than network communication overhead.

3.4.4 Evaluation Setup, Repeated-Run Protocol, and Performance Metrics

i) *Simulation Environment:* Evaluation was performed with THE ONE simulator. Identical mobility, radio, traffic, TTL, and buffer settings were used in all scenarios; only router behaviour and privacy-enabled record protection were varied. Table 1 provides the parameters used for simulating the proposed work.

The simulations were executed on a workstation running Windows 11 equipped with an Intel i7 and 32 GB RAM, using Java Development Kit 23.

ii) *Experimental Scenarios:* Four scenarios were implemented:

a) *Baseline Epidemic:* EpidemicRouter was assigned to all 100 nodes.

b) *Blackhole:* EpidemicRouter was assigned to 80 nodes and BlackholeRouter to 20 malicious nodes; non-destination messages were accepted and discarded.

c) *Markov:* MarkovReputationRouter was assigned to 80 nodes and BlackholeRouter to 20 nodes; forwarding was restricted by local reputation.

d) *Privacy-Markov:* MarkovReputationRouterPrivacy was assigned to 80 nodes and BlackholeRouter to 20 nodes; the same routing rule was retained while PID derivation, AES-128-GCM, HMAC-SHA256, and PID-bound AAD were added.

Table 1. Simulation parameters and scenario configuration

Parameter	Baseline Epidemic	Blackhole attack	Markov	Privacy-Markov
Simulator	THE ONE 1.6.0	THE ONE 1.6.0	THE ONE 1.6.0	THE ONE 1.6.0
Simulation duration	43,200 s	43,200 s	43,200 s	43,200 s
Warm-up period	1,000 s	1,000 s	1,000 s	1,000 s
Total nodes	100	100	100	100
Honest/routing nodes	100	80	80	80
Malicious nodes	0	20	20	20
Routing classes	EpidemicRouter	EpidemicRouter + BlackholeRouter	MarkovReputationRouter + BlackholeRouter	MarkovReputationRouterPrivacy + BlackholeRouter
Mobility model	ShortestPathMapBasedMovement	Same	Same	Same

Buffer capacity	5 MB	Same	Same	Same
Message size	500 kB–1 MB	Same	Same	Same
Message TTL	300 min	Same	Same	Same
Markov-transition seeds	Not applicable	Not applicable	1001–1030	1001–1030

3.4.5 Performance and Privacy Metrics

The following network metrics were extracted.

The delivery probability was calculated as

$$P_{delivery} = \frac{N_{delivered}}{N_{created}}, \qquad (8)$$

where $N_{delivered}$ is the number of messages successfully delivered to their destinations and $N_{created}$ is the total number of generated messages.

The routing-overhead ratio was calculated as

$$R_{overhead} = \frac{N_{relayed} - N_{delivered}}{N_{delivered}}, \qquad (9)$$

where $N_{relayed}$ is the total number of message-relay events.

End-to-end latency was measured from creation to delivery; hop count was averaged over delivered messages; and buffer residence was measured until delivery, dropping, or removal. Relayed, dropped, aborted, and delivered counts were also extracted. Buffer occupancy was sampled every 300 s and averaged across matched runs and simulation time.

The counters represent observed events only. Zero counts are not treated as proof against replay, substitution, PID swapping, collusion, or Sybil attacks because such attacks were not injected.

3.4.6 Statistical Analysis

Each seed was treated as a paired block. Means, medians, sample standard deviations, variances, ranges, interquartile ranges, standard errors, and 95% confidence intervals were calculated. Distributional assumptions were assessed on within-seed differences rather than on unpaired scenario values. For pairwise comparisons, the distribution of the within-seed differences was assessed using the Shapiro–Wilk test. A paired-samples t-test was used when the paired differences satisfied the normality criterion. Otherwise, the Wilcoxon signed-rank test was applied. The family-wise error rate for comparisons within each metric was controlled using the Holm adjustment, with statistical significance evaluated at $\alpha = 0.05$.

Paired effect sizes were reported using Cohen's d_z for paired t-tests and rank-biserial correlation for Wilcoxon tests.

4. Results

Results from 30 paired runs are presented for Baseline Epidemic, Blackhole, Markov, and Privacy-Markov. Means, sample standard deviations, 95% confidence intervals, paired tests, and privacy counters are reported.

4.1 Validation of the Repeated-Run Dataset

The validated dataset comprised 390 reports: 90 each for Baseline, Blackhole, and Markov, and 120 for Privacy-Markov.

Movement seeds 1-30 and Markov seeds 1001-1030 were covered. All runs ended near 43,200 s and generated 1,459 messages. No missing seeds, duplicates, malformed reports, or incomplete runs were detected.

Each seed was treated as a paired block; principal metrics are summarised in Table 2 as mean ± SD with 95% confidence intervals.

Table 2. Descriptive performance results across 30 paired simulation runs

Scenario	Delivery probability	Routing-overhead ratio	Average latency (s)	Average hop count	Buffer residence (s)
Baseline Epidemic	0.1735 ± 0.0069; 95% CI [0.1709, 0.1760]	73.605 ± 3.462; 95% CI [72.313, 74.898]	4483.3 ± 244.9; 95% CI [4391.8, 4574.7]	4.638 ± 0.223; 95% CI [4.554, 4.721]	1240.9 ± 25.5; 95% CI [1231.4, 1250.4]
Blackhole attack	0.2108 ± 0.0083; 95% CI [0.2077, 0.2139]	52.836 ± 2.562; 95% CI [51.879, 53.793]	4713.7 ± 189.5; 95% CI [4643.0, 4784.5]	3.200 ± 0.166; 95% CI [3.138, 3.263]	1256.1 ± 25.8; 95% CI [1246.4, 1265.7]
Markov	0.2829 ± 0.0112; 95% CI [0.2787, 0.2871]	2.171 ± 0.139; 95% CI [2.120, 2.223]	6296.5 ± 227.5; 95% CI [6211.6, 6381.5]	1.267 ± 0.031; 95% CI [1.256, 1.279]	8494.4 ± 267.6; 95% CI [8394.5, 8594.3]
Privacy-Markov	0.2829 ± 0.0112; 95% CI [0.2787, 0.2871]	2.171 ± 0.139; 95% CI [2.120, 2.223]	6296.5 ± 227.5; 95% CI [6211.6, 6381.5]	1.267 ± 0.031; 95% CI [1.256, 1.279]	8494.4 ± 267.6; 95% CI [8394.5, 8594.3]

Exact network-layer equivalence was observed between Markov and Privacy-Markov for every matched seed; this result is discussed in Section 5.

4.2 Delivery Performance

Mean delivery probabilities of 0.1735 ± 0.0069, 0.2108 ± 0.0083, and 0.2829 ± 0.0112 were obtained for Baseline, Blackhole, and both Markov scenarios, respectively, as shown in figure 2. Against Blackhole routing, a paired increase of 0.0721 (34.2%; bootstrap 95% CI: 0.0681-0.0759) was obtained for Privacy-Markov (Holm-adjusted $p<0.001$; Cohen's $dz=6.545$). Against Baseline, the increase was 0.1094 (63.1%).

A higher delivery probability was also observed for Blackhole than Baseline by 0.0373 (21.5%; Holm-adjusted $p<0.001$).

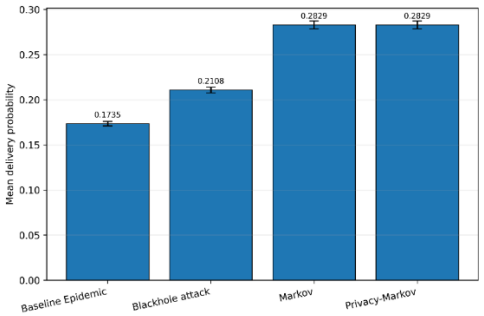


Fig. 2. Mean delivery probability across 30 paired simulation runs. Error bars represent 95% confidence intervals.

The congestion interpretation was supported quantitatively: overhead was reduced from 73.605 to 52.836 (28.2%), buffer occupancy from 81.49% to 75.41% (6.08 percentage points), and delivery was increased from 0.1735 to 0.2108. Thus, fewer replicas and lower contention were associated with higher aggregate delivery, without implying beneficial Blackhole behaviour.

4.3 Routing Overhead and Forwarding Activity

Mean overhead ratios of 73.605 ± 3.462, 52.836 ± 2.562, and 2.171 ± 0.139 were obtained for Baseline, Blackhole, and both Markov scenarios, respectively. Against Blackhole routing, a reduction of 50.665 units (95.9%; bootstrap 95% CI: 49.786-51.552; Holm-adjusted $p<0.001$) was obtained for Privacy-Markov. A logarithmic scale is used in Figure 3.

Mean relay counts were reduced from about 16,532 to 1,309 (92.1%), and drop counts from about 17,123 to 1,837 (89.3%), between Blackhole and Privacy-Markov.

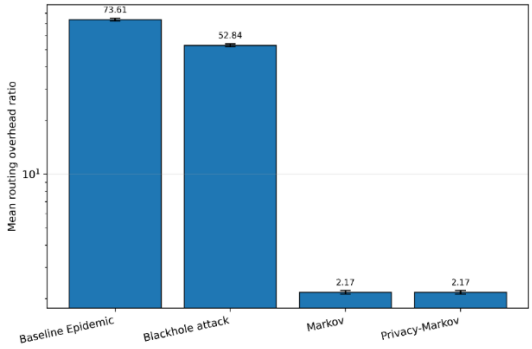


Fig. 3. Mean routing-overhead ratio across 30 paired simulation runs. Error bars represent 95% confidence intervals. The vertical axis is presented on a logarithmic scale.

4.4 Latency, Hop Count, and Buffer Residence

Mean latencies of 4483.3 ± 244.9 s, 4713.7 ± 189.5 s, and 6296.5 ± 227.5 s were obtained for Baseline, Blackhole, and both Markov scenarios. Against Blackhole routing, latency was increased by 1582.8 s (33.6%; bootstrap 95% CI: 1502.6-1670.1 s; Holm-adjusted $p<0.001$; Cohen's $dz=6.575$). Against Baseline, the increase was 1813.3 s.

The comparative end-to-end latency across the four scenarios is shown in Fig. 4.

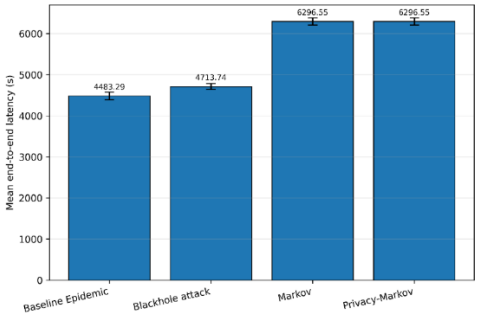


Fig. 4. Mean end-to-end latency across 30 paired simulation runs. Error bars represent 95% confidence intervals.

Mean hop counts of 4.638 ± 0.223, 3.200 ± 0.166, and 1.267 ± 0.031 were obtained for Baseline, Blackhole, and both Markov scenarios. Against Blackhole, a reduction of 1.933 hops (60.4%; bootstrap 95% CI: -1.995 to -1.873;

Holm-adjusted p<0.001; Cohen's dz=-11.053) was obtained. Fewer hops reflected restricted replication, not faster delivery.

The comparative average hop counts across the four scenarios are shown in Fig. 5.

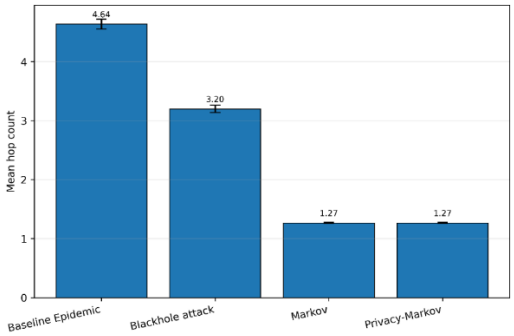


Fig. 5. Mean hop count across 30 paired simulation runs. Error bars represent 95% confidence intervals.

Mean buffer-residence times of 1240.9 ± 25.5 s, 1256.1 ± 25.8 s, and 8494.4 ± 267.6 s were obtained for Baseline, Blackhole, and both Markov scenarios. Against Blackhole, an increase of 7238.3 s was observed (bootstrap 95% CI: 7149.7-7331.3 s; Holm-adjusted p<0.001), showing prolonged waiting for eligible peers. 4.5 Delay Distribution and Buffer Occupancy

4.5.1 Conditional Delay Distribution

In the conditional delay ECDF, faster rises were observed for Baseline and Blackhole than for the Markov curve. Exact overlap was observed between Markov and Privacy-Markov. The right-shifted Markov curve was consistent with longer waiting for trusted contacts. Only delivered messages are represented by Fig. 6; overall delivery probability is reported separately in Section 4.2.

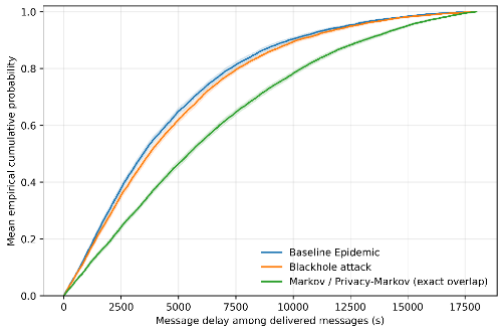


Fig. 6. Mean empirical cumulative distribution of end-to-end delay among successfully delivered messages across 30 runs. The Markov and Privacy-Markov curves overlap exactly.

4.5.2 Buffer Occupancy and Congestion-Related Evidence

Time-averaged occupancies of 81.49%, 75.41%, and 63.85% were obtained for Baseline, Blackhole, and both Markov scenarios. A 15.3% reduction relative to Blackhole was therefore obtained for Privacy-Markov. Peak mean occupancies were 87.78%, 83.29%, and 79.46%, respectively. Slower accumulation and lower occupancy were observed for the Markov variants as shown in Fig. 7.

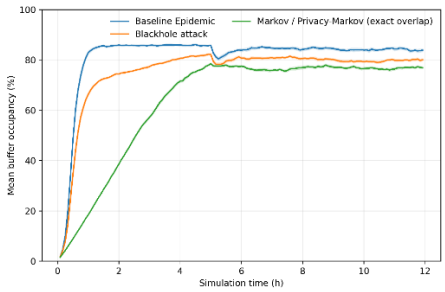


Fig. 7. Mean buffer occupancy over the 12-hour simulation period across 30 runs. Shaded regions represent 95% confidence intervals. Markov and Privacy-Markov overlap exactly.

Lower occupancy was consistent with reduced relaying and dropping. Fewer copies were stored, although individual messages could remain buffered longer while trusted contacts were awaited. Lower Blackhole occupancy than Baseline further supports the traffic-sink explanation: captured messages were not replicated, reducing epidemic pressure. Because congestion was not independently manipulated, association rather than isolated causation is claimed between selective forwarding, lower occupancy, and delivery. No identity-exposure, reputation-sharing, integrity-verification-failure, or AAD-authentication-failure events were recorded during normal execution as shown in Fig. 8.

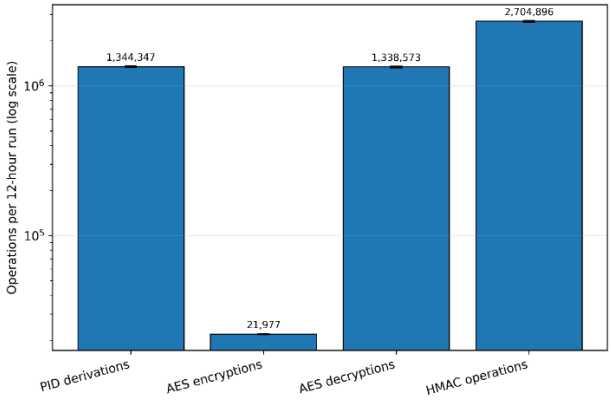


Fig. 8. Mean privacy-operation counts per 12-hour Privacy-Markov run across 30 simulations. Error bars represent 95% confidence intervals, and the vertical axis is logarithmic.

4.6 Contextual Cross-Study Comparison with Published Approaches

Because protocols, mobility, traffic, attacks, buffers, and aggregation differed, the comparison was treated as descriptive rather than controlled.

4.6.1 Reported Delivery Probability

A delivery probability of 0.2829 ± 0.0112 was obtained for Privacy-Markov. It was above the reported BiRep reputation-enabled value (about 0.2100) but below the reported T_CAFE, TBMOR, and BiRep no-reputation values in Fig. 9.

No universal ranking is inferred. Varying malicious-node proportions and multiple attacks were used in T_CAFE, whereas Prophet routing in a vehicular DTN was used for BiRep; a fixed 20% Blackhole proportion and map-based OIOT mobility were used here.

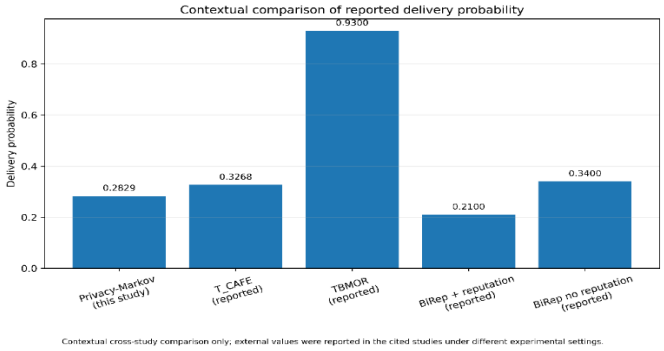


Fig. 9. Contextual comparison of delivery probability with values reported in previous studies.

4.6.2 Reported Latency

A latency of 6296.5 ± 227.5 s was obtained for Privacy-Markov, exceeding the reported T_CAFE, TBMOR, and BiRep values. The values are interpreted as different routing trade-offs, not a controlled ranking; T_CAFE, for example, reported about 2554.97 s under different conditions in Fig. 10.

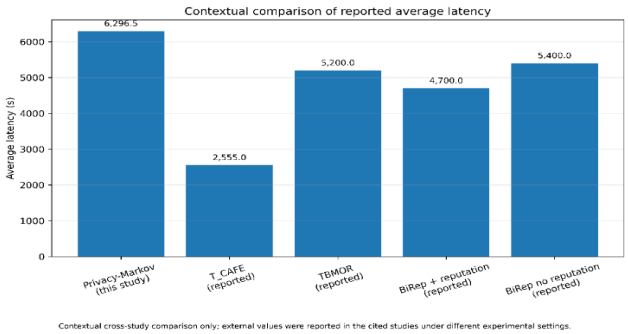


Fig. 10. Contextual comparison of average latency with values reported in previous studies.

4.6.3 Reported Routing Overhead

An overhead ratio of 2.171 ± 0.139 was obtained for Privacy-Markov, below the transcribed TBMOR and BiRep values in Fig. 11. Direct superiority is not claimed because metric definitions and experimental settings differed. Low overhead was demonstrated only within the present configuration.

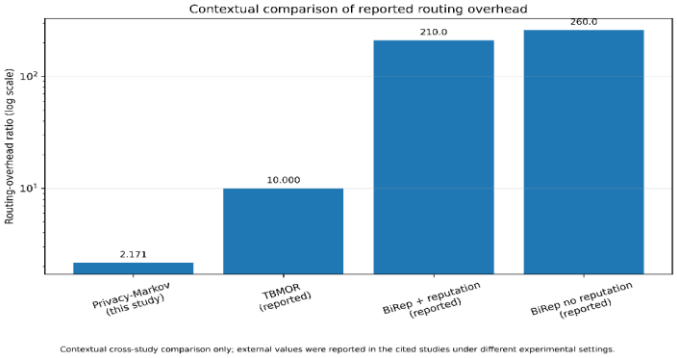


Fig. 11. Contextual comparison of routing-overhead ratios reported in previous studies.

5. Discussion

The 30 paired runs verified a delivery-overhead-latency trade-off in which Markov-based reputation-constrained forwarding reduced replication and resource consumption at the cost of longer waiting times. Relative to Blackhole routing, overhead decreased from 52.836 to 2.171, hop count from 3.200 to 1.267, and buffer occupancy from 75.41% to 63.85%, while latency increased from 4713.7 to 6296.5 s and buffer residence from 1256.1 to 8494.4 s. Higher delivery under Blackhole should not be interpreted as improved reliability because malicious nodes acted as traffic sinks, simultaneously reducing delivery and congestion. Privacy was strengthened through HMAC-derived PIDs, AES-128-GCM encryption, HMAC-SHA256 verification, and PID-bound AAD, although network-layer anonymity and resistance to Sybil, replay, collusion, and identity-splitting attacks were not evaluated.

Table 3 provides a qualitative comparison of the proposed framework with representative trust-based and privacy-aware routing approaches.

Table 3. Comparison of the Proposed Framework with Existing Trust-Based and Privacy-Aware Routing Approaches

Aspect	Representative trust-based schemes	Privacy-oriented schemes	Proposed framework
Identity–reputation association	Reputation may remain associated with persistent node or routing identifiers	Pseudonymity or anonymity may be provided, but trust enforcement can be absent	Reputation records are indexed using node-local HMAC-derived PIDs; simulator addresses remain in the underlying network model
Reputation handling	Trust values may be exchanged, aggregated, or reported	Reputation management may not be included	Reputation records remain local and are not globally exchanged
Storage confidentiality	Plaintext or identity-indexed trust storage may be used	Depends on the privacy mechanism	Local reputation records are encrypted using AES-128-GCM
Integrity protection	Protection against record replay or swapping varies among schemes	Often focused on identity or communication privacy	HMAC-SHA256 verification and PID-bound AES-GCM authentication are applied
Trust evolution	Deterministic, probabilistic, contextual, or reporting-based	Trust evolution may not be provided	Probabilistic Markov reputation updates
Main contribution	Malicious-node mitigation and trust-aware routing	Identity or communication privacy	Integration of local identity binding, reputation protection, and Markov routing
Evidence boundary	Depends on the original study	Depends on the original study	Evaluated under Blackhole forwarding; active privacy attacks and hardware costs were not tested

As shown in Table 3, mechanisms usually treated separately were combined; however, controlled baseline implementations were not replaced by this qualitative comparison.

6. Conclusion

This study presented a privacy-preserving Markov reputation framework for Opportunistic Internet of Things networks that integrates node-local pseudonymous identity binding, AES-128-GCM-encrypted reputation storage, HMAC-SHA256 verification, and PID-bound authenticated data into a probabilistic trust-aware routing lifecycle. Across 30 paired THE ONE runs, Privacy-Markov achieved a delivery probability of 0.2829 ± 0.0112, a routing-overhead ratio of 2.171 ± 0.139, an average latency of 6296.5 ± 227.5 s, and an average hop count of 1.267 ± 0.031. Relative to the Blackhole scenario, delivery increased by 34.2%, while overhead and hop count decreased by 95.9% and 60.4%, respectively; however, latency increased by 33.6%. The objectives were addressed as follows: G1 through HMAC-derived local PIDs; G2 through AES-GCM, HMAC verification, and PID-bound AAD; G3 through local-only reputation storage and zero recorded reputation-sharing events; G4 through identical Markov and Privacy-Markov network-layer results for all matched seeds; and G5 at the algorithmic level through constant-time local operations and linear storage growth.

Future work should implement identical-setting SOTA baselines, privacy-module ablations, replay protection using freshness counters or nonce windows, collusion detection through multi-observer evidence fusion, and lightweight identity validation using anonymous credentials or zero-knowledge proofs.

Data Availability Statement

The study-specific Java implementations, THE ONE simulation configuration files, random-seed definitions, raw 30-run simulation reports, statistical-analysis scripts, derived tables, and figures used in this study are publicly available at: https://github.com/abtet2020/privacy-markov-oiot-reproducibility. The underlying THE ONE 1.6.0 simulator is an external open-source platform and is not redistributed in the repository; it can be obtained separately from its official project website at https://akeranen.github.io/the-one/.

REFERENCES

1. A. I. A. Ahmed, S. H. Ab Hamid, A. Gani, and M. K. Khan, “Trust and reputation for Internet of Things: Fundamentals, taxonomy, and open research challenges,” J. Netw. Comput. Appl., vol. 145, Art. no. 102409, 2019.
2. Parveen, I. (2026). Artificial Intelligence-Based Intrusion Detection Using Graph Neural Networks for Internet of Things Networks. International Journal of Computer Information Systems and Industrial Management Applications, 18(14s), 600-625.
3. A. K. Subramanian, A. Samanta, S. Manickam, A. Kumar, S. Shiaeles, and A. Mahendran, "Linear Regression Trust Management System for IoT Systems," Cybernetics and Information Technologies, vol. 21, no. 4, pp. 15–27, 2021, doi: 10.2478/cait-2021-0040.
4. Rani, R. Y., Sar, A., Sree, P. H., Gupta, N., Sharma, G., & Anand, A. (2026). Instrument Development Using Structural Equation Modelling for Successful Factors for IOT Adoption: An Indian Manufacturing Perspective. International Journal of Computer Information Systems and Industrial Management Applications, 18(15s), 161-174.
5. A. Lohachab and A. Jangra, “A survey on security issues and challenges in opportunistic networks,” J. Inf. Secur. Appl., vol. 48, Art. no. 102341, 2019.
6. Qayyum, M., & Bari, M. A. (2026). A Comprehensive Review of Energy-Efficient Routing Protocols for Next-Generation IoT Networks. International Journal of Computer Information Systems and Industrial Management Applications, 18(12s), 438-466.
7. T. C. J. Jeaunita and V. Sarasvathi, "A Multi-Agent Reinforcement Learning-Based Optimized Routing for QoS in IoT," Cybernetics and Information Technologies, vol. 21, no. 4, pp. 45–61, 2021, doi: 10.2478/cait-2021-0042.
8. C. Nabais, P. R. Pereira, and N. Magaia, “BiRep: A reputation scheme to mitigate the effects of black-hole nodes in delay-tolerant Internet of Vehicles,” *Sensors*, vol. 21, no. 3, Art. no. 835, 2021, doi: 10.3390/s21030835.
9. R. Nigam, D. K. Sharma, S. Jain, K. K. Bhardwaj, and S. Banyal, “AI-enabled trust-based routing protocol for social opportunistic IoT networks,” *Trans. Emerg. Telecommun. Technol.*, vol. 35, no. 4, Art. no. e4330, 2024, doi: 10.1002/ett.4330.
10. F. S. Sadek, K. Belkadi, A. Abouaissa, and P. Lorenz, “Identifying misbehaving greedy nodes in IoT networks,” *Sensors*, vol. 21, no. 15, Art. no. 5127, 2021, doi: 10.3390/s21155127.
11. A. Sharma, N. Goyal, and K. Guleria, “Performance optimization in delay-tolerant networks using backtracking algorithm for fully credits distribution to contrast selfish nodes,” *J. Supercomput.*, vol. 77, pp. 6036–6055, 2021.
12. N. Gupta, J. Singh, S. K. Dhurandher, and Z. Han, “Contract-theory-based incentive design mechanism for opportunistic IoT networks,” *IEEE Internet Things J.*, vol. 10, no. 4, pp. 2881–2892, Feb. 2023, doi: 10.1109/JIOT.2021.3109162.
13. D. K. Bangotra, Y. Singh, A. Selwal, N. Kumar, and P. K. Singh, “A trust-based secure intelligent opportunistic routing protocol for wireless sensor networks,” *Wireless Pers. Commun.*, vol. 127, no. 2, pp. 1045–1066, 2022, doi: 10.1007/s11277-021-08564-3.

14. P. Garg, A. Dixit, and P. Sethi, “ML-Fresh: Novel routing protocol in opportunistic networks using machine learning,” *Comput. Syst. Sci. Eng.*, vol. 40, no. 2, pp. 703–717, 2022, doi: 10.32604/csse.2022.019557.
15. X. Li, H. Qi, and J. Wu, “Efficient path-sense transmission based on IoT system in opportunistic social networks,” *Peer-to-Peer Netw. Appl.*, vol. 15, no. 1, pp. 811–826, 2022.
16. Gurupriya K G, & Dr. A. S Aneeshkumar. (2026). Single- and Extended-Authority CP-ABE for E-Healthcare Security in Cloud-IoT. International Journal of Computer Information Systems and Industrial Management Applications, 18(21s), 752–772. https://doi.org/10.70917/ijcisim-2026-5342
17. Vineeta Shrivastava, Sandeep Kumar Agrawal, Pravesh Kumar Dwivedi, Jyoti Pasi, Arvind Sharma, & Yogita Verma. (2026). FLOW PATTERN CONSTRUCTION-BASED DECENTRALIZED INTRUSION DETECTION FRAMEWORK USING TL-IML3CSTM AND GAUSSIAN-Π FUZZY. International Journal of Computer Information Systems and Industrial Management Applications, 18(21s), 793–807. https://doi.org/10.70917/ijcisim-2026-5344
18. Rani and A. Malik, “A social relationship-based energy-efficient routing scheme for opportunistic Internet of Things,” *ICT Express*, vol. 9, no. 4, pp. 697–705, 2023, doi:10.1016/j.icte.2022.10.002..
19. G. Joshi and V. Sharma, “Hidden Markov trust for attenuation of selfish and malicious nodes in the IoT network,” *Wireless Pers. Commun.*, vol. 128, no. 2, pp. 1437–1469, 2023, doi: 10.1007/s11277-022-10007-6.
20. A. R. A. Moundounga, H. Satori, Y. Boutazart, and E. Abderrahim, “Malicious attack detection based on continuous hidden Markov models in wireless sensor networks,” *Microprocess. Microsyst.*, vol. 101, Art. no. 104888, 2023, doi: 10.1016/j.micpro.2023.104888.
21. B. Su and B. Zhu, “TBMOR: A lightweight trust-based model for secure routing of opportunistic networks,” *Egypt. Inform. J.*, vol. 24, no. 2, pp. 205–214, 2023, doi: 10.1016/j.eij.2023.02.002.
22. A. A. Hasan, X. Fang, S. Latif, and A. Iqbal, “Context-aware trust prediction for optimal routing in opportunistic IoT systems,” *Sensors*, vol. 25, no. 12, Art. no. 3672, 2025, doi: 10.3390/s25123672.
23. H. H. Mahmoud, J. Arshad, and A. Aneiba, “A systematic review of blockchain-based privacy-preserving reputation systems for IoT applications,” *Distrib. Ledger Technol. Res. Pract.*, vol. 3, no. 4, Art. no. 31, pp. 1–40, 2024, doi: 10.1145/3674156.
24. Ajit Karki, Vijay R, Harishchandra A. Akarte, & Abhijit Biswas. (2026). Software-Defined Networking Security: Architecture, Attack Surface, and Resilient Mechanisms. International Journal of Computer Information Systems and Industrial Management Applications, 18(21s), 731–743. https://doi.org/10.70917/ijcisim-2026-5340
25. A. Akli and K. Chougdali, “Towards tamper-proof trust evaluation of Internet of Things nodes leveraging IOTA ledger,” *Sensors*, vol. 25, no. 15, Art. no. 4697, 2025, doi: 10.3390/s25154697.