WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

TACS-VANET: Trust-Aware Clustering with Two-Dimensional Reputation and Ledger-Backed Revocation for Malicious-Vehicle Detection in Vehicular Ad Hoc Networks

Mr. Deepak Sain[1] ,Dr. Harsh Mathur[2]

Department of Computer Science and Engineering, Rabindranath Tagore University
Bhopal, (M.P.) India
[1] deepaksain.cse@satiengg.in, [2]harsh.mathur@rntu.ac.in

DOI : https://doi.org/10.55248/gengpi.07.1026.34078

ABSTRACT—

A malicious vehicle in a vehicular ad hoc network is an authenticated insider, which is precisely what makes it dangerous: cryptography establishes that its messages are genuine and says nothing about whether they are true. This paper presents TACS-VANET, a trust-aware clustering framework that turns the cluster structure itself into the detection mechanism, and evaluates it in simulation against a flat watchdog baseline. Three components carry the design. Cluster heads are elected by a weighted score combining behavioural reputation, residual energy, mobility stability and connectivity, so that only trustworthy and stable vehicles lead. Reputation is maintained in two independent Beta dimensions — behavioural trust from watchdog observations of relaying, and content trust from cluster-head plausibility and identity audits — and a vehicle is judged by its weaker dimension. Revocations are written to a ledger shared across roadside units so that exclusion is network-wide rather than local. Evaluation covers vehicle densities from 60 to 180, attacker populations from 0 to 40 per cent, and four attack types, each averaged over six independent runs. The framework raises the detection rate from 36.1% to 84.7% and the F1-score from 0.52 to 0.91 at 20 per cent attackers, while lowering end-to-end delay by 25 per cent and energy consumption by 9 per cent. The gains are not uniform: selective grayhole droppers are detected at only 34.7%, and the paper reports that limitation rather than averaging it away.

Index Terms—Vehicular ad hoc networks, malicious-vehicle detection, trust management, Beta reputation, clustering, cluster-head election, blockchain revocation, watchdog, intelligent transportation systems, network simulation.

I. INTRODUCTION

Vehicular ad hoc networks let vehicles exchange safety and traffic information with each other and with roadside infrastructure, and they are the communication substrate on which intelligent transportation systems and the Internet of Vehicles are built [1], [2]. The messages they carry are consequential in a way that ordinary network traffic is not: a collision warning that arrives late, or a congestion report that is false, translates directly into physical risk.

Most deployed protection is cryptographic. Certificateless and aggregate signature schemes authenticate messages efficiently, attribute-based encryption enforces fine-grained access control [2], and tree-based group key agreement secures platoon communication down to the antenna design [3]. These mechanisms answer one question well — did this message come from a credentialed vehicle and arrive unaltered — and are silent on a second question that matters just as much: is the vehicle behaving honestly. A malicious vehicle that holds valid credentials passes every cryptographic check while dropping the packets it agreed to relay, injecting fabricated congestion reports, or presenting forged identities to skew a vote [4], [5].

Two structural difficulties compound the problem. The first is scale: flat, network-wide security protocols incur overhead that grows with density and revoke slowly, exactly when traffic is heaviest. The second is mobility: at urban speeds the neighbour set of a vehicle turns over in seconds, so evidence about a neighbour's behaviour is gathered in fragments by observers who do not stay in contact long enough to accumulate a confident verdict individually. Clustering addresses both. Partitioning vehicles into locally managed groups bounds control traffic, and electing a cluster head to act as a local security manager gives the fragments somewhere to accumulate [9], [10].

Our earlier survey of this area concluded that the strongest published designs combine cryptographic identity, behavioural trust, and clustering, and sketched a conceptual framework on that basis [28]. What the survey could not establish is whether the combination actually works, or by how much. This paper closes that gap: it specifies the framework concretely enough to implement, implements it, and measures it.

The central design decision, and the one that turned out to matter most, concerns how reputation is represented. A single pooled trust score fails against attackers whose misbehaviour is orthogonal to their forwarding. A false-data injector relays every packet faithfully, so watchdog observations supply it with a stream of positive evidence that drowns out whatever the plausibility check contributes; in our own experiments a pooled score detected none of the false-data or Sybil attackers. Splitting reputation into behavioural and content dimensions, and judging a vehicle by its weaker dimension,

raised detection of those two classes to 96.5% and 90.3% respectively.

The contributions of this paper are fourfold.

- *A concrete trust-aware clustering protocol.* Cluster-head election is a weighted combination of behavioural reputation, residual energy, mobility stability and connectivity, with the weights and the election period fully specified, so that heads are both trustworthy and stable enough to remain heads.
- *Two-dimensional Beta reputation.* Behavioural and content evidence are accumulated in separate Beta distributions with exponential ageing, and a vehicle is judged by the minimum of the two. This is what makes detection of non-forwarding attacks possible at all, and the paper quantifies the difference.
- *Ledger-backed revocation with a measured contribution.* Revocation records are shared across roadside units so that exclusion is network-wide. The ablation isolates what this sharing is worth: without it, detection F1 falls from 0.91 to 0.43, because an excluded vehicle simply rejoins a different cluster.
- *A reproducible simulation study, including its negative results.* Five experiments cover density, attacker population, attack type, mechanism ablation and convergence, each averaged over six seeds. The grayhole result is reported at 34.7% rather than omitted, and the delivery-ratio cost of revocation is stated explicitly.

The remainder of the paper is organised as follows. Section II reviews recent work on VANET security, trust management, clustering and routing. Section III states the system and threat models. Section IV specifies the TACS-VANET protocol. Section V describes the simulation methodology, baselines and metrics. Section VI reports and discusses the results. Section VII states limitations, and Section VIII concludes.

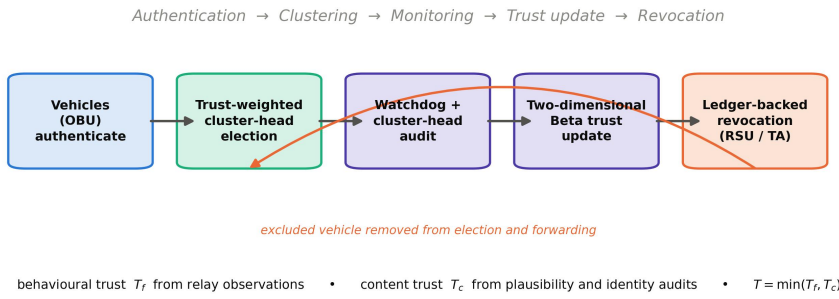


Fig. 1. The TACS-VANET closed loop. Authenticated vehicles form clusters; heads are elected on a trust-weighted score; watchdog observation and cluster-head audits feed two independent reputation dimensions; and a vehicle failing either dimension is revoked through a ledger shared by the roadside units, which removes it from both forwarding and future elections.

II. RELATED WORK

This review is organised around the four layers at which the surveyed literature intervenes — cryptographic identity, behavioural trust, cluster structure, and routing — followed by the physical-layer work that constrains how any of it should be simulated.

A. Cryptographic and Authentication Schemes

Cryptographic work establishes identity and integrity, and recent designs have made it efficient enough for vehicular timescales. Thuvva et al. present SAFE-V, arguing that conventional security systems scale poorly in vehicular settings and that blockchain alternatives impose a complex key-exchange burden [1]. Solanki et al. examine a recent attribute-based access-control scheme for VANETs, demonstrate concrete vulnerabilities including a collusion attack, and repair them [2] — a useful reminder that cryptographic schemes fail in ways that are found by analysis rather than assumed away. Rajasekaran et al. go further down the stack with TGKAV, a tree-based group key agreement scheme for vehicle platoons validated with a four-element planar antenna in the dedicated short-range communications band [3]. Hussain et al. compare elliptic-curve cryptography against RSA for the Internet of Drones and recommend the former on computational-cost grounds [22].

What none of these schemes attempts is behavioural judgement. Each authenticates the sender; none evaluates whether an authenticated sender is telling the truth. That boundary is the reason trust management exists as a separate layer.

B. Trust Management and Behavioural Evaluation

Trust systems judge behaviour rather than credentials. Zhang and Xu propose BDBAT, constructing Beta distributions from inter-node observations to assess trust in flying ad hoc networks, and report over 95 per cent detection accuracy for malicious drones [8]. The Beta formulation used in the present work follows the same lineage, with two differences: reputation is split across two evidence channels rather than pooled, and the evidence is aged so that a vehicle's distant past does not outweigh its recent conduct. Habelalmateen et al. combine an intelligent trust mechanism with collision-aware routing and particle-swarm optimisation, reporting a 10 per cent higher delivery ratio and 50 ms lower delay than their baselines [9].

Detection-oriented learning approaches attack the same problem from the data side. V. R et al. combine particle-swarm and bat-algorithm feature selection with an XGBoost classifier for distributed denial-of-service detection in VANETs, reporting 99.98 per cent accuracy [5]. Khenwar et al. apply a VGG16-based intrusion-detection system to black-hole attacks in ad hoc networks [6]. Subbarao et al. take a deliberately dataset-centric view of congestion prediction, benchmarking six model families under identical conditions and finding accuracy near 99 per cent across all of them — from

which they conclude, persuasively, that the result says more about redundancy in the feature space than about any model [18]. That paper is a caution worth carrying into any security evaluation: a very high number may be measuring the easiness of the benchmark rather than the strength of the method.

S. P. I et al. address black-hole mitigation by classifying nodes on forward rate and residual energy with dynamic thresholds, separating trusted from energy-constrained nodes [7]. Their coupling of energy state to trust is adopted here in the election weight, though we keep the two conceptually distinct: a vehicle low on energy is unsuitable as a head, which is not the same as being dishonest.

C. Clustering and Cluster-Head Selection

Clustering is the structural response to scale and mobility. Habelalmateen et al. propose a dynamic-mobility weighted cluster-head selection strategy assisted by unmanned aerial vehicles, observing that while clustering stabilises the network, an uncontrolled growth in cluster count increases energy congestion during communication [10]. Arivarasi et al. build clusters with type-II fuzzy logic and route between cluster heads using ant colony optimisation with a hybrid data-transmission mechanism to reduce the number of transmissions and thereby conserve energy [11]. The aggregation step in our own design serves the same purpose and, as Section VI shows, is a substantial part of why the clustered protocol consumes less energy despite carrying additional control traffic.

Where these schemes weight cluster-head selection by mobility, energy and connectivity, they do not generally weight it by trust. That omission has a specific consequence, which our ablation isolates: a Sybil attacker inflating its apparent degree is attractive to a degree-weighted election, and a compromised head is far more damaging than a compromised member because it audits its own cluster.

D. Routing, Resource Allocation and Energy

A second body of work optimises vehicular routing without a primary security focus, and supplies the performance baselines against which any secure protocol must be judged. Abdulsattar et al. propose maximum-lifetime particle-swarm routing, reporting a 7 per cent higher delivery ratio and 15 per cent better energy efficiency than their baselines [12], and separately develop a latency-reduction and traffic-management scheme for hybrid-surface VANETs reporting a 65 ms delay reduction [15]. Habelalmateen et al. report a 19 per cent higher delivery ratio and 130 ms lower delay for a mobility- and traffic-monitoring routing protocol using the Manhattan mobility model [13] — the same mobility model adopted here. Ihsan et al. address resource allocation by vehicle speed and overhead, reporting a 5 per cent higher delivery ratio [14]. Aerial assistance recurs as a theme: Habelalmateen et al. deploy multiple unmanned aerial vehicles with energy-fair resource allocation [16], and Sadiq et al. optimise peer selection in flying ad hoc networks under uncertain link quality using smell agent optimisation, reporting 25-35 per cent higher throughput [17].

Two observations follow. First, improvements in this literature are typically single digit to low double digit in delivery ratio, which sets a realistic expectation for what a protocol change can deliver. Second, these evaluations almost never include adversaries, so their delivery ratios are not directly comparable to those measured under attack.

E. Physical Layer, Propagation and Emerging Directions

Simulation fidelity is itself a research topic. S. K and Thilak analyse the impact of realistic radio propagation modelling on system-level VANET simulation, comparing 3GPP vehicle-to-vehicle channel models against simplistic assumptions and concluding that unrealistic propagation materially distorts application-level conclusions [19]. That finding directly shaped the channel model used here, which is a log-distance path loss with log-normal shadowing and a graded reception probability rather than a unit disc. Makarfi et al. analyse physical-layer security for reconfigurable intelligent surface-aided vehicular networks in terms of secrecy capacity and outage [20], and Jyostna et al. combine intelligent reflecting surfaces with reinforcement learning and hybrid beamforming [21]. These operate below the layer this paper addresses and are complementary to it: a physical-layer secrecy guarantee does not help against an insider who is a legitimate endpoint.

Broader integrations round out the picture. Shivhare and Raw examine how blockchain and artificial intelligence reinforce each other in autonomous vehicles, with blockchain supplying auditability and privacy [4]. Satyanarayana et al. apply an improved self-adaptive sailfish algorithm to secure routing and key management in mobile ad hoc networks [23]. G. S et al. integrate mobile ad hoc networks with 5G and a Q-learning forwarding scheme for disaster scenarios [24], and Boughariou et al. reduce energy consumption through mobile roadside units whose speed is predicted by a k-nearest neighbour model [25].

F. Positioning of This Work

Three gaps emerge. Cryptographic schemes secure identity but not behaviour. Trust schemes judge behaviour but generally pool all evidence into one score, which we show is the wrong representation when attacks are heterogeneous. Clustering schemes weight head selection by mobility and energy but rarely by trust, leaving the head — the most security-critical role in the structure — selected on criteria unrelated to honesty. TACS-VANET addresses all three at once and, unlike most of the work reviewed, reports the cases where it does not succeed.

III. SYSTEM AND THREAT MODEL

A. Network and Mobility Model

Vehicles move on a Manhattan street grid of 1500 m per side with 300 m blocks, following the mobility model used by Habelalmateen et al. [13]. Each vehicle travels at a constant speed drawn uniformly from 8 to 20 m/s and turns at a junction with probability 0.25, which produces the frequent topology changes characteristic of urban vehicular networks. Four roadside units are placed at the quarter points of the region; a vehicle within 200 m of one is considered associated with the infrastructure, and traffic is directed to the nearest unit.

The wireless channel follows a log-distance path loss with exponent 2.8 and log-normal shadowing of 4 dB, and the reception probability is a logistic function of the resulting signal-to-noise ratio rather than a hard threshold, so that links degrade gradually with distance. The choice follows the argument of S. K and Thilak that unrealistic propagation models distort system-level conclusions [19]. A transmission also contends with concurrent transmissions in the sender's neighbourhood, its success probability being scaled by $\exp(-\kappa c)$ where c counts the concurrent transmissions within range. This term is what makes a protocol's transmission count matter rather than being free.

B. Threat Model

All vehicles are assumed to have completed authentication and to hold valid credentials; the adversary is an insider, consistent with the class of threat that cryptography cannot address [1], [2]. A fraction of vehicles is malicious, each adopting one of four behaviours. A blackhole vehicle drops every packet it is asked to relay. A grayhole vehicle drops each packet with probability one half, which makes its behaviour statistically difficult to separate from ordinary channel loss. A false-data vehicle relays packets faithfully but corrupts their payload, so that a delivered packet carries incorrect information. A Sybil vehicle presents between three and seven forged identities, inflating its apparent connectivity and thereby its attractiveness in a degree-weighted cluster-head election.

The four are chosen because they stress different detection mechanisms. The first two are visible to a forwarding watchdog and invisible to a content check; the second two are exactly the reverse. Any scheme that reports a single aggregate detection rate over a mixed adversary population is concealing this structure, which is why Section VI reports the four separately.

IV. THE TACS-VANET PROTOCOL

A. Overview

Fig. 1 shows the closed loop. Authenticated vehicles self-organise into clusters; each cluster elects a head by a trust-weighted score; members and heads observe one another and audit reported content; observations update a two-dimensional reputation; and a vehicle whose reputation falls below threshold for consecutive evaluation rounds is revoked, with the revocation written to a ledger shared by the roadside units. A revoked vehicle is excluded both from forwarding and from future elections, so the loop closes back onto the clustering stage.

B. Trust-Weighted Cluster-Head Election

Elections run every 5 s. For each candidate vehicle i the election score combines four normalised terms as given in (1).

$$S_i = w_T T_i + w_E (E_i / E_0) + w_S \sigma_i + w_D \min(d_i / d_{ref}, 1) \qquad (1)$$

Here T_i is the composite trust of (4), E_i / E_0 the residual energy fraction, σ_i a mobility-stability term and d_i the neighbour count. The weights are (w_T, w_E, w_S, w_D) = (0.40, 0.20, 0.25, 0.15), giving trust the largest single share. Stability is defined in (2) from the mean relative speed between i and its neighbours, so that a vehicle moving with its neighbourhood scores higher than one cutting across it.

$$\sigma_i = \max(0,\ 1 - (1/|N_i|) \Sigma_{j \in N_i} \|v_i - v_j\| / 2v_{max}) \quad (2)$$

Vehicles are considered in descending score order; the highest unassigned candidate becomes a head and claims its unassigned neighbours as members. A vehicle left unclaimed heads a singleton cluster, which keeps the partition total. Revoked vehicles are excluded from candidacy entirely.

C. Two-Dimensional Beta Reputation

Each vehicle carries two independent Beta reputations. Behavioural evidence comes from watchdog observation in the manner of Marti et al. [26]: when a vehicle relays, up to three neighbours within earshot attempt to overhear the transmission, and each observation is itself subject to the observer's own link error rate. Content evidence comes from the cluster head, which checks reported events for plausibility against corroborating reports and checks identity claims for positional consistency. Following the Beta reputation formulation [27] and its vehicular adaptation [8], the two trust values are given in (3).

$$T_f = (\alpha_f + 1)/(\alpha_f + \beta_f + 2), \quad T_c = (\alpha_c + 1)/(\alpha_c + \beta_c + 2) \qquad (3)$$

The composite trust used everywhere else is the minimum of the two, as in (4).

$$T = \min(T_f, T_c) \qquad (4)$$

Taking the minimum rather than a weighted sum is deliberate and consequential. A weighted sum lets abundant good evidence in one dimension compensate for bad evidence in the other, which is exactly the compensation a false-data injector exploits: it forwards impeccably. Section VI-E quantifies what the pooled alternative costs.

All four evidence counters decay geometrically by a factor of 0.92 at each evaluation, so that reputation tracks recent conduct. A verdict is withheld until a vehicle has accumulated at least four units of evidence in some dimension, which prevents a newly encountered vehicle from being revoked on a single unlucky observation.

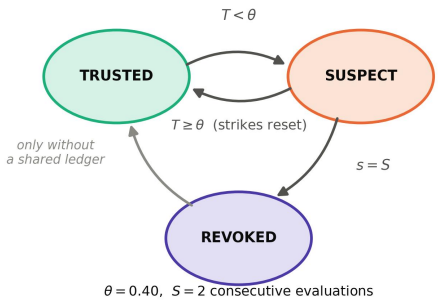


Fig. 2. Vehicle trust states. A vehicle enters SUSPECT when composite trust falls below the threshold and is revoked only after consecutive strikes; without a shared ledger the exclusion is local and the vehicle can rejoin elsewhere.

D. Revocation and the Shared Ledger

Fig. 2 shows the resulting state machine. A vehicle whose composite trust falls below $\theta = 0.40$ accumulates a strike; 2 consecutive strikes trigger revocation. The hysteresis matters because trust is estimated from a noisy channel, and a single depressed evaluation is a poor basis for permanent exclusion. On revocation the cluster head submits a record to a ledger replicated across the roadside units and the trusted authority, following the auditable-revocation pattern argued for by Shivhare and Raw [4] and by the blockchain-assisted schemes surveyed in [28]. Propagation is not instantaneous: a revocation becomes globally effective after a delay of 2 s.

Without the shared ledger a revocation is only meaningful inside the cluster that issued it, and the excluded vehicle can rejoin a neighbouring cluster whose head has no record of the offence. The ablation in Section VI-D models exactly this and measures the cost.

E. Routing and Aggregation

Members forward to their cluster head, which aggregates up to 4 member packets into a single frame before relaying towards the nearest roadside unit over a chain of cluster heads. Aggregation is the mechanism by which clustering reduces channel occupancy, and it follows the hybrid data-transmission argument of Arivarasi et al. [11]. A member whose head has drifted out of range, or whose head lies further from the infrastructure than the member itself, routes directly instead; without this rule the cluster structure imposes detours that cost more than aggregation saves.

Inter-cluster forwarding is greedy towards the destination roadside unit, with a bounded recovery mode permitting at most two non-improving hops. Pure greedy forwarding stalls whenever a city block lies between the carrier and the infrastructure, and in early experiments this single omission depressed delivery ratio by more than twenty points, which is worth stating because it is a property of the environment rather than of the security design.

V. SIMULATION METHODOLOGY

A. Simulation Environment and Parameters

The simulator is a discrete-time model written for this study in Python and released with the paper; the parameters are listed in Table I. Time advances in 0.5 s steps over a 180 s run. Every reported figure is the mean over 6 independent runs with different random seeds, and the shaded bands in the figures are one standard deviation across those runs.

Writing a purpose-built simulator rather than using NS-2, NS-3 or OMNeT++ is a trade-off and should be stated as one. What is gained is that every mechanism in Section IV — the two trust dimensions, the audit probabilities, the ledger delay — is directly instrumented, and the entire experiment reproduces from a single command. What is given up is the validated protocol stack those simulators provide, in particular a faithful IEEE 802.11p medium access layer. The contention term of Section III-A is an approximation of that layer, not a substitute for it, and Section VII treats this as the principal threat to validity.

B. Baseline and Metrics

The baseline is a flat protocol with the same channel, mobility, traffic and watchdog, differing only in that it has no cluster structure: vehicles forward greedily towards the infrastructure, and because there is no cluster head there is no plausibility or identity audit. The comparison therefore isolates what the cluster structure contributes, rather than comparing against a differently tuned protocol.

Network performance is reported as packet delivery ratio, the valid delivery ratio counting only uncorrupted packets, mean end-to-end delay, and total energy consumed. Detection performance is reported as detection rate, precision, false-positive rate and F1-score over the vehicle population, together with mean detection latency. The false-positive rate is reported throughout because a security mechanism that revokes honest vehicles degrades the network it is meant to protect.

TABLE I

Simulation parameters. Every value is read from the released configuration file.

Parameter	**Value**
Region	1500 m × 1500 m Manhattan grid
Block size	300 m
Vehicles	60 – 180
Speed	8 – 20 m/s, turn probability 0.25
Roadside units	4, association range 200 m
Radio range	300 m, log-distance ($n = 2.8$), shadowing 4 dB
Traffic	2 packets/vehicle/s, 4096 bit payload
Time step / duration	0.5 s / 180 s
Election period	5 s
Election weights (T, E, σ, d)	(0.40, 0.20, 0.25, 0.15)
Trust threshold θ / strikes	0.40 / 2
Evidence ageing	0.92 per evaluation
Aggregation factor	4 packets per frame
Ledger propagation delay	2 s

Parameter	Value
Malicious fraction	0 – 40 %
Independent runs per point	6

VI. RESULTS AND DISCUSSION

A. Detection Performance Across Vehicle Density

Fig. 3 shows detection rate and F1-score against vehicle density at a fixed 20 per cent attacker population. TACS-VANET detects 84.7% of malicious vehicles against 36.1% for the flat baseline at 120 vehicles, with F1 rising from 0.52 to 0.91; Table II reports the full metric set at that operating point. The margin holds across the whole density range, from 87.5% at 60 vehicles to 77.8% at 180, and the separation between the two curves exceeds their combined standard deviations everywhere.

The mechanism behind the gap is not subtle. The flat baseline has only the watchdog, which observes forwarding and nothing else, so it can detect the two attack types that manifest as dropped packets and is structurally blind to the other two. The clustered protocol adds the head's plausibility and identity audits, which cover precisely the complement. Detection is therefore not improved by better inference over the same evidence; it is improved by the cluster structure generating evidence that the flat protocol has no way to collect.

Both protocols maintain low false-positive rates — 0.00 for TACS-VANET and 0.00 for the baseline at 120 vehicles — which matters because the watchdog must distinguish a malicious drop from ordinary channel loss. Charging a relay only for transmissions it failed to make, rather than for onward losses outside its control, is what keeps this figure low.

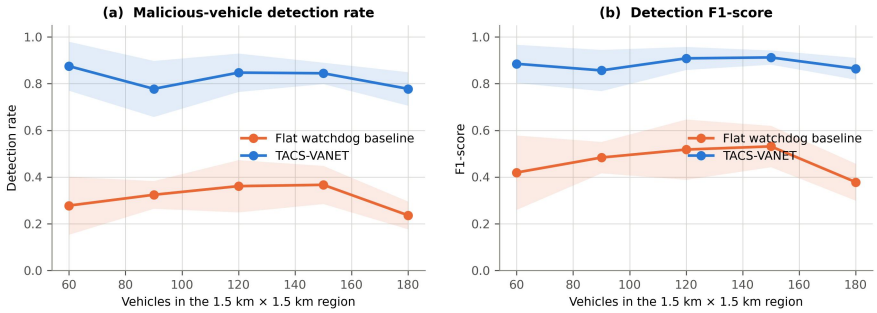


Fig. 3. Detection rate (a) and F1-score (b) against vehicle density at 20 % malicious vehicles. Shaded bands are one standard deviation over 6 runs.

TABLE II

Performance at 120 vehicles and 20 % malicious vehicles. Mean ± standard deviation over 6 independent runs.

Metric	Flat baseline	TACS-VANET
Detection rate	0.361 ± 0.112	0.847 ± 0.082
Precision	0.979 ± 0.047	0.984 ± 0.022
F1-score	0.518 ± 0.129	0.908 ± 0.049
False-positive rate	0.002 ± 0.004	0.003 ± 0.005
Detection latency (s)	15	31
Packet delivery ratio	0.394 ± 0.016	0.320 ± 0.025
Valid delivery ratio	0.383 ± 0.018	0.319 ± 0.025
End-to-end delay (ms)	21.4 ± 0.5	16.0 ± 0.2
Energy consumed (J)	9.3 ± 0.3	8.4 ± 0.2
Mean hop count	1.72	1.50

B. Network Cost

Fig. 4 reports delivery ratio, delay and energy against density, and here the result is mixed in a way worth stating plainly. TACS-VANET delivers 32.0% of packets against 39.4% for the flat baseline at 120 vehicles: clustering costs 7.3 points of raw delivery ratio. Two effects contribute. The uplink to a cluster head is an extra hop when the head is not itself on the path to the infrastructure, and — more interestingly — successful revocation removes vehicles from the relay pool, so a protocol that detects more attackers has fewer forwarders available. The security mechanism partly pays for itself in connectivity.

The picture changes when correctness rather than quantity is measured. Of the packets TACS-VANET delivers, 99.6 per cent are uncorrupted, against 97.3 per cent for the baseline, because false-data injectors are revoked rather than left in the forwarding path. For a safety application, a delivered packet carrying a fabricated hazard report is worse than no packet, so valid delivery ratio is the more meaningful measure, and by that measure the gap narrows to 6.4 points.

On delay and energy the clustered protocol is ahead: mean end-to-end delay falls from 21.4 ms to 16.0 ms, a reduction of 25 per cent, and energy consumption falls from 9.3 J to 8.4 J, a reduction of 9 per cent. Aggregation is responsible for both: combining up to 4 member packets into one frame reduces the number of transmissions, which lowers both the energy expended and the contention that transmissions cause each other.

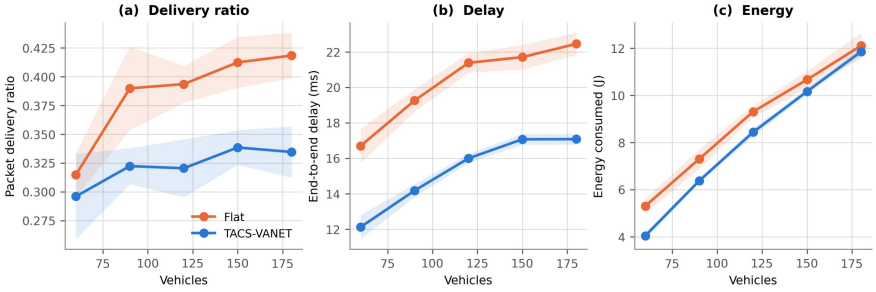


Fig. 4. Network cost against density: (a) packet delivery ratio, (b) end-to-end delay, (c) energy consumed. Clustering trades raw delivery ratio for lower delay and energy; part of the delivery gap is the cost of removing detected attackers from the relay pool.

C. Robustness to Attacker Population and Attack Type

Fig. 5 varies the malicious fraction from 0 to 40 per cent. With no attackers present both protocols correctly revoke almost nobody, the false-positive rate remaining at 0.00, which is the control condition a detection paper ought to report and frequently does not. As the attacker population grows, TACS-VANET holds detection between 74.1% and 84.7%, while the baseline stays below 36.1% throughout.

Fig. 6 and Table III decompose detection by attack type, and this is the most informative result in the paper. Blackhole vehicles are detected at 100.0%, false-data injectors at 96.5% and Sybil vehicles at 90.3%. Grayhole vehicles are detected at only 34.7%. The failure is intrinsic rather than incidental: a vehicle dropping half its packets produces an observation stream statistically similar to a vehicle on a poor link, and with the threshold set to keep the false-positive rate near 0.00 there is no operating point that separates them. Lowering the threshold would catch more grayholes at the cost of revoking honest vehicles in bad radio conditions, which is a worse failure. Detecting selective droppers requires evidence the watchdog does not supply — end-to-end acknowledgement or path-quality attribution — and that is the clearest direction for future work.

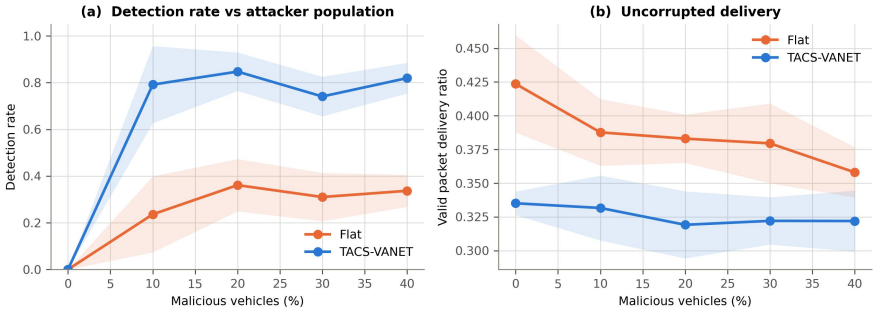


Fig. 5. Detection rate (a) and uncorrupted delivery ratio (b) against the malicious fraction, at 120 vehicles.

TABLE III

Detection by attack type at 120 vehicles, 20 % malicious, single-attack populations.

Attack	Detection	Precision	F1	FPR
Blackhole	1.000	0.993	0.997	0.002
Grayhole	0.347	0.936	0.501	0.005
Falsedata	0.965	0.987	0.975	0.003
Sybil	0.903	0.977	0.938	0.005

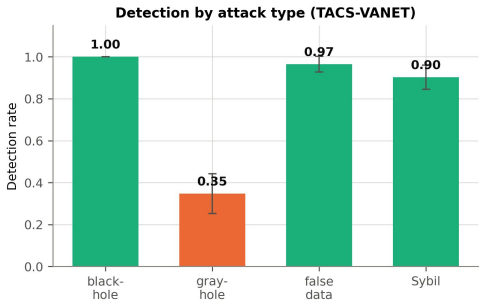


Fig. 6. Detection rate by attack type. Selective grayhole dropping is the one class the framework does not detect reliably.

D. Ablation of the Three Mechanisms

Fig. 7 and Table IV remove one mechanism at a time. Disabling trust management entirely reduces F1 from 0.91 to 0.00: with no reputation there is no revocation, and the cluster structure alone detects nothing. Removing the watchdog and retaining only the cluster-head audits gives 0.51, which corresponds to detecting the content attacks while missing the forwarding attacks — the mirror image of the flat baseline's failure, and a direct confirmation that the two evidence channels cover complementary threats.

Removing the shared ledger while keeping both evidence channels gives 0.43. This is the least obvious of the three results and the most practically important: the detection machinery still works, and the attackers are still identified, but because exclusion is only local a revoked vehicle rejoins a neighbouring cluster whose head holds no record of the offence. Detection without network-wide dissemination is close to no detection at all, which is a concrete argument for the ledger-backed revocation that the surveyed blockchain schemes propose [4], [28].

TABLE IV

Ablation at 120 vehicles and 20 % malicious vehicles.

Configuration	Detection	F1	FPR	PDR
Full TACS-VANET	0.847	0.908	0.003	0.320
No watchdog	0.340	0.506	0.000	0.325
No shared ledger	0.278	0.428	0.000	0.332
No trust management	0.000	0.000	0.000	0.321

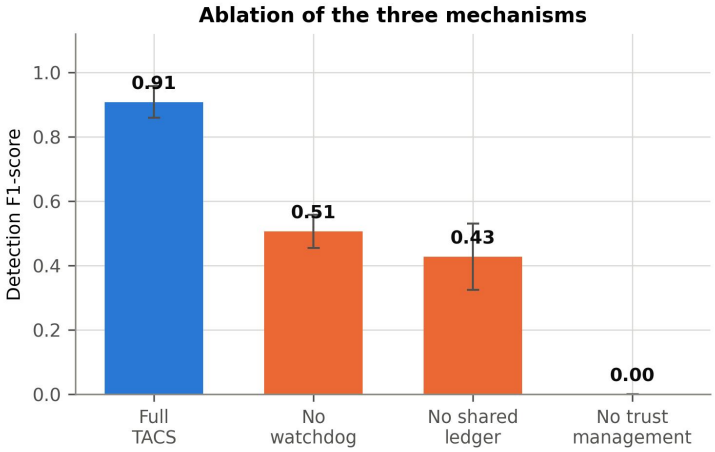


Fig. 7. Detection F1 with each mechanism removed in turn.

E. Positioning and the Cost of Pooled Trust

Table V places this study alongside the work reviewed in Section II. The comparison is presented with the caveat that most of those evaluations model no adversary at all, so their delivery ratios and delays are measured under conditions this paper does not share; the column recording what each study measured is therefore more informative than the numbers beside it. Among the studies that do model attackers, the Beta-trust scheme of Zhang and Xu reports over 95 per cent detection for malicious drones [8] and the learning-based detector of V. R et al. reports 99.98 per cent accuracy for distributed denial of service [5]. Neither figure is comparable to ours: the first is measured on a single behaviour class in an aerial network, and the second on a traffic classification task where attack and background traffic differ in volume by orders of magnitude. Our own aggregate is lower precisely because it averages over an adversary population that includes an attack class we do not solve.

Fig. 8 traces cumulative detection over simulation time. TACS-VANET separates from the baseline within the first 20 s and continues to climb as evidence accumulates, with mean detection latency of 31 s. The curve is bounded above by the grayhole population, which is never substantially detected, so the asymptote reflects the composition of the adversary rather than a limit of the learning process.

A final observation concerns the representation of trust itself. During development, an earlier version of this system pooled all evidence into a

single Beta reputation. Under that design the detection of blackhole and grayhole vehicles was unchanged, and the detection of false-data and Sybil vehicles was zero: every one of them escaped. The reason is arithmetical rather than mysterious. A false-data injector forwards faithfully, accumulating positive behavioural evidence at the rate of its traffic, while negative content evidence arrives only at the audit rate; the pooled ratio therefore never approaches the threshold. Splitting the dimensions and taking the minimum is a one-line change with a decisive effect, and we record it because the negative result is more instructive than the positive one.

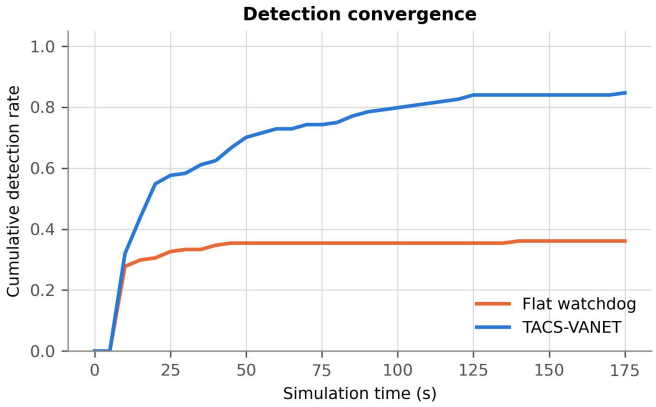


Fig. 8. Cumulative detection rate over simulation time, averaged over 6 runs.

TABLE V

Positioning against recent VANET and ad hoc network studies. Reported figures are as stated by the respective authors under their own conditions and are not directly comparable; the column records what each study measured.

Ref.	Focus	Adversary modelled	Reported outcome
[5]	DDoS detection, PSO-BAT + XGBoost	DDoS botnet	99.98 % accuracy
[7]	Blackhole mitigation by residual energy	Blackhole	not stated numerically
[8]	Beta trust for UAV networks	Malicious drones	> 95 % detection
[9]	Trust + collision-aware routing	Untrusted nodes	+10 % delivery ratio
[11]	Type-II fuzzy clustering + ACO	none	PDR, throughput, jitter
[12]	Max-lifetime PSO routing	none	+7 % delivery, +15 % energy
[13]	Mobility/traffic-aware routing	none	+19 % delivery, −130 ms
[14]	Speed/overhead resource allocation	none	+5 % delivery, −65 ms
[17]	FANET peer selection (SAO)	none	+25–35 % throughput
[18]	Congestion prediction, six models	none	≈99 % across all models
This work	Trust-aware clustering + ledger revocation	4 insider attack types	F1 0.91 vs 0.52 baseline

VII. LIMITATIONS AND THREATS TO VALIDITY

The most significant limitation is the simulation environment. A purpose-built model gives full instrumentation and exact reproducibility, but it lacks the validated IEEE 802.11p medium access layer that NS-3 or OMNeT++ provide. The contention term of Section III-A approximates carrier sensing and collision with a single exponential penalty; it does not model backoff, hidden terminals or the priority classes that safety traffic uses in practice. Absolute delivery ratios should therefore be read as internally comparable between the protocols evaluated here rather than as predictions of field performance. The argument of S. K and Thilak that propagation realism changes system-level conclusions [19] applies with equal force to medium access, and replicating these experiments on a standards-compliant stack is the necessary next step.

The audit model is parametric rather than derived. Plausibility and identity checks are represented by per-audit detection probabilities rather than by an implemented consistency algorithm operating on message content and position data. This is a reasonable abstraction for isolating the architectural question of where evidence is collected, but the probabilities are inputs rather than outputs, and a full evaluation would derive them from an implemented checker. The relative ordering of the ablation results does not depend on their exact values; the absolute detection rates do.

Several practical concerns are out of scope. Cryptographic authentication is assumed to have succeeded, so the overheads that schemes such as [3] and [22] quantify are absent from the energy figures. The ledger is a replicated append-only list with a propagation delay, not a consensus protocol with its own overhead and failure modes. Colluding attackers that vouch for one another are not modelled, and collusion is a known weakness of reputation systems. Finally the mobility model is a synthetic Manhattan grid rather than real trace data.

On the ethical side, a trust system that excludes vehicles from a safety network is a mechanism with real consequences for the excluded. A false positive removes an honest vehicle from the network that carries its collision warnings, which is why the false-positive rate is reported alongside every detection figure rather than relegated to a footnote, and why the strike-based hysteresis exists. Any deployment would need an appeal path and a bounded revocation lifetime rather than permanent exclusion on the basis of automated inference.

VIII. CONCLUSION

This paper has specified, implemented and evaluated TACS-VANET, a trust-aware clustering framework for detecting and isolating malicious vehicles in vehicular ad hoc networks. Cluster heads are elected by a score that weights behavioural reputation alongside energy, mobility stability and connectivity; reputation is maintained in separate behavioural and content dimensions and judged at its weaker point; and revocations are disseminated through a ledger shared across roadside units. Against a flat watchdog baseline sharing the same channel, mobility and traffic model, the framework raises detection from 36.1% to 84.7% and F1 from 0.52 to 0.91 at 20 per cent attackers, while reducing delay by 25 per cent and energy by 9 per cent.

Three findings deserve emphasis over the headline numbers. The two-dimensional representation of trust is not a refinement but a prerequisite: pooling the evidence reduced detection of false-data and Sybil attackers to zero, because faithful forwarding supplies enough positive evidence to mask an audit trail. Network-wide dissemination of revocations is worth nearly as much as detection itself, since local exclusion merely relocates the attacker. And selective grayhole droppers remain undetected at 34.7%, a limit imposed by the statistical similarity between selective dropping and ordinary channel loss rather than by any tuning choice.

Future work follows from the limitations. Replicating the study on a standards-compliant 802.11p stack would test whether the ordering of these results survives a faithful medium access layer. Replacing the parametric audits with implemented plausibility and identity-consistency checks would turn the audit probabilities from inputs into measured outputs. Adding end-to-end acknowledgement or path-quality attribution offers the most direct route to the grayhole problem. Finally, modelling colluding adversaries would test the reputation system against the attack it is most structurally vulnerable to.

ACKNOWLEDGEMENT

The authors thank the Department of Computer Science and Engineering, Rabindranath Tagore University, Bhopal, for providing the computing facilities used in this study.

REFERENCES

[1] A. Thuvva, R. Goyal and B. G. N, "SAFE-V: A Secure Framework for Vehicular Ad-Hoc Communications," in Proc. 2025 IEEE 1st Int. Conf. Smart Innovations in Systems, Infrastructure, Mechanical, Power, AI and Computing Technologies (SISIMPACT), Bhopal, India, 2025, pp. 1347-1352, doi: 10.1109/SISIMPACT67725.2025.11439808.

[2] R. Solanki, P. Chaudhari and M. L. Das, "Security of Vehicular Ad Hoc Networks with Fine-grained Access Control," in Proc. 2025 IEEE Int. Conf. Advanced Networks and Telecommunications Systems (ANTS), Delhi, India, 2025, pp. 1-6, doi: 10.1109/ANTS66931.2025.11429999.

[3] A. S. Rajasekaran, M. S. Obaidat, A. Maria, K. S. Kola, A. K. Das and Y. Park, "TGKAV: Tree-Based Group Key Agreement Scheme With Practical Antenna Implementation for Vehicle Platoon," IEEE Transactions on Intelligent Transportation Systems, vol. 26, no. 12, pp. 23342-23357, Dec. 2025, doi: 10.1109/TITS.2025.3610660.

[4] A. Shivhare and R. S. Raw, "VANET: Impact of Blockchain and Artificial Intelligence on Autonomous Vehicles," in Proc. 2025 Int. Conf. Networks and Cryptology (NETCRYPT), New Delhi, India, 2025, pp. 109-114, doi: 10.1109/NETCRYPT65877.2025.11102405.

[5] V. R, V. R, S. N, J. S, P. S K and B. R. K, "Machine Learning-Based DDoS Attack Detection in Vehicular Adhoc Networks," in Proc. 2025 Int. Conf. Intelligent Innovations in Engineering and Technology (ICIIET), Coimbatore, India, 2025, pp. 1-7, doi: 10.1109/ICIIET65921.2025.11377535.

[6] M. Khenwar, A. Sisodia, S. Vishnoi, S. Chugwani, R. Bhardwaj and N. Sharma, "Elevating Security Measures in Ad Hoc Networks: An Intrusion Detection Approach with Deep Learning," in Proc. 2025 Int. Conf. Advancements in Power, Communication and Intelligent Systems (APCI), Kannur, India, 2025, pp. 1-6, doi: 10.1109/APCI65531.2025.11137293.

[7] S. P. I, S. C and B. B, "Performance Analysis in Mitigation of Blackhole Attacker Considering Residual Energy," in Proc. 2025 IEEE 6th India Council Int. Subsections Conf. (INDISCON), Rourkela, India, 2025, pp. 1-6, doi: 10.1109/INDISCON66021.2025.11252238.

[8] D. Zhang and Q. Xu, "Beta Distribution-Based Adaptive Trust Management Scheme in UAV-Assisted IoTs," in Proc. 2025 IEEE Cyber Science and Technology Congress (CyberSciTech), Hakodate, Japan, 2025, pp. 610-614, doi: 10.1109/CyberSciTech68397.2025.00092.

[9] M. I. Habelalmateen, D. Jawad Hashim, A. S, A. A. Hameed, K. Aljawaheri and F. H. Alsalamy, "Active Evaluation and Intelligent Trust Mechanism with Collision Aware Routing in Vehicular Adhoc Networks," in Proc. 2025 11th Int. Conf. Communication and Signal Processing (ICCSP), Melmaruvathur, India, 2025, pp. 1548-1552, doi: 10.1109/ICCSP64183.2025.11089372.

[10] M. I. Habelalmateen, D. Jawad Hashim, R. Srivel, A. A. Hassan, K. Aljawaheri and F. H. Alsalamy, "Research on Dynamic Mobility Based Weighted CH Selection Strategy in UAVs Assisted Vehicular Adhoc Network," in Proc. 2025 11th Int. Conf. Communication and Signal Processing (ICCSP), Melmaruvathur, India, 2025, pp. 1135-1139, doi: 10.1109/ICCSP64183.2025.11089211.

[11] A. Arivarasi, V. S, S. Saravanan, V. Gokulakrishnan, R. S. Kumar and K. Madhan, "Improved Energy Efficiency in Vehicular ADHOC Networks Using T2FCATR Protocol," in Proc. 2026 Int. Conf. System, Computation, Automation and Networking (ICSCAN), Puducherry, India, 2026, pp. 1-6, doi: 10.1109/ICSCAN66520.2026.11588405.

[12] N. F. Abdulsattar et al., "Maximum Lifetime Based Optimized Particle Swarm Routing in Vehicular Adhoc Network," in Proc. 2025 11th Int. Conf. Communication and Signal Processing (ICCSP), Melmaruvathur, India, 2025, pp. 1124-1129, doi: 10.1109/ICCSP64183.2025.11088726.

[13] M. I. Habelalmateen, D. Jawad Hashim, B. Arunagiri, K. Aljawaheri, F. H. Alsalamy and I. Badi Khalaf, "An Effectual Approach for Mobility and Traffic Monitoring Based VANET's Routing Protocol in Urban Environment," in Proc. 2025 11th Int. Conf. Communication and Signal Processing (ICCSP), Melmaruvathur, India, 2025, pp. 1130-1134, doi: 10.1109/ICCSP64183.2025.11089478.

[14] M. Ihsan, H. M. Abed, E. Arasan, T. H. AlDaami, A. G. Abbood and F. Hashim, "Speed and Overhead Based Resource Allocation and Its Performance Analysis in V2X Communications," in Proc. 2025 11th Int. Conf. Communication and Signal Processing (ICCSP), Melmaruvathur, India, 2025, pp. 1558-1562, doi: 10.1109/ICCSP64183.2025.11088749.

[15] N. F. Abdulsattar, T. S. Madhi, Z. M. Alshukur, R. Al-Fatlawy, Z. H. Ward and F. H. Abbas, "Implementation of Latency Decrement and Traffic Management in Hybrid Surface Enabled Vehicular Adhoc Networks," in Proc. 2025 Int. Conf. Next Generation Computing Systems (ICNGCS), Coimbatore, India, 2025, pp. 1-6, doi: 10.1109/ICNGCS64900.2025.11182893.

[16] M. I. Habelalmateen, D. J. Hashim, S. U. Aswathy, F. H. Alsalamy, K. Aljawaheri and A. K. Zoori, "Energy Fairness Based Multiple UAVs Deployment with Resource Allocation in Vehicular Adhoc Networks," in Proc. 2025 Int. Conf. Next Generation Computing Systems (ICNGCS), Coimbatore, India, 2025, pp. 1-6, doi: 10.1109/ICNGCS64900.2025.11183227.

[17] B. O. Sadiq, M. D. Buhari, Y. I. Danjuma and A. Shabiluh, "Robust Peer-to-Peer Selection in FANETs With Uncertain Link Quality: A SAO-Based Approach for Improving Throughput, Fairness, and Energy Efficiency," IEEE Access, vol. 14, pp. 4435-4448, 2026, doi: 10.1109/ACCESS.2025.3650655.

[18] M. V. Subbarao, D. Ramesh Varma, P. D. Durga Parvathi, Y. Veera Sai Sindhu, D. G. Kumar and G. Challa Ram, "A Dataset-Centric Analysis of Traffic Congestion Prediction in Vehicular Ad Hoc Networks," in Proc. 2026 Int. Conf. Electronics and Renewable Systems (ICEARS), Tuticorin, India, 2026, pp. 706-711, doi: 10.1109/ICEARS67481.2026.11416533.

[19] S. K and A. S. Thilak, "The Impact of Realistic Radio Propagation Modeling on the System Level Performance of Vehicular Adhoc Network Simulation," in Proc. 2025 Int. Conf. Vehicular Technology and Transportation Systems (ICVTTS), Bangalore, India, 2025, pp. 1-6, doi: 10.1109/ICVTTS67119.2025.11296589.

[20] A. Makarfi et al., "Reconfigurable Intelligent Surface-Enabled Vehicular Networks: A Physical Layer Security Perspective," IEEE Open Journal of the Communications Society, vol. 7, pp. 239-252, 2026, doi: 10.1109/OJCOMS.2025.3642889.

[21] K. Jyostna, T. Ranganadh and B. S. Raj, "Performance Optimization of Reflective Surface Assisted Hybrid Beamforming in Multi User Vehicular ad Hoc Networks," in Proc. 2025 6th Int. Conf. Data Intelligence and Cognitive Informatics (ICDICI), Tirunelveli, India, 2025, pp. 1586-1591, doi: 10.1109/ICDICI66477.2025.11135360.

[22] S. F. M. Hussain, R. Karthikeyan, S. Ramamoorthi, V. Balaji, S. S. M. Gani and M. Sudhakar, "Enhanced Secure Transmission Using Cryptology in the Internet of Drones," in Proc. 2026 Int. Conf. Emerging Trends and Innovations in ICT (ICEI), Pune, India, 2026, pp. 1-6, doi: 10.1109/ICEI65890.2026.11447520.

[23] P. Satyanarayana, K. N. V. K. Rohitha, B. L. Aravind, D. Rajeswari, K. N. V. S. Varma and R. Ramprakash, "Optimization-Driven Network Security in MANET using an Improved Self-Adaptive Sailfish Algorithm (ISASA) for IoT Applications," in Proc. 2025 3rd Int. Conf. Intelligent Cyber Physical Systems and Internet of Things (ICoICI), Coimbatore, India, 2025, pp. 1380-1385, doi: 10.1109/ICoICI65217.2025.11252669.

[24] G. S, C. A and A. Roy, "Efficient Data Dissemination in a Challenging Environment Using a Hybrid Network," in Proc. 2025 IEEE 101st Vehicular Technology Conf. (VTC2025-Spring), Oslo, Norway, 2025, pp. 1-5, doi: 10.1109/VTC2025-Spring65109.2025.11174901.

[25] R. Boughariou, H. Marouane, M. Mosbah, H. Mnif and A. Meddeb-Makhlouf, "Optimization of Micromobility Integration in IoV Systems: Energy-Efficient Management Through Mobile RSUs and Machine-Learning Based Approaches," in Proc. 2025 28th Int. Symp. Real-Time Distributed Computing (ISORC), Toulouse, France, 2025, pp. 253-258, doi: 10.1109/ISORC65339.2025.00039.

[26] S. Marti, T. J. Giuli, K. Lai and M. Baker, "Mitigating Routing Misbehavior in Mobile Ad Hoc Networks," in Proc. 6th Annu. Int. Conf. Mobile Computing and Networking (MobiCom), Boston, MA, USA, 2000, pp. 255-265, doi: 10.1145/345910.345955.

[27] A. Jøsang and R. Ismail, "The Beta Reputation System," in Proc. 15th Bled Electronic Commerce Conf., Bled, Slovenia, 2002, pp. 41-55.

[28] D. Sain and H. Mathur, "Securing Vehicular Ad Hoc Networks Against Malicious Vehicles Using Advanced Clustering Techniques: A Survey of Cryptographic, Trust-Based, and Learning-Driven Defenses," International Research Journal of Modernization in Engineering Technology and Science (IRJMETS), vol. 8, no. 8, pp. 949-962, Aug. 2026, doi: 10.56726/IRJMETS102384. Available online: https://www.irjmets.com