Privacy-Preserving Pseudonym Management and Change Scheme in Fog-Based VANETs

Cong Zhao, Sheng Zhuang, Yikang Yang, Xinyang Deng, Xuan Ge, He Li*, Xiaopu Ma, Qinglei Qi, Wentao Li
School of Artificial Intelligence and Software Engineering, Nanyang Normal University, Nanyang, Henan, 473061, China
E-mail: zhaocong@nynu.edu.cn, zhuangsheng0302@163.com, yangyikang13@163.com, xinyang@dlnu.edu.cn, gexuan1013@gmail.com, * lihe@nynu.edu.cn, mapxiao@nynu.edu.cn, qiqinglei@nynu.edu.cn, wentaoli@hnu.edu.cn
*Corresponding author

Keywords: VANETs, privacy-preserving, fog computing, pseudonym change, pseudonym management

Received: April 21, 2026

Existing pseudonym schemes in Vehicular Ad-Hoc Networks (VANETs) incur significant communication and storage overhead. Moreover, while RSU-dependent schemes are constrained by the cost and scalability of infrastructure deployment, RSU-independent schemes frequently rely on centralized components that introduce both privacy vulnerabilities and performance bottlenecks. To address these issues, this paper proposes a fog-computing-based pseudonym-management architecture that utilizes distributed fog nodes to execute pseudonym generation, distribution, storage, and accountability. Within this framework, a vehicle entering a fog domain receives a pool of short-term pseudonyms from its nearest fog node for intra-domain communications. This architecture significantly reduces communication latency, alleviates the burden on the centralized certificate authority, and enhances system scalability. In addition, based on the proposed architecture, a vehicle-centric RSU-independent pseudonym-change mechanism is introduced. The mechanism dynamically determines the optimal timing for pseudonym changes by quantifying driving-behavior similarity and incorporating real-time vehicle states such as speed, direction, and relative position. Theoretical analysis shows that the proposed scheme provides effective privacy protection. Experimental results demonstrate that, compared with traditional RSU-based and centralized schemes, the proposed scheme significantly enhances anonymity while substantially reducing the risk of being tracked, offering a practical and scalable solution for privacy-preserving VANETs.

Povzetek: Predlagan je fog pristop za upravljanje in menjavo psevdonimov v VANET omrežjih, ki zmanjšuje zakasnitve in odvisnost od centralne infrastrukture. Rezultati kažejo boljšo anonimnost, manjšo sledljivost vozil in večjo razširljivost sistema.

1 Introduction

With the rapid development of information technology, the Internet of Things (IoT) has achieved widespread adoption across various industries. In recent years, the evolution of next-generation communication technologies has given rise to Vehicular Ad-Hoc Networks (VANETs), a key technology driving intelligent transportation systems. VANETs are revolutionizing the automotive industry, leading to profound changes that are already demonstrating immeasurable value and far-reaching impact [1]. By integrating state-of-the-art technologies such as wireless communications, vehicle positioning, data processing, and intelligent decision-making, VANETs have become an indispensable component of intelligent transportation systems. As the number of vehicles continues to increase, traffic safety has emerged as a critical concern. VANETs address this challenge by facilitating real-time information exchange, thereby significantly enhancing road safety and reducing accident rates [2]. This technology holds considerable potential for technological innovation, improving traffic safety and efficiency, while also contributing to environmental sustainability, commercial exploration, and policy formulation, making it a subject of increasing research attention.

A typical VANETs architecture consists of a Trusted Authority (TA), Roadside Units (RSUs), On-Board Units (OBUs), and other fundamental components [3]. The TA, which is managed by government agencies or vehicle manufacturers, is trusted by all other entities. TA is responsible for managing and maintaining system parameters, registration information, and authentication certificates. RSUs are infrastructure nodes deployed along roadways. They serve as intermediate nodes to facilitate communication and provide various services to vehicles within their jurisdiction [4]. Vehicles, as mobile communication nodes, are equipped with OBUs to send, receive, and process information to other entities via Dedicated Short Range Communications (DSRC) [5]. Vehicles periodically broadcast beacon messages every 100 to 300 ms. From a communication perspective, VANETs are primarily categorized into two modes: Vehicle-to-Vehicle (V2V) and

Vehicle-to-Infrastructure (V2I). In V2V communication, vehicles exchange critical real-time information such as traffic incidents, emergency braking alerts, driving conditions, weather conditions, and adverse road conditions [6]. In V2I communication, RSUs respond to vehicle requests, providing services such as traffic safety notifications, navigation, and other relevant services, while also offering infotainment services for drivers and passengers [7].

However, due to their inherent characteristics such as high node mobility and dynamic network topology, VANETs are highly susceptible to various attacks from both external and internal attackers [8]. The information leakage caused by broadcast beacons is a significant threat. This is because these beacons contain sensitive data like vehicle identity, GPS coordinates, speed, and direction [9], which may be exploited for malicious tracking, violating privacy and potentially compromising safety [10]. Moreover, due to the close association between drivers and vehicles, this linkage poses a significant threat to the confidentiality of drivers' locations in VANETs [11].

To address privacy concerns, vehicles use temporary authenticated identifiers during communication, commonly known as pseudonyms [12], preventing the leakage of real identity information [13]. While pseudonyms enhance anonymity, they also increase the complexity of inspecting and tracking. If a pseudonym is used for an extended period, attackers may infer the vehicle's true identity by analyzing trajectory, such as the destinations, driving patterns, and other contextual data [14]. However, frequent changes of pseudonyms result in additional storage and communication overhead for vehicles. Furthermore, even if only a small number of vehicles change their pseudonyms simultaneously, it may still be insufficient to achieve effective privacy protection. Therefore, the core challenge of pseudonym schemes lies in determining when and where to change pseudonyms in a reasonable manner to ensure unlinkability. These decisions directly impact the effectiveness of privacy protection. On one hand, in order to maximize anonymity, the number of vehicles within the pseudonym change area should be as large as possible. This would increase the difficulty for attackers to track a specific vehicle. On the other hand, to ensure unlinkability, it is preferable for all vehicles within the pseudonym change area to change their pseudonyms simultaneously, making it difficult for attackers to establish a continuous tracking chain. However, in the practice of pseudonym change, the primary principle of ensuring driving safety must also be considered [15]. Addressing how to protect privacy without impeding necessary communication is therefore a critical challenge [16].

Nowadays, many pseudonym schemes have been proposed. Widely adopted security standards, such as IEEE 1609.2 and ETSI 102941-v1.1.1, consider public keys certified by the Certificate Authority (CA) as pseudonyms. Vehicles change their pseudonyms based on the CA's instructions to facilitate location tracking in certain scenarios [17]. In urban environments, strategies based on hybrid regions have been widely implemented. These strategies tend to implement pseudonym changes at high-traffic locations such as busy intersections, parking lots, and gas stations [18]. However, in sparse traffic scenarios, such as highways and suburbs, the difficulty of forming effective hybrid regions weakens the anonymity effects [19]. To address this challenge, pseudonym change strategies based on hybrid contexts have emerged. These strategies dynamically determine the optimal timing for pseudonym changes based on real-time environmental factors such as traffic density, vehicle speed, and tracking probability. Through collaborative pseudonym changes or exchanges among vehicles, this approach can more effectively conceal their true identities [20]. This strategy provides stronger privacy protection, as it is suitable not only for dense urban areas but also for sparse environments, thus providing broader adaptability [21]. Compared to hybrid region-based schemes, hybrid context-based strategies are more flexible and efficient. However, existing solutions still face challenges such as high latency and significant overhead in pseudonym issuance, making it difficult to balance vehicle location privacy protection with network efficiency. Additionally, hybrid region-based strategies are often highly dependent on RSUs, which are expected to provide full coverage. In practice, however, this comprehensive coverage is not always feasible [22], which incurs high costs and limits their application. Furthermore, managing multiple pseudonyms demands considerable storage allocation on OBUs, which in turn complicates management for the TA and increases the storage and computation overhead on RSUs. Due to these issues, protecting vehicle location privacy both effectively and efficiently becomes a challenging task.

To address these challenges, this paper proposes a fog-based pseudonym management and change scheme that does not rely on RSU to protect privacy. The main contributions of this paper are summarized as follows:

(1) This paper proposes a fog-based pseudonym management scheme to reduce the communication and storage overhead associated with pseudonym distribution and accountability for RSUs. By leveraging the low-latency of fog computing, the scheme migrates the core management processes to the network's fog layer, thereby reducing service response time and alleviating the burden on the TA.

(2) A vehicle-centered, RSU-independent collaborative pseudonym-change mechanism is introduced. Similarity of driving behaviors is quantified using dynamic states such as driving direction, speed and position to identify and analyze vehicles that are about to enter the pseudonym change area. Subsequently, the selected vehicles invoke a proven-secure vehicular authentication and authenticated key-establishment protocol to establish secure communication for the cooperative pseudonym-change process.

(3) The simulation environment for our scheme is built upon the reliable vehicular network framework composed of Veins, OMNeT++, and SUMO. Leveraging a real-world map and actual traffic flow data, we conducted extensive simulations to evaluate and validate the proposed scheme.

The experimental results show that the proposed scheme significantly enhances anonymity while substantially reducing the risk of being tracked, offering a practical and scalable solution for VANETs.

The remainder of the paper is organized as follows. Section 2 introduces the related work. Section 3 presents the design of the pseudonym scheme. Then, Section 4 analyzes the security of the proposed scheme and Section 5 evaluates the performance. Finally, section 6 concludes the paper and discusses future work.

2 Related work

To protect the privacy and security of vehicle communications in VANETs, various pseudonym change strategies have been proposed. The primary goal of these strategies is to determine when and where vehicles change their pseudonyms, ensuring unlinkability between different pseudonyms and offering effective protection against pseudonym linking attacks [23]. This section provides an overview of these solutions for pseudonym change and management in VANETs.

2.1 Pseudonym change mechanisms dependent on RSU online services

Li et al. [24] employed a differential privacy index mechanism for pseudonym selection, enhancing anonymity but also increasing computational complexity. This added complexity negatively affects system performance. Additionally, the selection of differential privacy parameters is extremely sensitive, requiring careful adjustment. Chen et al. [25] proposed a scheme that enhances anonymity and reduces traceability. However, its effectiveness largely depends on the availability of RSUs and relies significantly on the accuracy of real-time traffic context prediction. Khodaei et al. [26]suggested that selecting relay vehicles to simulate non-existent relay vehicles can enhance user privacy. However, this method increases the communication overhead, and its effectiveness may be limited in areas with low vehicle density. Boualouache et al. [27] presented a scheme that can be easily deployed on widespread roadside infrastructure, offering secure pseudonym changing and management. However, this approach imposes stringent security requirements and heavily relies on extensive roadside infrastructure support. Li et al. [28] proposed a mix zone scheme by analyzing the frequently parked locations of vehicles. Within this zone, silent periods can be implemented to enhance privacy protection. However, broadcasting previous pseudonyms before the silent period may pose potential privacy leakage risks. I, Memon et al. [29] proposed that RSUs issue temporary keys to ensure communication among vehicles within the area. While this enhances security, it also increases communication delays and costs, bringing additional strain on RSUs.

The pseudonym change mechanisms dependent on RSU enhance vehicle privacy protection and security through various strategies. While these mechanisms improve anonymization effectiveness, they also come with several challenges. This dependence leads to increased communication delays, higher computational costs, and significant infrastructure requirements. Thus, while these mechanisms enhance privacy protection, they also face challenges related to cost, efficiency, and infrastructure development.

2.2 Pseudonym change mechanisms independent of RSU online service

Haider, D et al. [30] proposed a scheme that reduces reliance on external services, offering maximum location confidentiality. However, it is highly dependent on a central server, and reliance on it may pose security risks if the algorithm or the server is compromised. Zhang et al. [31] presented a scheme that alleviates the burden on the certificate authority, diminishes the delay in pseudonym requests, and elevates overall efficiency. However, security concerns in the fog computing environment require additional protective measures, and the reliability and credibility of fog nodes must be ensured. P. Mdee, M et al. [7] proposed a solution that operates independently of infrastructure, thereby enhancing the system's flexibility and scalability. However, frequent pseudonym changing increases system complexity and communication overhead. Hayat, Z et al. [32] introduced a method leveraging differential privacy to enhance data security and privacy protection. However, the incorporated noise impacts data accuracy, and in high-density scenarios, transmitting numerous perturbed messages escalates network load. To mitigate infrastructure dependency, A.P. Mdee et al. [33] proposed a method where pseudonyms are managed within the CA's pseudonym pool, with unused pseudonyms being cycled out for new ones to improve resource utilization. However, this approach increases the initial registration burden and the pseudonym usage threshold must be set appropriately to avoid compromising the timeliness and effectiveness of pseudonym replacement. Pranav Kumar Singh et al. [34] proposed that vehicles collaborate to exchange pseudonyms, preventing location tracking by service providers.However, the scheme must ensure sufficient randomness and unpredictability to avoid potential attacks. Moreover, it increases communication overhead.

The pseudonym change mechanism independent of RSUs offers an alternative solution to protecting vehicle privacy, but it also brings a series of challenges. These mechanisms often rely on a central server to manage pseudonyms, increasing security and computation demands on the central server and depending heavily on security communication and storage space. This makes privacy protection more difficult in sparsely populated traffic areas. While RSU-independent pseudonym change mechanism theoretically addresses privacy concerns, its practical implementation requires careful consideration of security, efficiency, resource consumption, and other factors.

The comparison among the existing pseudonym change schemes in VANETs is shown in Table 1.

3 The proposed scheme

This section provides a comprehensive overview of the system model and defines the fog-based pseudonym management strategy. Additionally, this section explores a vehicle-centered pseudonym-change scheme that is independent of RSUs. The relevant symbols and descriptions are presented in Table 2.

3.1 System model

The system model mainly consists of the user layer, fog layer, and cloud layer, as shown in Figure 1

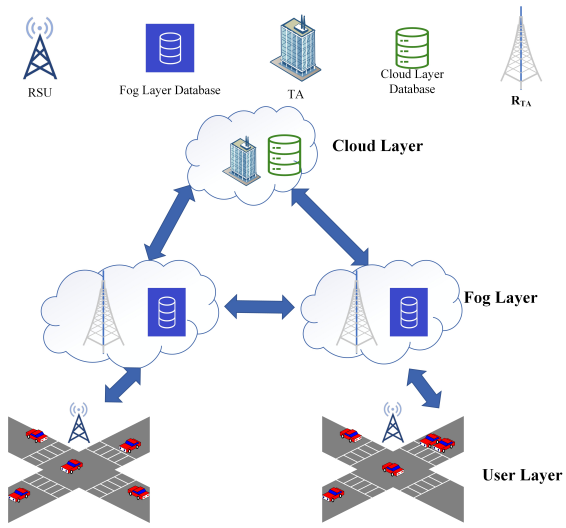

Figure 1: System model

Cloud Layer: The Trusted Authority (TA) is deployed on cloud servers within government departments and servers as the core of the cloud layer. The TA is assumed to be trustworthy. It consists of a long-term registration module and a pseudonym generation module. The registration module is responsible for verifying the real identity of vehicles. Once verified, the pseudonym generation module generates and assigns long-term pseudonyms to the verified vehicles. The TA records the initial mapping between vehicles' real identity and their long-term pseudonyms, securely storing this information in the vehicle-long-term pseudonym mapping database. In addition, the TA also continuously monitors all entities in VANETs. If the TA receives evidence of misbehavior from the Trusted Fog Node (FN) in the fog layer, the TA promptly performs detection and removes the malicious entity from VANETs to ensure the security and stability.

Fog Layer: The fog layer, situated between the user and cloud layers, is managed by the cloud layer. The Trusted Fog Node (FN) deployed on this layer is similar to the RSU in function, but with a broader communication range. The FN generates, stores, and distributes short-term pseudonyms to vehicles entering its coverage area for the first time. This mechanism significantly reduces the burden on the cloud layer in terms of pseudonym management and storage. The FN consists of three components: the short-term pseudonym pool, the local log database, and the mapping relationship database. The short-term pseudonym pool generates and securely stores short-term pseudonyms, which serve as temporary identifiers for vehicles to ensure anonymity while supporting system management and tracking. The local log database is responsible for recording misbehavior involving RSU and vehicles in the region, storing communication information with adjacent fog layers and enforcing mutual supervision. If misbehavior is detected, the database blacklists the offending vehicle or RSU and uploads the information to the cloud layer, where the vehicle's real identity is recovered for legal accountability and removal from the system. The mapping relationship database records the mapping between long-term and short-term pseudonyms, enabling rapid identification of vehicles when necessary.

User Layer: The user layer consists of vehicles, Roadside Units (RSUs), and eavesdropping devices deployed by attackers. The primary responsibilities of RSUs are to authenticate vehicles within their coverage area, broadcast service information, and upload the authentication information to the fog layer, and assist FN in the distribution of short-term pseudonyms. Vehicles register with the TA by submitting their real identities to obtain long-term pseudonyms. Each vehicle is equipped with electronic devices such as a Global Positioning System (GPS), an Onboard Unit (OBU) and various sensors. The OBU is used to store both short-term and long-term pseudonyms. Attackers deploy eavesdropping devices, primarily trackers. A tracker can be a server with significant resources and the ability to implement tracking algorithms. A tracker in the system poses various threats, including eavesdropping, tracking, and profile generation, even if it operates passively. In the proposed scheme, the tracker is global, meaning it covers the entire network and remains passive. This type of attacker is referred to as a Global Passive Attacker (GPA) in recent literature [29]. The GPA can eavesdrop on beacon messages but cannot modify the content of the captured messages.

Figure 2 shows the architectural components of the pseudonym system in VANETs.

3.2 Pseudonym management

Before joining the VANETs, a vehicle must register with the TA. Upon registration, the TA assigns the vehicle a private key SK, a public key PK, a long-term pseudonym PID_L. PID_L is associated with the vehicle's real identity (e.g., license plates, owner information), enabling the TA to accurately identify the vehicle based on PID_L.

RSUs regularly broadcast notifications about the provided services and distribution of short-term pseudonyms from FNs. When a vehicle enters the FN managed area for

Table 1: Pseudonym change schemes comparison

Scheme	RSU-independent	Pseudonym Management	Syntactic Linking Attack	Semantic Linking Attack	Radio Silence	Selfishness	Dynamic Decision
(24)	×	✓	×	×	×	×	✓
(25)	×	✓	×	×	✓	×	✓
(26)	×	✓	✓	✓	×	×	✓
(27)	×	✓	✓	✓	×	✓	✓
(28)	×	✓	✓	✓	✓	×	✓
(29)	×	✓	✓	✓	×	×	×
(30)	✓	✓	✓	✓	×	×	✓
(31)	✓	✓	×	×	×	×	✓
(7)	✓	✓	✓	✓	×	×	✓
(32)	✓	✓	✓	✓	×	×	✓
(33)	✓	×	✓	✓	×	×	✓
(34)	✓	✓	✓	✓	×	×	✓

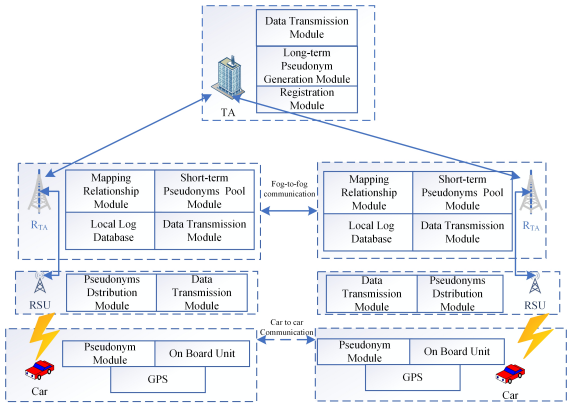

Figure 2: System architecture

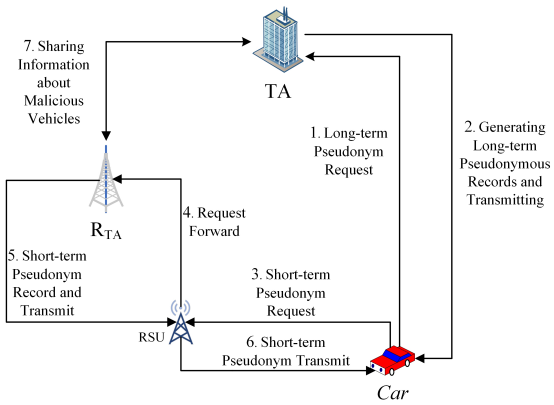

Figure 3: Pseudonym acquisition process

the first time and receives the broadcast announcement, it sends a message $msg_{V_a} = \{\text{Req}_i \parallel PID_L \parallel T\}$ to the nearest RSU, where Req_i is the request identifier for the short-term pseudonym, PID_L is the long-term pseudonym distributed by the TA, and T is the timestamp. After validating the PID_L, the RSU forwards msg_{V_a} to the nearest FN. The FN generates a set of short-term pseudonyms $PIDs = \{\text{PID}_1, \text{PID}_2, \text{PID}_3, \ldots, \text{PID}_i\}$ based on the pseudonym generation reward function described in [31]. These short-term pseudonyms are sent back to the registered vehicle via the RSU. Within this fog area, the vehicle uses the short-term pseudonyms for communication. The detailed procedure is shown in Figure 3.

If the vehicle leaves the current fog area, it must re-authenticate through a new RSU to a new FN to obtain a new set of short-term pseudonyms for subsequent communication. If a malicious vehicle is detected by a FN, relevant information is stored in the local log of the FN and uploaded to the TA. The real identity of the vehicle must be revealed by the TA based on the vehicle-long-term pseudonym mapping database. Subsequently, long-term pseudonyms of vehicle are then added to the blacklist and broadcast the blacklist to all FN in VANETs. Once FNs receive the blacklist, they no longer distribute pseudonyms to malicious vehicles.

3.3 Vehicle similarity calculation

Considering the difficulty in the widespread deployment of RSUs, this paper proposes an RSU-independent pseudonym change scheme. The scheme fully considers potential attacks and tracking capabilities of attackers and designs a dynamic area division strategy based on vehicle similarity.

In this scheme, the vehicle initiating the pseudonym change process is defined as Va. The set of all vehicles within the communication coverage range of Va is denoted

Table 2: Symbols and descriptions

Symbol	Description
V_{Range}	The set of vehicles within the communication range
PID_L, PID_i	The L-th long-term pseudonym and the i-th short-term pseudonym
T	Timestamp
SK_i	The private key of entity i
PK_i	The public key of entity i
$\Pi_{\text{Auth/AKE}}$	Secure authentication and authenticated key establishment
K_{ij}	Fresh session/group key
V_{change}	The group of vehicles undergoing a pseudonym change
C	Set of eligible candidate vehicles
$Sim_{D,P,L}$	Speed, Direction, Position and Comprehensive Similarity
N_t	Threshold for pseudonym-change participation
N_p	Actual number of eligible candidate vehicles, where $N_p = \|C\|$
v_i	The speed of vehicle i
d_i	The direction of vehicle i
(x_i, y_i)	The position of vehicle i
α	Short-term pseudonym lifetime threshold
w_1, w_2, w_3	Weights of speed, direction, and position similarity, respectively
s	Current road scene, where $s \in \{\text{highway}, \text{urban}\}$

as V_{Range}. Each vehicle in V_{Range} computes its driving similarity with Va using the vehicle's driving state, specifically focusing on position, speed, and driving direction. These indicators can be directly obtained from the beacon broadcast by the vehicle, enabling precise quantification of the similarity of vehicle driving states.

The driving similarity contains three dimensions: speed similarity, direction similarity, and position similarity. These three components are combined through a weighted summation to obtain the overall driving-state similarity. Using V_a and V_b as illustrative examples, the similarity measures are defined as follows:

(1) Speed similarity sim_v between V_a and V_b:

$$\begin{aligned} sim_v &= 1 - \frac{|v_a - v_b|}{\Delta v}, \\ \Delta v &= v_{\max} - v_{\min}\ (v_{\min}, v_{\max} \in V_{\text{Range}}). \end{aligned} \tag{1}$$

where v_a and v_b are the speeds of V_a and V_b, respectively. v_{max} and v_{min} are the maximum and minimum speeds in V_{Range}, respectively.

(2) Direction similarity sim_D between V_a and V_b:

$$\begin{aligned} sim_D &= 1 - \frac{|d_a - d_b|}{\Delta d}, \\ \Delta d &= d_{\max} - d_{\min}\ (d_{\min}, d_{\max} \in V_{\text{Range}}). \end{aligned} \tag{2}$$

where d_a and d_b are the movement directions of V_a and V_b along the horizontal coordinate, and Δd is the difference between the maximum d_{max} and minimum d_{min} values of the direction range of all vehicles, i.e., the total span of direction changes.

(3) Position similarity sim_L between V_a and V_b:

$$\begin{aligned} sim_L &= 1 - \frac{\sqrt{(x_a - x_b)^2 + (y_a - y_b)^2}}{\Delta L} \\ \Delta L &= \max_{V_a, V_b \in V_{\text{Range}}} \sqrt{(x_a - x_b)^2 + (y_a - y_b)^2}. \end{aligned} \tag{3}$$

where (x_a, y_a) and (x_b, y_b) are the position coordinates of V_a and V_b, and ΔL represents the maximum physical distance between any two vehicles in the set. It serves as a reference for evaluating the relative proximity of vehicles.

Based on the speed similarity, direction similarity, and position similarity between V_a and V_b, the proposed scheme designs a comprehensive measurement function to evaluate the similarity of their driving states. This measurement function aims to reflect the degree of proximity in the driving behaviors of the two vehicles. The more similar the driving states of the two vehicles, the more consistent their behavior patterns are, making it difficult for attackers to gather information through vehicle tracking and infer the relationship between the vehicles and their previous pseudonyms. The overall driving state similarity, denoted as sim, is calculated by combining the above three dimensions:

$$\begin{aligned} sim_c &= w_1 \times sim_v + w_2 \times sim_D + w_3 \times sim_L \\ w_1 &+ w_2 + w_3 = 1 \end{aligned} \tag{4}$$

The weights w_1, w_2, and w_3 are selected according to the current road and traffic conditions [24], allowing the comprehensive similarity to reflect the relative importance of speed, direction, and position. Accordingly, the weight tuple is selected from the following two preset vectors:

$$(w_1, w_2, w_3) = \begin{cases} (w_1^H, w_2^H, w_3^H), & s = \text{highway}, \\ (w_1^U, w_2^U, w_3^U), & s = \text{urban}. \end{cases}$$

For each $q \in \{H, U\}$, the preset vector satisfies $w_1^q + w_2^q + w_3^q = 1$ and $w_i^q > 0$. The values are selected offline and remain fixed during a pseudonym-change round. Thus, scene adaptation requires only a road-type lookup and a weight-vector selection, preserving the interpretability and low computational cost of the linear similarity function.

3.4 Short-term pseudonym change

After system initialization, each vehicle registers with the TA and receives a long-term pseudonym PID_L. When a vehicle enters a FN's coverage area and is successfully authenticated, it receives a short-term pseudonym pool $PIDs = \{PID_1, PID_2, PID_3, \ldots, PID_i\}$ issued by the FN. The vehicle initially uses PID_1. Each short-term pseudonym remains active until its elapsed usage time reaches the lifetime threshold α, at which point the vehicle initiates the cooperative pseudonym-change process and switches to a new unused short-term pseudonym.

When the elapsed usage time of the current short-term pseudonym reaches the lifetime threshold α, the vehicle initiates a cooperative pseudonym-change process. To establish an initiator-centric mixed area within the fog domain, surrounding vehicles are first evaluated according to their driving-state similarity. Vehicles with highly similar driving states are more suitable for simultaneous pseudonym changes because their trajectories are more difficult for an attacker to distinguish. After the eligible participating vehicles are determined, the initiator and the participating vehicles invoke the selected secure authentication and authenticated key-establishment protocol to establish a fresh shared group key $K_{V_a, V_{\text{change}}}$. Upon successful authentication and key establishment, the vehicles jointly establish the pseudonym-change area for the subsequent synchronized pseudonym change.

Figure 4 shows the timing diagram of the pseudonym change. At time t, vehicles are in the normal driving state, and the initiating vehicle broadcasts a request message. At time $t + 1$, vehicles receiving the request message validate the initiating vehicle's short-term pseudonym, quantitatively compute the comprehensive similarity (see the Vehicle Similarity Calculation section for details), and decide whether to participate. Vehicles willing to participate then respond to the initiating vehicle. At time $t + 2$, after receiving the responses, the initiating vehicle verifies the responding vehicles and recomputes their similarity values for comparison. The eligible vehicles are then selected according to their similarity values and the preset participation threshold N_t to form the pseudonym-change group V_{change}. The initiating vehicle and the selected participants subsequently invoke the selected secure authentication and key-establishment protocol to establish a fresh shared group key $K_{V_a, V_{\text{change}}}$ and jointly establish the pseudonym-change area. At time $t + 3$, the initiating vehicle and the participating vehicles simultaneously change their short-term pseudonyms. Then, all vehicles drive with the new short-term pseudonyms. This cooperative change mechanism significantly raises the difficulty for an attacker to track the real identities and locations of the vehicles.

(1) When the elapsed usage time of the current short-term pseudonym of vehicle V_a reaches the lifetime threshold α, V_a initiates a cooperative short-term pseudonym-change process. V_a first broadcasts a request message

$$msg_{V_a} = \{\text{Req}_i \parallel PID_i \parallel T\},$$

where Req_i is the pseudonym-change request, PID_i is the current short-term pseudonym of V_a, and T is the timestamp. This message indicates that V_a is actively seeking cooperation with other vehicles to collaboratively establish a pseudonym change mixed area to complete the pseudonym change process.

(2) When a vehicle in V_{Range} receives V_a's request, it validates PID_i. If validation succeeds, the vehicle selects (w_1, w_2, w_3) according to the current road scene s and calculates its driving similarity sim_c to V_a using Equations (1)–(4). Based on sim_c and the remaining time before its own current short-term pseudonym reaches the lifetime threshold α, it decides whether to participate. A responding vehicle $V_b \in V_{\text{Range}}$ sends the message $msg_{V_b} = \{sim_{ab} \parallel T \parallel PID_{V_b}\}$ back to V_a, where sim_{ab} is the driving similarity between V_a and V_b, T is the timestamp, and PID_{V_b} is the current short-term pseudonym used by vehicle V_b. If a vehicle does not respond, the message msg_{V_a} is discarded.

(3) Let the eligible-candidate set initially be $C = \varnothing$. After receiving a response, V_a verifies the legitimacy of the responding vehicle's short-term pseudonym. If verification is successful, V_a sets the received similarity as M_1, recomputes the similarity M_2 between itself and the responding vehicle, and compares M_1 with M_2. This comparison prevents vehicles from intentionally providing false similarity information to disrupt the normal pseudonym-change process. If $M_1 = M_2$, the responding vehicle is regarded as an eligible candidate and added to C.

(4) Let $N_p = |C|$ denote the actual number of eligible candidates in the candidate set C constructed in Step (3). The FN determines the participation threshold N_t according to service and environmental requirements. If $N_p \geq N_t$, a Top-N_t procedure selects the N_t vehicles with the highest similarity values from C to form V_{change}. If $0 < N_p < N_t$, V_{change} is set to C, and the k-anonymity compensation mechanism is applied with the target anonymity size $k = N_t$. A total of $N_t - N_p$ auxiliary pseudonym-change notifications are distributed among the N_p participants. These auxiliary notifications use the same observable format as the actual pseudonym-change notifications and

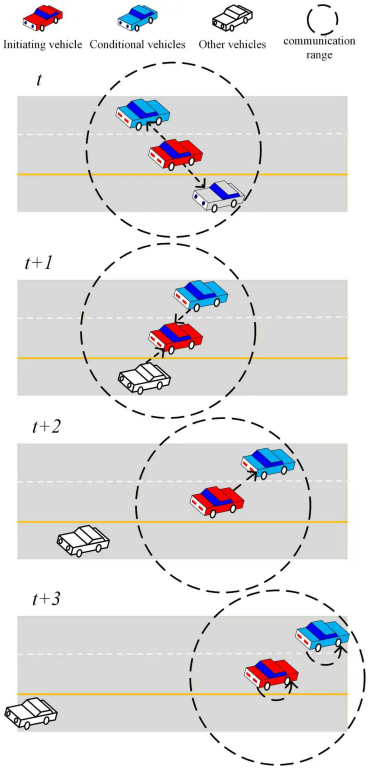

Figure 4: Vehicle pseudonym-change example

are transmitted within the same change interval. If $N_p = 0$, no cooperative pseudonym-change event is executed. This mechanism increases the difficulty for an attacker to distinguish the actual participants of a successfully executed pseudonym-change event.

(5) After V_a determines V_{change}, the initiating vehicle and the selected participants invoke a selected proven-secure vehicular authentication and authenticated key-establishment protocol. We abstract this replaceable protocol interface as

$$(accept, K_{V_a,V_{\text{change}}}) \leftarrow \Pi_{\text{Auth/AKE}}(V_a, V_{\text{change}}, PID, T). \quad (5)$$

Here, PID denotes the current valid short-term pseudonyms of the participating vehicles used as their authentication identifiers in the selected protocol, $accept$ denotes successful authentication and key establishment, and $K_{V_a,V_{\text{change}}}$ is the resulting fresh session/group key. If authentication or key establishment fails, the current cooperative pseudonym-change event is terminated.

(6) Once the selected protocol has successfully authenticated the eligible vehicles and established the shared key $K_{V_a,V_{\text{change}}}$, the initiating vehicle and the participating vehicles jointly establish the pseudonym-change area. Within this area, communications are encrypted using $K_{V_a,V_{\text{change}}}$. During the validity period of the mixed area, V_a and all vehicles in V_{change} simultaneously select new, unused short-term pseudonyms from their respective pseudonym pools, thereby achieving a synchronized pseudonym change and further protecting location privacy. The overall flowchart is shown in Figure 5.

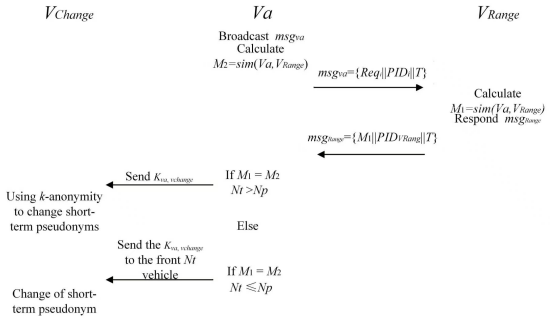

Figure 5: Pseudonym-change flowchart

4 Security analysis

4.1 Security assumptions and adversarial model

The TA and authorized FNs are trusted to correctly perform vehicle registration, pseudonym issuance, mapping maintenance, revocation, and accountability. Each legitimate vehicle is assigned a short-term pseudonym pool by the FN, whose pseudonyms are used sequentially with only one serving as the active communication identity at a time. Participating entities are assumed to be loosely time-synchronized for message-freshness verification.

We consider both external and internal security threats. An external network attacker may intercept, inject, modify, or replay messages over public wireless channels, but cannot compromise the trusted TA or authorized FNs or break the security guarantees of the selected proven-secure authentication and authenticated key-establishment protocol. A malicious registered vehicle may attempt to simultaneously misuse multiple short-term pseudonyms from its assigned pool to impersonate multiple vehicles, which is considered in the Sybil-attack analysis.

We also consider a Global Passive Adversary (GPA) that continuously observes publicly broadcast beacons and attempts to reconstruct vehicle trajectories. The GPA knows the proposed pseudonym-change strategy and can observe short-term pseudonyms, timestamps, positions, speeds, and directions, but cannot modify messages. Its objective is to link pseudonyms before and after a change and continuously track the target vehicle. Accordingly, the external attacker is considered in secure communication and replay-attack resistance, while the GPA is considered in location privacy, pseudonym unlinkability, and the formal privacy analysis.

4.2 Security properties

Traceability and accountability: Traceability and accountability ensure that the real identity of a malicious vehicle can be traced by authorized entities when misbehavior occurs, while legitimate vehicles remain anonymous during normal communication. The proposed scheme introduces a mapping mechanism between long-term and short-term pseudonyms at the fog layer, while establishing

Algorithm 1 Fog-assisted pseudonym-change algorithm

Input: Initiating vehicle V_a; neighboring set V_{Range}; short-term pseudonym lifetime threshold α; participation threshold N_t; current road scene $s \in \{\text{highway, urban}\}$; current timestamp T; fog-issued short-term pseudonym pools; protocol $\Pi_{\text{Auth/AKE}}$
Output: Pseudonym-change group V_{change} and a synchronized pseudonym-change result

The FN supplies the short-term pseudonym pool, the corresponding lifetime threshold α, and the regional participation threshold N_t.
if the elapsed usage time of the current short-term pseudonym of V_a has not reached α **then**
  return NoChange
end if
$(w_1, w_2, w_3) \leftarrow$ SelectWeights(s).
V_a broadcasts $msg_{V_a} = \{Req_i \parallel PID_i \parallel T\}$.
Initialize the eligible-response set $C \leftarrow \varnothing$.
for all $V_b \in V_{\text{Range}}$ receiving msg_{V_a} **do**
  if PID_i is legitimate **then**
    V_b calculates $sim_c(a, b)$ using Equations (1)–(4).
    if V_b intends to participate based on $sim_c(a, b)$ and the remaining time before its current pseudonym reaches α **then**
      V_b returns $msg_{V_b} = \{sim_{ab} \parallel T \parallel PID_{V_b}\}$.
      V_a sets $M_1 \leftarrow sim_{ab}$, verifies PID_{V_b}, and recomputes $M_2 \leftarrow sim_c(a, b)$.
      if $M_1 = M_2$ **then**
        $C \leftarrow C \cup \{V_b\}$.
      end if
    end if
  end if
end for
$N_p \leftarrow |C|$.
if $N_p \geq N_t$ **then**
  $V_{\text{change}} \leftarrow \text{Top}_{N_t}(C, sim_c)$ using a Top-N_t selection procedure.
else if $0 < N_p < N_t$ **then**
  $V_{\text{change}} \leftarrow C$ and set the target anonymity size to $k = N_t$.
  Distribute exactly $N_t - N_p$ auxiliary notifications among the N_p participants.
else
  return NoCooperativeChange
end if
$(accept, K_{V_a, V_{\text{change}}}) \leftarrow \Pi_{\text{Auth/AKE}}(V_a, V_{\text{change}}, PID, T)$.
if $accept =$ false **then**
  return NoCooperativeChange
end if
V_a and all vehicles in V_{change} simultaneously select new, unused short-term pseudonyms.
return $\langle V_{\text{change}}, \text{Success}\rangle$.

an association between vehicles' identities and long-term pseudonyms at the cloud layer. For instance, when surrounding vehicles report a malicious vehicle to the nearby RSU, the RSU forwards this report to a fog node. The fog node then uses the short-term pseudonym provided by the reporting vehicles to find the corresponding long-term pseudonym and forwards it to the TA. The TA, utilizing its records, retrieves the true identity associated with the long-term pseudonym, enabling it to hold the malicious vehicle accountable. Therefore, the proposed scheme provides traceability and accountability.

Location privacy: Location privacy ensures that an attacker cannot continuously track a vehicle and infer its real-time location and driving trajectory from periodically broadcast information. An attacker may eavesdrop on beacons broadcast by vehicles to track their locations and predict their future positions based on the vehicles' current speed and lane [1]. To defend against this threat, the initiating vehicle collaborates with surrounding vehicles to establish a pseudonym-change area. After the eligible vehicles are determined, the initiating vehicle and the participating vehicles invoke the selected secure authentication and key-establishment protocol to establish a fresh shared group key $K_{V_a, V_{\text{change}}}$. Within this area, communications are encrypted using $K_{V_a, V_{\text{change}}}$, preventing the attacker from obtaining valuable information exchanged during the pseudonym-change process. Furthermore, to increase the difficulty of tracking, all participating vehicles within this encrypted communication area change their short-term pseudonyms simultaneously. Moreover, the driving states of the participating vehicles are similar, and their behavior patterns are consistent, making it difficult for attackers to distinguish the vehicles and infer the relationship between their previous and new pseudonyms. The proposed scheme not only confuses attackers but also greatly protects the vehicles' driving trajectories and real-time location privacy.

Secure communication: Secure communication ensures that the messages exchanged during the cooperative pseudonym-change process are protected against unauthorized disclosure and modification [35]. In the proposed scheme, after the eligible vehicles are determined, the initiating vehicle and the participating vehicles invoke the selected proven-secure vehicular authentication and authenticated key-establishment protocol. As abstracted in Equation (5), successful execution of the protocol authenticates the participating entities and establishes a fresh shared group key $K_{V_a, V_{\text{change}}}$. The subsequent coordination messages within the pseudonym-change area are protected us-

ing the secure communication mechanism associated with the selected protocol and the established group key. Therefore, under the security guarantees of the selected protocol, an external network attacker cannot obtain the protected coordination information or successfully inject or modify such messages without being detected. Moreover, a fresh group key is established for each new cooperative pseudonym-change event, avoiding the reuse of the same session key across successive events. Therefore, the proposed scheme provides secure communication during the cooperative pseudonym-change process by relying on the security guarantees of the selected authentication and authenticated key-establishment protocol.

Vehicle anonymity: Vehicle anonymity ensures that the real identity of a vehicle cannot be directly obtained by other communicating entities or external observers during normal communication. In all communication exchanges, the sending vehicle uses only its pseudonym for transmitting information, which effectively ensures that the receiving party cannot obtain the real identity of the sender [36]. In the proposed scheme, the sending vehicle only uses its short-term pseudonym PID_i for communication, ensuring that the receiver cannot obtain the real identity of the sender. The receiver verifies the validity of the sender based on PID_i. The real identity of the vehicle is never disclosed during communication. The real identity of a vehicle can only be traced with the cooperation of both the TA and the FN, thereby ensuring vehicle anonymity.

Pseudonym unlinkability: Pseudonym unlinkability makes it difficult for an attacker to link messages sent by the same vehicle using two different pseudonyms [37]. In the proposed scheme, vehicles first acquire a long-term pseudonym, and when they enter a new fog area, they register through the nearest RSU as an intermediary to obtain short-term pseudonyms for communication within the fog area. Each short-term pseudonym is assigned the lifetime threshold α, and once its elapsed usage time reaches α, the vehicle switches to a new short-term pseudonym. To further enhance unlinkability, the proposed scheme designs a vehicle-centered pseudonym-change area. In this area, all participating vehicles simultaneously change to new short-term pseudonyms. Moreover, the driving states of the vehicles within the area are similar, and their behavior patterns are consistent, making it difficult for attackers to track the vehicles and deduce the correlation between their current and previous pseudonyms. The proposed scheme significantly reduces the risk of attackers linking two messages sent by the same vehicle under different pseudonyms, as the participating vehicles change their pseudonyms within nearly the same interval, making it difficult for attackers to continuously track and link their trajectories. Therefore, the proposed scheme achieves pseudonym unlinkability and reduces the risk of continuous vehicle tracking.

Replay-attack resistance: Replay attacks can undermine the integrity of the communication by allowing an attacker to capture and retransmit messages, which may lead to unauthorized actions or breach of privacy [15]. To effectively prevent replay attacks, the proposed scheme embeds a timestamp T in all transmitted messages and encrypts it along with the message. The proposed scheme assumes that the clocks of all entities are loosely synchronized. By embedding the timestamp T, the receiving entity checks the embedded timestamp T and determines whether the message is within the accepted freshness interval. Messages containing expired or previously processed timestamps are rejected and cannot be reused to trigger a new pseudonym-change operation. The proposed scheme enhances the security of the system and prevents attackers from successfully performing replay attacks by retransmitting old messages.

Sybil-attack resistance: A Sybil attack occurs when an attacker uses multiple pseudonyms to impersonate several vehicles simultaneously [22]. To defend against this threat, the proposed scheme restricts the simultaneous use of short-term pseudonyms assigned to each legitimate vehicle. The FN assigns a pool of unique short-term pseudonyms to each legitimate vehicle, and these pseudonyms are used sequentially, with only one serving as the vehicle's active communication identity at any given time. Before participating in a cooperative pseudonym-change process, the current short-term pseudonyms of the initiating and responding vehicles are validated. Once the current pseudonym is replaced by a new one, the previous pseudonym is no longer used as the vehicle's active communication identity. Therefore, a malicious registered vehicle cannot legitimately participate in the same communication state using multiple simultaneously active pseudonymous identities, thereby providing resistance to pseudonym-based Sybil attacks.

4.3 Formal anonymity and privacy analysis

The above analysis explains the privacy properties of the proposed scheme from the perspective of the pseudonym-change mechanism. To further quantify the anonymity and unlinkability achieved during a cooperative pseudonym change, this section analyzes the observable anonymity set and the linking probability of the GPA.

Let A_t denote the anonymity set observed by the GPA during a successfully executed pseudonym-change event at time t. When $N_p \geq N_t$, Algorithm 1 selects the N_t vehicles with the highest driving-state similarity to participate in the synchronized pseudonym change. When $0 < N_p < N_t$, the participating vehicles transmit a total of $N_t - N_p$ additional pseudonym-change notifications. These notifications use the same observable format and are transmitted within the same pseudonym-change interval. Therefore, from the perspective of the GPA, N_t plausible pseudonym-change records are observed in both cases, and the observable anonymity-set size is

$$|A_t| = \begin{cases} N_t, & N_p \geq N_t, \\ N_p + (N_t - N_p), & 0 < N_p < N_t. \end{cases} \tag{6}$$

Thus, both cases yield $|A_t| = N_t = k$. Accordingly, for each successfully executed cooperative pseudonym-change

event with $N_p > 0$, the observable anonymity-set size reaches the target value $k = N_t$. If $N_p = 0$, Algorithm 1 terminates the current round without executing a cooperative pseudonym change, and no anonymity guarantee is claimed for that round.

Let p_i denote the GPA's posterior probability that candidate $i \in A_t$ corresponds to the continuation of the target vehicle's previous pseudonym. The uncertainty of the GPA can be measured using the Shannon entropy H_t and the normalized degree of anonymity D_t:

$$H_t = -\sum_{i=1}^{|A_t|} p_i \log_2 p_i, \quad D_t = \frac{H_t}{\log_2 |A_t|}, \quad 0 \le D_t \le 1 \quad (7)$$

A larger D_t indicates greater uncertainty for the GPA in distinguishing the target vehicle from the other candidates. The driving-state similarity calculation in Equations (1)–(4) selects vehicles whose observable motion states are close to those of the initiating vehicle. This reduces the distinguishability among the participating vehicles and makes the GPA's posterior distribution closer to the uniform distribution.

In the ideal case where all candidates are equally indistinguishable to the GPA,

$$p_i = \frac{1}{|A_t|},$$

and the entropy reaches its maximum value $H_t = \log_2 |A_t|$, yielding $D_t = 1$. Under this maximum-anonymity condition, the GPA's optimal one-shot linking probability is

$$P_{\text{link}}(t) = \max_i p_i = \frac{1}{|A_t|} = \frac{1}{N_t} = \frac{1}{k}. \quad (8)$$

Therefore, increasing the target anonymity-set size k decreases the GPA's linking probability under the ideal indistinguishability condition. In practical traffic scenarios, the posterior distribution may deviate from the uniform distribution; therefore, the entropy-based anonymity degree D_t, together with the traceability and confusion metrics evaluated in Section 5, is used to characterize the actual privacy protection achieved by the proposed pseudonym-change mechanism.

5 Performance evaluation

5.1 Complexity and scalability

The computational overhead of the proposed scheme can be divided into vehicle-side pseudonym-change processing and fog-node pseudonym-management processing.

Vehicle-side computational complexity: Let d denote the number of neighboring vehicles within the communication range of an initiating vehicle. For each candidate vehicle, the proposed scheme evaluates speed, direction, and position similarities and combines them using a weighted linear function. Since each similarity calculation requires only a constant number of arithmetic operations, processing all neighboring candidates requires $O(d)$ time. After collecting the eligible responses, the initiating vehicle performs a Top-N_t candidate selection. With a fixed-size Top-N_t candidate set, this operation requires $O(d \log N_t)$ time. Since N_t is a preset participation threshold and does not scale with the number of neighboring vehicles, $\log N_t$ can be treated as a constant. Therefore, the overall vehicle-side pseudonym-change decision scales linearly with the number of neighboring candidates, i.e., $O(d)$. Hence, increasing vehicle density results in linear rather than quadratic growth in the dominant vehicle-side computation.

Message complexity: For each cooperative pseudonym-change event, the initiating vehicle broadcasts one request message and receives at most d response messages from neighboring vehicles. Subsequent coordination is restricted to the selected pseudonym-change group, whose size is bounded by N_t. Excluding the internal operations of the replaceable authentication and authenticated key-establishment protocol, the message complexity associated with the pseudonym-change decision is therefore $O(d + N_t)$, which reduces to $O(d)$ when N_t is fixed. More importantly, participating vehicles select new unused pseudonyms from the short-term pseudonym pools that have already been distributed by the FN. Therefore, an online pseudonym request to the FN or TA is not required for every pseudonym-change event, avoiding repeated infrastructure communication during normal pseudonym changes.

Fog-node storage and processing complexity: Let P_f denote the number of short-term pseudonym records maintained by one FN and q denote the number of short-term pseudonyms allocated in one pool request. The FN maintains the PID_L–PID_i mapping records together with expiration information, revocation states, the local pseudonym pool, and accountability records. With an indexed mapping structure, pseudonym lookup, revocation checking, and state update require $O(1)$ average time per record. Allocating a local pool containing q short-term pseudonyms requires $O(q)$ processing, which remains constant with respect to vehicle density when q is fixed. The FN storage requirement is $O(P_f)$ and therefore increases linearly with the number of short-term pseudonym records managed by the fog node.

Fog-node scalability evaluation: To evaluate the local processing performance of the FN under different pseudonym-record scales, a benchmark was conducted on the hardware platform listed in Table 4. For each value of P_f, 25,000 pseudonym-management requests were processed in each run, and the test was repeated 30 times. Each request included pseudonym lookup, revocation checking, state update, local pseudonym-pool allocation, and accountability recording. The Mean Request Time in Table 3 represents the average local processing time required for one request over the 30 runs, while the throughput represents the number of requests processed per second. Net-

Table 3: Fog-node processing performance

P_f	Mean Request Time (μs)	Throughput (req./s)
1,000	0.0290	3.4×10^7
5,000	0.0325	3.1×10^7
10,000	0.0333	3.0×10^7
50,000	0.0385	2.6×10^7
100,000	0.0376	2.7×10^7
250,000	0.0397	2.5×10^7
500,000	0.0512	2.0×10^7
1,000,000	0.0749	1.3×10^7

work transmission and cryptographic operations were excluded from the measurement.

The results are presented in Table 3. The first column gives the number of short-term pseudonym records managed by one FN. The second column reports the mean local processing time per composite FN request over 30 runs, while the third column reports the corresponding processing throughput. When the number of managed records increases from 10^3 to 10^6, the mean request-processing time increases from 0.0290 μs to 0.0749 μs, corresponding to an increase of approximately 2.58 times. Meanwhile, the throughput remains on the order of 10^7 requests per second throughout the tested range. These results are consistent with the average $O(1)$ complexity of the indexed lookup and state-maintenance operations, indicating that the local pseudonym-management workload does not become a rapidly growing computational bottleneck as the number of records managed by the FN increases.

5.2 Simulation

5.2.1 Simulation setup

The simulation in this paper utilizes the Privacy Extension version (PREXT) of the Veins VANETs simulator, which is based on the OMNET++ platform and integrates with the Simulation of Urban Mobility (SUMO) software to utilize the Veins framework, PREXT implements various privacy protection schemes and incorporates adversary modules capable of eavesdropping on beacon and tracking vehicle movements. These adversary modules are instrumental in evaluating the strength of privacy protection schemes. The attacker uploads the received beacons to a central entity known as the Vehicle Tracker [38]. The specific parameters used in the experiment and their corresponding values are listed in Table 4. In our experiment, The real roadmap of Munich city from SUMO is adopted in the experiment, as shown in Figure 6.

5.2.2 Results and discussion

We compared the proposed scheme with the schemes based on hybrid regions (Khodaei et al. [26], TCCM [8]) and hybrid contexts (CAPS [40], SLOW [42], LSHPC [4]) based on four dimensions: traceability, average pseudonym change per trace, average confusion per pseudonym change, and average confusion per trace.

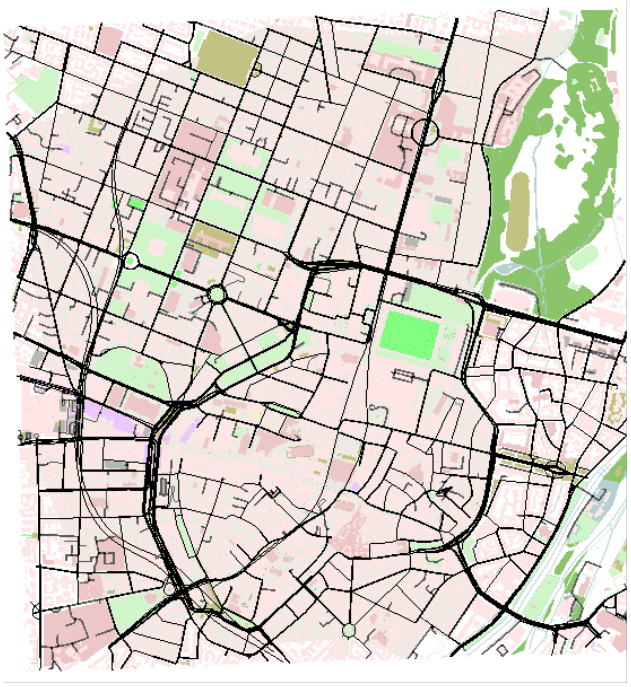

Figure 6: Sumo network based on munich city

Traceability: Traceability quantifies the probability that an attacker can successfully trace more than 90% of a vehicle's trajectory. The objective is to construct a comprehensive audit trail to enable precise identification and retrospective analysis of the target vehicle's historical trajectories and interaction locations [9]. The metric can be calculated using Formula 9:

$$P_i = \frac{1}{AS} \tag{9}$$

P_i denotes the probability that the i-th vehicle is tracked, AS denotes the anonymity-set size, i.e., the number of vehicles among which the target vehicle is indistinguishable to the attacker.

During the traceability measurement, the system automatically filters trajectory segments that have not undergone pseudonym changes. It is important to emphasize that for an attacker to successfully compromise vehicle privacy, continuous tracking of the target vehicle is required, and the effectiveness of such tracking heavily depends on the completeness of the reconstructed trajectory and the minimization of errors. The simulation results are illustrated in Figure 7. As illustrated in the figure, the x-axis represents the number of vehicles, while the y-axis represents traceability. The proposed scheme consistently achieves the lowest traceability among all compared schemes under different traffic densities, indicating stronger resistance to continuous vehicle tracking. From a quantitative standpoint, compared with the Khodaei et al., TCCM, CAPS, and SLOW schemes, the proposed scheme reduces the average traceability by approximately 94.3%, 94.4%, 94.2%, and 91.0%, respectively. Compared with LSHPC, which exhibits the best traceability performance among the five

Table 4: Simulation parameters

Module	Parameter	Values
Hardware Platform	System	13th Gen Intel(R) Core(TM) i5-13400 @ 2.50 GHz, 16 GB RAM
Operating System	Ubuntu	Version 16.04, 64-bit
Simulator	OMNeT++	Version 5.0
	PREXT	Version 1.0
Veins	Version	4.4
	MAC Layer	IEEE 802.11p
	Bit rate (Mbps)	18
	Beacon Rate	10 Hz
	Traffic Simulator	SUMO 0.25.0
	Map	Munich
	Coupling Using	TraCI
	Simulation Time (s)	300
	Number of Vehicles	50, 100, 150, 200
Adversary	Type	Global Passive (Eavesdropping)
	Eavesdropper Communication Range (m)	300
	Eavesdropper Overlap (m)	30
	Tracking Interval (s)	1

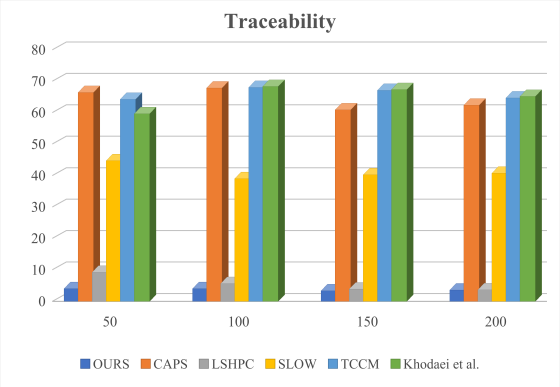

Figure 7: Traceability performance comparison

baseline schemes, the proposed scheme further reduces the average traceability by approximately 33.2%. These results demonstrate that the proposed scheme provides more effective protection against continuous trajectory tracking under different traffic densities.

Pseudonym change per trace: The average pseudonym change metric for each tracking indicates the frequency of a vehicle changing its pseudonym when being tracked by an adversary. A higher frequency of pseudonym changes complicates an adversary's efforts to trace the vehicle [32]. The metric can be calculated using Formula 10:

$$APCT = \frac{PseudonymChanges}{T} \tag{10}$$

where $TPseudonymChanges$ represents the cumulative number of pseudonym changes that occurred across all vehicles during the traces, and T represents the total number of traces or time intervals. A higher $APCT$ indicates a more frequent pseudonym-change strategy, which enhances privacy by reducing the linkability of vehicles' activities across different time intervals. This makes it more difficult for adversaries to trace or track the movements of vehicles over extended periods.

The results are illustrated in Figure 8, where the y-axis represents the pseudonym changes per trace. Experimental results indicate that LSHPC exhibits the highest pseudonym-change frequency, whereas the scheme proposed by Khodaei et al. shows the lowest frequency among all compared schemes. In contrast, the proposed scheme maintains a moderate and relatively stable pseudonym-change frequency under different traffic densities. It should be noted that pseudonym-change frequency is not a performance metric for which either a higher or a lower value is always preferable. More frequent pseudonym changes can reduce the linkability between consecutive pseudonyms, but excessive changes may introduce additional communication and pseudonym-management overhead. Conversely, an excessively low change frequency may increase the possibility of continuous trajectory linking. From a quantitative perspective, the proposed scheme requires approximately 32.5% fewer pseudonym changes than LSHPC. Although LSHPC changes pseudonyms more frequently, it still exhibits higher traceability and lower confusion per pseudonym change and average confusion per trace than the proposed scheme. These results indicate that the privacy advantage of the proposed scheme is not achieved by simply increasing the pseudonym-change frequency. Instead, it achieves stronger overall privacy protection with a moderate number of pseudonym changes, providing a better balance between pseudonym-change frequency and location-privacy protection.

Average confusion per pseudonym change: The average confusion per pseudonym change aims to quantify the average degree of confusion caused by pseudonym changes, thereby assessing the system's ability to protect user privacy [34]. A higher value of this indicator implies

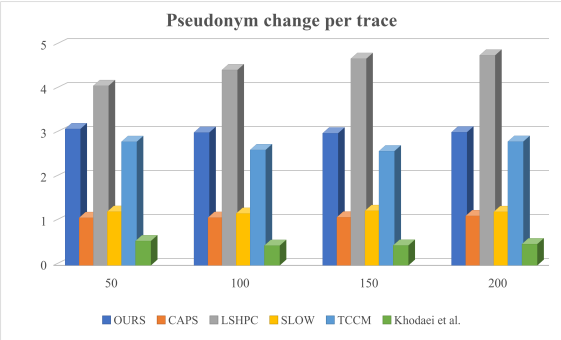

Figure 8: Pseudonym change per trace comparison

better performance of the scheme in protecting user privacy.

The experimental results are presented in Figure 9, where the y-axis represents the average confusion per pseudonym change. As shown in the figure, the proposed scheme consistently achieves the highest average confusion per pseudonym change under different traffic densities, while CAPS exhibits the lowest performance among the compared schemes. From a quantitative perspective, compared with CAPS, the proposed scheme achieves more than four times the average confusion per pseudonym change. Compared with SLOW, TCCM, LSHPC, and Khodaei et al., the proposed scheme achieves average improvements of approximately 112.6%, 110.7%, 98.0%, and 90.3%, respectively. It is worth noting that Khodaei et al. exhibits the best performance among the five baseline schemes for this metric, yet the proposed scheme still maintains a substantial advantage. The relatively lower confusion values of the compared schemes indicate their limited ability to confuse an adversary after pseudonym changes. In contrast, the proposed scheme produces a significantly higher degree of confusion, making it more difficult for the adversary to associate pseudonyms before and after a change and thereby providing stronger location-privacy protection.

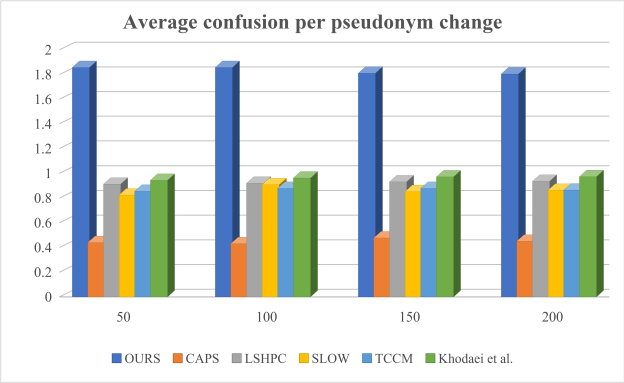

Figure 9: Average confusion per pseudonym change comparison

The advantage remains visible under different traffic densities, confirming that the improvement is not driven by a single traffic condition. Overall, the proposed scheme provides stronger privacy protection after each pseudonym change than the retained baselines.

Average confusion per trace: The average confusion per trace measures the level of confusion imposed on attackers by the scheme, where a higher value indicates stronger privacy protection performance [13]. A definition of confusion per trace is given in Formula 11:

$$Confusion = \frac{1}{|\tau|}\sum_{i=1}^{|\tau|} M\big(S(\tau_i), S'(\tau_i)\big) \qquad (11)$$

where $|\tau|$ is the number of traces, τ_i is the i-th trace, $S(\tau_i)$ is the set of actual vehicles present in τ_i, $S'(\tau_i)$ is the set of vehicles inferred from τ_i, and M quantifies the discrepancy between the real and inferred sets.

As illustrated in Figure 10, the y-axis represents the average confusion per trace. The experimental results demonstrate that the proposed scheme consistently achieves the highest average confusion per trace under different traffic densities. Among the five baseline schemes, LSHPC exhibits the best performance for this metric. From a quantitative perspective, compared with CAPS and Khodaei et al., the proposed scheme achieves more than eleven times the average confusion per trace. Compared with SLOW and TCCM, the proposed scheme achieves more than five times and more than twice the average confusion per trace, respectively. Even compared with LSHPC, which performs best among the five baseline schemes, the proposed scheme still achieves an average improvement of approximately 33.6%. These results demonstrate that the proposed scheme generates substantially greater overall confusion throughout a vehicle trajectory, making it more difficult for an adversary to continuously associate and reconstruct the target vehicle's movements. Together with the moderate pseudonym-change frequency shown in Figure 8, this result further indicates that the privacy improvement of the proposed scheme is achieved through more effective pseudonym changes rather than simply increasing the number of changes, thereby providing stronger and more reliable location-privacy protection.

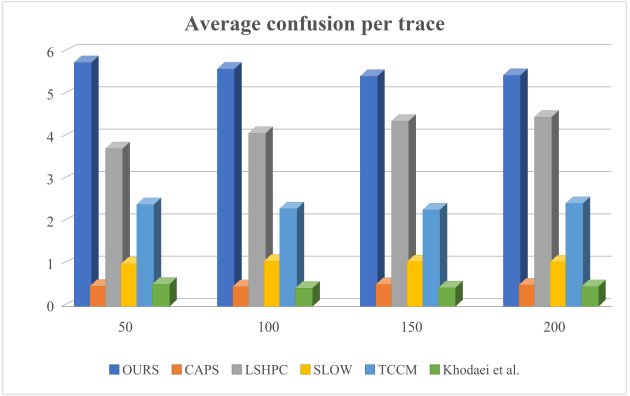

Figure 10: Average confusion per trace comparison

6 Conclusions

This paper proposed an improved pseudonym management scheme for VANETs based on fog computing. The primary objective is to mitigate the communication and storage overhead inherent in the lifecycle management of pseudonyms, including their generation, distribution, and revocation. The fog-based pseudonym system is geographically distributed, responsible for distributing multiple short-term pseudonyms to legitimate vehicles and for tracing malicious ones. By leveraging the distributed processing capabilities of the fog layer, the proposed system significantly reduces service-response latency, decreases communication-bandwidth consumption, and alleviates both computational and storage burdens on the TA. Additionally, this paper proposed a vehicle-centric cooperative pseudonym-change protocol that operates independently of RSUs. Vehicles select cooperative neighbours based on driving-state similarity and establish a dynamic mix-zone to change pseudonyms. Notably, a higher similarity in driving states leads to more uniform behavioral patterns among the cooperating vehicles, thereby significantly increasing the difficulty of vehicle tracking and pseudonym correlation. Comprehensive security, privacy, and performance analyses demonstrate the scheme's reliability, feasibility, and scalability, highlighting its high practicality for real-world deployment.

However, the proposed scheme has several limitations. First, its trust model relies on TA-authorized fog nodes and does not account for coordinated or cross-domain failures. Second, driving-state similarity is computed using fixed weight vectors for road scenes, whose adaptability under rapidly changing traffic conditions remains unverified. To address these issues, future work will develop adaptive fog-node trust and fault-tolerance mechanisms, refine similarity computation through dynamic weight adjustment and nonlinear scoring functions, and conduct large-scale multi-domain experiments that simultaneously evaluate resilience to positioning errors, cross-domain failures, and bandwidth constraints in realistic vehicular environments.

7 Acknowledgments

This work was supported by the Doctoral Research Funds for the Nanyang Normal University (GrantNo. 2024ZX009) and the Natural Science Foundation of Henan Province (Grant No 262300420703). and the Natural Science Foundation of Liaoning Province, China (Doctoral Research Start-up Project) (grant number: 2025-BS-0266) and Henan Province Key Scientific Research Projects for Higher Education Institutions (grant number:26B520033).

8 Author contributions

The authors confirm contribution to the paper as follows: study conception and design: Cong Zhao, Sheng Zhuang, Yikang Yang; data collection: Xinyang Deng, Xuan Ge, and He Li; analysis and interpretation of results: Xiaopu Ma, Qinglei Qi and Wentao Li; draft manuscript preparation: Cong Zhao, Yikang Yang and Xuan Ge. All authors reviewed the results and approved the final version of the manuscript.

9 Declaration of competing interests

The authors declare no financial or personal relationships that could be perceived as inappropriately influencing the content or conclusions of this paper. No external funding was received for this study that could lead to any bias or conflict of interest. The authors confirm that the work presented is their original research, free from any third-party influence or undisclosed collaborations.

References

[1] A. Alshaeri and M. Younis, "Certificate-Less Single-Use Pseudonym Scheme for Countering Trajectory Tracking Attacks in ITS," *IEEE Trans. Intell. Transp. Syst.*, vol. 26, no. 7, pp. 9214–9225, 2025. https://doi.org/10.1109/tits.2025.3568290

[2] J. Wang, Y. Sun, and C. Phillips, "Fake Beacon: A Pseudonym Changing Scheme for Low Vehicle Density in VANETs," in *Proc. IEEE Veh. Technol. Conf. (VTC-Spring)*, 2023, pp. 1–7. https://doi.org/10.1109/vtc2023-spring57618.2023.10199627

[3] S. Bartoletti *et al.*, "Integration of Sensing and Localization in V2X Sidelink Communications," *IEEE Commun. Mag.*, vol. 62, no. 8, pp. 185–191, 2024. https://doi.org/10.1109/mcom.001.2300748

[4] Y. Higashida and K. Sato, "Evaluation of a Pseudonym Change Scheme using LSH for Location Privacy in V2X Communication," *IEICE Commun. Express*, vol. 14, no. 3, pp. 115–118, 2025. https://doi.org/10.23919/comex.2024xbl0184

[5] S. Son *et al.*, "Design of Blockchain-Based Lightweight V2I Handover Authentication Protocol for VANET," *IEEE Trans. Netw. Sci. Eng.*, vol. 9, no. 3, pp. 1346–1358, 2022. https://doi.org/10.1109/tnse.2022.3142287

[6] J. Kang *et al.*, "Blockchain-Based Pseudonym Management for Vehicle Twin Migrations in Vehicular Edge Metaverse," *IEEE Internet Things J.*, vol. 11, no. 21, pp. 34254–34269, 2024. https://doi.org/10.1109/jiot.2024.3404559

[7] A. P. Mdee *et al.*, "Security Compliant and Cooperative Pseudonyms Swapping for Location Privacy Preservation in VANETs," *IEEE Trans. Veh. Technol.*, vol. 72, no. 8, pp. 10710–10723, 2023. https://doi.org/10.1109/tvt.2023.3254660

[8] Y. Wu, "TCCM: Trajectory Converged Chaff-Based Mix-Zone Strategy for Enhancing Location Privacy in VANET," *IEEE Access*, vol. 13, pp. 45436–45448, 2025. https://doi.org/10.1109/access.2025.3550442

[9] E. Khezri *et al.*, "Security Challenges in Internet of Vehicles (IoV) for ITS: A Survey," *Tsinghua Sci. Technol.*, vol. 30, no. 4, pp. 1700–1723, Aug. 2025. https://doi.org/10.26599/tst.2024.9010083

[10] S. K. Panigrahy *et al.*, "A Survey and Tutorial on Network Optimization for Intelligent Transport System Using the Internet of Vehicles," *Sensors*, vol. 23, no. 1, pp. 555–585, 2023. https://doi.org/10.3390/s23010555

[11] Z. Zhang *et al.*, "A Geo-Indistinguishable Context-Based Mix Strategy for Trajectory Protection in VANETs," *IEEE Trans. Veh. Technol.*, vol. 72, no. 12, pp. 16538–16552, 2023. https://doi.org/10.1109/tvt.2023.3293127

[12] A. Srivastava *et al.*, "Location Based Routing Protocols in VANET: Issues and Existing Solutions," *Veh. Commun.*, vol. 23, 2020. https://doi.org/10.1016/j.vehcom.2020.100231

[13] A. Hayat *et al.*, "A Novel Pseudonym Changing Scheme for Location Privacy Preservation in Sparse Traffic Areas," *IEEE Access*, vol. 11, pp. 89974–89985, 2023. https://doi.org/10.1109/access.2023.3303846

[14] X. Luo *et al.*, "Research on Data Privacy Protection of Internet of Vehicles Based on Differential Privacy," in *Proc. Int. Conf. Geo-Spatial Knowledge and Intelligence*, 2020, pp. 12007–12017. https://doi.org/10.1088/1755-1315/428/1/012007

[15] M. Al-Shalabi *et al.*, "Energy Efficient Multi-Hop Path in Wireless Sensor Networks Using an Enhanced Genetic Algorithm," *Inf. Sci.*, vol. 500, pp. 259–273, 2019. https://doi.org/10.1016/j.ins.2019.05.094

[16] L. Hou *et al.*, "Tracking Based Mix-Zone Location Privacy Evaluation in VANET," *IEEE Trans. Veh. Technol.*, vol. 70, no. 10, pp. 10957–10969, 2021. https://doi.org/10.1109/tvt.2021.3109065

[17] M. A. Saare *et al.*, "Relationships Between the Older Adult's Cognitive Decline and Quality of Life," *Int. J. Interact. Mobile Technol.*, vol. 13, no. 10, p. 42, 2019. https://doi.org/10.3991/ijim.v13i10.11288

[18] I. Saini *et al.*, "A Comprehensive Review of Pseudonym Changing Strategies in Vehicular Networks," *IJ Netw. Secur.*, vol. 21, no. 5, pp. 785–796, 2019. https://doi.org/10.6633/IJNS.201909_21(5).10

[19] S. Yogarayan *et al.*, "A Review of Routing Protocols for Vehicular Ad-Hoc Networks (VANETs)," in *Proc. Int. Conf. Inf. Commun. Technol. (ICoICT)*, 2020, pp. 1–7. https://doi.org/10.1109/icoict49345.2020.9166174

[20] T. Gao and L. Zhao, "Pseudonym Schemes Based on Location Privacy Protection in VANETs: A Survey," in *Proc. Int. Conf. Innov. Mobile Internet Services Ubiquitous Comput.*, 2020, pp. 597–605. https://doi.org/10.1007/978-3-030-50399-4_59

[21] W. Cheng *et al.*, "A Survey on Privacy-Security in Internet of Vehicles," in *Proc. IEEE DASC/PiCom/CBDCom/CyberSciTech*, 2021, pp. 644–650. https://doi.org/10.1109/dasc-picom-cbdcom-cyberscitech52372.2021.00109

[22] S. S. Moni and D. Manivannan, "CREASE: Certificateless and Reused-Pseudonym Based Authentication Scheme for Enabling Security and Privacy in VANETs," *Internet Things*, vol. 20, p. 100605, 2022. https://doi.org/10.1016/j.iot.2022.100605

[23] J. Guo *et al.*, "Improving Malicious Email Detection Through Novel Designated Deep-Learning Architectures Utilizing Entire E-Mail," *Digital Commun. Netw.*, vol. 157, pp. 257–279, 2023. https://doi.org/10.1016/j.neunet.2022.09.002

[24] X. Li *et al.*, "PAPU: Pseudonym Swap With Provable Unlinkability Based on Differential Privacy in VANETs," *IEEE Internet Things J.*, vol. 7, no. 12, pp. 11789–11802, 2020. https://doi.org/10.1109/jiot.2020.3001381

[25] Y. Li *et al.*, "A Secure Dynamic Mix Zone Pseudonym Changing Scheme Based on Traffic Context Prediction," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 7, pp. 9492–9505, 2022. https://doi.org/10.1109/tits.2021.3125744

[26] M. Khodaei and P. Papadimitratos, "Cooperative Location Privacy in Vehicular Networks: Why Simple Mix Zones Are Not Enough," *IEEE Internet Things J.*, vol. 8, no. 10, pp. 7985–8004, 2021. https://doi.org/10.1109/jiot.2020.3043640

[27] A. Boualouache *et al.*, "PRIVANET: An Efficient Pseudonym Changing and Management Framework for Vehicular Ad-Hoc Networks," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 8, pp. 3209–3218, 2020. https://doi.org/10.1109/tits.2019.2924856

[28] H. Li *et al.*, "Broadcast and Silence Period (BSP): A Pseudonym Change Strategy," *IEEE Trans. Veh. Technol.*, vol. 72, no. 10, pp. 13618–13630, 2023. https://doi.org/10.1109/tvt.2023.3279121

[29] I. Memon *et al.*, "Pseudonym Changing Strategy with Mix Zones Based Authentication Protocol for Location Privacy in Road Networks," *Wireless Pers. Commun.*, vol. 116, pp. 3309–3329, 2021. https://doi.org/10.1007/s11277-020-07854-6

[30] S. Haider *et al.*, "A Privacy Conserving Pseudonym Acquisition Scheme in Vehicular Communication Systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 9, pp. 15536–15545, 2022. https://doi.org/10.1109/tits.2022.3141744

[31] W. Zhang *et al.*, "DPSP: A Dynamic Pseudonym Swap Program Based Location Privacy Protection Algorithm for Internet of Vehicles," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 5, pp. 4525–4535, 2024. https://doi.org/10.1109/tnse.2024.3392709

[32] A. Hayat *et al.*, "A Novel Pseudonym Changing Scheme for Location Privacy Preservation in Sparse Traffic Areas," *IEEE Access*, vol. 11, pp. 89974–89985, 2023. https://doi.org/10.1109/access.2023.3303846

[33] A. P. Mdee *et al.*, "Infrastructure-Independent Pseudonym Swap Protocol for Vehicular Networks," in *Proc. IEEE ICUFN*, 2022, pp. 351–356. https://doi.org/10.1109/icufn55119.2022.9829588

[34] P. K. Singh *et al.*, "CPESP: Cooperative Pseudonym Exchange and Scheme Permutation to Preserve Location Privacy in VANETs," *Veh. Commun.*, vol. 20, p. 100183, 2019. https://doi.org/10.1016/j.vehcom.2019.100183

[35] T. Gao and X. Xin, "Location Privacy Protection Scheme Based on Random Encryption Period in VANETs," in *Proc. IMIS*, 2018. https://doi.org/10.1007/978-3-319-93554-6_34

[36] M. Wazid *et al.*, "Authentication in Cloud-Driven IoT-Based Big Data Environment: Survey and Outlook," *J. Syst. Archit.*, vol. 97, pp. 185–196, 2019. https://doi.org/10.1016/j.sysarc.2018.12.005

[37] F. Schaub *et al.*, "Privacy Requirements in Vehicular Communication Systems," in *Proc. IEEE PASSAT*, 2009. https://doi.org/10.1109/cse.2009.135

[38] K. Emara, "PREXT: Privacy Extension for Veins VANET Simulator," in *Proc. IEEE VNC*, 2016, pp. 1–2. https://doi.org/10.1109/vnc.2016.7835979

[39] L. Huang *et al.*, "Enhancing Wireless Location Privacy Using Silent Period," in *Proc. IEEE WCNC*, 2005, pp. 1187–1192. https://doi.org/10.1109/wcnc.2005.1424677

[40] K. Emara *et al.*, "CAPS: Context-Aware Privacy Scheme for VANET Safety Applications," in *Proc. ACM WiSec*, 2015, pp. 1–12. https://doi.org/10.1145/2766498.2766500

[41] Y. Pan and J. Li, "Cooperative Pseudonym Change Scheme Based on the Number of Neighbors in VANETs," *J. Netw. Comput. Appl.*, vol. 36, no. 6, pp. 1599–1609, 2013. https://doi.org/10.1016/j.jnca.2013.02.003

[42] L. Buttyan *et al.*, "SLOW: A Practical Pseudonym Changing Scheme for Location Privacy in VANETs," in *Proc. IEEE VNC*, 2009, pp. 1–8. https://doi.org/10.1109/vnc.2009.5416380