Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

Future generation Security for smart and connected Cars - FuSeCar

Deliverable D4.3: Cooperative misbehavior detection algorithm

WP4: Misbehavior detection for vehicular communication protocols and architectures

Authors:
Giovanni Gambigliani Zoccoli[1], Mattia Trabucco[2], Mauro Andreolini[2], Luca Ferretti[2], and Mirco Marchetti[1]
{name.surname}@unimore.it

[1]Department of Engineering "Enzo Ferrari"
[2]Department of Physics, Informatics and Mathematics
University of Modena and Reggio Emilia

Current revision: R1.1
Delivery date: February 19th, 2026

Revision history

Authors	Changes	Date	Revision
Giovanni Gambigliani Zoccoli Mattia Trabucco Mirco Marchetti	Creation of the document, tentative structure	February 3rd, 2025	R0.1
Mattia Trabucco Mauro Andreolini Luca Ferretti Mirco Marchetti	First draft of Section 2	April 11th, 2025	R0.2
Giovanni Gambigliani Zoccoli Luca Ferretti Mauro Andreolini Mirco Marchetti	First draft of Section 3	May 28th, 2025	R0.3
Luca Ferretti Mauro Andreolini Mirco Marchetti	Refinement of the draft of Section 3	July 17th, 2025	R0.4
Giovanni Gambigliani Zoccoli Mattia Trabucco Mirco Marchetti	First draft of Section 4	September 21th, 2025	R0.5
Mattia Trabucco Luca Ferretti Mauro Andreolini Mirco Marchetti	Refinement of the draft of Section 4	October 18th, 2025	R0.6
Mattia Trabucco Luca Ferretti Mauro Andreolini Mirco Marchetti	First draft of Section 5	December 14th, 2025	R0.8
Luca Ferretti Mauro Andreolini Mirco Marchetti	First draft of complete document	January 20th, 2026	R1.0
Giovanni Gambigliani Zoccoli Mattia Trabucco Luca Ferretti Mauro Andreolini Mirco Marchetti	Revision of document and minor fixes	February 19th, 2026	R1.1

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

Contents

1 Introduction

The automotive ecosystem has undergone a profound transformation with the advent of Cooperative Intelligent Transportation Systems (C-ITS), where vehicles, infrastructure, and network services cooperate to enhance road safety, traffic efficiency, and environmental sustainability. As discussed in Deliverable D4.1, the transition from isolated vehicular systems to interconnected Vehicular Ad-hoc Networks (VANETs) represents a paradigm shift in intelligent mobility architectures [1, 6]. In such systems, vehicles are no longer independent entities but active nodes in a distributed cyber-physical environment.

C-ITS relies on continuous communication among heterogeneous entities, including vehicles, Road Side Units (RSUs), Intelligent Traffic Lights (ITLs), and cloud services. This interconnected architecture enables real-time exchange of information regarding position, speed, heading, acceleration, and road events through standardized Vehicle-to-Everything (V2X) messages [11, 21]. While this connectivity enhances situational awareness and cooperative decision-making, it also introduces significant security challenges.

Security in VANETs is traditionally ensured through Public Key Infrastructure (PKI), which guarantees authentication, integrity, and non-repudiation of V2X messages [8, 17]. However, as highlighted in D4.1, PKI only ensures the authenticity of the sender and does not validate the semantic correctness of the transmitted data. A malicious vehicle equipped with valid credentials can broadcast incorrect yet syntactically valid information, thereby acting as an internal attacker [26, 25].

Recent studies have shown that misbehavior detection systems relying solely on plausibility and consistency checks over V2X messages struggle against stealthy trajectory-manipulation attacks such as SixPack $v2$ [10]. In such attacks, the adversary carefully reconstructs realistic vehicle trajectories that remain coherent with surrounding traffic conditions, thereby evading traditional detection mechanisms based only on communication analysis.

Building upon these findings, this deliverable proposes a novel *infrastructure-centric cooperative misbehavior detection algorithm*. Instead of relying exclusively on vehicle-based perception, the proposed approach leverages infrastructure entities, specifically RSUs and Intelligent Traffic Lights (ITLs), as trusted physical observers of the traffic environment.

Infrastructure nodes offer several advantages compared to mobile detectors used on D4.1:

- They are stationary and calibrated, providing stable observation perspectives.
- They can be equipped with high-quality sensing systems such as cameras [20, 5].
- They operate under municipal or trusted authority control, reducing the likelihood of compromise.
- They can monitor critical areas such as intersections, where a high concentration of vehicles and safety-critical events occurs.

The core idea of the proposed approach is to fuse two complementary data sources:

1. **Real data**: Absolute vehicle positions reconstructed from infrastructure perceptual systems (e.g., multi-object tracking via calibrated ITL cameras or radar-equipped RSUs).
2. **Virtual data**: Digital representations of vehicles obtained through V2X communication.

By comparing the digital representation broadcasted by vehicles with the physical representation reconstructed by infrastructure sensors, inconsistencies can be detected and attributed to potential misbehavior. This perception validation extends traditional communication-based misbehavior detection by introducing an external trusted reference from the physical world.

Moreover, infrastructure-based detection enables cooperative and potentially centralized reasoning. Multiple RSUs and ITLs can forward anomaly reports to a municipal Management Authority (MA), where global correlation and aggregation techniques can be applied to improve detection robustness and reduce false positives.

This deliverable therefore extends the local perception-based detection paradigm introduced in D4.1 to a cooperative infrastructure-based detection architecture, maintaining compatibility with ETSI ITS-G5 and IEEE WAVE standards while strengthening resilience against sophisticated internal attacks and potential byzantine behavior.

The remainder of this document is structured as follows. Section 1.1 revisits the fundamentals of C-ITS and VANET architectures. Section 1.2 describes V2X communication mechanisms and message formats. Section **??** provides the security background, including PKI, internal attacker models such as SixPack, and analysis of the current state-of-the-art anomaly detection for VANETs. Section 3 presents the design of the proposed Cooperative Infrastructure Detection Algorithm, while Section 3.3 details the detection procedure adapted to RSUs and ITLs. Section 4 presents the performance evaluation of the system. Lastly, Section 5 present the conclusions of this work.

1.1 C-ITS and VANETs

Cooperative Intelligent Transportation Systems (C-ITS) are built upon Vehicular Ad-hoc Networks (VANETs), which provide the communication substrate enabling vehicles and infrastructure entities to exchange information in real time [1, 6]. VANETs exhibit distinctive characteristics compared to traditional wireless networks: nodes are highly mobile, network topology changes rapidly, connectivity may be intermittent, and safety-critical applications impose strict latency and reliability constraints.

C-ITS deployments support multiple interaction paradigms, most notably Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I), which enable cooperative safety and traffic-efficiency services. Through V2V communication, vehicles periodically broadcast awareness information that allows surrounding entities to construct a local dynamic map of the environment. Through V2I communication, vehicles interact with roadside infrastructure to obtain contextual information such as signal phase and timing at intersections, speed recommendations, or hazard notifications [11, 21]. The ability to establish a shared situational awareness among heterogeneous actors is a cornerstone of connected and automated mobility.

From a protocol perspective, vehicular communication is commonly based on the IEEE 802.11p standard, which represents an amendment of IEEE 802.11 specifically designed for low-latency communication in high-mobility environments. IEEE 802.11p operates in the 5.9 GHz band and introduces enhancements to support rapid link establishment without the need for traditional association procedures. This capability is essential in vehicular environments, where communication opportunities between nodes may last only a few seconds.

Above the physical (PHY) and medium access control (MAC) layers defined by IEEE 802.11p, two main protocol families are adopted depending on the geographic deployment context. In North America, the IEEE 1609 Wireless Access in Vehicular Environments (WAVE) stack defines higher-layer services, including networking, transport, and security mechanisms. In Europe, the ETSI ITS-G5 architecture builds upon IEEE 802.11p and specifies a layered stack composed of:

- Access layer (IEEE 802.11p PHY/MAC),
- Networking and transport layer (GeoNetworking and BTP),
- Facilities layer (defining CAM, DENM, MAP, SPAT messages),
- Security layer (based on Public Key Infrastructure).

The facilities layer is particularly relevant because it standardizes the structure and semantics of cooperative awareness messages. These standardized formats ensure interoperability between vehicles and infrastructure

manufactured by different vendors. The security layer, typically compliant with ETSI TS 103 097, guarantees message authenticity and integrity through digital signatures and pseudonym certificates.

Infrastructure entities such as RSUs and Intelligent Traffic Lights (ITLs) participate in the same protocol stack as vehicles. They implement IEEE 802.11p at the access layer and process CAMs or BSMs at the facilities layer, enabling seamless integration within the VANET ecosystem. However, unlike vehicles, infrastructure nodes are stationary and can combine communication capabilities with persistent sensing.

The stationarity of RSUs and ITLs introduces a fundamental architectural advantage. Since their position is fixed and known, they can be calibrated with respect to road geometry and can continuously monitor a defined geographic area. When equipped with sensing systems such as calibrated cameras infrastructure nodes can reconstruct physical vehicle trajectories and positions with temporal stability [19, 12]. This persistent observation capability is particularly valuable for misbehavior detection, as it reduces uncertainty caused by temporary occlusions and perspective changes typical of vehicle-based sensing.

This deliverable adopts an infrastructure-integrated view of VANETs. Rather than treating infrastructure as a passive communication relay, RSUs and ITLs are considered trusted anchors that observe the physical world and support cooperative reasoning. By combining IEEE 802.11p-based communication with infrastructure-grounded perception, the proposed architecture introduces a hybrid model in which decentralized message exchange is complemented by infrastructure-backed validation.

1.2 V2X Messages

At the core of C-ITS communication are V2X messages, which enable vehicles and infrastructure entities to exchange structured information describing their dynamic state and relevant environmental events. These messages allow each participating node to construct a digital representation of its surroundings, forming the basis for cooperative awareness and safety applications.

In the European ITS-G5 framework, two primary message types are defined at the facilities layer: Cooperative Awareness Messages (CAMs) and Decentralized Environmental Notification Messages (DENMs). In the IEEE WAVE stack adopted in North America, the equivalent periodic awareness message is known as the Basic Safety Message (BSM) [11].

CAMs and BSMs are periodic broadcast messages generated by vehicles at a frequency typically ranging between 1 and 10 Hz, depending on vehicle dynamics. These messages include essential state parameters such as:

- Absolute position (latitude and longitude),
- Speed,
- Heading,
- Acceleration,
- Vehicle dimensions and classification.

The generation of CAMs follows event-driven triggering conditions defined by ETSI standards: a new message is transmitted when significant changes in position, speed, or heading occur, or when a maximum time interval since the last transmission is reached. This adaptive transmission mechanism balances communication efficiency with awareness freshness.

DENMs, on the other hand, are event-driven messages used to report specific hazardous situations such as accidents, road works, or emergency braking events. Unlike CAMs, DENMs are disseminated when a particular event is detected and may be forwarded through GeoNetworking mechanisms to reach vehicles beyond one-hop communication range.

From a functional perspective, V2X messages enable each vehicle to maintain a Local Dynamic Map (LDM), which represents the positions and states of surrounding entities based on received broadcasts. Each received CAM or BSM contributes to updating this digital map, allowing vehicles to anticipate potential conflicts or dangerous maneuvers.

Infrastructure entities, such as RSUs and ITLs, process V2X messages in the same manner as vehicles. They receive periodic broadcasts and can construct their own digital representation of the monitored area. However, unlike vehicles, infrastructure nodes can complement this digital representation with physical measurements obtained from onboard sensing systems.

It is important to emphasize that V2X communication is inherently broadcast-based and does not require session establishment. Messages are transmitted using IEEE 802.11p in ad-hoc mode without prior association, meaning that every node within communication range can receive the broadcast. While this mechanism ensures low latency and scalability, it also introduces security challenges: all nodes rely on the correctness of self-reported data contained in the messages.

Although the security layer ensures message authenticity and integrity through digital signatures and pseudonym certificates, it does not guarantee the correctness of the transmitted information. A malicious vehicle can broadcast syntactically valid CAMs or BSMs containing incorrect position or velocity data while still possessing legitimate credentials. As demonstrated in previous work on stealthy trajectory-manipulation attacks such as SixPack $v2$ [10], attackers can carefully craft message sequences that remain statistically plausible and consistent with traffic flow.

The digital representation constructed from V2X messages therefore reflects a *claimed state* of the vehicle rather than its verified physical state. This distinction is fundamental for the detection methodology proposed in this deliverable.

In the infrastructure-centric approach presented in this work, V2X messages represent the **virtual data source**, while RSU and ITL sensing systems provide the **real data source**. By comparing the claimed position and motion parameters contained in CAMs or BSMs with the physically observed vehicle state reconstructed through infrastructure perception, inconsistencies can be identified.

Consequently, V2X messages are not discarded nor replaced in the proposed architecture; instead, they are treated as hypotheses that must be validated against trusted physical observations. This hybrid validation paradigm forms the foundation of the Cooperative Infrastructure Detection Algorithm described in the following sections.

2 Background

This section provides the necessary background for understanding the security challenges and the detection approach proposed in this deliverable. We begin by discussing the role of PKI in ensuring secure communication in vehicular networks. We then review existing misbehavior detection systems like F^2MD, which are designed to identify anomalies caused by internal attackers. Next, we introduce the SixPack attack, which can evade detection in the F^2MD framework. Finally, we discuss Byzantine faults and their implications for detection in cooperative systems.

2.1 Public Key Infrastructure in VANETs

PKI is a cornerstone of security in C-ITS. In the context of VANETs networks, PKI ensures the authenticity, integrity, and non-repudiation of V2X messages exchanged between vehicles and infrastructure entities. The system relies on the use of digital certificates and asymmetric cryptography, where each vehicle and roadside unit (RSU) is issued a public-private key pair. The public key is used to verify the authenticity of messages, while the private key is used for signing the messages.

In practice, V2X messages such as CAMs or BSMs are signed by the sender using its private key. Upon receiving a message, the recipient uses the sender's public key (included in the message's certificate) to validate the message's authenticity and integrity. This cryptographic process ensures that the message was sent by a legitimate entity and has not been altered during transmission.

The primary role of PKI in VANETs is to protect against external attacks, such as spoofing or man-in-the-middle attacks, where unauthorized entities attempt to inject false messages into the network. By verifying the sender's identity through certificates, PKI prevents unauthorized communication from non-registered vehicles or infrastructure nodes, thus maintaining the integrity of the V2X communication system [15, 2].

However, PKI does not address the problem of internal attackers, who already possess valid certificates but may send false information to mislead other nodes in the network. While PKI guarantees the authenticity of the source, it does not guarantee the truthfulness of the data being transmitted. As vehicles and infrastructure entities are part of the trusted system, internal attackers can exploit their access to broadcast incorrect position data, speed, or other critical information without triggering PKI-based defenses.

This limitation highlights the need for additional layers of security and misbehavior detection, especially in scenarios where trusted entities in the system, such as internal vehicles, become malicious or compromised. Detecting such attacks requires mechanisms that can validate the content of V2X messages beyond simple authentication, which is the focus of the misbehavior detection systems discussed in the following sections.

2.2 Misbehavior Detection Systems: F^2MD and Internal Attacks

Misbehavior Detection Systems (MDS) are designed to identify and mitigate malicious behavior in VANETs, ensuring the integrity and safety of Cooperative Intelligent Transportation Systems (C-ITS). These systems monitor the messages exchanged between vehicles and infrastructure nodes, comparing the received data against expected or plausible values based on vehicle dynamics, traffic flow, and environmental conditions. One of the most widely used frameworks for misbehavior detection is the Framework for Misbehavior Detection (F^2MD), which implements a series of plausibility and consistency checks to evaluate the validity of V2X messages [14].

F^2MD operates on two levels: local detection at the vehicle level and global detection at the infrastructure level. In the local detection phase, each vehicle analyzes its own V2X messages, as well as those received from nearby vehicles, to check for discrepancies such as implausible positions, speed violations, or inconsistent accelerations. The local detection can flag a message as suspicious if, for instance, the reported position is too far from the vehicle's last known location or the speed exceeds a plausible limit for the given road conditions.

The system also verifies whether the vehicle's motion is consistent with its previous states, ensuring that sudden jumps or abnormal changes in direction are flagged as anomalies.

The global detection phase extends the analysis by aggregating misbehavior reports from multiple vehicles or infrastructure nodes. While MDS, like F^2MD, are designed to detect anomalies and misbehavior for internal attackers. These attackers have valid cryptographic credentials and can pass the PKI-based authentication checks. However, they can still broadcast false position, speed, or event information that is consistent with normal behavior, evading many of the plausibility checks used in systems like F^2MD. In particular, dynamic attacks like the SixPack attack exploit this limitation by simulating realistic vehicle behavior while deviating from the actual trajectory. These attacks are difficult to detect because they do not exhibit obvious inconsistencies in the data being transmitted.

F^2MD relies heavily on the assumption that data consistency and plausibility checks can effectively identify misbehavior. However, the SixPack attack, which was presented in detail in D4.1, manipulates the vehicle's reported position, speed and acceleration to blend seamlessly with the surrounding traffic. For example, during the FakeBrake phase, the attacker sends a forged emergency braking message, while the vehicle continues without braking. As F^2MD checks only the plausibility of the broadcasted data (i.e., whether the data fits within expected parameters), it cannot detect this discrepancy. Similarly, during the Recovery and Rejoin phases, where the attacker simulates a gradual recovery of the forged position, the vehicle appears to move normally, passing F^2MD's consistency checks.

Thus, while F^2MD is a powerful tool for detecting many types of misbehavior, its effectiveness is limited when dealing with stealthy internal attacks that manipulate V2X data without violating basic plausibility constraints. This limitation is critical in environments where internal vehicles may be compromised, as their behavior will appear normal to detection systems relying solely on message content analysis.

Moreover, while the approach proposed in D4.1 has demonstrated promising results in detecting the SixPack attack, it does not account for Byzantine actors. These actors, who may behave arbitrarily despite being part of the trusted system, pose a potential threat that could undermine the effectiveness of detection systems. By leveraging infrastructure-based validation, the detection framework can better address the challenges posed by both byzantine actors, providing a more resilient solution for C-ITS security.

3 Design of the Cooperative Detection Algorithm

The design of the Cooperative Detection Algorithm introduces a novel approach to misbehavior detection in C-ITS by shifting the detection responsibility from individual vehicles (as proposed in D4.1) to infrastructure nodes, such as RSUs and ITLs. These infrastructure nodes are equipped with high-fidelity sensors and are placed at strategic, elevated positions, offering an unobstructed view of the surrounding traffic environment. By integrating V2X communication data with real-world sensor data from infrastructure nodes, the algorithm improves detection providing a more reliable solution for identifying stealthy attacks like SixPack. This section outlines the design of the detection algorithm which is described in Section 3.1, while in Section 3.2 are presented the advantages and disadvantages of this solution compared to the vehicle-based proposed in D4.1. Lastly, the detection procedure is detailed in Section 3.3.

3.1 Infrastructure design

The infrastructure design is built upon a hybrid approach that integrates **virtual data** from V2X messages with **real-world sensor data** from infrastructure nodes. This approach leverages the power of both vehicle communication and infrastructure sensors to improve the detection of misbehaving entities in a C-ITS. The infrastructure detection method is designed to be deployed on RSUs participating in the detection task within the C-ITS framework. In our scenario, ITLs are considered as our reference RSUs since they are equipped with cameras and other sensors capable of identifying traffic violations and providing critical data for detecting C-ITS entities.

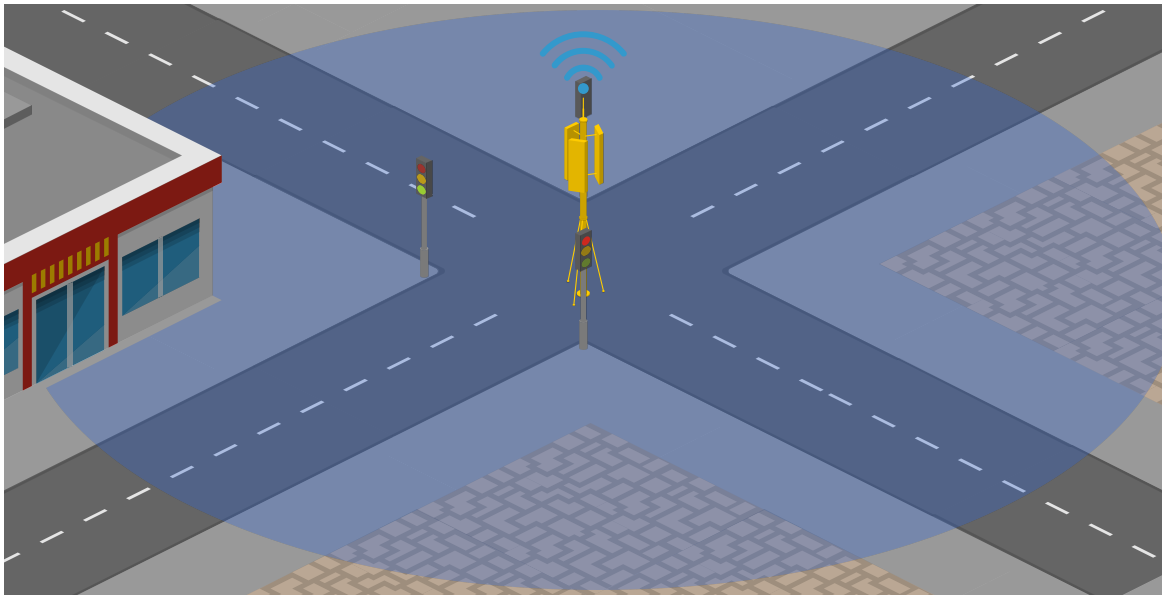

Figure 1: Representation of the circular-shaped coverage area surrounding the ITL

The coverage area of the ITL is modeled as a circular region centered around the pole of the ITL, as depicted in Figure 1. In this model, the radius r defines the extent of the detection range of the ITL. Therefore, the system verifies if each vehicle's GPS coordinates (x_p, y_p) fall within the coverage area of the ITL by checking whether the following condition is satisfied:

$$(x_p - x_c)^2 + (y_p - y_c)^2 < r^2 \tag{1}$$

where (x_c, y_c) are the coordinates of the ITL, and r is the radius of the coverage area. This equation ensures that only vehicles within the effective detection range of the ITL are considered for validation.

The infrastructure-based detection method benefits from the stationary nature of ITLs, allowing them to consistently monitor a well-defined area without the limitations faced by vehicle-based systems, such as occlusions or sensor range limitations. This makes the ITL-based detection method particularly effective in scenarios where high accuracy and continuous monitoring are critical, such as at road intersections or traffic hubs.

3.2 Advantages and disadvantages of cooperative detection

The cooperative detection algorithm offers several key advantages over traditional local detection systems presented in D4.1 where the detection mechanism relies on vehicle-based sensors to monitor surrounding traffic. While this approach is effective in some scenarios, it suffers from significant limitations due to occlusions and field-of-view constraints.

The cooperative detection approach offers several significant advantages over traditional vehicle-based detection systems, primarily due to the use of infrastructure nodes such as RSUs and ITLs. One of the key strengths of this approach is the unobstructed view provided by infrastructure nodes, which are typically positioned elevated and stationary. Unlike vehicles, which can be easily occluded by surrounding traffic, infrastructure nodes offer a clear line of sight over a larger area, such as intersections or complex road networks. This allows them to continuously monitor and validate vehicle positions and movements with much higher accuracy and reliability, reducing the likelihood of misbehavior going undetected due to sensor occlusions or limited field-of-view issues faced by vehicles. Moreover, infrastructure nodes have the advantage of broader coverage, as they can monitor entire intersections or road sections simultaneously, whereas vehicles are limited by their sensor ranges and can only observe a small portion of the surrounding traffic at any given time.

Furthermore, the infrastructure-based approach improves scalability and real-time performance. Infrastructure nodes, such as ITLs, can cover large urban areas and dense traffic zones, offering continuous and consistent monitoring, which is particularly important in urban environments with frequent traffic disruptions. However, the infrastructure-based detection approach also presents certain challenges. One of the main disadvantages is the high cost associated with deploying and maintaining infrastructure nodes. Installing RSUs and ITLs requires significant investment in both physical infrastructure and sensor technology, and the cost of ongoing maintenance and calibration can add up over time. Additionally, low-density or rural areas may struggle to justify such investments, as infrastructure nodes might be underutilized, leading to inefficiencies.

Another challenge lies in the privacy concerns associated with the continuous monitoring of traffic. Infrastructure nodes have the potential to collect large amounts of sensitive data, such as vehicle positions, speeds, and movements. While data privacy measures like pseudonyms and encryption can mitigate some of these concerns, ensuring that data is only used for security purposes and not misused is critical.

Despite these limitations, the cooperative detection approach offers a highly effective solution for misbehavior detection, particularly in urban environments where complex traffic scenarios present challenges for traditional vehicle-based systems. By leveraging infrastructure sensors with a broader field of view, the algorithm can detect stealthy attacks, such as SixPack, that might evade vehicle-based detection systems. Moreover, another significant advantage of the infrastructure-based approach is that RSUs and ITLs can be considered as fully trusted entities within the C-ITS ecosystem. Unlike vehicles, which are mobile and potentially compromised by internal attacks, infrastructure nodes are typically controlled and maintained by trusted authorities such as municipalities or road operators. This means that RSUs and ITLs do not require mechanisms to account for potential Byzantine faults.

3.3 Detection Procedure

The Cooperative Detection Algorithm operates in a sequence of well-defined phases, ensuring that V2X communication data from vehicles is validated through real-world sensor data obtained from infrastructure nodes. The process begins upon receiving a V2X message, which could be either a `BSM` (Basic Safety Message) in the DSRC/WAVE standard or a `CAM` (Cooperative Awareness Message) in the ETSI ITS-G5 standard. Both `BSM`s and `CAM`s contain three mandatory parameters:

- `Position`: The GPS coordinates of the vehicle;
- `Timestamp`: The timestamp associated with the last `BSM` or `CAM` message containing position information;
- `Pseudonym`: An integer used to uniquely identify the vehicle.

It is important to note that the `pseudonym` is part of the mandatory component of V2X communication messages, referred to as `TemporaryID` in [16] and `authorization tickets` in [9]. The pseudonym is used to enhance the privacy of communications within C-ITS networks, as discussed in [3].

The detection algorithm discards any messages where the `coordinates` fall outside the coverage area of the perceptual system used by the ITL. If the GPS coordinates extracted from the message are within the coverage range, the system proceeds by analyzing the output from the `sensor fusion system` to identify the vehicle corresponding to those coordinates. The algorithm accounts for possible `approximation errors` resulting from the sensor fusion process. If the location derived from the message does not match any of the identified entities in the perceptual system, an `anomaly` is flagged.

A diagram of the detection overview is depicted in Figure 2, where the three phases (V2X message evaluation, perceptual system representation, and anomaly detection) are highlighted in yellow, blue, and red, respectively, demonstrating the flow of information and data processing steps.

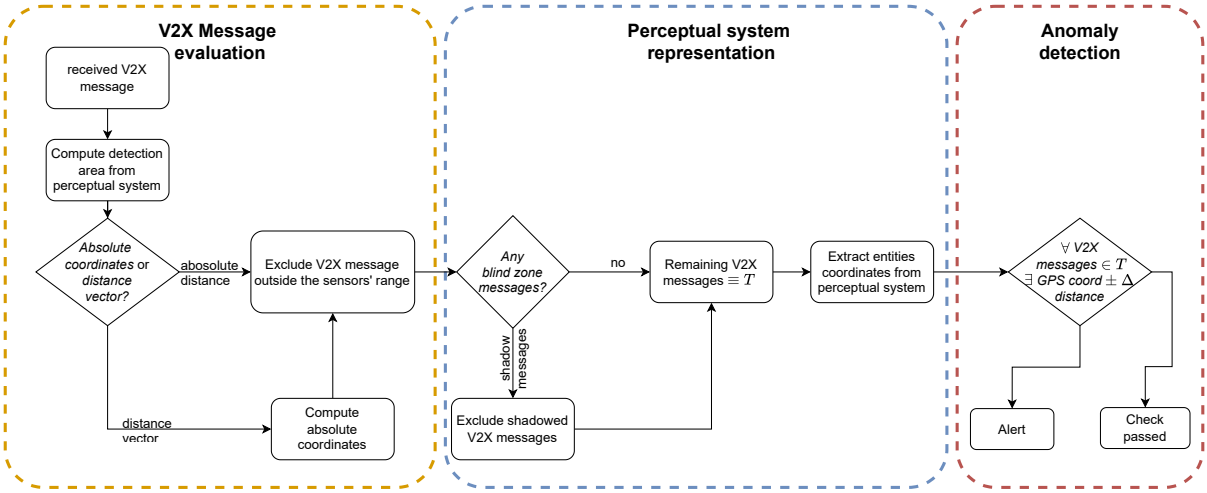

Figure 2: Flowchart of the detection procedure

We define the detection process as consisting of three main phases: *V2X message evaluation*, *perceptual system representation*, and *Anomaly detection*.

1. **V2X Message evaluation**: V2X messages (such as CAMs and BSMs) are broadcast by vehicles and received by infrastructure nodes like RSUs and ITLs. These messages contain essential information about the vehicle, including its position, speed, heading, and pseudonym. The first step in the detection process is to evaluate this data and create a digital representation of the vehicle's state. The RSUs and ITLs analyze

the structure of the V2X messages, ensuring that the data is both syntactically valid and conforms to the expected message format. At this stage, the system checks for basic issues like message integrity and timestamp validity. While this phase primarily handles the verification of the message format, it lays the groundwork for the subsequent validation steps.

2. **Perceptual System representation**: Infrastructure nodes, equipped with high-fidelity sensors such as cameras, continuously monitor the surrounding environment. These sensors capture real-world data about the positions and movements of vehicles within the detection area. Unlike vehicle-based systems, which are constrained by their field of view, infrastructure nodes have the advantage of a broader, unobstructed view. This enables them to track vehicles even in dense traffic or situations where other vehicles might block a vehicle's sensors. In this phase, the infrastructure nodes use sensor data to reconstruct the positions of nearby vehicles, forming an accurate physical representation of the environment. This real-time perceptual system representation is essential for validating the claimed positions reported in the V2X messages.

3. **Anomaly Detection**: After receiving the V2X messages and gathering the real-world data from the perceptual system, the algorithm fuses these two data sources to detect discrepancies. The fusion process compares the digital claims from the vehicle (as reported in the V2X messages) with the observed data from the infrastructure sensors. An anomaly is raised if the digital representation of an entity does not align with any entity seen from the perceptual system with a tolerance of ± 2 meters to account for any sensor's approximation errors. We set this value considering a tolerance higher than the ones presented in [4], where the authors experimentally demonstrate that different automotive perception sensors have an approximation error up to 1 meter. In particular, we increase the approximation error by 1 meter to consider any additional approximation error of the GPS antenna in the presence of obstruction, which is not evaluated in [4]. We remark that this additional approximation error on the GPS antenna is extremely conservative, and that in the presence of high buildings (i.e., the GPS satellites are not visible at road level) the GPS system could provide a position that falls between 1 and 5 meters [24].

   Additionally, we implemented a time-based filter to overcome the limitation of the simulation environment, where the data (i.e., V2X messages and perception systems output) are only available after the simulation is completed. In particular, each detection entity compares the position of each entity extracted by the *perceptual system representation* with the last V2X message received for each entity up to a maximum of T_b, which is the beacon interval of the simulation. In our experiments, T_b is set equal to 1 second. We remark that this threshold value in the simulation environment allows the detection entity to use only the last message sent by the visible C-ITS entities, and that in a real-world application C-ITS entities could have different beacon period, thus requiring an additional filtering step.

This method allows for robust detection by validating the digital claims made by vehicles against physical data obtained from infrastructure nodes. By doing so, the algorithm significantly improves the detection of stealthy attacks, such as SixPack, where the attacker manipulates V2X messages to simulate realistic vehicle behavior. The combination of V2X message evaluation, perceptual system validation, and anomaly detection provides a more trustworthy and accurate system for identifying misbehavior in C-ITS. The cooperative nature of the system, where multiple RSUs and ITLs can share data and collaborate on detection, further enhances its effectiveness in challenging environments.

The entire detection procedure phases are reported in Algorithm 1 starting from the reception of a V2X message to the last signaling phase. We remark that, the `check detection area` function of line 3 use the formula described in Section 3.1 to verify if the V2X message fall within the coverage area of the ITL.

Algorithm 1: Detection Procedure

```
1  Function COMPARISON_CHECKS(msg)
2      (x_c, y_c) ← EXTRACT_POSITION(msg);
3      if CHECK_DETECTION_AREA(x_c, y_c) then
4          find ← False;
5          entities ← GET_SURROUNDING_ENTITIES();
6          for e in entities do
7              e_pos ← GET_POSITION(e);
8              if e_pos is a distance vector then
9                  e_pos ← CALCULATE_COORD(e_pos);
10             S ← CALCULATE_S_AREA(v_pos, e_pos);
11             if m_pos in (C_a − S) and m_pos ≈ e_pos then
12                 find ← True;
13                 BREAK;
14         if not find then
15             sender ← EXTRACT_SENDER(msg);
16             RAISE_ANOMALY();
17     else
18         return;                                  // detection procedure ends
```

4 Experimental evaluation

In this section we evaluate the detection performance of our novel detection methodology against the F^2MD detection framework in Section 4.2. The simulation setup, its parameters and performance indexes used as reference for the evaluation process are presented in Section 4.1.

4.1 Simulation setup

The testing environment relies on two main components: the VEINS simulator [22] and the F^2MD evaluation framework [14]. VEINS is an open source platform used for modeling network communications based on the *OMNet++* [23] and the *SUMO* [18] simulators. We configured the F^2MD framework to use the *ExperiCheck* as the only plausibility check in the detection task to minimize the false positives, as demonstrated in [10]. Despite this configuration might seem to introduce a bias in the detection performance of F^2MD, we experimentally evaluated that it is the best configuration for the F^2MD framework to detect the SixPack attack, hence we chose to compare the detection evasion capabilities of our attack against the worst case scenario. Our tests are based on the Luxembourg SUMO Traffic minified (LuSTMini) scenario [7], publicly released by the VehicularLab at the University of Luxembourg. The LustMini scenario is designed to cover a restricted area of the city of Luxembourg, spanning close to 84 kilometers of roads with 300 intersections, 43 semaphores, and a total of 536 vehicles over 20 minutes of simulation, with a maximum of 105 vehicles simultaneously present in the simulated environment. In all of our tests, we configured 27 vehicles (equal to 5% of the overall number of vehicles) to activate the attack, with the remaining 509 vehicles that can participate in the detection or not.

For the performance evaluation task we use *precision*, *recall*, and $\mathcal{F}_1$, a set of widely used indexes to measure the quality of the detection.

$$Precision = \frac{TP}{TP + FP} \tag{2}$$

$$Recall = \frac{TP}{TP + FN} \tag{3}$$

$$\mathcal{F}_1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{4}$$

The *precision* metric (Formula 2) measures the proportion of correctly identified anomalies among all detected anomalies, while the *recall* metric (Formula 3) reflects the proportion of true anomalies successfully identified out of the total number of actual anomalies. In this context, TP, FP, and FN represent the counts of true positives (anomalous messages correctly detected), false positives (legitimate messages wrongly flagged as anomalies), and false negatives (anomalous messages incorrectly classified as normal), respectively. The $\mathcal{F}_1$ (Formula 4) is calculated as the harmonic mean of precision and recall, providing a balanced summary where both metrics are equally weighted. Each of these performance measures ranges from 0 to 1, with values near 0 indicating poor detection capability (correspondingly, the ability of the attack to evade detection) and values close to 1 indicating almost perfect detection performance (i.e., the inability of the attack to evade detection).

4.2 Detection performance of the Cooperative Detection algorithm

In this section, we compare the detection performance of the *cooperative* detection method with the performance of F^2MD against SixPack $v2$. The setup of the simulation environment is the same as already described in Section 4.1.

For the comparison of the detection performance of the *cooperative* detection method and F^2MD, we execute multiple simulations activating a different number of detection entities in the simulation scenario, being either *ITLs* (for the *cooperative* detection method) or vehicles (for the F^2MD detection framework). Following the previous test scenarios, we tested 10 different detector ratios (from 10% to 100% of the overall number of available detecting entities), replicating each test 100 times to remove any bias in the simulation.

Since the *cooperative* detection method is based on data gathered from *ITLs* centered in a road intersection (with a total of 43 traffic light intersections), we prioritized the activation of *ITLs* for detection purposes starting from the most important ones and gradually activating the lesser important ones. The importance of a road intersection is based on the number of lanes converging to the intersection, with higher numbers identifying more important intersections. In the LuSTMini scenario, there are 2 traffic light intersections with more than 14 lanes, followed by 3 intersections with 12 lanes, 6 with 10 lanes, 5 with 8 lanes, 7 with 7 lanes, 6 with 6 lanes, 10 with 5 lanes, while the remaining 4 intersections have at most 4 lanes. In our tests, we set the radius r of the coverage area for the *ITLs* of the *cooperative* scenario equal to 100 meters [13, 4].

The results of this experimental evaluation are presented with the same metrics used in 4.1, being the *precision* (Figure 3), *recall* (Figure 4), and $\mathcal{F}_1$ (Figure 6). The x-axis of each figure showcases the percentage value of detecting entities (*ITLs* for the *cooperative* detection method and vehicles participating in the detection for F^2MD), while the y-axis shows the value of the performance index. In each figure, results depicted in blue are related to the *cooperative* detection method, while results depicted in red are related to the F^2MD detection framework.

The *precision* of both *cooperative* and F^2MD are compared in Figure 3, which highlights that our novel detection method outperforms the F^2MD performance by always reaching the maximum value of 1.0 regardless of the number of *ITLs* used for the detection. This implies that our approach does not raise any false positives, thus being extremely effective in the identification of an ongoing instance of SixPack $v2$. On the other hand, the median *precision* value of the F^2MD framework has a decreasing trend inversely proportional to the percentage of vehicles contributing to the detection process, starting with a value slightly higher than 0.4 with 10% of detectors down to a value close to 0.23 with 100% of vehicles participating in the detection process. The decrease in both the median and variance of the precision for F^2MD occurs because, as more vehicles contribute reports, the MA aggregates a larger amount of partially redundant information, increasing the likelihood of false positives. This results in more consistent (lower variance) but slightly less precise (lower median) outcomes. When only a few vehicles participate, the system relies on a smaller and more heterogeneous set of reports, leading to greater variability and occasional high-precision results when the contributing vehicles have favorable coverage.

The results of the *recall* performance index presented in Figure 4 show that, once again, our novel detection approach outperforms F^2MD in all different scenarios, improving its overall detection performance by increasing the number of *ITLs* used in the detection task. To evaluate how the radius values influence the *recall* of the *cooperative* detection method, we repeated the same experimental procedure used in the previous tests (variable percentage of detectors and 100 tests for each percentage). In these additional tests, we modified the radius values of the coverage area of the *cooperative* detection area and compared the default value equal to 100 meters with two reduced ranges of 50 and 25 meters, respectively.

The results are depicted in Figure 5 where, as in the other Figures, the x-axis shows the percentage of detectors, while the y-axis reports the value of the *recall*. The three different radius values are showcased using three shades starting from the original blue color (100 meters) used in the other Figures to represent the performance of the *cooperative* detection method. The *recall* values of Figure 5 indicate that decreasing the radius does not lead to a substantial degradation of the *recall* metric. Notably, the 50 meters configuration achieves *recall* values comparable to - and in some cases exceeding - those obtained by the F^2MD framework, while the 25 meters radius exhibits a lower *recall*, as expected, but still maintains reasonable performance given

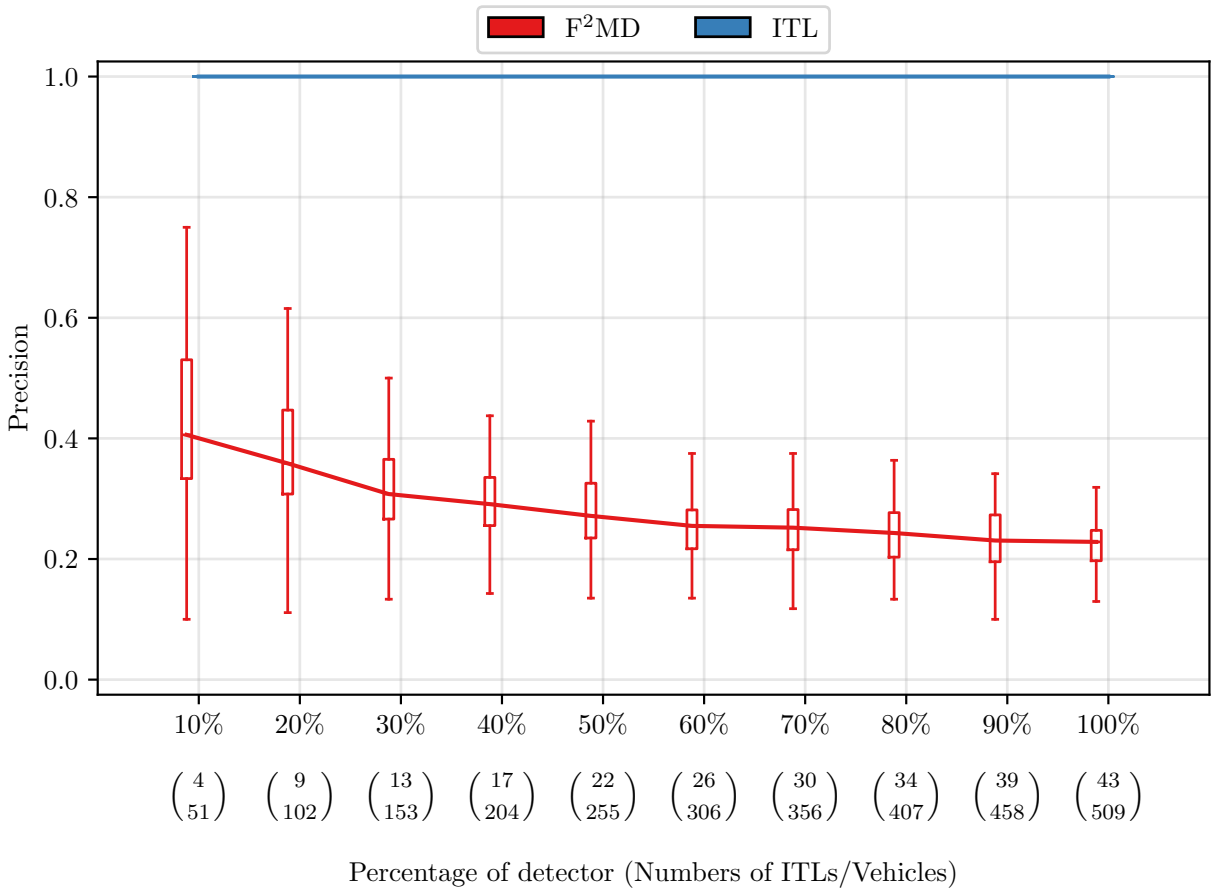

Figure 3: Comparison of the *precision* index of F^2MD and the *cooperative* detection methods against the SixPack $v2$ attack.

its limited coverage.

The overall detection performance of both *cooperative* and F^2MD against SixPack $v2$ is summarized by means of $\mathcal{F}_1$ in Figure 6, which outlines that our novel *cooperative* detection method is able to achieve higher $\mathcal{F}_1$ than F^2MD. We can observe that the overall results of F^2MD are the same as those presented in D4.1, being close to the same value regardless of the number of vehicles participating in the detection task. On the other hand, by focusing on the detection performance of the *cooperative* detection method, we can observe that its $\mathcal{F}_1$ is directly related to the number of participating *ITLs* for the detection task, being able to reach a maximum median value of 0.9411.

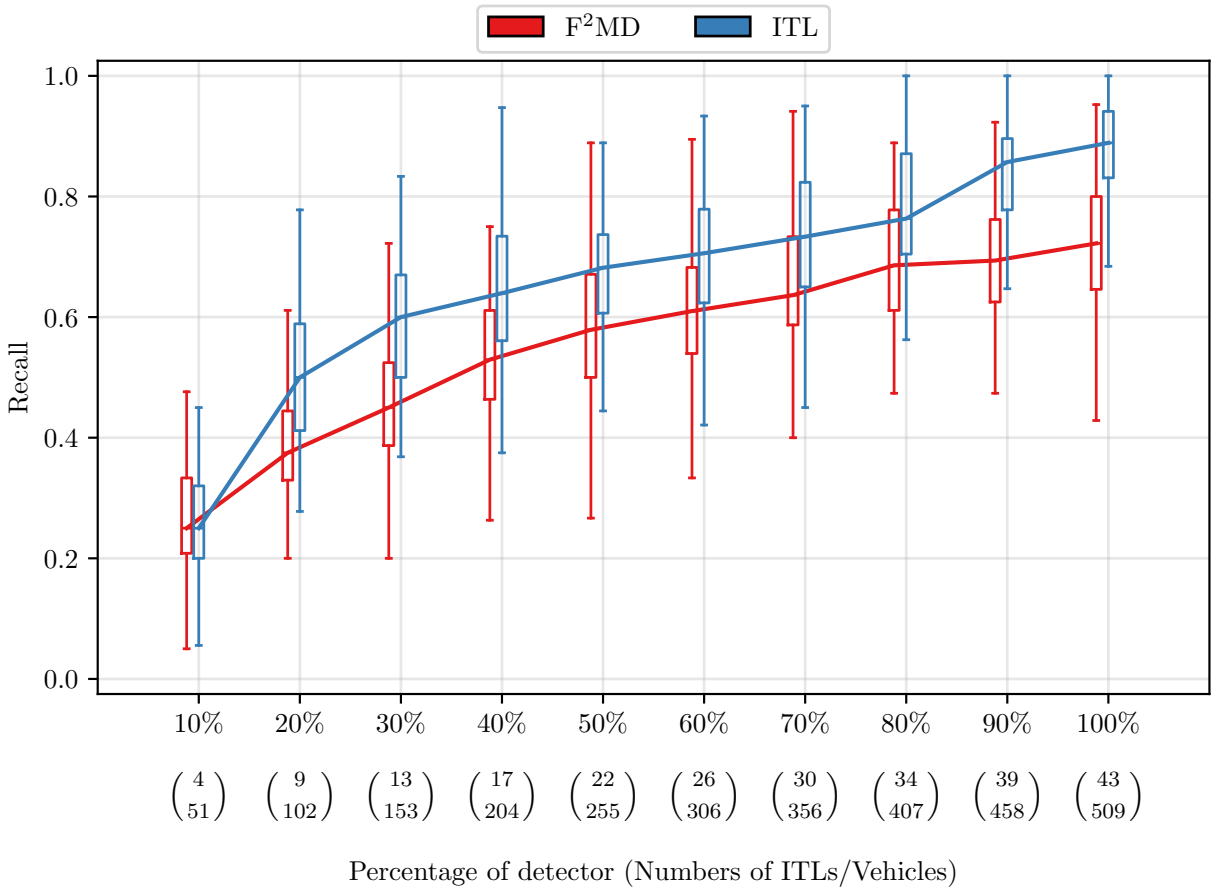

Figure 4: Comparison of the *recall* index of F^2MD and the *cooperative* detection methods against the SixPack $v2$ attack.

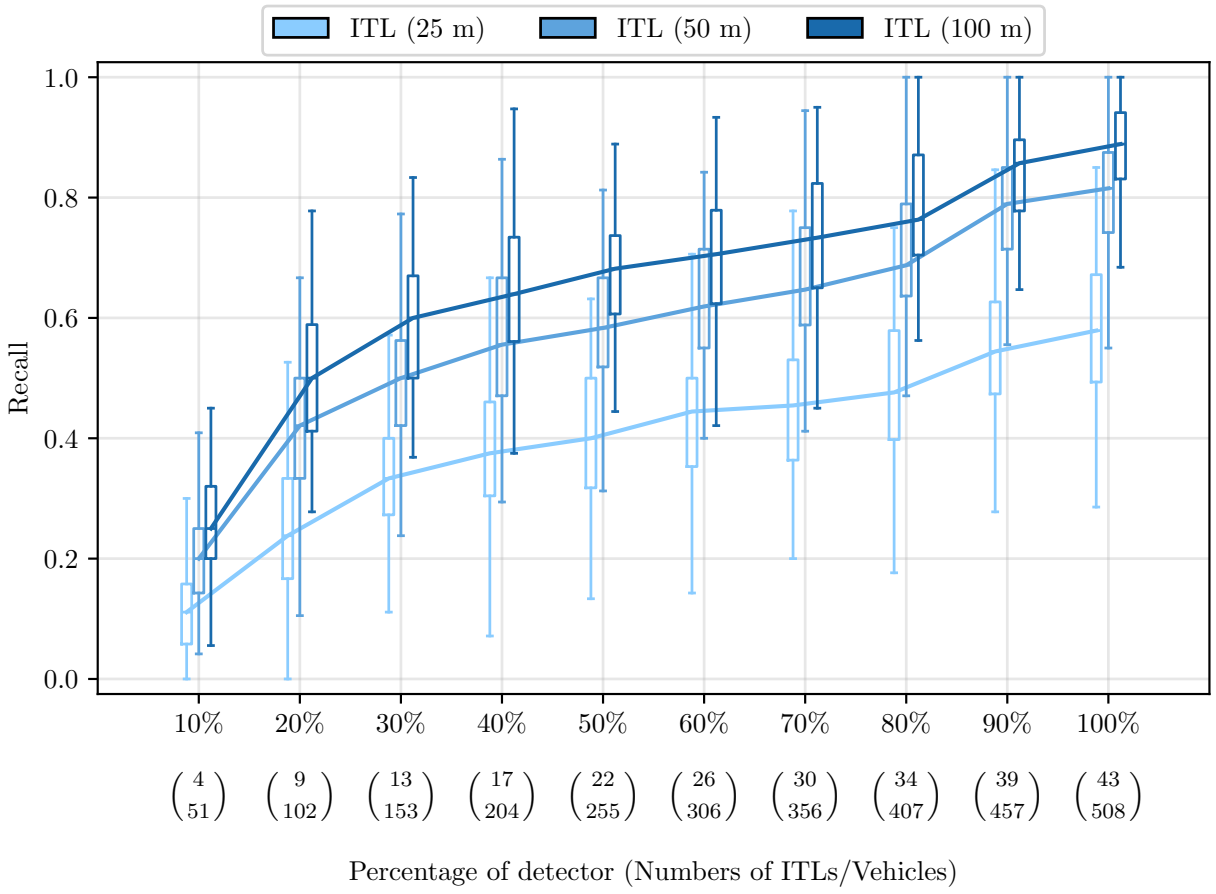

Figure 5: Comparison of the *recall* value considering different *radius* values of the *cooperative* detection method.

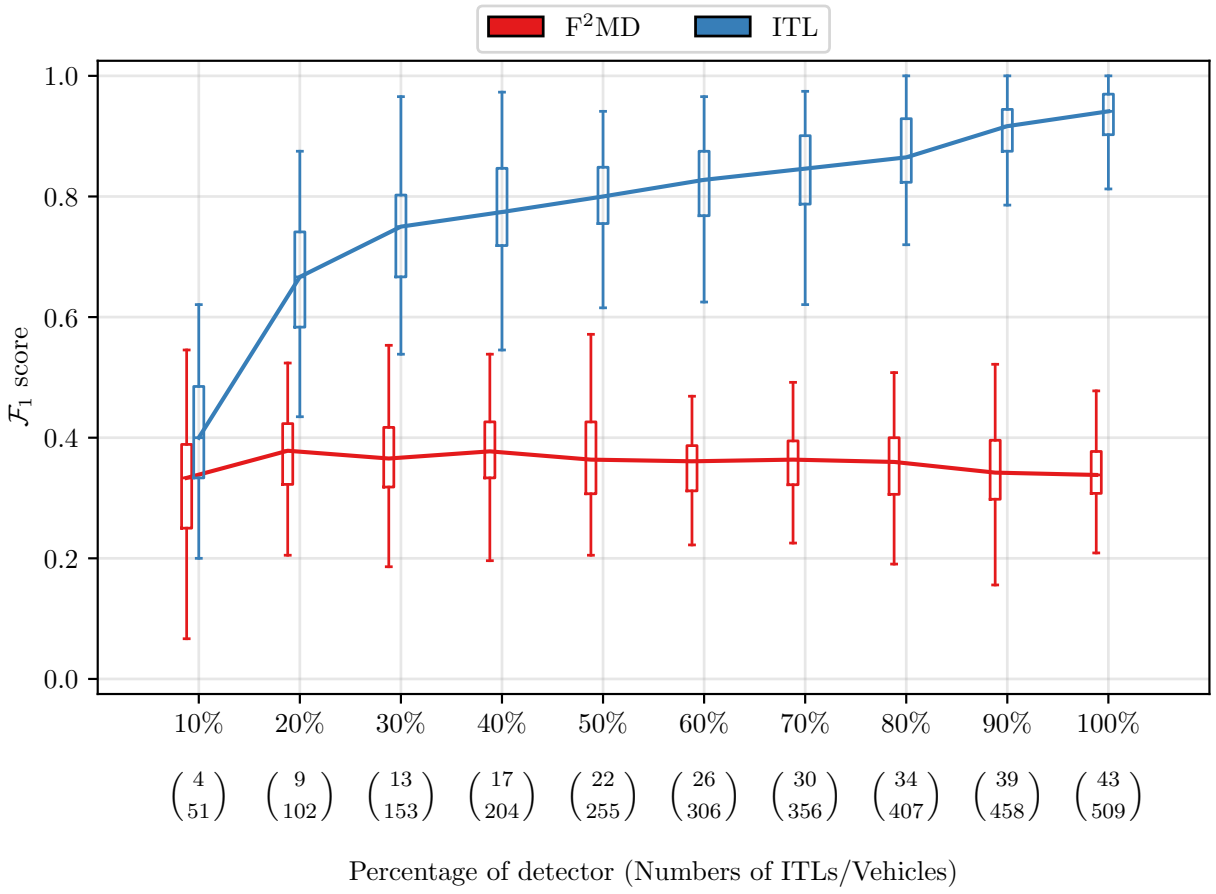

Figure 6: Comparison of the $\mathcal{F}_1$ index of F^2MD and the *cooperative* detection methods against the SixPack $v2$ attack.

5 Conclusions

In this work, we present a novel cooperative misbehavior detection method based on the comparison of data extracted from V2X messages and the data gathered by commonly deployed perceptual systems to improve misbehavior detection against stealthy attacks on V2X communications.

Our novel detection approach utilizes perceptual systems configurations that can be found on modern infrastructure, such as calibrated video cameras. The detection method compares real-world data gathered from *infrastructure nodes* (RSUs and ITLs) with the claimed states in the V2X messages, thus offering more reliable verification of vehicle positions and behaviors.

The experimental results demonstrate that our novel detection method achieves superior detection performance compared to the current state-of-the-art. The highest $\mathcal{F}_1$ scores were achieved when multiple *ITLs* participated in the detection task. Additionally, our approach limits false positives across all tested scenarios, making it especially suitable for situations where *precision* is paramount in detection algorithms. Moreover, in this work we evaluated how different percentage of detector (RSUs and ITLs) impact the anomaly detections accuracy and precision, providing a sensitive analysis which can be adopted by a municipality to decide how many ITLs/RSUs are necessary to guarantee that the anomalies are detected.

References

[1] Ahmad Abuashour and Michel Kadoch. Vehicular ad-hoc networks: Architecture, applications and challenges. *International Journal of Computer Science & Network Security*, 20(2):26–36, 2020.

[2] M. Babaghayou, N. Labraoui, A. A. A. Ari, N. Lagraa, and M. A. Ferrag. Pseudonym change-based privacy-preserving schemes in vehicular ad-hoc networks: A survey. *Journal of Information Security and Applications*, 2020:1–17, 2020.

[3] Messaoud Babaghayou, Nabila Labraoui, Ado Adamou Abba Ari, Nasreddine Lagraa, and Mohamed Amine Ferrag. Pseudonym change-based privacy-preserving schemes in vehicular ad-hoc networks: A survey. *Journal of Information Security and Applications*, 2020.

[4] Minjin Baek, Donggi Jeong, Dongho Choi, and Sangsun Lee. Vehicle trajectory prediction and collision warning via fusion of multisensors and wireless vehicular communications. *Sensors*, 20(1):288, 2020.

[5] Celso A. R. L. Brennand, Geraldo P. Rocha Filho, Guilherme Maia, Felipe Cunha, Daniel L. Guidoni, and Leandro A. Villas. Towards a fog-enabled intelligent transportation system to reduce traffic jam. *Sensors*, 19(18), 2019.

[6] Claudia Campolo, Antonella Molinaro, and Riccardo Scopigno. *Vehicular Ad Hoc Networks: Standards, Solutions, and Research*. Springer, 2015.

[7] Lara Codecá, Raphaël Frank, Sébastien Faye, and Thomas Engel. Luxembourg SUMO traffic (LuST) scenario: Traffic demand evaluation. *IEEE Intelligent Transportation Systems Magazine*, 9(2):52–63, 2017.

[8] ETSI. Intelligent transport systems (ITS); security; security header and certificate formats; release 2. Technical Specification ETSI TS 103 097 V2.1.1, European Telecommunications Standards Institute, Sophia Antipolis Cedex, France, October 2021.

[9] European Telecommunications Standards Institute (ETSI). Intelligent Transport Systems (ITS); Security; ITS communications security architecture and security management. Technical Report TS 102 940 V2.1.1, ETSI, 2021. Accessed: November 2025.

[10] Giovanni Gambigliani Zoccoli, Francesco Pollicino, Dario Stabili, and Mirco Marchetti. Sixpack v2: Enhancing sixpack to avoid last generation misbehavior detectors in VANETs. In *2022 IEEE 21st International Symposium on Network Computing and Applications (NCA)*, pages 243–249. IEEE, 2022.

[11] Sohan Gyawali, Shengjie Xu, Yi Qian, and Rose Qingyang Hu. Challenges and solutions for cellular based V2X communications. *IEEE Communications Surveys & Tutorials*, 23(1):222–255, 2021.

[12] Henry Alexander Ignatious, Manzoor Khan, et al. An overview of sensors in autonomous vehicles. *Procedia Computer Science*, 198:736–741, 2022.

[13] Henry Alexander Ignatious, Manzoor Khan, et al. An overview of sensors in autonomous vehicles. *Procedia Computer Science*, 198:736–741, 2022.

[14] J. Kamel, M. R. Ansari, J. Petit, A. Kaiser, I. Ben-Jemaa, and P. Urien. Simulation framework for misbehavior detection in vehicular networks. *IEEE Transactions on Vehicular Technology*, 69(4):3401–3414, 2020.

[15] John B. Kenney. Dedicated short-range communications (dsrc) standards in the united states. *Proceedings of the IEEE 99.7*, 2011.

[16] John B. Kenney. Dedicated short-range communications (dsrc) standards in the united states. *Proceedings of the IEEE*, 99(7):1162–1182, 2011.

[17] Salabat Khan, Fei Luo, Zijian Zhang, Mussadiq Abdul Rahim, Mubashir Ahmad, and Kaishun Wu. Survey on issues and recent advances in vehicular public-key infrastructure (VPKI). *IEEE Communications Surveys & Tutorials*, 24(3):1574–1601, 2022.

[18] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lücken, Johannes Rummel, Peter Wagner, and Evamarie WieBner. Microscopic traffic simulation using sumo. In *The 21st IEEE International Conference on Intelligent Transportation Systems*, 2018.

[19] Enrique Marti, Miguel Angel De Miguel, Fernando Garcia, and Joshue Perez. A review of sensor technologies for perception in automated driving. *IEEE Intelligent Transportation Systems Magazine*, 11(4):94–108, 2019.

[20] Mükremin Özkul, Ilir Capuni, and Elton Domnori. Context-aware intelligent traffic light control through secure messaging. *Journal of Advanced Transportation*, 2018(1):1–10, 2018.

[21] Miguel Sepulcre, Javier Gozalvez, and Gokulnath Thandavarayan. On the potential of V2X message compression for vehicular networks. *IEEE Access*, 8:20826–20842, 2020.

[22] C. Sommer, R. German, and F. Dressler. Bidirectionally coupled network and road traffic simulation for improved ivc analysis. *IEEE Transactions on Mobile Computing*, 2011.

[23] András Varga and Rudolf Hornig. An overview of the omnet++ simulation environment. In *Proceedings of the 1st international conference on Simulation tools and techniques for communications, networks and systems & workshops*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

[24] Duojie Weng, Baoguo Yu, Jingbo Zhao, Shuo Li, Hangyu Zhou, and Ying Xu. Augmenting vehicle gnss positioning through cross-street measurements in urban canyons. *Measurement*, 257:118603, 2026.

[25] X. Xu, Y. Wang, and P. Wang. Comprehensive review on misbehavior detection for vehicular ad hoc networks. *Journal of Advanced Transportation*, 2022:1–27, 2022.

[26] Y. Zhang et al. Security in vehicular ad hoc networks: Challenges and solutions. *IEEE Transactions on Intelligent Transportation Systems*, 18(9):2549–2561, 2017.