Computer Networks 244 (2024) 110332

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

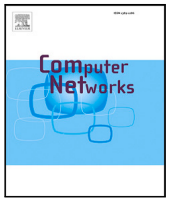

Review article

DIVA: A DID-based reputation system for secure transmission in VANETs using IOTA

Angelo Feraudo *, Nicolò Romandini, Carlo Mazzocca, Rebecca Montanari, Paolo Bellavista

Department of Computer Science and Engineering, University of Bologna, Bologna, Italy

ARTICLE INFO

Keywords:
VANETs
V2V
Decentralized identifiers
DAG
IOTA
Authentication
Security

ABSTRACT

Today's advancement in Vehicular Ad-hoc Networks (VANET) constitutes a cornerstone in ensuring traffic safety in Intelligent Transportation Systems (ITS). In this context, vehicle-to-vehicle (V2V) communications are a pivotal enabler for road safety, traffic optimization, and pedestrian protection. However, V2V communications lack effective and efficient security solutions that can adequately ensure the trustworthiness of the source of the transmitted content. In this work, we originally propose DIVA, i.e., a Decentralized Identifier-based reputation system for secure transmission in VAnets. In particular, we claim the suitability of utilizing IOTA, a Direct Acyclic Graph (DAG)-based ledger, to securely store reputation scores and of leveraging Decentralized Identifiers (DIDs) to identify participating vehicles. DIVA also incorporates and implements a reputation algorithm that computes reputation scores by analyzing both safety and non-safety messages, exchanged among vehicles and Road Side Units (RSUs) in compliance with the related European Telecommunications Standards Institute (ETSI) standards. Thus, DIVA can effectively identify malicious contributors and decrease their reputation scores. The reported experimental results clearly show the feasibility and effectiveness of DIVA, by working on an extended and comprehensive dataset of realistic V2V messages; the dataset has been made openly accessible to the research community, also to increase result reproducibility.

1. Introduction

In recent years, Vehicular Ad-hoc NETworks (VANETs) have received increasing interest from both academia and industry. These networks are essential components of Intelligent Transportation Systems (ITS), enabling the exchange of road-related messages. VANETs facilitate the broadcasting of this information, enhancing vehicle awareness and contributing to the overall safety of the driving environment.

The content of these messages can differ due to varying specifications in different regions. For example, the European Telecommunications Standards Institute (ETSI) has made significant advancements in the field of ITS with the publication of several documents aiming to constitute a set of standards for the development of Cooperative Intelligent Transport Systems (C-ITS) [1]. Specifically, ETSI defines message structures for both periodically exchanged messages (non-safety messages) and messages exchanged in exceptional situations (safety messages) [2,3]. The former is designed to enhance awareness among vehicles within the VANET, while the latter is a valuable tool in the effort to reduce traffic accidents.

Given the content of these messages and their pivotal role in maintaining traffic safety, it is imperative to guarantee the trustworthiness of the conveyed information. Moreover, VANET scenarios present a multitude of new security challenges [4–6] that require utmost attention and proactive measures. For example, an attacker may exploit vulnerabilities to forge messages containing false or misleading information, potentially leading to severe consequences such as car accidents or misrouted vehicles. Additionally, passive attacks, such as packet eavesdropping, can be carried out to gather sensitive information about vehicles and traffic in specific areas.

Therefore, it is of crucial importance to establish mechanisms that guarantee the authenticity and integrity of messages within the VANET, while simultaneously safeguarding the privacy of end-users. In this direction, numerous solutions have been proposed, many of which rely on reputation systems [7–9]. These systems typically assign a numerical value to each vehicle based on the whole history of message interactions, thus summarizing the vehicle's level of trust. By utilizing suitable detection algorithms, it becomes feasible to determine the authenticity and truthfulness of the transmitted messages, which will be used to update reputation values. However, most existing reputation-based solutions for VANET do not take into account the truthfulness of

* Corresponding author.
E-mail addresses: angelo.feraudo@unibo.it (A. Feraudo), nicolo.romandini@unibo.it (N. Romandini), carlo.mazzocca@unibo.it (C. Mazzocca), rebecca.montanari@unibo.it (R. Montanari), paolo.bellavista@unibo.it (P. Bellavista).

https://doi.org/10.1016/j.comnet.2024.110332
Received 19 September 2023; Received in revised form 29 January 2024; Accepted 12 March 2024
Available online 16 March 2024

the information contained within the standard-defined message structures. Additionally, the employed strategies perform score computation directly on vehicles; thus, requiring the processing of the information included in the communications and generating non-negligible delays during message exchanges.

Besides these limitations, reputation-based schemes must tackle various technical challenges stemming from the highly dynamic nature of the VANET environment. In particular, the use of temporary identifiers (such as pseudonyms, as referenced in [10]) for vehicle identification, combined with the frequent changes in network topology within a VANET, pose challenges in the definition of reputation scores and in the collection of adequate vehicle information. Furthermore, the computation and propagation of reputation scores require time, which can be a significant concern for VANET-based applications with latency-sensitive requirements.

To deal with the aforementioned challenges, this paper proposes DIVA, an innovative Decentralized Identifiers (DID)-based reputation system for secure transmission in VAnets. The primary objective of DIVA is to offer a reliable solution to identify malicious messages by leveraging a reputation mechanism based on Distributed Ledger Technologies (DLTs). DIVA requires every vehicle to have a unique identifier and share road information with other entities within the network: for this purpose, involved vehicles are identified with DIDs, while their participation is regulated using Verifiable Credentials (VCs). Initially, each DID is associated with a default reputation score, which is then adjusted based on the history of exchanged message contents and their truthfulness. Reputation scores are stored in the IOTA Tangle, a Direct Acyclic Graph (DAG)-based ledger, which is spread across multiple edge nodes covering a specific geographical area, e.g., a city. A vehicle can obtain the reputation of other participants only if its reputation overcomes a given threshold. The computation of reputation scores occurs at edge nodes, without incurring in computation overhead and delays on the participating vehicles. Let us note that the DIVA reputation computation originally employs the message format defined by ETSI, thus ensuring standard compliance and full compatibility with the emerging road infrastructures.

To show the feasibility and effectiveness of DIVA, we have also generated an extensive dataset including non-safety and safety messages that are exchanged by vehicles in V2V communications. To the best of our knowledge, this represents the first effort to provide an ETSI-compliant message dataset, thus favoring the adoption of the standard in the VANET research community. The dataset was curated by capturing messages within the context outlined in [11]: building upon this foundation, we further manipulated the dataset by introducing controlled instances of malicious messages. This modification was crucial for evaluating how DIVA reacts when a vehicle misbehaves. By exploiting this extensive dataset, we have performed a wide set of experiments to test the feasibility and efficiency of DIVA. Experimental results reported in the paper show that DIVA accurately detects malicious contributions in VANET environments, by achieving an accuracy of approximately 90% across various scenarios. Moreover, we have made the DIVA implementation openly accessible to the community [12], serving as a valuable solution to conduct thorough comparisons on this topic. By summarizing, the primary original contributions of the paper are as follows:

- We propose a novel DID-based reputation scheme for VANET, with full compliance with ETSI standards;
- We have generated the first extensive ETSI-compliant communication dataset and made it openly available to the research community in the field;
- We have implemented and extensively evaluated DIVA against that dataset and also in 5G-enabled deployment scenarios, by showcasing the DIVA effectiveness in handling malicious contributions.

Table 1
Table of acronyms.

Abbreviation	Definition
CAM	Cooperative Awareness Message
DAG	Direct Acyclic Graph
DENM	Decentralized Environmental Notification Message
DID	Decentralized Identifier
DLT	Distributed Ledger Technology
OBU	On Board Unit
RSU	Road Side Unit
VANET	Vehicular Ad-hoc NETwork
VC	Verifiable Credential
VP	Verifiable Presentation
V2V	Vehicle-to-Vehicle
V2I	Vehicle-to-Infrastructure
V2X	Vehicle-to-Everything

The remainder of the paper is organized as follows. In Section 2, we concisely present the needed background about VANET and DLT. Section 3 presents the DIVA system model, while Section 4 is dedicated to presenting the originally defined and implemented DIVA protocols. Section 5 offers an extensive performance evaluation of DIVA under different circumstances, while Section 6 provides the readers with the security analysis of the proposed reputation scheme. The description of related work employing DLT and some conclusive remarks end the paper.

2. Background

This section offers the needed background to understand the key elements of our proposal. Additionally, for the sake of clarity, we summarized the acronyms used in this paper in Table 1.

2.1. VANET

VANETs are self-organizing networks that rely on vehicles to function as routers or mobile nodes with a wide communication range. To dynamically establish the network, vehicles are equipped with an On Board Unit (OBU), which enables communication with neighbors. In traditional VANET scenarios, Roadside Units (RSUs) are strategically positioned alongside roads to provide service access and facilitate wide-range communication between vehicles.

Currently, two primary technologies promote VANET communications: IEEE 802.11p [13] and Cellular V2X (C-V2X) [14,15]. The 802.11p standard operates at the physical and medium access control (MAC) layers, while it utilizes the 802.11 wireless technology to enable wireless V2V and vehicle-to-infrastructure (V2I) communications. However, despite the good performance of this standard, it does not assure satisfactory reliability of message delivery. In 2019, the Internet Task Force initiated the development of a novel V2X communication standard, known as IEEE 802.11bd. This standard is based on the IEEE 802.11ac (i.e., Wi-Fi 5) which makes it more powerful than its predecessor IEEE 802.11p. Specifically, it should guarantee twice the performance of IEEE 802.11p [16] in terms of throughput, latency, reliability, and communication range while guaranteeing backward compatibility with the previous standard. Furthermore, the use of Low-density Parity-check Coding (LDPC) and Midambles enables IEEE 802.11bd to achieve better reliability of message delivery compared to IEEE 802.11p, i.e., about 88% vs 75% [17].

The message specification for this technology may vary across regions. For example, in the United States V2V communications are governed by the Wireless Access in Vehicular Environments (WAVE) standard [18], while in Europe, the ETSI group has established the ITS-G5 standard [19]. Both variants include messages with comparable information but different encoding, such as the Basic Safety Message (BSM) in WAVE and the Cooperative Awareness Message (CAM) in ITS-G5. An example of a protocol facilitating the dissemination of

these messages is the GeoNetworking Protocol [20]. This protocol leverages geographical positions for data packet distribution, relying on a geographical routing process to direct messages to specific geographic areas.

On the other hand, C-V2X, developed by the 3rd Generation Partnership Project (3GPP), leverages the existing cellular network infrastructure to offer a unified solution for V2V, V2I, and vehicle-to-pedestrian (V2P) communications. Initially introduced in Release 14 (Long-Term Evolution V2X) [14], C-V2X supported only broadcast communications. However, with the introduction of New Radio V2X in Release-16 [15], support was extended to unicast and groupcast communications. In such a scenario, vehicles exploit the PC5 interface for direct V2V communications independent of cellular networks, while the Uu interface is used to facilitate traditional cellular communications, enabling vehicles to receive information about road and traffic conditions.

2.2. Decentralized identifiers and verifiable credentials

DIDs [21] have transformed identifier systems in decentralized identity frameworks by uniquely identifying DID Subjects, whether human or non-human entities [22]. A DID consists of three parts: a Uniform Resource Identifier (URI), a specific DID method identifier, and a method-specific DID identifier. Each DID resolves to a machine-readable JSON-LD document known as a DID Document, containing cryptographic public keys, service endpoints, authentication parameters, timestamps, and metadata. DIDs eliminate the need for identity providers and centralized authorities, allowing entities to prove ownership by using private keys corresponding to the embedded public keys in the DID document. Verification is achieved by accessing the public DID Document, typically shared through a verifiable data registry often implemented via DLTs.

VC [23] is another specification formalized by the W3C that provides a standardized data structure for representing cryptographically verifiable and tamper-proof claims. VCs play a key role in an ecosystem comprising holders, issuers, verifiers, and a verifiable data registry. Holders are entities controlling one or more VCs, while issuers create new VCs. In a vehicular context, the issuer could be the the Ministry of Transport or the Department of Motor Vehicles. Verifiers leverage VCs to establish trust, i.e., an RSU collecting data from a vehicle. The verifiable data registry manages the creation and verification of identifiers, keys, credential schemas, and related data. A VC includes elements like the subject URI, issuer's URI, unique credential identifiers, claim expiration conditions, and cryptographic signatures.

Furthermore, the W3C Verifiable Credentials Working Group introduced verifiable presentations (VPs), specifying methods for signing and presenting VCs by holders to prove their ownership.

2.3. Distributed ledger technologies

DLTs are decentralized data sources that eliminate centralized storage and administration. The information stored in DLTs is replicated across multiple nodes, ensuring that each node maintains a coherent copy of the data. The collaborative nature of DLTs allows various entities to contribute without relying on centralized or third-party intermediaries. Unlike traditional databases, data stored in DLTs cannot be tampered with due to their append-only data structures. These systems rely on decentralized Peer-to-Peer (P2P) networks, avoiding the vulnerabilities associated with a centralized control entity and the risks of a single point of failure. To ensure network synchronization, DLTs employ consensus protocols based on strong cryptographic principles that guarantee trust among participants. In this work, we use DLTs to maintain DID Documents and the reputation values of each vehicle.

DLTs can be categorized according to the data structure used to store transactions. The most popular options are blockchains and DAGs. A blockchain stores information in blocks linked together by hash pointers. This design enables the detection of data tampering, as altering a block would disrupt the entire chain. In contrast, a DAG is a data structure that does not follow the traditional linked list of blocks since it is organized as a directed graph without cycles. Moreover, DLTs can also be devised based on their access model. Permissionless ledgers are open and accessible to the public, allowing anyone to participate. It operates with complete decentralization involving unknown parties. On the other hand, permissioned ledgers restrict participants' ability to write data, read data, or both. This setup generally leads to partial decentralization. In this work, we employ a permissioned DAG-based ledger, wherein only a restricted set of nodes is authorized to both publish and read transactions.

Table 2
Example of edge node internal table fields.

Field name	Type	Description
$vehicle\_did$	String	DID uniquely assigned to a vehicle.
$repScore$	Double	Reputation Score associated with the DID identifying the vehicle.

2.3.1. IOTA

IOTA [24] is a next-generation DLT designed to address the scalability issues of blockchain technology, while maintaining the same security features. It utilizes a unique data structure called Tangle, which is a DAG-based ledger made up of interconnected nodes that store transactions. Transactions are validated by the nodes they are connected to, allowing for fast performance without the need for middlemen such as miners or validators. The Tangle also allows for zero-value transactions, which do not require validation by network participants since they do not involve any transfer of value. As a result, they can be attached to the Tangle without the risk of double spending, significantly reducing the time required to share information. As blockchain, also IOTA networks can be deployed as private or public. In addition, IOTA distinguishes clients and nodes. A client is any entity (i.e., human or not) that submits transactions to a node. A node is responsible for verifying their correctness and attaching them to the Tangle.

Finally, an IOTA network comprises additional node types named *Coordinator* and *Permanode*. In each IOTA network, there is a unique Coordinator that regularly produces *milestones*, trusted signed transactions used by nodes to confirm transactions. The signature guarantees that nobody can fake the signatures on milestones, thus, they are always legit. In particular, a transaction is confirmed only when directly or indirectly referenced by a milestone that nodes have validated. It is worth noting that the use of the Coordinator is temporary since it will be removed in incoming updates. On the other hand, Permanodes are responsible for keeping the history of all the transactions that occurred. Such a component is particularly relevant in specific scenarios since nodes may be constrained devices that cannot memorize the entire Tangle, thus periodically deleting recorded transactions through a pruning operation.

3. The DIVA architecture and model

In this section, we outline the key components of DIVA. We first describe the elements encompassed within our architecture along with the underlying assumptions. Then, we present the communication requirements and a comprehensive overview of the adversarial model.

3.1. General overview

Our scenario comprises a set of edge nodes, which are interconnected through a 5G core network, and considers the IEEE 802.11p standard for both V2V and V2I communications. In this context, DIVA

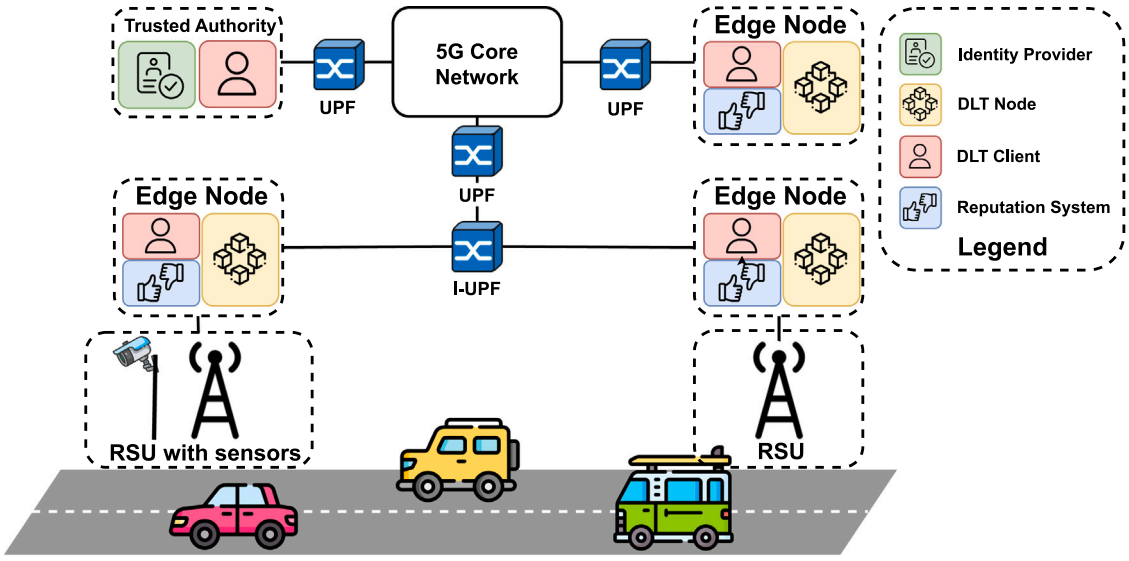

Fig. 1. System Model.

offers a reputation system that stores reputation-based information on IOTA, which is maintained across multiple edge nodes. An edge node serves as the entity providing computational resources in the form of a small data center located at the network edge, e.g., close to car drivers [25–27]. The precise location of these nodes might be determined by the network infrastructure deployment, as discussed in [28].

In a VANET scenario, edge nodes may act as collection points for messages from vehicles and perform computations to determine the trustworthiness of each participant. Edge nodes are responsible for updating the vehicles' reputation through transactions stored in the ledger. Furthermore, we assume that they are designed to be resilient against tampering or unauthorized access. Each edge node also maintains an *internal table* that holds the reputation associated with the vehicles, enabling them to quickly respond to reputation requests without querying the ledger. This organization minimizes the computation overhead on the vehicle side since participants in the VANET can directly leverage these data to assess the trustworthiness of other vehicles. Table 2 illustrates the fields maintained within the internal table.

Vehicles are uniquely identified through DIDs, which are assigned and registered on the ledger by a Trusted Authority (TA). This entity plays a crucial role in ensuring the integrity of vehicles and granting them access to the network. As illustrated in Fig. 1, the TA consists of an Identity Provider, such as the Ministry of Transport in Italy or an authorized inspection center, which is responsible for periodically verifying the vehicles' compliance with legal requirements by detecting any physical tampering or modifications. Once the inspections are completed, the Identity Provider issues a VC for the vehicle's DID, certifying its eligibility.

Finally, RSUs, assumed to be integral and trusted components of the infrastructure, play a central role within the system to verify the message reliability. The presence of an RSU, providing road-related data, may help in the identification of malicious vehicles, as their contributions inherently deviate from that of the RSU. However, it is worth noting that not all RSUs are equipped with sensors for monitoring road-related events. Some RSUs may solely work as relay nodes, extending vehicles' communication range and facilitating communication with the infrastructure, such as edge nodes or cloud services. Hence, RSUs providing trusted road data are not uniformly distributed across all roadways. In instances where RSU sensors are absent, RSU functioning as relay nodes, the system operates under the assumption that the majority of vehicles are benign.

3.2. Communication requirements

DIVA strives to enhance the safety and efficiency of vehicular communications by adhering to a comprehensive set of requirements commonly used in the literature [29,30]:

1. *Confidentiality*: information exchanged between vehicles and RSUs has to be protected from unauthorized access or eavesdropping attempts.
2. *Integrity*: tampering or alteration of messages has to be detected, thereby assuring that the received information is trustworthy and unaltered.
3. *Authentication*: interactions in the VANET have to be authenticated, ensuring that communication is only established with trusted entities.
4. *Privacy*: no entity in the VANET can infer the real identity of the user vehicle;
5. *Traceability*: messages have to be traced back to their origin, enabling accountability and facilitating legal actions against malicious participants. Traceability also serves as a deterrent, discouraging potential attackers from engaging in malicious activities within the network.
6. *Non-repudiation*: vehicle cannot disown or repudiate its role as the sender of a particular message.

3.3. Adversarial model

The security of communications heavily depends on the robustness of the proposed scheme against potential attacks. In the assessment of the security, presented in Section 6.2, we use the Dolev–Yao model adversary [31], which allows for a comprehensive evaluation of the system's resilience. The Dolev–Yao model assumes that the attacker can intercept any message passing through the network, initiate communication with any entity, and even act as the intended receiver of any transmission. We consider that a VANET can be susceptible to the following attacks:

1. *Eavesdropping Attack*: By compromising the confidentiality of the information, the attacker gains unauthorized access to sensitive data, such as location information, personal details, or messages exchanged between vehicles.

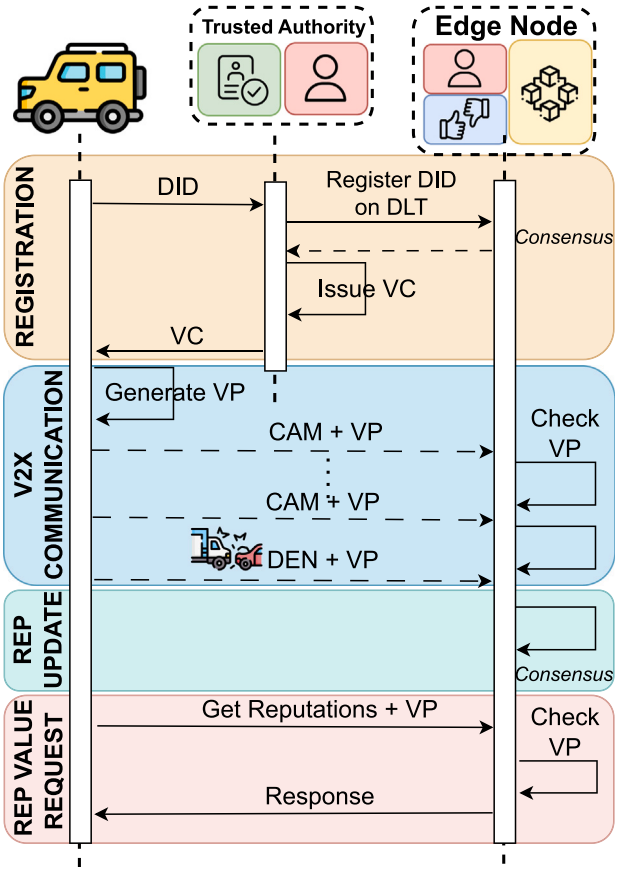

Fig. 2. DIVA workflow.

2. *Replay Attack*: The adversary captures and subsequently retransmits a previously recorded message to mislead other vehicles. This can lead to unexpected situations, such as traffic management, collision avoidance, and cooperative driving.
3. *Forgery Attack*: The attacker attempts to impersonate legitimate users or entities within the VANET. By forging messages, the attacker can mislead neighboring vehicles and manipulate their behavior, leading to unintended consequences or hazardous situations.
4. *Sybil Attack*: The adversary creates multiple fake identities or vehicles to deceive other participants within the VANET. This can be exploited to disrupt routing protocols, manipulate traffic patterns, or spread false information within the network.

4. DIVA at work

This section gives the design details about how the DIVA entities interact at runtime, whose workflow is depicted in Fig. 2. The DIVA operations supported in our current implementation can be classified into: registration, V2X communication, and reputation update.

4.1. Registration

To become a DIVA participant, a vehicle must first register a DID on the ledger and have a VC issued by a TA such as the Department of Motor Vehicles (DMV) of a city. The TA is responsible for verifying that the vehicle is compliant with the law and has not been stolen or altered. Once the vehicle successfully meets the required checks, it generates a DID and submits it to the TA. This ensures that only the holder of the DID retains complete control over its associated information. The TA interacts with the ledger to record the DID and the corresponding DID Document. Then, it issues a VC that confirms the DID's successful registration and returns it to the vehicle. Given the VC, the vehicle can generate a VP, which will be appended to each communication. The VP serves as evidence of the vehicle's identity and its ownership of the issued VC.

DIDs and VCs are completely anonymous, as they do not contain any sensitive information that can be traced back to the identity of the vehicle or its user. This ensures the privacy of the vehicle owner and helps prevent any leaks.

4.2. V2X communications

As described in Section 2.1, vehicles may use different technologies to communicate. Within the scope of this paper, we focus on V2V and V2I communications compliant with the ETSI standard and based on the 802.11p technology.

We identify two main categories of messages between vehicles: *safety messages* and *non-safety messages*. The former indicates exceptional situations such as accidents or road closures, which can have a substantial impact on traffic conditions. These messages are of utmost importance as they provide critical traffic information, playing a vital role in ensuring traffic safety. On the other hand, vehicles engage in periodic exchange of non-safety messages to share details regarding their speed and position. This continuous exchange fosters a comprehensive awareness of the surrounding environment among the vehicles. In this direction, the ETSI standards define two main message types to cover the aforementioned cases: CAM [2] and Decentralized Environmental Notification Message (DENM) [3]. CAM is used by vehicles to create awareness among one another. It includes status and attribute information, like vehicle position, speed, and activated systems, and is exchanged periodically among V2V and between vehicles and infrastructure V2I. On the other side, the DEN messages are exchanged during exceptional situations to notify road users of a detected event to improve road safety and traffic efficiency. The information exchanged may include the device originating the event, situation description, detection time, and event location.

In our system scenario, message dissemination is facilitated by the GeoNetworking Protocol [20], briefly introduced in Section 2.1. Specifically, the forwarding scheme adopted in cooperative awareness basic service is Single-Hop Broadcasting, where messages are sent only to one-hop neighborhoods. Conversely, the decentralized environmental notification basic service employs GeoBroadcasting for packet forwarding. In this case, packets are forwarded hop-by-hop until they reach the destination area specified within the packet. Once the packet reaches the destination area, nodes within the area start packet rebroadcasting to ensure its dissemination.

4.3. Reputation update

In DIVA, vehicles communicate with surrounding entities, including RSUs, via Single-Hop and GeoBroadcasting protocols. Therefore, we assume that each edge node can intercept the transmitted messages within the VANET. Fig. 2 shows the edge node behavior upon receiving a DENM. After verifying the VP, the edge node initiates the process of evaluating the received information and proceeds to refresh the reputations of the involved vehicles.

4.3.1. Calculating reputation scores

The DIVA reputation system consistently treats information originating from RSUs equipped with traffic-monitoring sensors as inherently accurate. The inclusion of these RSUs enables the establishment of a definitive reference point for verifying the authenticity of vehicle-generated messages and subsequently assigning reputation scores. Nevertheless, there exist certain scenarios where RSUs lack sensor equipment, making them incapable of generating messages on

Algorithm 1 Reputation Score Computation.

Input: $\tau_{eventType}, \tau_{sitQuality}, \tau_{cam}, \omega_{rsu}, \omega_{msg}, defScore, m$
Output: $repScores$
Require: $m_{age} \geq \tau_{eventType}; m_{sitQuality} \geq \tau_{sitQuality}$
$rsuScore \leftarrow (\omega_{rsu} * defScore)$
$msgCohScore \leftarrow (\omega_{msg} * defScore)$
if $initialReputationExists(m_{source})$ **then**
  $repScore \leftarrow loadInitalRepScore(m_{source})$
else
  $repScore \leftarrow defaultRepScore$
end if
if not $messageInsideEdgeArea(m)$ **then**
  return
end if
if not $messageCoherentWithRSU(m)$ **then**
  $repScore \leftarrow repScore - rsuScore$
end if
$camCoh \leftarrow computeCAMCoherency(m, \tau_{cam})$
if $camCoh < 10$ **then**
  $repScore \leftarrow repScore - msgCohScore$
else
  if $10 \leq cam\_coh \leq 30$ **then**
    $repScore \leftarrow repScore$
  else
    $repScore \leftarrow repScore + msgCohScore$
  end if
end if
$similarEvents \leftarrow findSimilarEvents(m, \tau_{eventType}, \delta_{eventType})$
if $len(similarEvents) > 2$ **then**
  $repScore \leftarrow repScore - msgCohScore$
else
  $repScore \leftarrow repScore + msgCohScore$
end if
$updateRepScore(did_m, repScore)$ ▷ Applying Eq. (1)

Algorithm 2 CAM coherency.

Input: $denm, \tau_{cam}$
Output: $camCoherency$
function COMPUTECAMCOHERENCY($denm, \tau_{cam}$)
  $cams \leftarrow loadCam(denm_{source}, \tau_{cam})$
  $camCoh \leftarrow 0$
  for all cam **in** $cams$ **do**
    $distance \leftarrow getDistance(denm_{evPos}, cam_{refPos})$
    if $distance \leq \delta_{eventType}$ **then**
      $camCoh \leftarrow camCoh + 1$
    end if
  end for
  return
  return $(\frac{camCoh}{\text{len}(cams)}) \times 100$
end function

traffic conditions. To address this challenge, DIVA provides an outlier detection algorithm that identifies messages generated by vehicles that remarkably deviate from the average. Following this evaluation, a reputation score is assigned to the DID responsible for generating that message.

Each DID is associated with a numeric value $r_{DID} \in [0, 1]$ representing the vehicle reputation. The outlier detection algorithm, outlined in Algorithm 1, requires the freshness of the messages $\tau_{eventType}$, which depends on the type of event leading to its dissemination, and the situation information quality above a certain threshold, $\tau_{sitQuality}$. These

Algorithm 3 DENM coherency.

Input: $denm, \tau_{eventType}, \delta_{eventType}$
Output: $similarEvents$
function FINDSIMILAREVENTS($denm, \tau_{eventType}, \delta_{eventType}$)
  $centroidsDenms \leftarrow loadCentroids()$
  for all el **in** $centroidsDenms$ **do**
    $eventType \leftarrow el[eventType]$
    $\tau_{centroid} \leftarrow el[time]$
    $\delta_{centroid} \leftarrow el[space]$
    if $eventType = denm_{eventType}$ **and** $abs(\tau_{centroid} - denm_{detTime} \leq \tau_{eventType}$ **and** $distance(\delta_{centroid} - denm_{eventPos}) \leq \delta_{eventType}$ **then**
      $el[time] \leftarrow \frac{\tau_{centroid} + denm_{detTime}}{2}$
      $el[space] \leftarrow centroid(\delta_{centroid}, denm_{eventPos})$
      return $element[denms]$
    end if
  end for
end function

data are extracted directly from the message content. For example, in the DEN message structure, information quality is a value ranging from 0 to 7, indicating the quality level of information provided by specific vehicles. This may be influenced by the condition of the sensor that gathers that information. Moreover, both the CAM and DENM structures contain information about when the message was generated, allowing for the identification of its age. Hence, messages identified as excessively dated or of low quality are excluded from the reputation computation, since they are usually ignored by vehicles within VANETs.

After confirming the quality and freshness of the message, the algorithm verifies whether the event position indicated by the message m falls within the area managed by the respective edge node through the function $messageInsideEdgeArea(m)$. Then, it compares the content of the message with the information provided by the RSU and assesses its coherence with the CAMs generated by the same source. It is worth noting that we only consider CAMs falling within a time window defined by a threshold τ_{cam}. The time window can be adjusted according to the dynamic nature of the VANET. For instance, in a highway scenario where vehicles move at higher speeds, older CAMs may not accurately reflect the current state of the environment. In such cases, a smaller window is more appropriate. As shown in Algorithm 2, the coherency assessment relies on the Euclidean distance between the event location indicated in the DENM and the vehicle positions specified in the corresponding CAMs.

Furthermore, the algorithm evaluates the similarity of the event provided in the DENM with those already received using a centroid-based approach, as outlined in Algorithm 3. This involves a progressive computation of two centroids: a time centroid (denoted as $\tau_{centroid}$) and a distance centroid (denoted as $\delta_{centroid}$). Upon receiving the first DENM, they are initialized with DENM detection time and position. Subsequent DENMs trigger updates to the centroids by averaging the old centroids with the detection time and position contained in the received message. The algorithm uses the time centroid to determine a time window, $|\tau_{centroid} - cadDen_{detTime}| \leq \tau_{eventType}$, in which the detection time has to fall to be considered similar in temporal terms. The spatial aspect undergoes a similar process, employing $\delta_{centroid}$ along with the specified radius $\delta_{eventType}$ to establish an area, where the message's position is expected to fall for similarity assessment. This process produces a list of similar DENMs employed to assess the accuracy of the current DENM.

All these evaluations are utilized to compute a reputation score, denoted as $repScore$, which is then used to update the reputation of the vehicle responsible for generating the DENM.

4.3.2. Updating reputation scores

The edge node updates the reputation of the vehicles according to the following equation:

$$(r_{DID})_t = \alpha \times (r_{DID})_{t-1} + \beta \times ((r_{DID})_{t-1} + repScore) \quad (1)$$

where $(r_{DID})_t$ represents the updated reputation score of the source DID, while $(r_{DID})_{t-1}$ refers to the reputation score from the previous time step $t-1$. The α and β parameters dictate the relative impact of each contribution. Their combined value is restricted to be equal to 1, indicating that they jointly shape the overall influence on the updated reputation score. Thus, the new reputation score is determined through a combination of the previous reputation and the score computed by Algorithm 1. It should be noted that a variation in the reputation score is contingent on the type of misbehavior exhibited by the vehicle. Each misbehavior results in a degradation of the reputation score determined by the corresponding weight (e.g., ω_{msg} for message coherency). Additionally, these weights can be dynamically adjusted based on the geographical area managed by the respective edge node.

4.3.3. Storing reputation scores

The final value of the reputation score is stored in the Tangle through a zero-value transaction. However, this requires that there can be no conflicts, i.e., two edge nodes that contemporary update the reputation of the same vehicle. In VANETs, we can assume that the edge nodes are positioned distant enough, thus preventing a vehicle from communicating simultaneously with two different nodes. Indeed, in real-world deployments, edge nodes are located either at the network edge, near the Radio Network node (e.g., RSU), or at network aggregation points, where they manage multiple Radio Network nodes [28]. Hence, by defining the area managed by the edge nodes as the coverage of their associated radio network nodes, the intersection of areas supervised by distinct edge nodes becomes negligible. Furthermore, the time needed for a vehicle to move from one node's coverage to another one exceeds the time required to update the vehicle's reputation.

Each vehicle can now request reputations related to the area managed by that edge node. To obtain the latest updated reputations, a vehicle must explicitly request them to provide its VP. This ensures that only authorized users can access the reputation data, excluding contributions coming from malicious vehicles.

5. DIVA evaluation

To evaluate the performance of the DIVA prototype, we run extensive simulations running on realistic deployment settings. We first created an extensive dataset with the CAMs and DENMs exchanged between vehicles in the scenario presented in [11]. Subsequently, we manipulated the dataset containing DENMs to create some malicious instances.

We built our simulation network environment by connecting 5 edge nodes, using the OMNeT++ simulation environment along with the Artery framework [32] and Simu5G communication library [33]. Artery is a VANET simulator that operates within the OMNeT++ environment and incorporates SUMO [34] for road traffic modeling. Simu5 g realistically models both the core network and the Radio Access Network (RAN) of a 5G network. We employed Artery to collect V2V messages that we originally generated in full compliance with ETSI standards. On the other hand, we used Simu5G to simulate network-related behaviors. This comprehensive approach allowed us to rigorously evaluate the performance and effectiveness of DIVA in a controlled, yet realistic, experimental setup.

5.1. V2V communications dataset

To evaluate the effectiveness of our approach in detecting malicious messages, we created a dataset gathering ETSI-compliant messages exchanged during a road hazard event. This dataset has been made publicly accessible,[1] facilitating further research and evaluation. The dataset was generated considering the scenario analyzed in [11], where the authors extended the Artery tool [35] to enable dynamic VANET scenarios and present a Decentralized Environmental Notification use case.

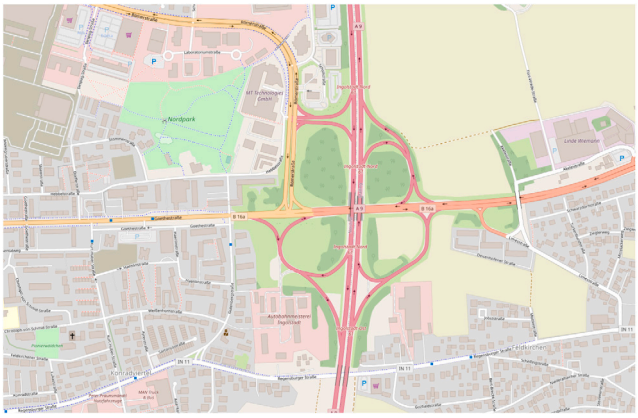

Fig. 3. Ingolstadt North highway junction.

This scenario involves the modeling of a limited sight zone, such as a foggy area where entering vehicles suddenly reduces their speed. Fig. 3 illustrates the region of interest under study, which corresponds to the area near a highway junction *Ingolstadt Nord* in Ingolstadt, Germany. In this area, a car at the Ingolstadt Nord junction performs an emergency stop, thus causing a collision risk that triggers the broadcasting of DENMs. Generally, vehicles detecting a critical situation with high crash probability generate a DENM containing event-specific details such as the *detection time*, *cause code*, *event position*, and more. These messages are collected in our dataset as they play a crucial role in malicious behavior detection. Moreover, we collect CAMs generated within the designated area, as they constitute an integral part of the procedure for validating the reliability of information within DENMs in the DIVA system. This encompasses identifying any inconsistencies in detection time and vehicle positions. In Table 3 we provide an overview of the key components of the DENM and CAM datasets with their relative features.

Additionally, to evaluate DIVA's capacity to identify DENM malicious content, we manipulated the DEN dataset by introducing specific malicious instances. Initially, we select a subset of sources to undergo malicious transformations. The approach involves introducing distortions to values by injecting noise into messages sent by these sources, specifically targeting some fields such as detection time, longitude, latitude, and altitude. However, it is not sufficient to simply add random noise to those values, as it might lead to nonsensical alterations for the respective columns, such as negative values. To address this, a normal distribution is created for each column, centered in zero, with a standard deviation tailored to match the specific deviation of that column. Each value in the messages from the selected sources is then adjusted by adding a randomly sampled value from the corresponding distribution. This process ensures that, despite the modifications, the resulting values remain realistic and fall within the acceptable range for each column. Each value undergoes independent modification, by introducing a varying degree of distortion even for messages from the same source, thereby by adding complexity to the task of identifying malicious sources.

[1] https://github.com/MMw-Unibo/ETSI-V2V-Dataset

Table 3
Summary of dataset characteristics.

Dataset name	Description	Illustrative features
DENM dataset	It comprises a subset of the information contained in DEN messages exchanged during the execution of the simulated scenario detailed in [11].	source; situation_eventType; detection_time; simulation_time; eventPos_lat; eventPos_long; eventPos_alt.
CAM dataset	It includes a selection of data extracted from CAMs exchanged within the context of the simulated scenario outlined in [11].	source; referencePositionLat; referencePositionLong; referencePositionAlt; simulationTime.

Table 4
Threshold Study with 20% of malicious vehicles.

Threshold	Event type	$\alpha = \beta$	TPR	TNR	FPR	FNR
Mode	dangerousEndOfQueue	0.5	100	49.51	50.49	0
	collisionRisk	0.5	87.23	92.8	7.20	12.77
	trafficCondition	0.5	100	78.48	21.52	0
	Total	0.5	99.93	78.14	21.86	0.07
Median	dangerousEndOfQueue	0.5	100	49.51	50.49	0
	collisionRisk	0.5	78.72	92.80	7.20	21.28
	trafficCondition	0.5	100	78.48	21.52	0
	Total	0.5	99.89	78.14	21.86	0.11
Mean	dangerousEndOfQueue	0.5	100	100	0	0
	collisionRisk	0.5	61.70	100	0	38.30
	trafficCondition	0.5	100	100	0	0
	Total	0.5	99.80	100	0	0.20

5.2. Experiments

In this section, we explore the parameter selection and DIVA performance in detecting malicious messages. We implemented DIVA[2] and assessed its performance using the dataset detailed in Section 5.1. The algorithm has been implemented through a Python script that processes every single message characterizing the DEN-based dataset. All experiments were conducted on a Linux virtual machine equipped with 16 CPUs and 32 GB of RAM.

5.2.1. Thresholds selection

This study aims to determine the thresholds that most effectively enable DIVA to distinguish between malicious and non-malicious messages. To find their optimal values, we employed three distinct aggregation metrics, namely *mode*, *median*, and *mean*. Such metrics consider the distances from the vehicle generating the event and the event position itself, obtained by combining CAMs content with information provided in the current DENM. The computation of these thresholds relies on the benign datasets described in Section 5.1. Concerning the message age, we used the thresholds described in different reference documents (e.g., [36]). To determine which CAMs are considered for coherency computation (Algorithm 2), we defined a predetermined value of τ_{cam}, set at 600 s. This choice specifically addresses the scenario examined in this manuscript, to identify Cooperative Awareness messages in proximity to a specific event. It is important to note that, as mentioned earlier, the outlier algorithm provides flexibility to fine-tune this value for the specific geographical area under consideration.

5.2.2. DIVA performance

We conducted several experiments varying the values of α and β to determine the optimal parameter combination. These experiments employed the manipulated versions of the DENM dataset, considering different percentages of malicious sources.

Initially, we run DIVA to compute the reputation of sources contained in the manipulated versions of DENM dataset. Then, we reprocessed the dataset to evaluate the reputation score computed in the previous step, to identify malicious messages generated by sources with a reputation score below a defined threshold denoted as $\tau_{repScore}$. This procedure was repeated ten times, incrementing β values by 0.1 each time within the range of 0 to 1.

5.3. Results and discussions

This section presents the results of our experiments aimed at identifying the most effective thresholds and evaluating the performance of DIVA in detecting and discarding malicious messages.

5.3.1. Threshold selection

To ensure a uniform evaluation of the factors outlined in Eq. (1), we analyzed the three selected thresholds, by employing identical values for both α and β. Table 4 shows the results obtained for each event type within the DENM dataset when α and β are both 0.5. It provides details on the percentage of accurately classified messages, i.e., True Positive Rate (TPR) for malicious messages and True Negative Rate (TNR) for non-malicious ones. Additionally, it indicates instances of misclassification, including False Positive Rate (FPR) and False Negative Rate (FNR) for malicious and non-malicious messages respectively. For the sake of briefness, we only report the metrics for 20% of malicious sources within the dataset. It is worth mentioning that the majority of the events refer to *traffic condition*, as it constitutes the most frequently transmitted event in DENMs. Consequently, the event type *Total*, aggregating results regardless of the situation contained in the message, is notably influenced by their outcomes.

Furthermore, the table shows that the *mode*, i.e., the most common distance value, as $\delta_{eventType}$ leads to the highest performance in detecting malicious messages, achieving a TPR of 99.93%. However, as for the *median*, when employing this threshold, the algorithm tends to act conservatively, resulting in the rejection of around 20% of non-malicious messages. On the contrary, using the *mean* ensures the accurate identification of all non-malicious messages achieving a TNR of 100%, albeit it may leave certain events, like the *collision risk* event (TPR 61.70%), unprotected. Based on this analysis, in the following section, we employ both *mean* and *mode* as $\delta_{eventType}$ to delve deeper into DIVA's performance while varying the values of α and β, as well as the percentage of the malicious sources.

5.3.2. DIVA performance

Figs. 4 and 5 illustrate the performance obtained in identifying the nature of vehicles when using *mean* and *mode* as thresholds. By comparing the graphs, DIVA demonstrates better performance when using the *mean* as the threshold. This is particularly evident, in Fig. 4(a), for values of β ranging from 0.2 to 0.6: the incidence of incorrectly identified messages approaches zero. For lower values of β, the performance of DIVA degrades as the influence of the new reputation score in Eq. (1) is minimal. Indeed, the percentage of incorrectly identified messages increases to approximately 20% for these values of β in all the scenarios considered. On the other hand, when increasing the number of malicious sources, DIVA maintains a high accuracy for values of β above 0.3, achieving around 94% with 30% of malicious sources (Fig. 4(b)), and 89% with 40% of malicious sources (Fig. 4(c)).

When considering *mode* as the space-related threshold, DIVA acts conservatively. Although it tends to accurately detect malicious messages, it also ends up discarding approximately 40% of the total received messages, including non-malicious messages. As illustrated in

[2] https://github.com/MMw-Unibo/DIVA

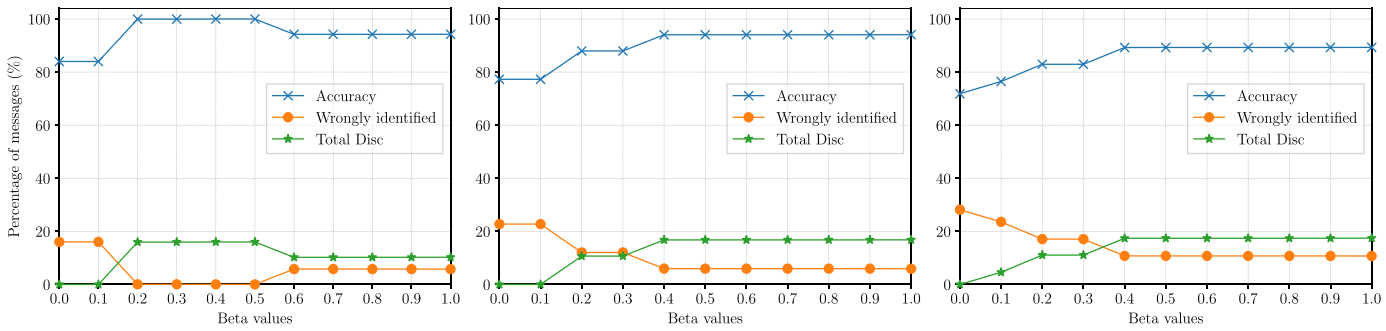

Fig. 4. Performance of DIVA using *mean* as threshold.

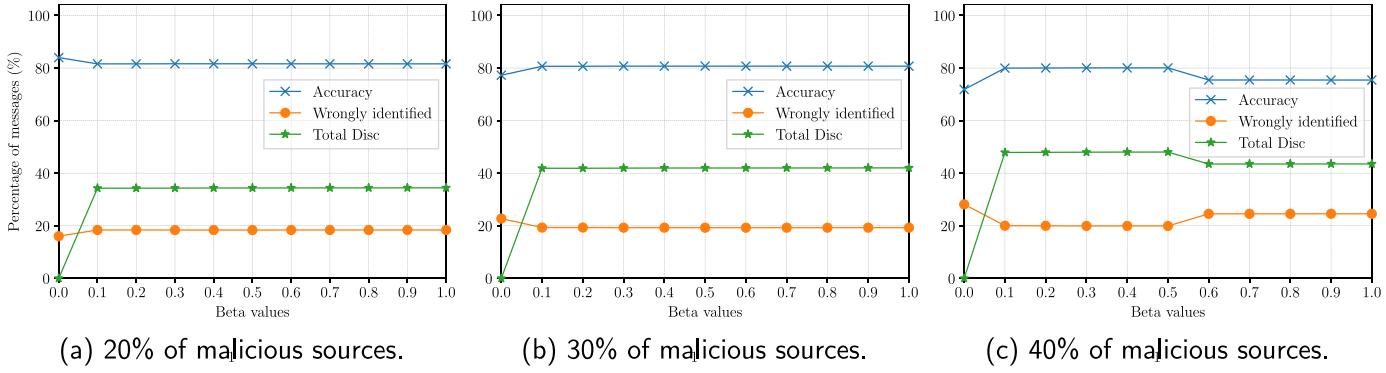

Fig. 5. Performance of DIVA using *mode* as threshold.

Fig. 5(c), this number increases with the highest number of malicious sources to approximately 50%.

These findings demonstrate DIVA's proficiency in distinguishing malicious sources, using values of β ranging from 0.4 to 0.6. By assigning them low reputation values, DIVA effectively enables vehicles with benign intent to filter out these potentially harmful sources from V2V communications. This capability enhances the overall security and trustworthiness of the VANET environment, creating a safer network for all participating vehicles.

Finally, we carried out tests on execution latency. As mentioned earlier, transactions published on the Tangle do not necessitate validation by other nodes. Therefore, the overall latency for propagating new reputations by a node is equivalent to the latency of a broadcast message. Utilizing the simulation tool, we measured this latency for 5 edge nodes, which, however, proved to be negligible, i.e. a few microseconds. It is worth noting that the number of nodes does not impact the above latency indicator since it specifically measures one-to-one communication. Furthermore, the total number of nodes is constrained by the scenario of interest, which revolves around a city setting, totaling approximately a few hundred nodes. Thus, DIVA can be deployed in a real-world scenario since the overhead introduced by the Tangle is minimal.

6. Security analysis

In this section, we present a comprehensive analysis of how DIVA effectively fulfills the security requirements outlined in Section 3.2 and offers robust protection against the adversary model discussed in Section 3.3.

6.1. Communication requirements

Confidentiality. All V2V communications are geographically broadcast in the VANET. Usually, they do not contain information that can be leveraged by the attacker. Indeed, since we are referring to a cooperative awareness system, the main objective is to enable communication and collaboration between vehicles. Encrypting road information would be pointless due to the public nature of the shared information. However, if sensitive data related to the vehicle (e.g., vehicle's location or direction) are exchanged, the communication can be encrypted with the receiver's public key, which is included in the DID Document stored in the Tangle.

Integrity. Each vehicle signs messages through its private key, making it computationally infeasible for the attacker to compute a valid private key and forge digital signatures. Any messages altered by an attack can be detected by the receivers since the signature cannot be faked. Moreover, the integrity of valuable information such as the vehicle's reputation is guaranteed by storing these data in the Tangle, which is tamper-proof by design.

Authentication. The authenticity of messages is guaranteed by the VPs and reputation score. On the one hand, VPs ensure that the sender is provided with a valid VC issued by the TA and that the messages have been signed with the corresponding sender's private key. On the other hand, the reputation score provides awareness about the historical quality of the information provided. If the reputation drops under a specific threshold, messages are automatically discarded.

Privacy. To preserve privacy, the adversary cannot extract or infer the unique identity of the vehicle user. In DIVA, all the interactions are performed through DIDs, which act as pseudonyms and cannot be mapped to the user's identity. Thus, preserving the privacy of the drivers.

Traceability. To discourage malicious actions such as intentionally broadcasting incorrect road conditions, all the information is processed and evaluated to compute the vehicle's reputation. RSUs and vehicles use reputation to filter all communications with participants whose reputations drop under a certain threshold, corresponding vehicles are marked as malicious and will not receive further updates.

Non-repudiation. Non-repudiation is another security requirement aiming to deter incorrect behaviors. Vehicles cannot deny their involvement in the operations performed or the message's authenticity. In DIVA, all the communications are signed with the sender's private key, and public keys are included within the DID Documents shared through the Tangle. The private key is only owned by the vehicle sending the data, preventing later repudiating their involvement.

6.2. Adversarial model

Resistance to Eavesdropping Attack. V2V communications are geographically broadcast and involve either public information (e.g., road conditions) and/or sensitive data, such as the location and direction of vehicles. Listening to public information does not directly bring any advantages to potential attackers. On the other hand, when sensitive data are transmitted, the communication can be secured through end-to-end encryption using the receiver's public key (shared in the DID Document).

Additionally, along with vehicular information, a vehicle also attaches its VP to prove the trustworthiness of the provided data. Therefore, there is a concern that the adversary could intercept valid VPs and use them to transmit false or misleading information, disrupting the reputation of the associated vehicle or deceiving other participants in the VANET. DIVA prevents this type of attack by leveraging digital signatures. To attack our reputation system or deceive other vehicles, the attacker would require access to the private key of the sender to sign the incorrect information.

Resistance to Replay Attack. When a message is received, the edge node checks its timestamp and discards if it is not fresh enough. This helps prevent the acceptance of outdated or reused messages. Furthermore, the sender embeds a challenge - a unique random string - within its VP. This challenge is used as a safeguard against attackers attempting to reuse the VP with another verifier. It is worth noting that the VP primarily serves the purpose of verifying the sender's authorization to contribute to cooperative awareness. The security of the system lies in the fact that the attacker does not possess the private key of the sender. Without it, it is virtually impossible for the attacker to successfully spoof the sender's identity and inject falsified data into the system.

Resistance to Forgery Attack. In DIVA, each vehicle is assigned a unique DID that is linked with a key pair. To inject malicious data into the system, the adversary would need to gain access to the private key of an authorized vehicle. The private key is securely held by the respective vehicle and is not easily accessible to unauthorized individuals. Alternatively, the attacker would need to compute the secret key associated with a valid entity within the system, which is computationally infeasible due to the key generation scheme based on Elliptic Curve Cryptography (ECC).

Resistance to Sybil Attack. Each vehicle is associated with only one DID. To perform multiple registrations, the adversary should own as many vehicles as the number of identities. It is highly improbable that an adversary may buy several cars to attack a cooperative awareness system. Furthermore, once the reputations of the vehicles drop under a given threshold, all their message are automatically discarded by other participants. Therefore, also in the extreme case where the adversary may own or be able to control a large number of vehicles, its attack window is extremely narrowed to the time needed to go under the chosen threshold.

7. Related work

This section aims to provide a literature review of reputation schemes based on blockchain and DAG — the two main types of DLTs that have been employed in vehicular scenarios.

7.1. Blockchain-based

Reputation schemes based on blockchain have been widely explored in the context of vehicular communication networks. The immutable and transparent nature of blockchain can be used to build trust and incentivize good behavior among participants. Li et al. [29] introduced a novel mechanism named BDRA, which enables secure registration and authentication using a double-layer blockchain and DIDs. The first blockchain layer consists of all the authorized RSUs. While the second one devises vehicles into areas according to the coverage of the corresponding RSU. In each area, the RSU and the vehicles collaborate to form a consortium blockchain. DIDs are leveraged to reduce reliance on third-party intermediaries for the registration phase and implement the reputation feedback mechanism. This mechanism updates reputations relying on messages sent by vehicles. There is no description of whether or how the truthfulness of these messages is verified. In a scenario like the one described in the article, where all vehicles might be malicious, this mechanism could be easily exploited.

Similarly, Fernandes et al. [30] presented a decentralized reputation system based on a consortium blockchain and smart contracts called BRS4VANET. RSUs are responsible for calculating vehicles' reputations and storing them on a blockchain. The reputation is estimated by the RSUs collecting feedback from other vehicles. To preserve the privacy of vehicles, they have a pseudo-anonymous certificate, which will be revoked if the reputation falls under a given threshold. However, the authors do not specify how the certificates are revoked.

Lu et al. [37] proposed BARS, a blockchain-based anonymous reputation system to increase the security of VANETs by preventing the propagation of malicious messages. BARS uses pseudonyms obtained through public keys to preserve the privacy of users. Additionally, the authors designed a reputation system that considers both direct and indirect interactions among vehicles, utilizing blockchain technology to securely and immutably store information. Each vehicle's reputation is recorded in a certificate, along with its public key, on the blockchain. If a vehicle's reputation score drops to zero, its public key is revoked, thereby prohibiting it from interacting with others. While the reputation evaluation algorithm is explained in the paper, there is no mention of how the system verifies the truthfulness of messages.

Yang et al. [38] presented a trust management system for vehicular networks based on blockchain. Each vehicle generates a reputation score for each received message based on a Bayesian Inference Model. The score is then sent to the RSU, which aggregates all the reputation values of vehicles in the vicinity to create a block to add to the blockchain. RSUs maintain the blockchain and compete to add blocks. The authors demonstrate how the system is effective in establishing the veracity of messages and is resilient against the compromise of vehicles or RSUs. However, vehicle identity numbers are used for identification, potentially compromising privacy.

7.2. DAG-based

Reputation schemes based on DAG are a relatively new area of research. DAG-based systems offer advantages such as scalability, fast transactions, and low transaction fees. Recently, Li et al. [39] proposed a partitioned DAG-based ledger with local consistency for managing vehicular reputation in partitioned VANETs. The designed reputation scheme is based on the idea that the truthfulness of information is valuable only for nearby vehicles and their usefulness is limited over time. To establish trust, a node extracts interaction information from transactions and uses it to calculate the reputation of other nodes according to the situation. However, since the number of transactions grows over time, is unfeasible to use them to estimate reputations. As a possible solution, the authors introduced an additional data structure to store intermediate reputation calculations, leading to further overhead that impacts the overall scalability. The presented approach mainly

focuses on how to estimate and update the reputation of vehicles, neglecting their identification and communication which are key concerns to enhance the security of communication in VANETs.

Li et al. [40] presented a DAG-based reputation mechanism that aims to realize the authenticity, immutability, and accessibility of all vehicle reputations while preserving their privacy. The reputations of vehicles are used to determine the degree of protection of their privacy when a peer captures and uploads an image of a traffic accident on the DAG for mutual supervision. Although reputation should be the main theme of their proposal, the authors mainly focus on privacy concerns. As a result, they do not provide in-depth details on how the reputation score is estimated or updated. To extend the tamper-proof capability of blockchain to large-scale vehicle networks, Du et al. [41] proposed an information-sharing approach based on a DAG. To resist chain attacks, the authors designed a reputation-based rate control strategy. The reputation is used to select reliable vehicles that have to vote to reach a consensus and to limit the number of transactions published.

8. Conclusion

VANETs hold the promise of revolutionizing V2V communications, enabling novel vehicular services, and significantly contributing to road safety. The integrity and reliability of transmitted messages are crucial to ensure that vehicles can make informed and secure decisions. To address this concern, this paper introduces DIVA, a novel DID-based reputation system for ensuring secure message transmission in VANETs using IOTA. The system operates by establishing anonymous identities for vehicles via DIDs and assigning them reputation scores, which are recorded on IOTA. Within the VANET ecosystem, messages from other vehicles are evaluated based on their source vehicle's reputation score.

To validate and assess the efficacy of DIVA, we implemented and evaluated it using a dataset compliant with the ETSI standards. This dataset was generated by ourselves and was made openly accessible to the wider research community. Experimental findings conclusively demonstrate that DIVA succeeds in identifying and mitigating malicious contributions, achieving an impressive accuracy of around 99% when specific thresholds are employed. This significantly enhances trustworthiness and security among previously unknown vehicles within VANETs.

In future work, we plan to use DIVA as a labeling algorithm to identify malicious instances and use the labeled dataset to train a machine learning (ML) model that can detect other misbehaving vehicles.

CRediT authorship contribution statement

Angelo Feraudo: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Nicolò Romandini:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Carlo Mazzocca:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Rebecca Montanari:** Resources, Supervision, Writing – review & editing. **Paolo Bellavista:** Resources, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

We have shared the link to the dataset and code used in the manuscript.

Acknowledgements

This work was partially supported by the project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan program funded by the European Union - NextGenerationEU. All authors have read and agreed to the published version of the manuscript.

References

[1] T. ETSI, Intelligent Transport Systems (ITS); Cooperative ITS (C-ITS);Release 1 , Technical Report ETSI TR (101 607), 2020, pp. 448–451.

[2] T. ETSI, Intelligent transport systems (its); Vehicular communications; Basic set of applications; part 2: Specification of cooperative awareness basic service, Draft ETSI TS 20 (2011) (2011) 448–451.

[3] E. ETSI, 302 637-3 "Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 3: Specification of Descentralized Environmental Notification Basic Service", ETSI, 2014, IEEE Open (2020).

[4] I.A. Sumra, I. Ahmad, H. Hasbullah, J.-l. bin Ab Manan, Classes of attacks in VANET, in: 2011 Saudi International Electronics, Communications and Photonics Conference, SIECPC, 2011, pp. 1–5, http://dx.doi.org/10.1109/SIECPC.2011.5876939.

[5] F. Sakiz, S. Sen, A survey of attacks and detection mechanisms on intelligent transportation systems: VANETs and IoV, Ad Hoc Netw. 61 (2017) 33–50, http://dx.doi.org/10.1016/j.adhoc.2017.03.006, URL https://www.sciencedirect.com/science/article/pii/S1570870517300562.

[6] A. Bujari, M. Conti, C. De Francesco, C.E. Palazzi, Fast multi-hop broadcast of alert messages in VANETs: An analytical model, Ad Hoc Netw. 82 (2019) 126–133, http://dx.doi.org/10.1016/j.adhoc.2018.07.024.

[7] Q. Li, A. Malip, K.M. Martin, S.-L. Ng, J. Zhang, A reputation-based announcement scheme for VANETs, IEEE Trans. Veh. Technol. 61 (9) (2012) 4095–4108, http://dx.doi.org/10.1109/TVT.2012.2209903.

[8] X. Yao, X. Zhang, H. Ning, P. Li, Using trust model to ensure reliable data acquisition in VANETs, Ad Hoc Netw. 55 (2017) 107–118, http://dx.doi.org/10.1016/j.adhoc.2016.10.011, URL https://www.sciencedirect.com/science/article/pii/S1570870516302943. Self-organizing and Smart Protocols for Heterogeneous Ad hoc Networks.

[9] H. Sedjelmaci, S.M. Senouci, An accurate and efficient collaborative intrusion detection framework to secure vehicular networks, Comput. Electr. Eng. 43 (2015) 33–47.

[10] I. ETSI, Intelligent Transport Systems (ITS); Security; Pre-Standardization Study on Pseudonym Change Management, Technical Report ETSI TR 103 415 V1. 1.1 (2018-04), 2018.

[11] C. Obermaier, R. Riebl, C. Facchi, Dynamic scenario control for VANET simulations, in: 2017 5th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems, MT-ITS, 2017, pp. 681–686, http://dx.doi.org/10.1109/MTITS.2017.8005599.

[12] A. Feraudo, N. Romandini, C. Mazzocca, DIVA-ReputationAlgorithm, 2024, http://dx.doi.org/10.5281/zenodo.10522096.

[13] IEEE standard for information technology– Local and metropolitan area networks– Specific requirements– Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications amendment 6: Wireless access in vehicular environments, 2010, pp. 1–51, http://dx.doi.org/10.1109/IEEESTD.2010.5514475, IEEE Std 802.11p-2010 (Amendment to IEEE Std 802.11-2007 as amended by IEEE Std 802.11k-2008, IEEE Std 802.11r-2008, IEEE Std 802.11y-2008, IEEE Std 802.11n-2009, and IEEE Std 802.11w-2009).

[14] 3rd Generation Partnership Project, 3GPP TR 21.914 V14. 0.0 (2018-05) 3rd generation partnership project; technical specification group services and system aspects; release 14 description; summary of rel-14 work items (release 14), 2018.

[15] J. Meredith, F. Firmin, M. Pope, Release 16 Description; Summary of Rel-16 Work Items, 3rd Generation Partnership Project (3GPP), Technical report (TR) 21, 2022.

[16] B. Sun, H. Zhang, 802.11 NGV proposed PAR, in: Proc. IEEE NGV Meeting, 2018, pp. 2–3.

[17] B.Y. Yacheur, T. Ahmed, M. Mosbah, Analysis and comparison of IEEE 802.11p and IEEE 802.11bd, in: F. Krief, H. Aniss, L. Mendiboure, S. Chaumette, M. Berbineau (Eds.), Communication Technologies for Vehicles, Springer International Publishing, Cham, 2020, pp. 55–65.

[18] IEEE Guide for Wireless Access in Vehicular Environments (WAVE) Architecture, IEEE Std 1609.0-2019 (Revision of IEEE Std 1609.0-2013), 2019, pp. 1–106, http://dx.doi.org/10.1109/IEEESTD.2019.8686445.

[19] E. ETSI, ETSI EN 302 663 V1. 3.11, intelligent transport systems (ITS); ITS-G5 access layer specification for intelligent transport systems operating in the 5 GHz frequency bands, 2020, ITS-G5 Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency bands.
[20] T. ETSI, 102 636-4-1 V1. 1.1, Intell. Transp. Syst. (ITS). Geonetwork. (2011).
[21] W3 Recommendation, Decentralized Identifiers (DIDs) v1.0, 2022, URL https://www.w3.org/TR/did-core/.
[22] C. Mazzocca, A. Acar, S. Uluagac, R. Montanari, P. Bellavista, M. Conti, A survey on decentralized identifiers and verifiable credentials, 2024, arXiv preprint arXiv:2402.02455.
[23] W3 Recommendation, Verifiable Credentials Data Model v1.1, 2022, URL https://www.w3.org/TR/vc-data-model/.
[24] S. Popov, The tangle, White Pap. 1 (3) (2018) 30.
[25] W.Z. Khan, E. Ahmed, S. Hakak, I. Yaqoob, A. Ahmed, Edge computing: A survey, Future Gener. Comput. Syst. 97 (2019) 219–235, http://dx.doi.org/10.1016/j.future.2019.02.050, URL https://www.sciencedirect.com/science/article/pii/S0167739X18319903.
[26] ETSI, GS MEC 003 - V3.1.1, 2022, URL https://www.etsi.org/deliver/etsi_gs/MEC/001_099/003/03.01.01_60/gs_mec003v030101p.pdf.
[27] A. Feraudo, A. Calvio, A. Bujari, P. Bellavista, A novel design for advanced 5G deployment environments with virtualized resources at vehicular and MEC nodes, in: 2023 IEEE Vehicular Networking Conference, VNC, 2023, pp. 97–103, http://dx.doi.org/10.1109/VNC57357.2023.10136327.
[28] S. Kekki, W. Featherstone, Y. Fang, P. Kuure, A. Li, A. Ranjan, D. Purkayastha, F. Jiangping, D. Frydman, G. Verin, et al., MEC in 5G networks, ETSI White Pap. 28 (2018) (2018) 1–28.
[29] X. Li, T. Jing, R. Li, H. Li, X. Wang, D. Shen, BDRA: Blockchain and decentralized identifiers assisted secure registration and authentication for VANETs, IEEE Internet Things J. (2022) 1, http://dx.doi.org/10.1109/JIOT.2022.3164147.
[30] C.P. Fernandes, C. Montez, D.D. Adriano, A. Boukerche, M.S. Wangham, A blockchain-based reputation system for trusted VANET nodes, Ad Hoc Netw. 140 (2023) 103071, http://dx.doi.org/10.1016/j.adhoc.2022.103071.
[31] D. Dolev, A. Yao, On the security of public key protocols, IEEE Trans. Inform. Theory 29 (2) (1983) 198–208, http://dx.doi.org/10.1109/TIT.1983.1056650.
[32] A. Hegde, A. Festag, Artery-C: An OMNeT++ based discrete event simulation framework for cellular V2X, in: Proceedings of the 23rd International ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems, 2020, pp. 47–51.
[33] G. Nardini, D. Sabella, G. Stea, P. Thakkar, A. Virdis, Simu5G–An OMNeT++ library for end-to-end performance evaluation of 5G networks, IEEE Access 8 (2020) 181176–181191, http://dx.doi.org/10.1109/ACCESS.2020.3028550.
[34] P.A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, E. Wiessner, Microscopic traffic simulation using SUMO, in: 2018 21st International Conference on Intelligent Transportation Systems, ITSC, 2018, pp. 2575–2582, http://dx.doi.org/10.1109/ITSC.2018.8569938.
[35] R. Riebl, H.-J. Günther, C. Facchi, L. Wolf, Artery: Extending veins for VANET applications, in: 2015 International Conference on Models and Technologies for Intelligent Transportation Systems, MT-ITS, 2015, pp. 450–456, http://dx.doi.org/10.1109/MTITS.2015.7223293.
[36] CAR 2 CAR, CAR 2 CAR communication consortium triggering conditions and data quality traffic jam CAR 2 CAR communication consortium about the C2C-CC, 2022, URL https://www.car-2-car.org/fileadmin/documents/Basic_System_Profile/Release_1.6.2/C2CCC_RS_2007_TrafficJam.pdf.
[37] Z. Lu, Q. Wang, G. Qu, Z. Liu, BARS: A blockchain-based anonymous reputation system for trust management in VANETs, in: 2018 17th IEEE International Conference on Trust, Security and Privacy in Computing and Communications/12th IEEE International Conference on Big Data Science and Engineering, TrustCom/BigDataSE, 2018, pp. 98–103, http://dx.doi.org/10.1109/TrustCom/BigDataSE.2018.00025.
[38] Z. Yang, K. Yang, L. Lei, K. Zheng, V.C.M. Leung, Blockchain-based decentralized trust management in vehicular networks, IEEE Internet Things J. 6 (2) (2019) 1495–1505, http://dx.doi.org/10.1109/JIOT.2018.2836144.
[39] N. Li, Y. Guo, Y. Chen, J. Chai, A partitioned DAG distributed ledger with local consistency for vehicular reputation management, Wirel. Commun. Mob. Comput. 2022 (2022) 6833535, http://dx.doi.org/10.1155/2022/6833535.
[40] Y. Li, X. Tao, X. Zhang, J. Xu, Y. Wang, W. Xia, A DAG-Based reputation mechanism for preventing peer disclosure in SIoV, IEEE Internet Things J. 9 (23) (2022) 24095–24106, http://dx.doi.org/10.1109/JIOT.2022.3189108.
[41] G. Du, Y. Cao, J. Li, Y. Zhuang, Secure information sharing approach for internet of vehicles based on DAG-Enabled blockchain, Electronics 12 (8) (2023) http://dx.doi.org/10.3390/electronics12081780.

Angelo Feraudo received the master's degree in computer science engineering from the University of Bologna, Italy, in 2020. Currently, he is a Ph.D. student working at the Department of Computer Science and Engineering at the University of Bologna. His research interests include vehicular networks, next-generation networks and vehicular computing integration in cloud-continuum spectrum.

Nicolò Romandini graduated from the University of Bologna, Italy, where he received M.Sc. degree in computer science engineering. He is currently a Ph.D. student at the Department of Computer Science and Engineering at the University of Bologna. His research focuses mainly on blockchain, cybersecurity and machine learning, and how to integrate them into IoT domains.

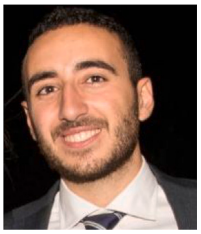

Carlo Mazzocca received his M.Sc. and B.Sc. degrees in Computer Engineering in 2018 and 2020, respectively, both from the University of Naples Federico II, Italy. He is currently a Ph.D. student in Computer Science and Engineering at the University of Bologna, Bologna, Italy. His research interests mainly include authentication and authorization solutions for the cloud-to-thing continuum.

Rebecca Montanari full professor at the University of Bologna since 2020 carries out her research in the area of information security and of the design/development of middleware solutions for the provision of services in mobile and IoT systems. Her research is currently focused on blockchain technologies to support various supply chains, including agrifood, manufacturing and fashion and on security systems for Industry 4.0.

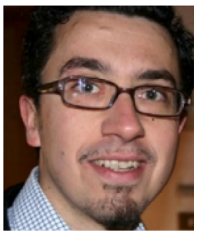

Paolo Bellavista received the Ph.D. degree in computer science engineering from the University of Bologna, Italy, in 2001. He is currently a Full Professor with the University of Bologna. His research interests include middleware for mobile computing, QoS management in the cloud continuum, infrastructures for big data processing in industrial environments, and performance optimization in wide-scale and latency-sensitive deployment environments. He serves on the Editorial Boards of IEEE COMMUNICATIONS SURVEYS AND TUTORIALS, IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT, IEEE TRANSACTIONS ON SERVICES COMPUTING, ACM CSUR, ACM TIOT, and PMC (Elsevier). He is the Scientific Coordinator of the H2020 IoTwins Project (https://www.iotwins.eu).