From Panoptic to Secure and Privacy-Preserving Lateral Surveillance in Vehicle Restricted Areas

Carles Anglés-Tafalla, Alexandre Viejo, Rolando Trujillo-Rasua, and Jordi Castellà-Roca, *Member, IEEE*

***Abstract*— In recent years, the implementation of vehicle restricted areas (e.g., Low Emission Zones or Congestion Charge Zones) has proven effective in addressing urban traffic congestion and environmental pollution. Traditionally, these zones have been enforced using the highly centralized and infrastructure-expensive panoptic model. However, recent advancements in smart vehicle technology have enabled a shift towards decentralized, infrastructure-free approaches such as the lateral surveillance model. While offering significant advantages, this novel model also presents relevant security and privacy challenges, including potential misuse and widespread surveillance. This paper introduces a novel autonomous, decentralized, and infrastructure-free solution aligned with the lateral surveillance model. Leveraging Blockchain technology and privacy preservation measures, the system ensures user anonymity during fee pricing and payment processes, and it is capable of detecting and sanctioning dishonest drivers. To assess the applicability of the new scheme in real-world scenarios, it was implemented and tested in controlled environments as well as a low-traffic street. The evaluation included an assessment of the gas costs associated with the smart contracts and a detailed scalability analysis. The theoretical and experimental results demonstrate the feasibility of the proposed solution.**

***Index Terms*— Lateral surveillance, privacy-preserving systems, blockchain, smart contracts, vehicle restricted areas.**

I. Introduction

THE number of cars worldwide is growing year by year. This trend is driven by various factors including population growth, urbanization, rising incomes in many parts of the world leading to increased car ownership, and improvements in accessibility to automobiles [1]. As a direct result of this increment, in the recent years, traffic congestion has become a growing problem in large metropolitan areas. This issue, in turn, has pushed the affected cities to high levels of environmental pollution that often surpass the safe health thresholds set by the World Health Organization (*WHO*) [2].

Received 10 June 2024; revised 14 January 2025 and 3 April 2025; accepted 5 May 2025. Date of publication 21 May 2025; date of current version 21 October 2025. This work was supported in part by the Project "HERMES" funded by INCIBE and European Union NextGenerationEU/PRTR; in part by the Project PID2021-125962OB-C32 "SECURING/-DATA" funded by MCIN/AEI/10.13039/501100011033/FEDER, UE; and in part by the Government of Catalonia under Grant 2021SGR 00115. The work of Rolando Trujillo-Rasua was supported the Ramon y Cajal Grant from Spanish Ministry of Science and Innovation and European Union under Grant REF: RYC2020-028954-I. The Associate Editor for this article was M. Shojafar. *(Corresponding author: Alexandre Viejo.)*

The authors are with the Departament d'Enginyeria Informàtica i Matemàtiques, Universitat Rovira i Virgili, 43007 Tarragona, Spain (e-mail: carles.angles@urv.cat; alexandre.viejo@urv.cat; rolando.trujillo@urv.cat; jordi.castella@urv.cat).

Digital Object Identifier 10.1109/TITS.2025.3568143

The enforcement of vehicle restricted zones has emerged as an effective solution that eases both problems mentioned above. These designated areas can be named "Low Emission Zones (*LEZ*)s" or "Congestion Charge Zones (*CC*)s" depending on the purpose behind the applied restrictions. In particular, LEZs are sectors where access restrictions to vehicles are enforced based on the environmental pollution, while CCs are zones in which restrictions are applied according to the traffic density. The significant proliferation of LEZs and CCs in European countries,[1] along with the intention of certain countries like Spain[2] to legislate in their favor, serves as evidence of their importance and impact.

The adoption of automated enforcement control systems in major urban environments has become the method of choice to comply with the restrictions imposed within these controlled areas. Specifically, according to the Transport Decarbonisation Alliance,[3] as well as the organizations POLIS[4] and C40,[5] the most widespread systems for enforcing those access restriction schemes are based on automatic number-plate recognition (ANPR) camera networks [3] that indiscriminately recognize the license plate of vehicles circulating within the restricted areas so that a central entity can accordingly determine and charge the applicable fees. London [4], Barcelona,[6] or Stockholm[7] serve as representative examples of this kind of approach, implementing around 1400 and 150 cameras in a network layout for the two first cases respectively, and 18 cameras control points in cordon layout for the latter.

The automated enforcement procedure stated above is aligned with the *Panoptic surveillance* model [5], which refers to surveillance activities via the use of video cameras carried out by the authorities. This approach has been proved to be feasible and effective, however, this model is inherently intrusive since all vehicles are identified by the charging entity every time they approach a control infrastructure. This situation jeopardizes the privacy of the drivers and underscores the necessity for alternative systems that enforce access restrictions in privacy-preserving manner.

In addition to the aforementioned privacy concerns, the panoptic model also delivers two significant downsides that

[1] Urban Access Regulations, http://urbanaccessregulations.eu/userhome/map

[2] Law on Climate Change and Energy Transition, https://www.boe.es/eli/es/l/2021/05/20/7

[3] Transport Decarbonisation Alliance (TDA), www.tda-mobility.org

[4] POLIS - Cities for Transport Innovation, www.polisnetwork.eu

[5] C40 cities network, www.c40.org

[6] Barcelona Metropolitan Area - LEZ, https://www.zbe.barcelona

[7] Stockholm charging scheme, https://miljobarometern.stockholm.se/trafik

require attention: i) the entity in charge of the vehicular restricted area is required to arrange a substantial infrastructure, resulting in significant implementation, deploy, and management costs; and ii) given that this entity is responsible for overseeing the entire system, including the vehicles and drivers involved, it is essential that it be fully trusted, however, this assumption may be too strong in many practical scenarios.

The panoptic model and its associated limitations are closely linked to traditional vehicles that lack smart features. Nevertheless, society is moving into the era of the Internet of Vehicles (IoV), where next-generation vehicles are equipped with sensors, cameras, computing resources, and vehicle-to-everything (V2X) communications [6]. These advancements enable smart vehicles to assume new roles, such as autonomous driving, platooning, and traffic optimization. This evolution paves the way for an effective transition from the dominant panoptic paradigm to the innovative *lateral surveillance* model [7].

The lateral surveillance model is characterized by individuals using surveillance tools to monitor and track one another and gather information, rather than relying on agents of public or private institutions. Within this paradigm, new-generation vehicles can leverage their advanced detection and communication systems to act as enforcement entities within restricted areas. This approach reduces the need for costly control infrastructures supervised by a centralized authority, alleviating the burdens of deployment and operation that hinder scalability, as well as addressing the trust issues inherent to the panoptic model.

Despite the undeniable benefits, the implementation of the lateral surveillance model also comes with relevant challenges. In particular, this paradigm may lead to indiscriminate and systematic tracking of drivers, jeopardizing their privacy; dishonest drivers may try to take advantage of a system without a central trusted authority to gain an advantage over others; and some drivers might lack the motivation to employ their own resources for monitoring others, consequently diminishing the effectiveness of the system.

Given the aforementioned points, additional research is necessary to develop a secure and privacy-preserving lateral surveillance system that mitigates these risks and ensures robust control over vehicles in different restricted areas.

A. Related Work

In the past decade, significant research efforts have been undertaken within both panoptic and lateral surveillance paradigms, focusing extensively on addressing security and privacy concerns associated with managing vehicle restricted areas.

Initially, panoptic-based proposals got the spotlight by means of works such as [8], [9], [10], [11], [12], and [13], which addressed security issues and relied on a centralized Service Provider (SP) for fee computation using location and time data gathered by the vehicles' On-Board Units (OBUs). However, these schemes largely neglected privacy threats, resulting in an omnipresent SP equipped with a vast network of cameras capable of fully tracking all vehicles in their vicinity. Since then, the scientific community has diverged into two distinct branches. One branch has embraced advancements such as AI-based surveillance techniques, which improve vehicle identification through feature extraction [14], [15], and edge computing traffic monitoring systems, which process data locally near the point of capture and transmit only processed information to central servers [16], [17]. These technologies further empower the centralized SP to enhance real-time analysis and tracking capabilities. The other branch, driven by growing societal concerns over privacy, has focused on mitigating these threats by proposing solutions that limit surveillance scope and emphasize data minimization, thereby addressing the critical privacy issues inherent in panoptic-based approaches

Within this latter branch, several works [18], [19], [20], [21] have endorsed a refined approach within the panoptic paradigm. In particular, they advocated for the targeted use of surveillance cameras, specifically aimed at drivers involved in fraudulent activities, rather than employing indiscriminate monitoring. These four proposals share a common underlying approach but differ primarily in terms of computational and communication costs, or in the specific electronic payment system utilized.

In [22], the authors introduce the latest and pivotal panoptic-based solution in the literature. This work reconceptualizes the payment landscape, addressing inherent structural challenges observed in earlier panoptic approaches linked to centralized entities. Specifically, the authors employ blockchain technology and smart contracts to decentralize their architecture. This approach transforms vehicle accesses to restricted areas into blockchain transactions, enabling decentralized pricing and fee management seamlessly. Though this work marks a significant stride toward decentralization within the examined scenario, being panoptic in nature, it remains tied to a set of centralized entities tasked with implementing, overseeing, and maintaining street infrastructures alongside their associated expenses.

Unlike the panoptic approach, the *lateral surveillance* model emerged unburdened by the aforementioned costs, owing to its inherently decentralized nature. The emergence of this state-of-the-art paradigm became feasible, first, through the increased integration of advanced sensors and cameras in vehicles, and, next, by means of actively sharing the gathered data with the other drivers. Pioneering applications of this concept, exemplified in works like [23], [24], and [25], leverage crowd-sourced data and GPS readings to provide real-time information about parking space availability and road incidents such as traffic congestion or accidents.

While initially designed to assist drivers, sensors and cameras installed on vehicles have transcended their original purpose. In this way, works such as [26], [27], and [28] have shifted from aiding users to delving into the realm of surveillance. Specifically, these proposals pave the way for utilizing onboard vehicle cameras to surveil other vehicles passing nearby, eliminating the need for centralized and costly infrastructures. However, their focus is solely on utilizing this technology for reporting traffic violations and tracking designated vehicles.

In the specific context of applying the lateral surveillance paradigm to enforce vehicle restricted zones, an initial comprehensive approach to this issue was introduced in [29]. This study outlines a strategy where vehicles, equipped with embedded video cameras and utilizing automatic number plate recognition techniques, collaborate through Vehicle-to-Vehicle (V2V) communications to establish surveillance networks. These networks play a pivotal role in monitoring and reporting suspicious vehicles accessing or traversing designated zones to nearby law enforcement forces. However, this solution incurs in the indiscriminate collection of license plate recordings, posing a significant privacy concern previously identified in some of the panoptic-based approaches mentioned earlier.

Finally, the work in [30] presents an Electronic Road Pricing (ERP) control system that employs lateral surveillance to detect fraudulent drivers attempting to evade tolls. This envisioned system integrates radars and V2V communications to identify drivers who fail to transmit their pseudonymized position to the road side units (RSUs). Following the model proposed in [19], the onboard cameras of vehicles capture license plate photos exclusively in punishable circumstances. In the final step, collaboratively signed evidence is dispatched to a centralized toll server. However, despite leveraging vehicles as a decentralized surveillance network, the system requires an entire network of RSUs to serve as location proof for tolling. This poses a significant drawback, as the proposed solution still relies on a centralized tolling entity with a costly fixed infrastructure, amplifying the system complexity through the utilization of a vehicle-based surveillance network.

B. Contribution and Plan of This Paper

In recent years, the lateral surveillance paradigm has garnered growing interest as a viable method for the infrastructure challenges associated with the well-established panoptic approach. However, as previously exemplified, a significant gap persists in the literature concerning an infrastructure-free procedure that ensures secure and privacy-preserving access to vehicle-restricted areas.

Accordingly, this article presents an autonomous and fully decentralized solution for managing vehicle access to controlled zones. The proposed system leverages cutting-edge detection and communication systems embedded in next-generation vehicles, thereby obviating the need for deploying and maintaining costly control infrastructures. Furthermore, it employs smart contracts and blockchain technology to decentralize vehicle access management, pricing, and payment processes, effectively eliminating the involvement of centralized third-party service providers.

The proposed scheme achieves the following properties, which are not jointly addressed by existing works in the literature:

- *Infrastructure-free:* The system eliminates the necessity for fixed access control infrastructure within vehicle-restricted areas.
- *Anonymous payment:* The system ensures user anonymity throughout the fee pricing and payment processes by leveraging blockchain technology and privacy-preservation measures.
- *Decentralized architecture:* The system decentralizes the management of tolling, payment, and fraud control for vehicles entering restricted areas. It replaces the need for centralized third-party involvement with a distributed network founded on smart contracts and blockchain technology, complemented by advanced detection systems embedded within next-generation vehicles.
- *Revocable anonymity:* The system preserves the privacy of compliant drivers while maintaining the ability to identify and sanction dishonest users when necessary.

The rest of the paper is organized as follows: Section II provides a general overview of the proposed scheme; Section III formalizes the protocols that sustain the proposal; Section IV examines the system's ability to identify drivers who evade payment when passing through restricted zones; Section V details the implementation and assesses the system's performance in a practical environment; Section VI evaluates the scalability of the solution; Section VII examines the gas costs associated with the key methods implemented in the governing smart contract; Section VIII analyzes the system's security and privacy guarantees; and, finally, Section IX presents the conclusions.

II. General Overview

This section provides a general overview of the proposed system, introducing actors, functionality, integration with blockchain technology, and considerations regarding security and privacy. First, it outlines the key actors involved. Next, it offers a high-level understanding of the system's functioning, reserving intricate details about the underlying communication protocols and verification processes for the next section. Following this, it examines the incorporation of blockchain technology into the framework. Lastly, it highlights the functional, security, and privacy requirements taken into account during the design phase.

A. Actors

The proposed system involves the following five actors:

- *Restricted Area Administrator (RAA)*: This entity manages the vehicle restricted area and enforces the corresponding constraints imposed on vehicles. It is tasked with issuing credentials to other entities, deploying the smart contract governing the designated restricted area, and, when needed, defining the emission categories for vehicles.
- *Vehicles (Vs)*: They are the entities that access and drive through the restricted areas overseen by the RAA. It is assumed that each V is equipped with a tamper-proof on-board unit with cryptographic capabilities, GPS technology, 4G and short range communication systems (e.g. Bluetooth, Zigbee or DSRC).
- *Access Control Vehicles (ACVs)*: These entities constitute a subset of V entities endowed with the capability to verify whether other vehicles hold a correct access ticket. To fulfill this function, each ACV must be equipped with an Advanced Driving Assistance System (ADAS), encompassing proximity sensors and cameras in both front and rear sections.

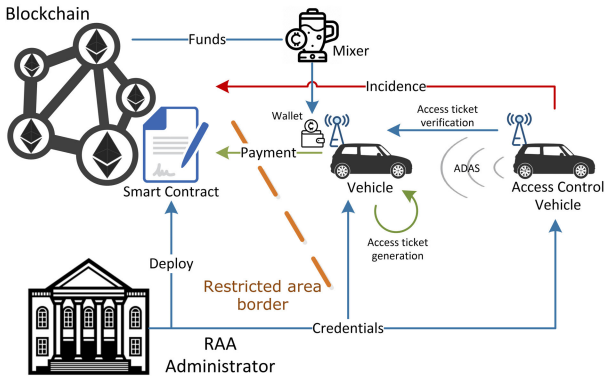

Fig. 1. General overview of the proposed scheme.

- *Smart Contract (SC)*: This is a specific programmed transaction protocol to publish the restricted area vehicular access details in the blockchain. This piece of software can automatically verify, price and charge, in terms of digital currencies, the vehicles driving through a certain restricted area without needing a central authority or external supervision.
- *Cryptocurrency Mixing Service (M)*: This is an independent for-profit entity (e.g., ETH-Mixer[8] or Tornado-Cash[9]) responsible for obfuscating the traceability of cryptocurrency transactions, so that the transferred funds cannot be trailed back to the source digital wallet. The services offered by M are expected to be used when a V is transferring digital currencies between two of its own wallets, as in the *wallet renewal* process or during the *wallet filling* protocol. Both procedures will be explained in detail thereafter.

B. The Proposed System in a Nutshell

Figure 1 shows a general perspective of the proposed scheme, showcasing the main interactions among the aforementioned actors. Before gaining entry to the restricted area, Vs and ACVs need to acquire valid credentials from the RAA to interact with other actors. Moreover, they must generate a set of digital wallets to interact with the smart contract that manages the payment system. To settle the restricted area pricing taxes and cover the blockchain network fees, Vs must purchase cryptocurrency coins through an online exchange service. This process involves the intermediary services of M to prevent the linkage of Vs' wallets from their owners' bank accounts.

In the depicted scenario, the envisioned protocol commences as a V prepares to enter the restricted area. In that moment, V generates its own cryptographically-signed *access ticket* encompassing an access identifier (i.e., $access\_ID$) and certain entrance details related to fees. Once within the restricted area, any proximate ACV issues a challenge to V to verify its self-generated *access ticket*. The verification challenge starts the moment an ACV detects an adjacent V through its sensors, followed by an attempt to initiate a challenge through short-range communication system with the detected vehicle. As a result of this secure communication process, ACV obtains an a digitally signed access ticket from V as an evidence of the vehicle's access parameters to the restricted area.

Throughout this procedure, user anonymity is preserved through the use of a vehicle temporary alias, which can be altered at the user's discretion to prevent other entities from linking different access tickets with the same V. However, if V engages in misconduct during the process, and the access ticket cannot be obtained nor verified, the ACV can use its front/rear camera to capture a photo of the adjacent vehicle's license plate. This evidence is then reported to the RAA as a potential fraud detection. To prevent unwarranted accusations, the proposed system incorporates a threshold mechanism that triggers sanctions and anonymity revocation solely when a specific number of ACVs report the same misbehaving V.

Upon exiting the restricted area, V must fulfill the associated access fees. To achieve this, V initiates a remote call to the payment method offered by the deployed smart contract, submitting its entry and exit data as parameters. Then, the smart contract's logic subsequently runs two key functions: i) it calculates the payment amount by leveraging the supplied data and the established prices on the blockchain; and ii) it autonomously transfers the determined payment in digital currency from V's digital wallet to the digital wallet of RAA.

At a later reasonable time, any ACV that has challenged V and acquired its signed access ticket may verify, by querying the blockchain, whether the access data and the corresponding fee payment have been properly processed. If any irregularity is detected, the ACV can report an incidence by calling the designated smart contract's method and sharing its own copy of the access ticket, which is digitally signed by V. It is important to note that reporting on V can only occur after a predefined time interval has elapsed. This time interval is set by RAA and published on the blockchain. Through the publication of access ticket information on the blockchain, only RAA possesses the capability to identify V's owner and implement punitive measures if deemed necessary.

C. Handling Non-Compliant Vehicles

In scenarios where a vehicle lacks the required on-board unit with compatible connectivity, or when the unit is damaged or otherwise unusable, the proposed system faces challenges in maintaining seamless operation. To address such situations, the RAA can provide a supplementary mobile or web application that allows drivers to manually report their presence within the restricted area by submitting their vehicle's license plate number and completing the process through a suitable payment method.

While the technical details of this mechanism are beyond the scope of this paper, it serves as a fallback solution, ensuring that all vehicles, regardless of their technological capabilities, can access restricted areas without issue. However, this solution comes at the cost of forfeiting privacy entirely, as the vehicle's activity within the restricted area is directly linked to the driver's identity. This setup is analogous to parking

[8]ETH-Mixer, https://eth-mixer.com
[9]Tornado Cash, https://defirate.com/tornado-cash

reservation systems[10] in urban areas, where drivers provide identifying information in exchange for access to a service.

D. Blockchain Integration

The blockchain is essentially a digital ledger that records transactions in a secure and transparent manner, allowing parties to transact without the need for intermediaries. Its underlying principle, known as *Proof-of-Work (PoW)*, involves providing equal decentralized trust to any node capable of solving computational challenges as a mechanism to reach consensus [31]. Originally conceived as a means to decentralize financial transactions, this concept, as explained in [32], was further expanded with Ethereum. In this extension, programmable logic programs, known as smart contracts, operate on the blockchain, facilitating the decentralized settlement of transactions that encompass different resources.

The PoW principle faces notable challenges including high processing times, limited scalability, and the necessity for transaction fee payments, rendering it suboptimal for scenarios involving diverse and typically lightweight devices, such as the Internet of Things (IoT) or the Internet of Vehicles (IoV). To overcome these challenges and incorporate the logic of Ethereum Virtual Machine (EVM) smart contracts into these domains, researchers have suggested alternative consensus mechanisms to PoW. In this way, Ethereum 2.0 introduced the *Proof-of-Stake (PoS)*, wherein nodes possessing the highest stake or deposit in the network are chosen as block validators. Similarly, [33] proposed the *Randomized Delegated Proof of Stake (Roll-DPoS)*, a PoS variant tailored for IoT scenarios and currently implemented in the IoTeX[11] blockchain network.

In contrast, IOTA [34] takes a different approach to solve the aforementioned blockchain drawbacks in IoT scenarios. In particular, it uses a Hash Directed Acyclic Graph (DAG), known as *The Tangle*. The Tangle employs the Fast Probabilistic Consensus (FPC) as a consensus protocol, combined with *Mana*, a delegated proof of token ownership used as a transferable reputation metric to prevent dishonest nodes from disrupting the validation process.

While IOTA addresses scalability through its unique DAG structure and consensus mechanisms, researchers have also explored layer-2 blockchains as an even more scalable solution to address efficiency challenges in IoT scenarios. Layer-2 blockchains operate as secondary frameworks built on top of main blockchains (layer-1), significantly reducing the computational and storage burden of the primary chain while enhancing throughput, latency, and cost-efficiency. One prominent example of a layer-2 blockchain addressing IoT challenges is Polygon [35] (formerly Matic Network), which employs sidechains and rollups to dramatically increase transaction throughput and reduce costs. Polygon's architecture has been effectively leveraged in IoT use cases where lightweight and highly scalable interactions are critical.

Despite differences in their internal architectures, IoTeX, IOTA, and Polygon are distributed ledgers tailored for IoT environments that support EVM compatibility. This feature enables the native implementation of Ethereum smart contracts on these networks, providing access to the full Ethereum development and testing ecosystem and significantly simplifying the integration of smart contracts in IoT and IoV scenarios. Leveraging this capability, the proposed system utilizes the programmable logic embedded in EVM's smart contracts to seamlessly incorporate vehicle access data into blockchain transactions. This integration automates the pricing of access fees and facilitates the subsequent transfer of corresponding digital currencies. Through this process, blockchain network nodes autonomously validate vehicle access to restricted areas, enabling a fully distributed scheme free from centralized third-party supervision. Additionally, it is envisaged that end-users within the system, i.e., Vs and ACVs, may serve as miners or validators, helping offset a portion of the expenses incurred while driving through restricted areas.

E. Functional, Security, and Privacy Requirements

To fulfill the operational objectives of the proposed system, it is designed to meet the following *functional requirements*:

F1) *Detection reliability*: Ensure the accurate detection of vehicles that pass through a restricted area without making the required payment.

F2) *Lightweight V2V protocols*: Design protocols that satisfy the Real-Time Computing (RTC) constraints of real-world vehicular scenarios.

F3) *Scalability*: Maintain system performance as the number of vehicles and traffic density increases.

F4) *Efficient decentralized payment system*: Implement a smart contract-based system that is affordable in terms of gas consumption.

F5) *Exculpability*: Prevent honest drivers, i.e., those who follow the system and protocol specifications, from being incorrectly fined.

Following this, we elaborate on the four *security requirements* addressed by the new scheme. The first three requirements (S1, S2, and S3) are directly mapped to the four core security properties: confidentiality, integrity, authentication, and non-repudiation. These core properties ensure the foundational security of the interactions between ACVs and Vs within the system. The fourth requirement (S4) focuses on the broader security risks inherent to blockchain technology and their potential impact on the payment protocol.

Certain assumptions serve as the foundation for these requirements: i) the RAA is trusted but curious; ii) credentials are correctly generated and distributed by the RAA; iii) vehicles can establish a secure connection with the RAA; and iv) vehicles' onboard units remain untampered with. The requirements are as follows:

S1) *Secure access ticket verification*: Whenever an ACV receives an access ticket supposedly from a driver X, it must be guaranteed that X generated and sent the access ticket during the protocol execution. This requirement establishes vehicle authentication and ensures the integrity of access tickets.

S2) *Non-repudiable access ticket generation*: Ensure that a driver cannot deny having generated an access ticket.

[10]EasyPark - https://www.easypark.com

[11]IoTeX, https://iotex.io

S3) *Confidential V2V communication*: Ensure that sensitive information, such as access tickets and location data, remains confidential for honest drivers.
S4) *Resilience to blockchain security risks*: Ensure that the proposed system remains robust against inherent security vulnerabilities of decentralized blockchain technology, thereby safeguarding the payment protocol.

Focusing now on the *privacy requirements* and acknowledging that no digital system can prevent drivers from using their cameras to record fellow drivers, the proposed system strives to minimize the transmission of private information to the central authority. Additionally, it aims to pose difficulties for potential attackers attempting to trace other vehicles. Consequently, the following privacy requirements are established:
P1) *Data minimization*: Avoids or minimizes the disclosure of an honest driver's picture to the RAA.
P2) *Protocol anonymity*: Ensures that, based solely on the messages exchanged within the system, it is not possible to track an honest driver. Specifically, the communication protocol between drivers and ACVs must not disclose identifying or traceable information about drivers. Additionally, the payment system must preserve driver anonymity and prevent traceability.

Note that location privacy is not a requirement of our system. This is because malicious ACVs can always use their cameras to track other drivers, much like bystanders can use their phone cameras to track others. However, we do require our system to avoid increasing the risk of location tracking by ensuring that the access ticket and payment protocols satisfy untraceability.

Additionally, our system does not guarantee the anonymity of drivers who have been reported as potential rule breakers. This is an intentional feature rather than a limitation, as these drivers will undergo an external investigation that necessitates identifying information.

III. The Proposed System in Detail

The lifecycle of the proposed system is divided into eight sub-protocols or phases:

- *On-board unit set-up*: It describes the registration process that Vs and ACVs should complete with the RAA to obtain the credentials which are required to interact with the system's actors.
- *Wallet filling*: It outlines the steps that V must follow to anonymously purchase cryptocurrency coins, which serve as the currency in the proposed system.
- *Access ticket generation*: It describes the steps that V should follow to self-generate its access ticket.
- *Access ticket verification*: It details the sequence of interactions by which an ACV obtains and verifies the access ticket from a V traveling within a restricted area
- *Payment*: It outlines the steps that V should follow to price and pay the corresponding access fee by means of smart contract interaction.
- *Payment Verification*: It describes the countermeasures that ACVs may apply to report a V who tries to commit fraud during the payment phase.
- *Fraud prevention*: It outlines the countermeasures that ACVs may employ to report a V attempting to bypass the access ticket verification protocol.
- *Vehicle code renewal*: It describes the process by which a V renews its credentials and certificates as a measure to safeguard its privacy

In the following sections, those eight protocols are described, giving enough detail to allow their implementation and verification.

A. On-Board Unit Set-up

This initial protocol involves configuring the On-Board Units (OBU) of vehicles, both Vs and ACVs, to obtain the necessary credentials for secure interaction with other vehicles in the system. To achieve this, the OBU of a V establishes a secure communication channel with the RAA and submits relevant vehicle information, including license plate number, car manufacturer, model, etc. It is assumed that the OBU cannot be tampered with to provide inaccurate information, and that the RAA (a governmental entity, such as a city council) possesses the capability to validate the vehicle's data and retrieve its owner's information. Once this data is obtained by the RAA, it proceeds with the following tasks:

a) RAA verifies whether the vehicle's data corresponds to its license plate number, and subsequently retrieves the owner's details, including name, address, email, and phone number.
b) RAA creates a pseudo-random temporary alias, denoted as β, and binds it with the owner of V.
c) RAA transmits β to the owner of V through an alternative channel, such as email, phone, or a public administration electronic identification and authentication system.[12]

Once the owner of V obtains β and inserts it into V's OBU, this device proceeds with the following steps:

a) V generates a key pair (sk_V, pk_V).
b) V computes a certificate signing request $CSR(pk_V)$ for the generated public key, containing the vehicle's temporary alias β in the *CommonName* field instead of V's personal data.
c) V sends the certificate request $CSR(pk_V)$ back to RAA.

When RAA receives a valid $CSR(pk_V)$, it performs the following operations:

a) RAA verifies the validity of the code β contained in $CSR(pk_V)$. If this verification fails, the vehicle registration is aborted; otherwise, the process continues.
b) RAA recovers the vehicle's information bound to β.
c) RAA issues the certificate $Cert(pk_V)$, including β in the *CommonName* field and the vehicle emissions category cat as a certificate extension. Note that, RAA may include other extensions (e.g. a LEZ or CC residence proof).
d) RAA sends the generated certificate $Cert(pk_V)$ to V.

Finally, in order to complete the registration process, V proceeds with the following operations:

[12]Cl@ve - Electronic Identity for the Administration - https://clave.gob.es

a) V verifies the validity of the certificate $Cert(pk_V)$ through $Cert(pk_{RAA})$. If this verification fails, the procedure is aborted; otherwise, it continues.
b) V securely stores $Cert(pk_V)$ and (sk_V, pk_V).

After V's credentials have been stored in the OBU, V generates an *EVM digital wallet* W_V or, depending on its privacy preferences, a set of them $W_V^1, \ldots, W_V^n$. To do so, the following steps are undertaken:

a) V generates, for each wallet, a private key sk of 256-bit, a public key pk of 512-bit, and it derives its address according to the EVM key specs.
b) V securely stores the set of digital wallets $W_V^1, \ldots, W_V^n$.

B. Wallet Filling

In order to settle LEZ or CC road pricing taxes and the blockchain network's fees, Vs purchase cryptocurrency coins, which serve as the native currency within our system. To accomplish this, Vs uses an online service where crypto coins can be purchased and directly transferred to their wallets,[13] by means of classic payment mechanisms (e.g., credit card or wire transfer). It is assumed that RAA offers a service of this nature. However, the process of purchasing coins carries the potential risk of linking Vs' digital wallets to their respective bank accounts, thereby compromising their privacy. To prevent this issue, V undertakes the following steps:

a) V creates a temporary wallet W_V^T.
b) V accesses an online service and purchases *crypto coins*. These coins are transferred to the temporary wallet W_V^T.
c) V requests a *cryptocurrency mixing service*, M, to transfer the purchased crypto coins from the temporary wallet W_V^T to V's set of wallets $W_V^1, \ldots, W_V^n$. Through a mixing process, M obscures the link between the source and destination wallets during the funds transfer, thereby preventing the selling entity from associating any payment information with V's transactions on the blockchain.
d) Once the coins have been transferred, V discards the temporary wallet W_V^T.

Although the purchaser's identity cannot be linked to her wallet after this process, transactions associated with the same wallet address can still be correlated. Different strategies can be employed to tackle this problem. A straightforward approach is for V to utilize a single wallet at a time, generating a new one whenever it wants to break the link between its transactions. Any remaining funds in the previous wallet can then be transferred to the new one via M, which offers a transaction obfuscation service.

C. Access Ticket Generation

Once a vehicle finds itself at the edge of a restricted area, it initiates the *access ticket generation* protocol, which generates a valid ticket for ACVs to verify. The protocol consists of the following steps:

a) The vehicle V retrieves spatial and temporal access parameters, namely the entrance position pos_V and the current entry timestamp t_V^{in}. The superscript of t_V^{in} refers to *in* a restricted area. Later, we will use a timestamp t_V^{out} to mean *out* of the restricted area.
b) V generates a 128-bit access ID δ by hashing the vehicle temporary alias β, the access parameters, the vehicle category cat, and a randomly-generated nonce n_V, giving $\delta = \mathrm{h}(n_V \| \beta \| pos_V \| t_V^{in} \| cat)$. This ID will be used to identify the blockchain transaction that corresponds to the payment of the access fee.
c) V generates an access ticket $at = \{\delta, n_V, \beta, pos_V, t_V^{in}, cat\}$. Notice that δ is the hash of the access data.

[13] My Ether Wallet - https://ccswap.myetherwallet.com

D. Access Ticket Verification

This sub-protocol, outlined in Figure 2, starts when a vehicle operating as an ACV detects a nearby vehicle V using its front or rear sensors of the ADAS system. The goal of the ACV is to obtain a valid access ticket and a non-repudiable proof of encounter with V. The ACV executes the following steps to initiate a handshake with V:

a) ACV obtains its GPS position and determine the restricted area identifier, denoted RA_{ID}, where the detection has occurred.
b) ACV retrieves its current timestamp, denoted t_{ACV}.
c) ACV generates a *detection announcement* $da = \{RA_{ID}, n_{ACV}, t_{ACV}\}$, where n_{ACV} is a nonce.
d) ACV signs the da, denoted by $\sigma_{ACV}(da)$.
e) ACV broadcasts da and $\sigma_{ACV}(da)$ to all nearby Vs.

Upon reception of the detection announcement da and signature $\sigma_{ACV}(da)$, a V assesses the freshness of da by examining t_{ACV} and whether RA_{ID} corresponds to V's present location. Optionally, utilizing its front/rear sensors, V can also verify the proximity of vehicles that may potentially be the ACV that issued da.

If all validations are successful, V accepts the detection announcement and thereby creates a secure communication channel via TLS with the ACV. This process involves robust bilateral authentication, with both parties exchanging certificates, i.e. $Cert(pk_V)$ and $Cert(pk_{ACV})$. The TLS secure channel is used to transmit a non-repudiable proof of the generation of a valid access ticket. This is achieved as follows:

a) V retrieves its access ticket at which contains the associated access data $\{\delta, n_V, \beta, pos_V, t_V^{in}, cat\}$.
b) V obtains the current timestamp t_V^c.
c) V generates an encounter proof $(ep, \sigma_V(ep))$ where $ep = (at, t_V^c, \sigma_{ACV}(da))$. The encounter proof contains the access ticket, the current timestamp, and the signature of the detection announcement.
d) V send over the secure channel the encounter proof $(ep, \sigma_V(ep))$ to ACV.
e) Upon reception of the encounter proof $(ep, \sigma_V(ep))$, the ACV uses $Cert(pk_V)$ to verify that $\sigma_V(ep)$ is valid. The ACV also validates the correctness of β and cat, which are part of the access ticket, by looking into V's certificate extensions. Lastly, the ACV checks that the access ticket timestamp t_V^{in} is lower than both t_{ACV} and t_V^c.
f) If these validations are successful, ACV locally stores the tuple $(t_{ACV}, ep, \sigma_V(ep))$, which establishes that at

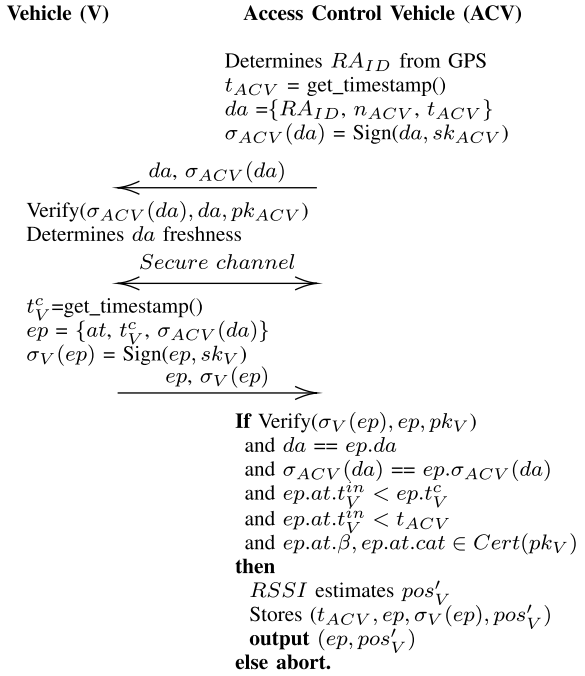

Fig. 2. Access ticket verification protocol.

time t_{ACV} a vehicle V sent an encounter proof ep which contains a valid access ticket. This information will be used later to verify whether V pays for the access ticket.

g) Additionally, ACV estimates the relative position of V based on the physical properties of the communication channel, e.g. by employing a received signal strength indicator system (e.g., RSSI [36]) on the transmitted data as introduced in [30] for ERP systems.

h) ACV thus stores pos'_V together with the tuple $(t_{ACV}, ep, \sigma_V(ep))$ as the output of the access ticket verification protocol.

i) This sub-protocol aborts, i.e. gives no output, if any of the steps above fails.

We let the ACV execute the access ticket verification protocol with all vehicles that respond to the detection announcement, widening the scope of the verification to many vehicles. Therefore, the output of the access ticket verification protocol is a list of encounter proofs sent by Vs together with their relative positions. Formally, given a set of vehicles $\{V_1, \ldots, V_n\}$ that successfully finished the access ticket verification protocol with ACV, we use $\{(ep_{V_1}, pos'_{V_1}), \ldots, (ep_{V_1}, pos'_{V_n})\}$ to denote their output.

E. Fraud Prevention

The goal of the fraud prevention mechanism is to check whether the detected vehicle V has sent an encounter proof. If not, the ACV deems V fraudulent and thus takes a photo with its front/rear camera as a visual proof of the presence of V in the restricted zone. In short, either the ACV receives a cryptographic proof of encounter or it seeks a visual one. The fraud prevention mechanism works as follows.

- Let $\{(ep_{V_1}, pos'_{V_1}), \ldots, (ep_{V_1}, pos'_{V_n})\}$ be the the output of the access ticket verification sub-protocol and pos_V the position of the vehicle V detected with the front or rear sensor of the ACV. Note that pos_V can be calculated straightforwardly from pos_{ACV}.
- If, for every $i\{1, \ldots, n\}$, pos'_{V_i} does not match pos_V, then the ACV captures a photo of the sensors-detected vehicle. The photo is sent to the RAA together with cryptographically signed detection announcement da.

Over time, RAA accumulates the fraud evidences sent by ACVs. Once the number of collected evidences concerning a single V, issued by different ACVs within a similar time frame, surpasses a predetermined threshold (K), the RAA initiates sanctioning measures against the dishonest V. If the incident concludes with V being unable to provide a valid access ticket at, which is subsequently published and paid for on the blockchain, the owner of V is sanctioned, and the contributing ACVs are appropriately compensated.

F. Payment

Following V's departure from the restricted area, there is a specified time-frame for settling the access fees. Note that the duration of this timeframe is determined by the RAA and may vary from several hours to several days. The entire process, illustrated in Figure 3, is governed by smart contracts supported by blockchain distributed ledger technology, eliminating the need for any central authority's involvement.

To initiate the payment protocol, V invokes the *register_access* method of the SC, providing its fee-related data from the access ticket at and departure information as parameters. These parameters include the access ID δ, entry timestamp t_V^{in}, exit timestamp t_V^{out}, vehicle category cat, and restricted area RA_{ID}. Like t_V^{in}, the departure timestamp t_V^{out} is self-generated. Nonetheless, our system ensures that t_V^{out} is higher than or equal to the timestamp of the the latest encounter proof transmitted to an ACV, deterring fraudulent vehicles from cheating on their calculation of t_V^{out}.

Upon calling this method, the SC runs the following on-chain operations:

a) Based on t_V^{in} and t_V^{out}, the SC calculates the duration for which the V remained within the restricted area.

b) The SC retrieves the applicable prices for the restricted area RA_{ID} from the blockchain, corresponding to the vehicle category cat, and computes the corresponding fee.

c) The SC transfers the corresponding fee amount in cryptocurrency from V to the digital wallet of RAA.

d) If V's wallet does not have enough funds, all coins in V's wallet are transferred, and the transaction debt is updated accordingly. V can repeat this process, using several wallets, until the remaining debt is settled, at which point the transaction is set as $paid$.

e) The SC stores transaction δ along with the following attributes: t_V^{in}, t_V^{out}, cat, and the remaining $debt$. Regardless of the outcome of the *register_access* execution, the transaction is added to the blockchain.

It is worth noting that the RAA, which is expected to be the owner of the SC, has the authority to modify the pricing table used by the SC to calculate vehicle fees within

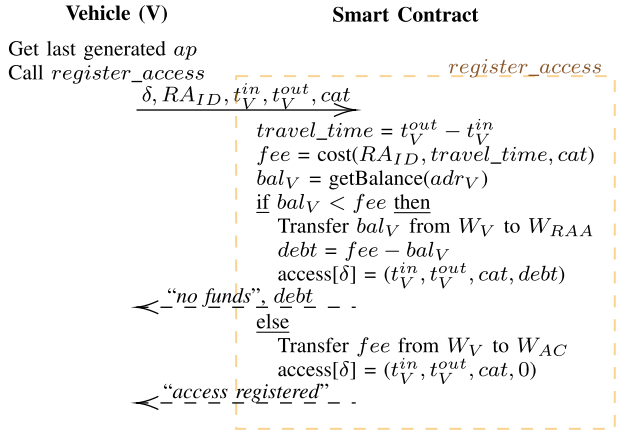

Fig. 3. Payment phase.

the restricted area. This assumes that the SC incorporates a method restricted to the owner for this specific task.

G. Payment Verification

After the timeframe set by the RAA for validating and publishing the access transaction δ to the blockchain expires, ACVs that challenged and acquired an access proof ep from V can verify its payment status by undertaking the following steps:

a) ACV hashes the access ID δ from the ep data. Specifically, $\delta = \text{h}(n_V \| \beta \| pos_V \| t_V^{in} \| cat)$.
b) ACV retrieves the current transaction δ data by invoking the SC's *get_access_status* method.
c) ACV verifies whether the transaction with δ ID exists, and ensures that its current debt is set to "0".
d) ACV additionally verifies the consistency of the published data with its own ep copy. Additionally, the ACV checks whether the exit timestamp t_V^{out} is larger than the timestamp t_V^c within the encounter proof.
e) If any of the aforementioned conditions are not satisfied, ACV publishes an incidence on the blockchain using the method *payment_incidence* of the SC. During this procedure, the encounter proof $(ep, \sigma_V(ep))$, which contains the vehicle's temporary alias β and V's signature, are sent as parameters.
f) On the contrary, if the verification proceeds smoothly, ACV deletes the local copy the encounter proof ep. This copy is no longer necessary, either because the payment has been successfully completed or because the access data has been recorded on the blockchain as an incident.

In turn, the logic of the SC verifies whether the time constraints for publishing an incident are met, using timestamps from ep as reference to determine if the designated timeframe has elapsed since the encounter proof was generated. It is important to note that the RAA, as the SC owner, is responsible for setting this timeframe. Once $(ep, \sigma_V(ap))$, which contains β, is published on the blockchain, the RAA can initiate sanctioning measures against V. This is possible because the access incident data uploaded to the blockchain includes V's temporary alias, i.e. β, along with V's signature, which can be verified using V's certificate $Cert(pk_V)$. Since the RAA knows the link between β and V's data, it is the sole entity capable of identifying the vehicle owner.

Once RAA resolves the incidence, the SC automatically rewards the ACV that published it on the blockchain, transferring a certain amount of cryptocurrency to their wallet. This transaction enables the SC to generate revenue from its services. It is assumed that the SC utilizes an owner-restricted method for this specific purpose.

H. Vehicle Code Renewal

Despite the concealment of each V's identity behind a vehicle temporary alias β or an Ethereum wallet address, these elements, if used in different interactions, may be successfully correlated, leading to the disclosure of V's owner identity.

To prevent this situation, any V can request a new code β^* from the RAA at any time. This action serves to effectively thwart attempts by other entities to link independent accesses. The change of a vehicle temporary alias β implies that new cryptographic keys (sk_V, pk_V) and the corresponding certificate $Cert(pk_V)$ have to be generated. This is because the code β is embedded in the certificate's common name field.

Regarding the case where V uses a single Ethereum wallet configuration for its payments, it is advisable to create a new Ethereum wallet W_V to utilize a fresh address when publishing transactions on the blockchain. To carry out this renewal process, the same operations outlined in III-A are followed. Once the new wallet W_V^* is generated, a method for anonymously transferring funds between the old wallet W_V and the new one W_V^* becomes necessary. For this purpose, V should send a request to a mixing service M for transferring the remaining cryptocurrencies from the old wallet to the new one, thereby obscuring the connection between the source and destination addresses.

IV. Fraud Detection Evaluation

In this section, we evaluate the detection rate of our system, which corresponds to functional requirement *F1* as defined in Section II-E. The evaluation of the remaining requirements, including security and privacy requirements, is given in the sections that follows.

Given the decentralized nature of the proposed solution and the lack of fixed infrastructure control points, the effectiveness of fraud detection and enforcement measures directly depend on the quantity of operational ACVs navigating the restricted area. In this context, assessing the fraud detection ratio across different traffic loads within the restricted area emerges as a pivotal factor in evaluating the viability of the proposed solution, a task accomplished through a methodical set of simulations, which is next described and discussed.

A. Simulation Scenario

To conduct these simulations, we used a combination of SUMO [37], OMNET++ [38], and VEINS [39] for urban traffic simulation and the underlying vehicle-to-vehicle (V2V) communication network. SUMO, a versatile traffic simulation

tool, is specially suited for modeling and simulating transportation systems, encompassing vehicles, public transportation systems, and pedestrians. It eases the analysis of various aspects such as traffic flow, congestion, and emissions. On the other hand, OMNeT++ is used for simulating network data transmission, providing insights into network performance, quality of service, and other network-related metrics. Lastly, Veins serves as a simulation framework that seamlessly integrates SUMO and OMNET++, facilitating communication between the two simulators. This framework is instrumental in simulating vehicular networks and conducting a comprehensive analysis of their performance.

The Luxembourg SUMO traffic (LuST) scenario [40] is the simulated urban area used to depict vehicular traffic that closely resembles real-world conditions. This scenario consists of 203 signalized intersections, a total of 2247 nodes, and 5779 road segments, covering an area of almost 156 km^2 with a total of 930 km of roads. Moreover, it includes real traffic demand and mobility patterns of 295,979 vehicles and 224 traffic jams over 24 hours, simulating a mid-sized European city with varying traffic patterns throughout a typical day. These patterns include high density during rush hours, moderate density during the daytime, and low density during nighttime.

Next, we configured the VEINS framework to simulate the movement of three distinct vehicular actors within the LuST scenario, presupposing the operational deployment of the proposed system. These vehicular actors are:

- *ACVs*: Upon identifying a nearby vehicle within a 10-meter range, these vehicles activate the access ticket verification protocol. Notably, the 9-10 meter distance corresponds with the safety guidelines endorsed by the Spanish Directorate-General for Traffic (DGT) for maintaining appropriate separation between vehicles at the recommended speed in an urban setting.[14]
- *Honest Vs*: These vehicles follow the protocol by consistently responding to the challenges presented by ACVs.
- *Fraudulent Vs*: Vehicles that can either ignore the challenges posed by ACVs or evade payments if no challenges are received during their transit through the controlled restricted area.

The configured simulation model aims to assess the probability of Vs moving undetected within the restricted area. This evaluation relies on the traffic parameters of the LuST scenario observed over a 24-hour period, taking into account different proportions of the aforementioned vehicular actors. To this end, we conducted multiple simulations using the 24-hour vehicular data from the previously mentioned LuST scenario. We systematically adjusted the proportion of vehicles serving as ACVs, specifically at 10%, 20%, 40% and 80% of the total 295,979 vehicles, in order to analyze their impact on the feasibility of the proposed system.

Additionally, the designed set of simulations also allows to investigate the influence of the threshold detection parameter K, namely, the minimum number of license plate photos from distinct ACVs necessary to identify and penalize deceptive non-communicative Vs. Note that, for the sake of concreteness, non-communicative Vs are those vehicles that ignore any incoming challenge from an ACV. In this evaluation, we examined the probability of Vs traveling undetected for K threshold values set at 1, 3, 5, and 8 detections, assuming a 20% proportion of ACV.

[14]DGT - Road Safety Tips and Rules - https://www.dgt.es

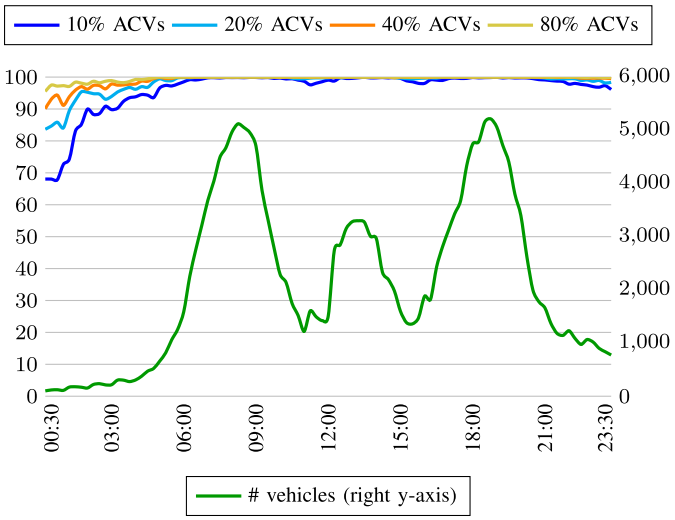

Fig. 4. SUMO simulation. Proportion of detected vehicles for 10%, 20%, 40% and 80% of ACVs.

B. Simulation Results

Figure 4 depicts the number of detected vehicles in the LuST 24-hour simulation for proportions of ACVs comprising 10%, 20%, 40%, and 80% of the total vehicles.

As these results shown, each simulation exhibits a consistent trend, revealing an elevated vehicle detection rate during peak hours that progressively diminishes with decreasing traffic density during off-peak hours. More specifically, it is noticeable that during the morning and evening peaks (6:30-10:00 and 16:45-20:30), the detection ratio reaches 100% in simulations with a high proportion of ACVs (40% and 80%). In contrast, in simulations with a lower ACV proportion (i.e., 10%) a few vehicles manage to go undetected during the same time interval, although their number never exceeds 0.50%.

A similar pattern is noted during periods of moderate vehicle density throughout the daytime (10:00-16:45), where simulations featuring substantial proportions of ACV (i.e., 40% and 80%) consistently maintain a 100% vehicle detection rate. However, for lower proportions of ACVs (i.e., 10% and 20%), the percentage of undetected vehicles slightly rises to 0.85% and 0.22% respectively, reaching peaks during the morning post-peak (11:00) at 2.43% and 0.56% of undetected vehicles.

Finally, during the night's off-peak hours (2:30-6:30 and 20:30-0:00), when the surcharges in the restricted area are relaxed, the number of undetected vehicles rises across all simulations. In particular, they reach their peaks from 2:00 to 3:00 in the morning with 10.20%, 6.94%, 3.69% and 1.79% of vehicles for proportions of 10%, 20%, 40%, and 80% ACVs respectively. This trend aligns with periods of lower vehicle density in the simulations. It is important to note that the

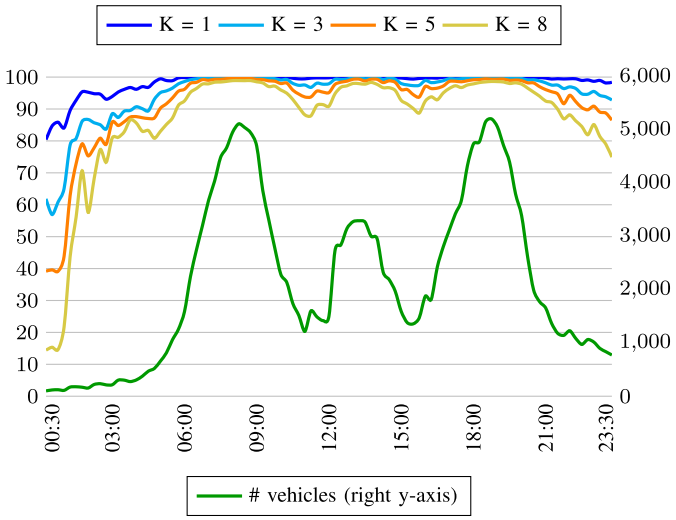

Fig. 5. SUMO simulation. Proportion of detected vehicles for thresholds K of 1, 3, 5 and 8 detections for 20% of ACVs.

intervals from 0:30 to 2:00 are excluded, as this represents the simulation's initialization phase and reflects an unusually low number of vehicles. This initialization effect is evident when comparing the number of vehicles at the simulation's end (23:59), which is 775 vehicles, to the beginning (0:15), which is 112 vehicles.

Despite the increase in undetected vehicles during the night's off-peak hours, the proposed system achieves high detection rates of vehicles: 99.09%, 99.67%, 99.88% and 99.95% for simulations with ACV proportions of 10%, 20%, 40%, and 80%, respectively. Moreover, the results indicate that regardless of the ACV proportion, achieving 100% detection is likely when more than 944 ACVs are circulating in the simulation. This implies that, given the 930 km of roads in the LuST scenario, a density of 1.02 ACVs per kilometer is required to detect any vehicle in the simulation at least once.

Upon examination of the detection rates per ACV proportion depicted above, it becomes apparent that doubling the number of ACVs from 10% to 80% does not result in a proportional increase on the total number of detected vehicles. In that matter, increasing the proportion of ACVs from 10% to 20%, 40%, and 80% yields detection increases of 0.58%, 0.21%, and 0.07% of the total vehicles in the simulation, respectively. Consequently, it can be inferred that escalating the number ACVs beyond 20% does not offer a favorable trade-off in terms of vehicle detection, considering the potential unrealistic assumption of a high proportion of ACVs driving through the designated restricted area. Based on this observation, we fix the proportion of ACVs at 20% for the remainder of this study.

In this regard, Figure 5 illustrates the impact of the threshold detection parameter K on the vehicle detection rate within a simulation featuring the aforementioned 20% proportion of ACVs. In this scenario, K values of 1, 3, 5 and 8 were tested. The results achieved under these conditions show average detection ratios of 99.67%, 98.57%, 97.19%, and 94.79% for K values of 1, 3, 5, and 8, respectively, over a 24-hour period in the LuST scenario. As it can be appreciated, the same detection patterns that were depicted in Figure 4 persist throughout various times of the day, with a consistently higher vehicle detection rate during peak hours that gradually decreases during off-peak hours. However, the requirement of detecting each V by multiple distinct ACVs shows a more notable influence on the detection ratio than observed during the assessment of varied ACV proportions.

Going into the details and distinguishing between different daytime periods, it can be observed that, during morning and evening peaks (i.e., 6:30-10:00 and 16:45-20:30), setting K at 3, 5, and 8 results in average detection rates of 99.68%, 99.14%, and 98.01% for vehicles, respectively. It is worth mentioning the 100% detection rate achieved by the simulation with $K = 3$ during rush hours. Focusing on the moderate vehicle density periods (i.e., 10:00-16:45), these detection ratios decrease to 98.59%, 96.92% and 93.90% for K at 3, 5, and 8, respectively. Finally, putting the spotlight on late and early morning off-peak hours (i.e., 2:30-6:30 and 20:30-0:00), the average detection ratios decrease to 94.06%, 90.70% and 85.73% when detecting Vs at least 3, 5 and 8 times, reaching peaks of 26.68% undetected vehicles for $K = 8$ during periods of lower vehicle density.

In summary, the conducted simulations have demonstrated that, with a 20% proportion of ACVs, the proposed solution attains an average detection ratio of 99.67% when only one license plate photo is necessary (i.e., $K = 1$), and 94.79% when a minimum of 8 ACVs must have identified the rogue vehicle (i.e., $K = 8$). While the detection ratio may drop to 73.32% during certain early morning traffic off-peaks with $K = 8$, it is important to note that the parameter K is dynamically configured by the RAA in charge of the restricted area. Consequently, it is anticipated that RAAs will employ distinct detection thresholds K based on traffic density, aiming to optimize the overall detection rate.

V. Implementation and Performance Assessment

The section evaluates functional requirement *F2*, aiming to assess the feasibility of the proposed system in a scenario aligned with Technology Readiness Level 5 (TRL5), as defined by the European Commission.

To achieve this, we begin by detailing the implementation of the system prototypes, including the hardware and software specifications of the V and ACV entities. Following the implementation, we conduct a series of experiments in two distinct settings. First, a controlled laboratory environment is used to evaluate the protocol's performance under optimal conditions, allowing for precise measurement of computation and communication costs. Second, field experiments on a low-traffic street are performed to validate the system's performance in a realistic scenario, considering the dynamic nature of vehicular movement and potential interference.

Finally, we contextualize the experimental results by comparing the performance of the proposed system with similar approaches in the literature. This comparative analysis highlights the computational and communication efficiencies of our solution while demonstrating its advantages in implementing the lateral surveillance model without the need for costly infrastructure.

Fig. 6. In-vehicle set up (top). Detection ultrasonic sensor (bottom).

A. Implementation Details

To evaluate the feasibility of the proposed system in a realistic setting, we implemented two prototypes representing the V and ACV entities involved in the *access ticket verification* protocol. These prototypes were designed as standalone compact LCD devices equipped with short-range communication capabilities to simulate the On-Board Units (OBUs) used by vehicles. The testbed was configured as follows:

- The OBUs of V and the ACV (Figure 6) are implemented on an APU2.E4 single board featuring a 1 GHz quad-core AMD G-Series GX-412TC CPU, 2GB RAM DDR3-1333Mhz, 16GB M-Sata SSD, Debian OS, and powered by 12V AC batteries. Each OBUs is equipped with a Compex WLE200NX miniPCI-e card, configured to operate in accordance with the DSRC (Dedicated short range communication) specifications, specifically utilizing the 5.9 GHz frequency band and operating in Out of the Context of a Base-station (OCB) mode. Moreover, the ACV is enhanced with an ultrasonic sensor designed to detect nearby vehicles and automatically initiate the *access ticket verification* protocol.
- The protocols running on both ACV and V's OBUs were developed using Java7 (openjdk-1.7). The communication between these entities was facilitated through IEEE 802.11p/DSRC, employing the Java.net Sockets library.
- The ECDSA cryptosystem with 256-bit key size was employed for computing the signatures and encryptions needed in the designed protocols. Furthermore, within the *access ticket verification* protocol, the AES symmetric encryption scheme with 256-bit key size was utilized for managing the session keys.
- The smart contract SC managing the *Payment* protocol was implemented in Solidity, an object-oriented programming language specifically designed for the Ethereum Virtual Machine (EVM). Solidity facilitates the deployment of smart contracts on EVM-compatible networks, such as Ethereum, Binance Smart Chain, IoTeX, and IOTA. For our experiments, we opted for the Ethereum ecosystem due to its maturity and the availability of advanced smart contract tools for testing, debugging, and conducting performance analysis. To simplify the evaluation process, the SC was deployed on the Sepolia Test network,[15] a testnet blockchain designed to mimic the operating environment of an EVM "mainnet". Through the use of Sepolia, we can assess the smart contract's performance without incurring actual economic costs, as would be the case on the main Ethereum network.
- To run cryptographic operations within the digital wallets, we followed the EVM key specifications. Specifically, we employed ECDSA with a 256-bit private key and a point on the secp256k1 ECDSA curve, represented as an (x, y) point, for the corresponding public key. The public Ethereum addresses were generated by taking the lower 160 bits of the Keccak-256 (also known as SHA-3) digest of the public key.

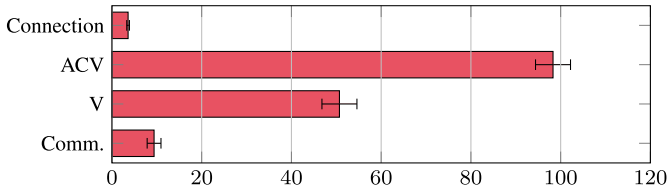

Fig. 7. Protocol execution and deviation time in milliseconds.

B. Performance Assessment in a Laboratory Environment

In order to analyze the performance of the communication protocol phases under real-time constraints (i.e, functional requirement *F2*), a set of experiments was conducted in a *controlled laboratory environment*. This setting allowed for the assessment of the new scheme's performance under optimal conditions. Both V and ACV prototypes were positioned in close proximity, ensuring a clear line of sight and no interference in their communication frequency. Twenty executions of the *access ticket verification* protocol were carried out to establish a confidence interval.

The results highlight an average completion time of 162 milliseconds for the *access ticket verification* protocol, measured from the moment the ACV requests a challenge until the nearby V receives a signed access proof. Noteworthy, observed completion times extremes range from 153 and 173 milliseconds. At the worst completion time, a V traveling at the urban speed limit of 50 km/h could travel approximately 2.40 meters away from the ACV. This suggests that the communication range between vehicles should be at least this distance to ensure reliable protocol execution.

Taking a closer look at the test results, Figure 7 examines the obtained average time, revealing key components of the

[15]Ethereum Sepolia Faucet - https://sepoliafaucet.com

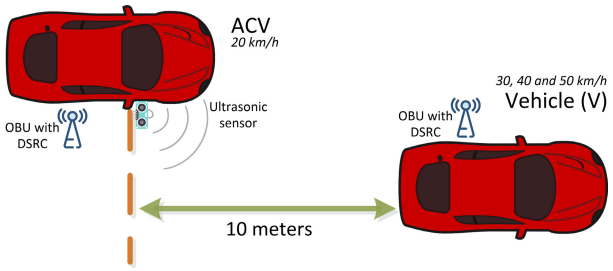

Fig. 8. Scenario for experiments on a low-traffic street.

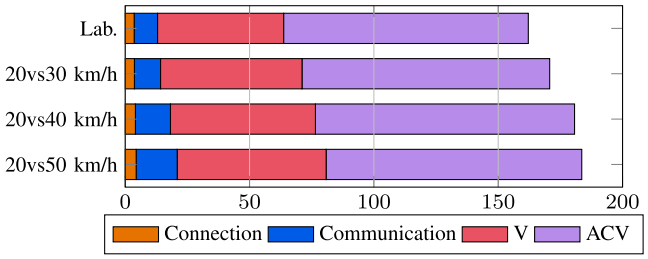

Fig. 9. Average protocol completion time (milliseconds) at various speeds.

access ticket verification protocol. This breakdown includes computation costs, communication overheads, and connection establishment time. It is worth noting that the protocol cost times incorporate the associated efforts in establishing a robust bilateral authentication between V and ACV.

As it can be appreciated in that figure, 91.95% of the total protocol runtime is allocated to computation. This is distributed between 98.3 milliseconds on the ACV side and 50.7 milliseconds on the V side. The remaining 13 milliseconds are attributed to communication overhead, encompassing 3.6 milliseconds for establishing the connection between ACV and V's OBUs, and the remaining 9.4 milliseconds for transmitting 1869 bytes of data.

It is worth mentioning that the use of DSRC, a vehicular-specialized communication technology, results in significantly low establishing and transmission times. Despite a variability coefficient of 16.4% in data transmission, the absolute impact on the protocol time remains limited to a range of ± 3.5 milliseconds.

In summary, the obtained results indicate that the proposed system successfully addresses the real-time constraints imposed upon it, demonstrating its feasibility in an interference-free scenario.

C. Performance Assessment in a Low Traffic Street

In this section, we extend the validation of the functional requirement *F2*. Following the analysis of communication protocol performance under real-time constraints in a controlled laboratory setting, we now conduct a series of experiments in a practical scenario to affirm the TRL5 maturity level of the novel proposal.

For this purpose, first, we deployed the implemented prototypes in two separate cars as follows: i) The V's OBU device was placed in the co-driver's seat, ready to receive incoming detection announcement; and ii) The ACV's OBU device was situated in the co-driver's seat, equipped with an ultrasonic sensor attached to the vehicle's rear-view mirror to measure the distance to the V's vehicle. Subsequently, a series of tests were conducted on a low-traffic street situated in an industrial area. During these tests, the ACV maintained a constant speed of 20 km/h, while the V traveled at speeds of 30, 40, and 50 km/h, overtaking the ACV. For each proposed speed, the ACV initiated the *access ticket verification* protocol whenever its sensor detected the V within a 10-meter range from the vehicle's front-end. The protocol was run ten times for each considered speed. The scenario of this set of tests is depicted in Figure 8.

The outcomes of this experiment, presented in Figure 9, reveal that the *access ticket verification* protocol takes 171, 181, and 184 milliseconds for completion when the V exceeds the ACV speed by 10, 20, or 30 km/h, respectively. These figures indicate an increase ranging from 9 to 21 milliseconds compared to the experiments conducted in the controlled laboratory environment. Such variations in time are understandable in the context of a realistic scenario, influenced by factors like movement, distance between devices, and potential signal interference arising from the vehicles' bodies.

Table I provides a numerical representation of the protocol performance illustrated in Figure 9, examining the same protocol components analyzed in the controlled laboratory experiments. Notably, these results reveal an ascending pattern in all communication-related processes as the speed difference between the ACV and V increases. Specifically, DSRC connection times increase by 18.9%, and data transfer overheads surge by 57.1% between the 20-30 km/h and 20-50 km/h experiments. This upward trend in communication-related overheads can be attributed to the growing distance between V and ACV due to the difference in vehicle speeds during the protocol run.

In terms of the time required to complete the *access ticket verification* protocol between V and ACV, the outdoor experiments yielded results indicating that this process takes only 184 milliseconds in the worst-case scenario (i.e., 20 vs. 50 km/h). With a speed difference of 30 km/h, V and ACV separate by 1.53 meters during the protocol runtime. Considering the 10-meter condition required to initiate the protocol, this results in a total distance of 11.53 meters between both vehicles at the protocol's completion. Such a distance can be easily covered by the OBUs' WLE200NX cards implementing IEEE 802.11p/DSRC V2X communication technology.

Based on the aforementioned considerations, it can be stated that the proposed system is viable within a relevant scenario, aligning with Technology Readiness Level 5 according to the European Commission standards.

D. Comparing Our Proposal With Similar Alternatives

In the preceding sections, we have shown that the phases of our solution undergoing real-time constraints are sufficiently lightweight for a TRL 5 scenario. However, a fully distributed approach like the one we propose may introduce additional overheads to secure operations, which centralized

TABLE I
LOW TRAFFIC STREET RESULTS (IN MILLISECONDS)

		Conn.	Protocol run time			Total
			V	ACV	Comm.	
20 vs 30	Avg.	3.7	56.9	99.6	10.5	170.8
	Dev.	0.1	14.8	5.2	1.9	15.4
20 vs 40	Avg.	4.1	58.3	104.2	14.1	180.9
	Dev.	0.4	1.7	5.2	1.5	7.7
20 vs 50	Avg.	4.4	59.9	102.8	16.5	183.6
	Dev.	0.4	3.7	8.6	2.6	7.0

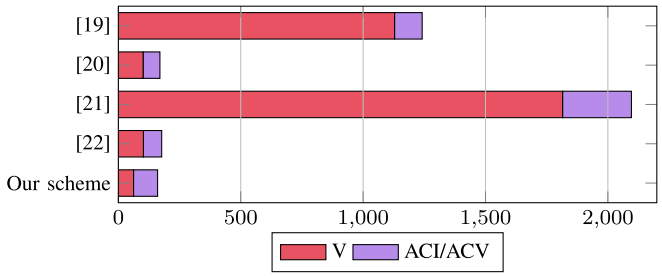

Fig. 10. Computational costs (milliseconds) in the Access phase.

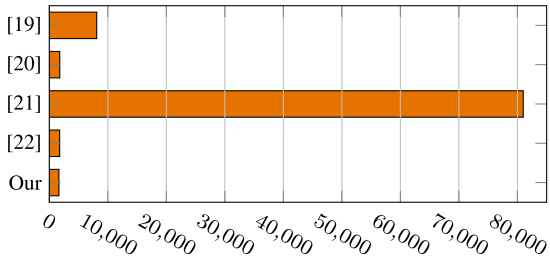

Fig. 11. Bytes transferred during the Access control phase.

infrastructures inherently offer. With this premise, in this section, we aim to compare the computation and communication costs of our proposal with similar schemes in the literature that heavily rely on infrastructure. The objective of this comparison is to illustrate that our new proposal eliminates the need for deploying costly infrastructures in line with the lateral surveillance model, while still maintaining a lightweight profile similar to recent panoptic-based approaches.

To fulfill this goal, we implemented and tested the related schemes [19], [20], [21], [22], featuring real-time constraints on access control phases with similar privacy requirements. Specifically, we replicated the vehicle-to-infrastructure access control phases from these works and conducted tests within our controlled laboratory environment. For consistency, the Access Control Infraestructures (ACI), which serve as verifiers in the panoptic-based proposals, were emulated on the same APU2.E4 single board utilized for the ACV verifiers in our proposal.

The aim of these tests was to assess the computational workload and data transmission overhead encountered by each participant during the vehicle access procedure. It is worth mentioning that we excluded the proposal presented in [30] from our comparative analysis, as its privacy requirements are notably more relaxed compared to the other schemes being analyzed.

Focusing on the computational burden, Figure 10 illustrates the computational requirements of each scheme during their respective *access control* phases, alongside the equivalent stage in our proposal, namely the *Access ticket verification* sub-protocol. These numbers highlight the computational efficiency of the new system in comparison to [19], [20], [21], and [22].

In particular, our proposal demonstrates notable lightweight characteristics at V's side, requiring less computational effort than the other studied schemes. These differences are very significant in the case of [19] and [21], where the former requires a smart card for securely storing V's keys and generating signatures, while the latter employs costly cryptographic primitives before each interaction with the deployed infrastructure. In contrast, concerning [20], [22], our new proposal exhibits slightly superior performance due to the reduced steps involved in access control, wherein V merely needs to generate and dispatch a signed access proof upon receiving a challenge, without requiring acknowledgment from the verifier.

On the other hand, from the perspective of the verifier, (i.e., ACI/ACV), Figure 10 depicts that our new proposal entails a higher computational load for this entity compared to the structure-dependent schemes outlined in [20] and [22]. This situation arises due to the additional detection and challenge engagement steps inherent in our *Access ticket verification* sub-protocol. Nevertheless, despite this increased burden on the verifier's side, when combined with the computational load experienced by V, our proposal still yields a more favorable overall computational cost compared to the alternative schemes under consideration.

Shifting the focus to communication costs, Figure 11 shows the data transfer volumes during the *access control* phase for each analyzed system and our corresponding *Access ticket verification* sub-protocol. Notably, [21] incurs in a very high data exchange volume compared with the other schemes. This is attributed to the use of zero-knowledge proofs to ensure driver anonymity during debt accumulation wallet transactions. Conversely, [19] exhibits lower data exchange, albeit still transmitting 8,102 bytes between V and the deployed infrastructure, primarily due to the transmission of signed price information during each interaction. The remaining schemes experience similar volumes of transferred data, as they undertake a comparable set of steps.

In summary, our conducted tests demonstrate that the computational workload and communication overhead incurred by the new system align with the most lightweight approaches in the literature. This achievement is realized while implementing the lateral surveillance model, which allows the new scheme to get rid of expensive dedicated infrastructures.

VI. SYSTEM SCALABILITY

In this section, we demonstrate our system's compliance with the functional requirement *F3* (*Scalability*), as defined in Section II-E. To achieve this, we identified and evaluated potential bottlenecks that could arise under heavy traffic conditions. Specifically, we analyzed the system's response in two scenarios associated with high vehicle density: i) a heavy load

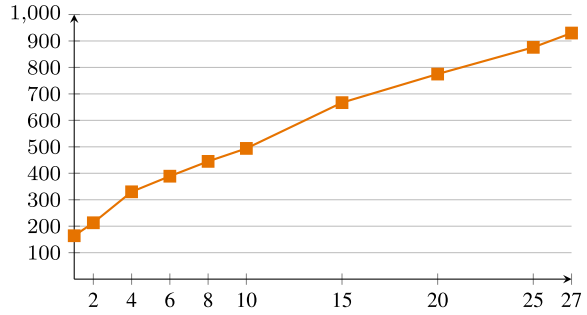

Fig. 12. Access ticket verification time in milliseconds for varying numbers of simultaneous requests.

of access verification requests on ACVs, and ii) a large volume of payment transactions on the blockchain network.

A. Access Verification Load on ACVs

Under conditions of heavy traffic and high vehicle density, ACVs, acting as validators, may face a significant number of access ticket verification requests. This influx could potentially overwhelm their capacity, jeopardizing the system's viability in real-time scenarios.

To evaluate the capability of ACVs to handle such situations, we first conducted a series of tests on our OBU prototype to assess the response times of ACVs when processing an increasing number of concurrent access ticket verification requests, simulating various high-traffic conditions. Figure 12 illustrates the average completion times (in milliseconds) for the access ticket verification protocol under workloads ranging from 1 to 27 simultaneous requests.

We then analyzed the combined impact of handling multiple access ticket requests and varying vehicle speeds on the ability of ACVs to detect and respond to dishonest Vs attempting to evade the system. This analysis is based on the following considerations:

- The maximum recognition range of commonly available ANPR mobile cameras is 25 meters.[16] Therefore, 25 meters is the maximum distance a V can travel away from an ACV before the access ticket verification request must be completed. If a dishonest V exceeds this distance before detection, it will not be properly identified.
- The analysis considers a scenario where V and ACV travel in opposite directions at speeds of 30 km/h, 50 km/h (the maximum allowable urban speed), and 70 km/h (significantly exceeding the urban speed limit).

Under the considered speed conditions, an ACV handling 10 simultaneous verifications completes the protocol in 494 milliseconds, allowing a V to travel 8.24, 13.72, and 19.21 meters at speeds of 30 km/h, 50 km/h, and 70 km/h, respectively, well within the 25-meter recognition range of ANPR cameras.

With 15 simultaneous verifications, the protocol duration increases to 667 milliseconds, during which a V can travel 11.12 meters at 30 km/h, 18.52 meters at 50 km/h, and 25.93 meters at 70 km/h. At 70 km/h, which exceeds the urban speed limit, this distance slightly surpasses the recognition range, potentially preventing the V from being detected if the protocol is not completed in time.

[16]Raiper ANPR mobile camera - https://www.anprcameras.com/wp-content/uploads/2021/03/Rapier-35P-Datasheet-2024-ITS-1.pdf

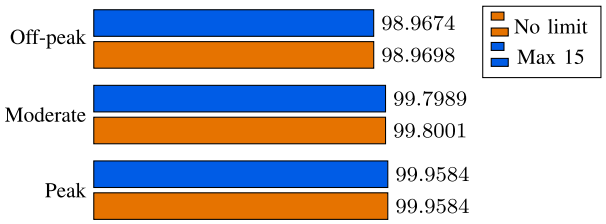

Fig. 13. Percentage of detected Vs in the RA when ACVs are limited to 15 verification requests per second vs. no request limit.

Finally, with 27 simultaneous verifications, the ACV requires 930 milliseconds per request, during which a V can travel 15.50 meters at 30 km/h, 25.84 meters at 50 km/h, and 36.18 meters at 70 km/h. At 50 km/h, this represents the first scenario where Vs adhering to the urban speed limit may bypass the protocol, as the camera's recognition range is exceeded.

These results indicate that high volumes of requests can increase protocol response time, potentially compromising service quality. According to Alfréd Rényi's parking constant [41], a two-lane, two-way, 20-meter street can accommodate a maximum of 15–16 vehicles in dense traffic. Limiting each ACV to verifying the 15 closest vehicles ensures that even at speeds above legal urban limits of up to 70 km/h, in-range verifications can be reliably completed without degrading service quality.

After establishing a limit of 15 simultaneous access ticket verifications per ACV, we evaluated the impact of this limitation on the system's overall performance. Specifically, we analyzed how this constraint affects the ability of ACVs to detect Vs under various traffic conditions. To this end, we replayed the simulation shown in Figure 5 (ACVs = 20%, K = 1), restricting ACVs to a maximum of 15 simultaneous access ticket verifications per second.

Figure 13 illustrates the percentage of Vs detected in the restricted area (RA) by ACVs with and without this limitation under off-peak (2:30–6:30 and 20:30–0:00), moderate (10:00–16:45), and peak (6:30–10:00 and 16:45–20:30) traffic conditions. As observed, limiting each ACV to 15 simultaneous verifications has no impact during peak traffic hours, with the same number of Vs detected as in the unrestricted scenario. During moderate and off-peak traffic conditions, the effect is minimal, resulting in only 1 or 2 fewer Vs being detected in absolute terms.

These findings demonstrate that the proposed system can effectively handle the load on ACVs, confirming its scalability even under varying traffic conditions.

B. Payment Transaction Load on the Blockchain

Blockchain networks can experience congestion during periods of high transaction volume, when transaction demand exceeds the blockchain's processing capacity. This congestion can lead to delays in transaction validation and inclusion in

TABLE II
SMART CONTRACT OPERATION COSTS IN GAS AND DOLLARS

Method	Gas	Binance Smart Chain
deploy	741,914	0.0041 bnb (≃0.92 $)
register_access	50,551	0.0003 bnb (≃0.06 $)
payment_incidence	182,909	0.0001 bnb (≃0.22 $)

the blockchain, with the delay duration often influenced by associated gas fees and the network's priority rules.

In our proposal, the time required for transaction validation determines the speed at which a V's access payment is verified by ACVs. To mitigate this issue, our scheme is intentionally designed to operate independently of delays caused by blockchain network congestion. Specifically, protocols involving interactions with smart contracts are decoupled from real-time operations. In this way, the contract owner (i.e., the RAA) can define a specific time interval within which Vs must validate their payment transactions on the blockchain, preventing ACVs from submitting incident claims during this period.

This design introduces a flexible time window, ranging from hours to days, allowing the procedure to adapt to varying blockchain transaction validation times. This flexibility offers two key advantages: i) it ensures the system remains unaffected by network congestion caused by blockchain scalability issues; and ii) it enables the use of low-priority gas prices, due to the extended time available for transaction validation.

Finally, for scenarios requiring enhanced throughput and reduced latency, the smart contract can be deployed on a layer-2 network, such as Polygon (as discussed in Section II-D), further improving the scalability and efficiency of smart contract-based protocols.

In summary, the described measures, including the flexible time window for transaction validation and the potential use of layer-2 networks like Polygon, demonstrate that the payment transaction load on the blockchain does not significantly affect the scalability of the system. These features ensure that the proposed solution remains efficient and reliable, even under varying network conditions.

VII. SMART CONTRACT GAS COST

In this section, we validate the functional requirement *F4* (*Efficient decentralized payment system*) by conducting a series of experiments to assess the gas costs associated with the key methods implemented in the smart contract. These experiments also evaluate the monetary costs incurred by the system's actors when invoking these functionalities. Specifically, we analyze:

- The *register_access* method, invoked by Vs to validate, price, and pay for their accesses to the restricted area.
- The *payment_incidence* method, invoked by ACVs to report incidents involving dishonest drivers.
- The *deployment* of the smart contract, carried out by the RAA.

Table II presents the gas costs and their monetary equivalents in dollars for the three mentioned smart contract operations. The conversion to dollars was calculated based on the deployment of the smart contract on the BNB Chain, an EVM-compatible network, considering the average gas price in October 2023[17] (i.e., 5.5 Gwei) and the Binance Coin (BNB) exchange rate[18] (i.e., 226.42$).

These figures reveal that the *deployment* of the smart contract carries the highest gas cost, amounting to 741,914. Next, the *payment_incidence* method incurs a cost of 182,909. Finally, the *register_access* method results in a cost of 50,551. To contextualize these costs, it is noteworthy that a basic transfer of Ether (ETH) or Binance Coin (BNB) on their respective native networks consumes 21,000 gas. Thus, all smart contract operations are expected to consume, at a minimum, this amount of gas. In light of this, it can be inferred that the smart contract operations utilized in the new scheme, particularly the user-end calls, do not incur significant gas expenses.

The final column in Table II quantifies the gas consumption in dollars. As the two most expensive operations, *deployment* and *payment_incidence*, are executed by for-profit entities and, therefore, do not impose a substantial burden on the system. Notably, the smart contract deployment is a one-time operation performed by RAA, which predominantly receives access fee incomes. Similarly, the *payment_incidence*, costing 0.22 dollars, is invoked by ACVs whenever a payment irregularity is detected, leading to a more substantial reward upon incidence validation.

On the other hand, the validation cost of the *register_access* operation is assumed by Vs, the end-users of the proposal, and thus warrants special consideration. In this way, running the *register_access* method for an average gas price costs 0.06 dollars, a value considered affordable for real-world applications.

As discussed in Section VI-B, the proposed scheme is not constrained by fixed transaction validation times. This flexibility offers two significant benefits to end-users regarding smart contract costs: i) they can opt for low-priority gas prices (e.g., 3 Gwei) in exchange for longer validation times; and ii) they are unaffected by occasional network congestion that could otherwise increase smart contract operation costs. Both aspects help ensure that the costs incurred by end-users during the *Payment* phase remain low.

VIII. SECURITY AND PRIVACY ANALYSIS

This section is dedicated to the analysis of the security and privacy requirements of our system, i.e., S1, S2, S3, S4, P1 and P2, and to the analysis of the functional requirement F5. Whenever possible, we have resorted into formal security proofs, otherwise providing arguments of security. Formal proofs can be identified by the keyword *Proof* right after the security lemma, while the keyword *Proof sketch* is used for security arguments.

A. Threat Model

Because vehicles use a wireless communication channel, we assume a Dolev-Yao attacker in full control of the network.

[17]BSC Average Gas Price - https://bscscan.com/chart/gasprice - 01/11/2023

[18]BNB price - https://coinmarketcap.com/es/currencies/binance-coin - 01/11/2023

That is, an attacker that can inject, block and modify messages sent to the network. In addition to a network presence, attackers may have a physical presence by moving around in their vehicles. The attacker, however, cannot tamper with the onboard-units of vehicles that have been registered by the RAA. This security feature can be accomplished by means of secure hardware and standard attestation procedures. Other assumptions that we use in our analysis are as follows:

- The RAA is trusted but curious: It operates correctly according to the system's protocol but may attempt to infer additional information from the data it processes.
- Credentials are correctly generated and distributed by the RAA: The generation and distribution processes are secure, ensuring the authenticity and integrity of credentials.
- Vehicles' On-Board Units (OBUs) remain untampered with and do not provide inaccurate information: OBUs are assumed to operate as designed, free from unauthorized modifications or compromises that could affect the accuracy or reliability of the data they provide.
- Vehicles can establish a secure connection with the RAA: Secure communication channels between vehicles and the RAA are assumed to be resistant to eavesdropping and tampering.
- ACVs can determine the proximity of Vs that successfully complete the *access ticket verification* protocol: This can be achieved through mechanisms such as distance bounding, signal strength analysis, spatio-temporal verification, or similar techniques.
- A consolidated blockchain network, such as Ethereum or Binance Smart Chain, is used: The blockchain network provides robust security features, including resistance to majority control attacks, and ensures high reliability for executing smart contracts.

In our analysis, it is often useful to refer to some vehicles as *honest* or *dishonest*. Honest vehicles behave as expected by the system and protocol specifications, while dishonest vehicles can behave in arbitrary ways.

The remainder of this section is dedicated to the analysis of the requirements stated in section II-E against the threat model we have just introduced.

B. Security Requirements

In this section, we address the security requirements of the proposed system. The first lemma focuses on requirements S1, S2, and S3, which are directly mapped to the traditional security properties (i.e., confidentiality, integrity, authentication, and non-repudiation). The second lemma addresses S4, which focuses on the payment protocol's resilience against systemic risks inherent in blockchain technology.

Security Lemma 1 (Secure Access Ticket Verification and Confidential Communication): Whenever a verifier, an ACV, completes a run of the *access ticket verification* protocol, apparently with a vehicle V, this indicates that V had previously engaged in the protocol with the ACV, and both parties agreed on the encounter proof $(ep, \sigma_V(ep))$ during the protocol run. Moreover, the access ticket remains secret.

Proof: We provide a formal security proof using the protocol verification tool TAMARIN.[19] The TAMARIN specification is available at https://github.com/rolandotr/access-ticket-verification/. Below, we list the two security properties that we verified, namely non-injective-agreement and secrecy. They can be found towards the end of the Tamarin specification file.

```
lemma noninjective_agreement:
    "All acv v ep #i.
        Commit(acv,v,ep) @i
            ==> (Ex #j. Running(v,acv,ep) @j)"
```

In this context, $Commit(v, acv, ticket)$ is an action fact issued by the vehicle establishing the expectation of agreeing on the content of the ticket with an ACV. $Running(acv, v, ticket)$ is an agreement claim made by an ACV.

```
lemma secrecy:
  "All x #i.
    Secret(x) @i ==> not (Ex #j. K(x)@j)"
```

The statement above states that all terms that are marked as secret, via the action fact $Secret$, are not leaked to the adversary, where $K(x)@j$ denotes that the adversary knows x at the jth step of the execution.

It is worth remarking that, in our specification, we abstracted the notion of time and location to global constants. This means that this security property is satisfied even when the time and location of vehicles are known. □

The preceding lemma demonstrates that our system ensures the integrity, authenticity, and confidentiality of encounter proofs collected by ACVs. These properties are essential for identifying drivers passing through restricted zones while protecting the privacy of honest users. The non-repudiation requirement is satisfied by the cryptographic signature embedded in the encounter proof and the assurance that secret keys remain secure within the vehicles' OBUs.

Security Lemma 2 (Resilience of the payment protocol to blockchain-specific security risks): The payment protocol's integration with decentralized blockchain technology is resilient to inherent blockchain security risks, ensuring the protocol's security remains intact.

proof sketch As outlined in the assumptions linked to the threat model, the proposed system operates within a consolidated blockchain network, such as Ethereum or Binance Smart Chain (BSC). These networks implement foundational security principles, including robust consensus mechanisms (e.g., PoW and PoS), diverse peer participation, and economic disincentives for malicious behavior. These features effectively mitigate inherent blockchain risks, such as majority control attacks (51% attacks), forks, Sybil attacks, and double-spending.

While these networks provide robust security, they often encounter scalability challenges, leading to delays in transaction validation. Nevertheless, as discussed in Section VI-B, our system is designed to tolerate high validation latency without compromising its functionality. This design ensures that the proposal is well-suited for deployment within consolidated

[19] https://tamarin-prover.com

blockchain networks, leveraging their security strengths while mitigating the inherent risks associated with decentralized technology.

The lemma above justifies that deploying the smart contract in a consolidated blockchain network mitigates the inherent security risks associated with decentralized technology. Consequently, the security requirement S4 is successfully fulfilled.

C. Privacy Requirements

We first address the privacy requirement P1, that of minimizing the probability of disclosure of the picture of an honest driver. We prove next that such probability is smaller than the probability of a network failure.

Privacy Lemma 1: When the network is reliable, no picture of an honest vehicle is disclosed to the RAA.

Proof: Because on-board units do not leak secrets, reports to the RAA must come from on-board units that can establish a secure channel with the RAA. An ACV takes and sends a picture to the RAA only if it senses a vehicle in front and it doesn't receive a correct encounter proof. The latter only happens if there is a network problem, as honest vehicle always self-generate a correct access ticket and Lemma 1 establishes that encounter proofs from honest vehicles arrive integral to verifiers. □

Next, we analyze the anonymity of honest drivers (i.e., privacy requirement P2).

Privacy Lemma 2: The access ticket verification protocol do not disclose identifying or traceable information about vehicles, except for the vehicle's certificate.

Proof: By inspecting the protocol description, one can notice that during the handshake both V and ACV exchange certificates to establish a secure communication channel. Any other information sent by V remains confidential, as was already proven in Lemma 1. □

Corollary 1: An honest vehicle can remain untraceable by frequently updating its certificate.

Privacy Lemma 3: The payment system let honest vehicles remain anonymous. They can remain untraceable by not using the same certificate twice.

Proof sketch: Certificates in our system are anonymous, but not untraceable. We conclude that, if a vehicle wishes to remain untraceable, it needs to change certificates after each use. This includes the credentials use to publish in the blockchain. Beside credentials, the other source of identification is the data published in the blockchain. The information available in the blockchain is: the transaction identifier δ, the entrance and departure timestamps t_V^{in} and t_V^{out}, respectively, the restricted zone ID and the vehicle category. The transaction identifier is anonymous. And, the timestamps and zone identifier are weak identification attributes. Therefore, the payment process cannot be used to de-anonymized an honest vehicle.

We conclude that our system lets honest vehicles remain anonymous. As far as untreaceability is concerned, our system offers a trade-off. The more a vehicle changes credentials, the less traceable it becomes, which comes at the cost of a higher computational cost.

D. Functional Requirement F5: Exculpability Validation

We end this theoretical analysis by arguing that our system also satisfies the functional requirement (F5) as established in Section II-E. The other functional requirements have already been evaluated in previous sections.

Functional Lemma 1: Honest drivers, i.e. those who follow the system and protocol specifications, are not incorrectly fined by our system.

Proof sketch: There exist two types of accusations against vehicles, and both are triggered by ACVs. The first type occurs when an ACV notices that an access ticket has not been paid. In this case, the ACV can execute the corresponding smart-contract to log the incident. The logic of the smart-contract, however, prevents ACVs from successfully logging an incident on a ticket that has been paid or that it is still within the payment expiration time window. Therefore, an honest driver, i.e. a driver that correctly self-generates and pays its tickets, cannot be falsely accused of not paying.

The second type of accusation occurs when enough ACVs report a vehicle V for not sending a correct encounter proof. This means that the RAA has gathered sufficient visual evidences about the presence of V in a restricted. If the number of reports, notably from different ACVs, is larger than a predefined threshold, then the RAA opens an investigation. During that investigation the RAA will establish a secure channel with V to request proofs of a paid access ticket that is consistent with the visual evidence reported by the ACVs. V can prove having paid the access to the restricted zone by pinpointing to the transaction in the payment system which the RAA can verify to match the vehicle's pseudo-identifier. Because V is honest, the area identifier as well as the entry/exit timestamps will be consistent with those reported by the ACVs.

Note that, for this proof, we are assuming that the payment system and vehicle OBU's are reliable. That is to say, we are assuming that vehicles are able pay and self-generate access tickets.

As expected, our system is conservative in the sense that it lets dishonest drivers to get away without paying with some probability, but it does not fine honest drivers.

IX. Concluding Remarks

In recent years, the enforcement of vehicle restricted areas (e.g., LEZs and CCs) has emerged as an effective solution to the traffic congestion and high environmental pollution problems prevalent in large metropolitan areas. Traditionally, the enforcement of these designated zones has been accomplished using the highly centralized and infrastructure-expensive *panoptic model*. Nowadays, the rise of smart and sensing-capable vehicles has allowed the transition to the decentralized and infrastructure-free *lateral surveillance model*.

While the *lateral surveillance model* offers important benefits, its implementation also presents significant challenges. These include the indiscriminate and systematic tracking of drivers, potential system misuse by dishonest drivers exploiting the absence of a central trusted authority, and drivers'

lack of motivation to invest resources in monitoring others, consequently reducing the system's effectiveness.

This paper introduces a novel autonomous, decentralized, and infrastructure-free solution for managing vehicle access to vehicular restricted areas, in line with the innovative *lateral surveillance model*. This new scheme, first, guarantees user anonymity during fee pricing and payment processes through Blockchain technology and privacy preservation measures. Second, it can detect dishonest drivers, revoke their anonymity, and impose suitable sanctions.

The new proposal has been implemented and evaluated in both a controlled environment and a low-traffic street. Results from these experiments indicate that the decentralized solution is lightweight and viable in a relevant scenario consistent with the Technology Readiness Level 5 according to the European Commission standards. Our assessment of the smart contract's gas cost reveals that this factor does not significantly impact the overall feasibility of the system.

Additionally, the scalability of the system has been analyzed, considering two critical factors: the access verification load on ACVs and the payment transaction load on the blockchain. The results demonstrate that the proposed system remains scalable even under varying traffic conditions, effectively handling high volumes of simultaneous access verifications and adapting to potential blockchain network congestion without compromising performance or reliability.

REFERENCES

[1] K. Gwilliam, "Cities on the move—Ten years after," *Res. Transp. Econ.*, vol. 40, no. 1, pp. 3–18, Apr. 2013.

[2] World Health Organization and UNAIDS, *Air Quality Guidelines: Global Update 2005*. Geneva, Switzerland: World Health Organization, 2006.

[3] POLIS, Transport Decarbonisation Alliance and C40. (2020). *Zero-Emission Zones: Don't Wait to Start With Freight!* Accessed: Sep. 28, 2023. [Online]. Available: https://www.polisnetwork.eu/news/how-to-guide-zero-emission-zones-freight/

[4] G. Santos, "Urban congestion charging: A comparison between London and Singapore," *Transp. Rev.*, vol. 25, no. 5, pp. 511–534, Sep. 2005.

[5] D. Lyon, *Surveillance Society: Monitoring Everyday Life*. New York, NY, USA: McGraw-Hill, 2001.

[6] X. Shen, R. Fantacci, and S. Chen, "Internet of Vehicles [scanning the issue]," *Proc. IEEE*, vol. 108, no. 2, pp. 242–245, Feb. 2020.

[7] M. Andrejevic, "The work of watching one another: Lateral surveillance, risk, and governance," *Surveill. Soc.*, vol. 2, no. 4, Sep. 2002.

[8] R. A. Popa, H. Balakrishnan, and A. J. Blumberg, "VPriv: Protecting privacy in location-based vehicular services," in *Proc. 18th USENIX Secur. Symp.*, Aug. 2009, pp. 335–350.

[9] X. Chen, G. Lenzini, S. Mauw, and J. Pang, "A group signature based electronic toll pricing system," in *Proc. 7th Int. Conf. Availability, Rel. Secur.*, Aug. 2012, pp. 85–93.

[10] J. Balasch, A. Rial, C. Troncoso, B. Preneel, I. Verbauwhede, and C. Geuens, "Pretp: Privacy-preserving electronic toll pricing," in *Proc. USENIX Secur. Symp.*, vol. 10, 2010, pp. 63–78.

[11] S. Meiklejohn, K. Mowery, S. Checkoway, and H. Shacham, "The phantom tollbooth: Privacy-preserving electronic toll collection in the presence of driver collusion," in *Proc. USENIX Secur. Symp.*, vol. 201, 2011, pp. 1–16.

[12] J. Day, Y. Huang, E. Knapp, and I. Goldberg, "SPEcTRe: Spot-checked private ecash tolling at roadside," in *Proc. 10th Annu. ACM Workshop Privacy Electron. Soc.*, Oct. 2011, pp. 61–68.

[13] F. D. Garcia, E. R. Verheul, and B. Jacobs, "Cell-based privacy-friendly roadpricing," *Comput. Math. Appl.*, vol. 65, no. 5, pp. 774–785, Mar. 2013.

[14] H. K. Leung, X.-Z. Chen, C.-W. Yu, H.-Y. Liang, J.-Y. Wu, and Y.-L. Chen, "A deep-learning-based vehicle detection approach for insufficient and nighttime illumination conditions," *Appl. Sci.*, vol. 9, no. 22, p. 4769, Nov. 2019.

[15] I. Lashkov, R. Yuan, and G. Zhang, "Edge-computing-facilitated nighttime vehicle detection investigations with CLAHE-enhanced images," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 11, pp. 13370–13383, Nov. 2023.

[16] B. Wang, H. Zheng, K. Qian, X. Zhan, and J. Wang, "Edge computing and AI-driven intelligent traffic monitoring and optimization," *Appl. Comput. Eng.*, vol. 67, no. 1, pp. 41–46, Jun. 2024.

[17] H. Yang, J. Cai, C. Liu, R. Ke, and Y. Wang, "Cooperative multi-camera vehicle tracking and traffic surveillance with edge artificial intelligence and representation learning," *Transp. Res. C, Emerg. Technol.*, vol. 148, Mar. 2023, Art. no. 103982.

[18] R. Borges, F. Sebé, and M. Valls, "An anonymous and unlinkable electronic toll collection system," *Int. J. Inf. Secur.*, vol. 21, no. 5, pp. 1151–1162, Oct. 2022.

[19] R. Jardí-Cedó, M. Mut-Puigserver, M. M. Payeras-Capellà, J. Castellà-Roca, and A. Viejo, "Time-based low emission zones preserving drivers' privacy," *Future Gener. Comput. Syst.*, vol. 80, pp. 558–571, Mar. 2018.

[20] C. Anglès–Tafalla, J. Castellà-Roca, M. Mut–Puigserver, M. M. Payeras–Capellà, and A. Viejo, "Secure and privacy-preserving lightweight access control system for low emission zones," *Comput. Netw.*, vol. 145, pp. 13–26, Nov. 2018.

[21] V. Fetzer, M. Hoffmann, M. Nagel, A. Rupp, and R. Schwerdt, "P4TC—Provably-secure yet practical privacy-preserving toll collection," *Proc. Privacy Enhancing Technol.*, vol. 2020, no. 3, pp. 62–152, Jul. 2020.

[22] C. Anglés-Tafalla, A. Viejo, J. Castellà-Roca, M. Mut-Puigserver, and M. M. Payeras-Capellà, "Security and privacy in a blockchain-powered access control system for low emission zones," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 1, pp. 580–595, Jan. 2023.

[23] S. Mathur et al., "ParkNet: Drive-by sensing of road-side parking statistics," in *Proc. 8th Int. Conf. Mobile Syst., Appl., Services*, 2010, pp. 123–136.

[24] Waze Mobile. (2015). *Waze Outsmarting Traffic, Together*. [Online]. Available: http://www.waze.com

[25] Q. Liu, T. Han, J. L. Xie, and B. Kim, "Real-time dynamic map with crowdsourcing vehicles in edge computing," *IEEE Trans. Intell. Vehicles*, vol. 8, no. 4, pp. 2810–2820, Apr. 2023.

[26] M. Jameela, H. Afzal, K. Khurshid, and A. W. Malik, "Crowdsourced system to report traffic violations," in *Proc. 4th Int. Conf. Vehicle Technol. Intell. Transport Syst. (VEHITS)*, vol. 1, Funchal, Portugal, 2018, pp. 315–322.

[27] H. Chen, B. Guo, Z. Yu, and Q. Han, "CrowdTracking: Real-time vehicle tracking through mobile crowdsensing," *IEEE Internet Things J.*, vol. 6, no. 5, pp. 7570–7583, Oct. 2019.

[28] Z. Jiang et al., "CrowdPatrol: A mobile crowdsensing framework for traffic violation hotspot patrolling," *IEEE Trans. Mobile Comput.*, vol. 22, no. 3, pp. 1401–1416, Mar. 2023.

[29] L.-W. Chen, Y.-C. Tseng, and K.-Z. Syue, "Surveillance on-the-road: Vehicular tracking and reporting by V2V communications," *Comput. Netw.*, vol. 67, pp. 154–163, Jul. 2014.

[30] S. Bouchelaghem and M. Omar, "Reliable and secure distributed smart road pricing system for smart cities," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 5, pp. 1592–1603, May 2018.

[31] S. Nakamoto. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. [Online]. Available: https://bitcoin.org/bitcoin.pdf

[32] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum Project Yellow Paper*, vol. 151, pp. 1–32, Jan. 2014.

[33] X. Fan and Q. Chai, "Roll-DPoS: A randomized delegated proof of stake scheme for scalable blockchain-based Internet of Things systems," in *Proc. 15th EAI Int. Conf. Mobile Ubiquitous Syst. Comput. Netw. Services*, 2018, pp. 482–484.

[34] S. Popov et al., "The coordicide," *Accessed Jan*, vol. 12, pp. 1–30, Jan. 2020.

[35] J. Kanani, S. Nailwal, and A. Arjun, "Matic whitepaper," Sep. 2021. [Online]. Available: https://github.com/maticnetwork/whitepaper

[36] R. S. Yokoyama, B. Y. L. Kimura, L. A. Villas, and E. D. S. Moreira, "Measuring distances with RSSI from vehicular short-range communications," in *Proc. IEEE Int. Conf. Comput. Inf. Technol.; Ubiquitous Comput. Commun.; Dependable, Autonomic Secure Comput.; Pervasive Intell. Comput.*, Oct. 2015, pp. 100–107.

[37] SUMO. (2023). *Sumo—Simulation of Urban Mobility*. Accessed: Apr. 12, 2023. [Online]. Available: http://sumo.sourceforge.net/
[38] OMNET++. (2023). *OMNET++—Network Simulation Framework*. Accessed: Apr. 12, 2023. [Online]. Available: https://omnetpp.org/
[39] VEINS. (2023). *The Open Source Vehicular Network Simulation Framework*. Accessed: Apr. 12, 2023. [Online]. Available: http://veins.car2x.org/
[40] L. Codeca, R. Frank, and T. Engel, "Luxembourg SUMO traffic (LuST) scenario: 24 hours of mobility for vehicular networking research," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Dec. 2015, pp. 1–8.
[41] A. Rényi, "On a one-dimensional problem concerning random space-filling," *Publications Math. Inst. Hung. Acad. Sci.*, vol. 3, pp. 109–127, Jan. 1958.

Carles Anglés-Tafalla received the Ph.D. degree in computer science from Universitat Rovira i Virgili, Tarragona, Spain, in 2020. He is currently a Post-Doctoral Researcher of the CRISES Research Group, Universitat Rovira i Virgili University. He has participated in several national funded research projects and authored several papers and conference contributions. His interests and line of research include data privacy, data security, and cryptographic protocols applied to vehicular scenarios.

Alexandre Viejo received the Ph.D. degree in computer science from Universitat Rovira i Virgili, Tarragona, Spain, in 2008. In 2009, he was a Researcher at Humboldt-Universität zu Berlin, Berlin, Germany. He is currently an Associate Professor at Universitat Rovira i Virgili. He has authored several papers and conference contributions. His fields of activity are data privacy, data security, and cryptographic protocols.

Rolando Trujillo-Rasua is currently a Ramon y Cajal Researcher at Rovira i Virgili University (URV). Prior to joining the URV, he held a senior lecturer position at Deakin University, Australia, and a post-doctoral position at the University of Luxembourg. His work appears in the most important publication venues in computer security, such as ACM CCS, S&P, ESORICS, and CSF, and his research interests span the areas of formal methods, computer security, and privacy protection.

Jordi Castellà-Roca (Member, IEEE) received the Engineering degree in computer systems from the University of Lleida in 1998, the Engineering degree in computer science from Rovira i Virgili University in 2000, and the Ph.D. degree in computer science from the Autonomous University of Barcelona in 2005. He is currently an Associate Professor at Rovira i Virgili University. He has published over 70 works and is the co-author of seven patents. His research interests include the fields of cryptographic protocols and privacy.