Private and Verifiable Storage of User Vehicles' Privacy Settings through Decentralisation

Gianpietro Castiglione[1,*], Sergio Esposito[1], Daniele Francesco Santamaria[1] and Giampaolo Bella[2]

[1]Department of Mathematics and Computer Science, University of Catania
[2]Department of Political and Social Sciences, University of Catania

Abstract

Processing personal data and finding effective storage instruments requires certain security and privacy properties must be guaranteed. There is no exception in the automotive realm where various services are involved and potentially sensitive information is at risk. The privacy preferences the users set across these services can reveal their personal traits and may not always be honoured by service providers. Thus, the user should have access to a system that enables them to concurrently store their preferences and legally demonstrate that they have made specific choices, thereby discouraging service providers from the improper enforcement of data processing. Simultaneously, it is crucial to distinguish a user's true identity from their choices, to prevent direct inferences.

In this contribution, we investigate the problem of preserving users' privacy preferences set aboard modern vehicles, by guaranteeing the security and privacy properties of pseudonymity, public verifiability, and integrity. The proposed design is partially built on the Hyperledger Fabric blockchain.

Keywords

User-centric, Privacy, Blockchain, Hyperledger, Privacy Enhancing Technology

1. Introduction

People's control over their personal data is worrisomely restricted in today's sophisticated automotive world. Modern vehicles, since equipped with numerous sensors and connectivity features, continuously gather extensive data, from location tracking to driver behaviour and vehicle media usage. While the proliferation of data enhances driving experience and vehicle functionalities, it significantly raises privacy risks. Correlated with this technological advancement, a systemic concern exists: when managing compliance with the EU General Data Protection Regulation (GDPR) almost all brands do not appear to be entirely aligned [1]. Although GDPR sets stringent standards for personal data protection, demanding transparency, user consent, and user control over their data, a recent study [2] analysed that many car manufacturers do not allow proper controls over the data collection and processing practices that GDPR requires. This is particularly relevant when dealing with sensitive information, often submitted by the users when choosing their preferences over the various services, sometimes mandatory, that can dangerously reveal the user's traits. Therefore, storing the privacy preferences in the automotive scenario needs to be fixed according to the regulatory requirements. The users should be able to demonstrate, in legal proceedings, that they have set specific preferences.

In light of these observations, the following research question arises:
RQ. How can we privately and verifiably store the user privacy preferences set aboard modern cars? The proposed solution within this work has been designed to ensure three privacy and security properties: **pseudonymity**, **public verifiability**, and **integrity**. Embedding these properties while finding a storing solution means that, from the user point of view, with pseudonymity, the preference would not allow one to trace back any physical person, as it is linked to a pseudonym rather than an actual identity; with public verifiability in the event of a dispute, the user can publicly prove that they

DLT2025: 7th Distributed Ledger Technology Workshop, June, 12-14 2025 - Pizzo, Italy
*Corresponding author.
gianpietro.castiglione@phd.unict.it (G. Castiglione); sergio.esposito@unict.it (S. Esposito); daniele.santamaria@unict.it (D. F. Santamaria); giampaolo.bella@unict.it (G. Bella)

selected certain preferences and that in no way can a third party decide to apply contrary ones; finally, with the integrity of the preferences, the user can demonstrate which privacy preference was set at a certain time.

In this work, we chose to adopt a decentralised solution such as blockchain, and in particular, Hyperledger Fabric. With Hyperledger we can a priori guarantee the integrity and public verifiability properties, together with management properties resulting suitable for the automotive context. Therefore, we propose a simple design that allows the enforcement of the pseudonymity property. In this way, the user can register their preferences in the blockchain without being part of it, not ultimately clashing with the permissioned nature of Hyperledger.

Paper Outline Section 2 discusses why and which security properties the proposed approach guarantee; Section 3 provides an overview of Hyperledger Fabric, as the chosen blockchain for the implementation of the solution; Section 4 illustrates the proposed general design of the solution; Section 5 illustrates the operational flow; Section 6 investigates possible attack scenarios and how could be intrinsically mitigated; Section 7 explores the state of the art; Section 8 concludes the paper by summarising the proposed approach and presenting some future research directions.

2. How to Meet the Security Properties

At the core of the design presented in this paper are the security and privacy properties, henceforth referred to as *properties*, which we aim to achieve. The properties we aim to fulfil regarding user-set privacy preferences include ***pseudonymity***, ***integrity*** and ***public verifiability***. The properties are accurately chosen considering what GDPR mandates, what rights a user might have, and the consequent actions a user might take over their data, protecting the user over a legal debate. The properties are entirely independent of the platform or tool that will be used to ensure them, and they are contingent only upon the automotive-legal environment we are considering.

Ultimately, we assert that blockchain is the appropriate solution for storing users' privacy preferences, as intrinsically allows us to obtain integrity and public verifiability, while to obtain pseudonymity we will refer to the proposed solution.

The properties' motivations are described in the following subsections.

2.1. Pseudonymity

In an ideal security scenario, it would be arguably desirable from a privacy standpoint to obfuscate the user's privacy preferences set aboard modern cars. But while applying obfuscation could be useful in avoiding direct inferences on the data, e.g., improper training or protecting from data breaches, obfuscation cannot entirely guarantee the presented security properties. They mainly fall in the legal domain and are particularly intended to protect a user from any legal issues, so not necessarily technological.

In the case of user preferences, a property that may suffice rather than actual obfuscation as interpreted above is the **pseudonymity**. With pseudonymity, the user's preferences would remain transparent and easily identifiable, while being linked solely to a pseudonym rather than their actual identity. The pseudonymity ensures that the preferences can be accessed and understood without revealing personal information, thereby enhancing privacy and protecting the user's true identity from potential misuse or unwanted exposure.

2.2. Public Verifiability

With pseudonymity, we obfuscate who is the user's identity to which the preferences are related, leaving the preferences themselves clear. But why not use obfuscation for the privacy preferences? Here comes the second property of **public verifiability**; the final fulfilment is to have privacy preferences publicly readable so that we can ensure these two main properties:

1. **Dispute Resolution**: In cases where there are disputes regarding data usage or privacy preferences, public verifiability can provide evidence to support claims made by users. This can facilitate resolution processes and enhance fairness;
2. **Independent Verification**: Stakeholders, including regulators, can independently verify that organisations are adhering to the privacy preferences set by users. This can help ensure compliance with data protection laws and regulations, such as the GDPR.

The crucial property is that the user should be entitled to publicly demonstrate that they chose a specific preference in case of dispute, should they feel that the provider was not compliant.

2.3. Integrity (Immutability)

Contextually, a user should be able to prove *many* preferences they set, hence proving how the preferences have changed over time, allowing them to prove also *past* settings. Guaranteeing such a property means guaranteeing the integrity of the preferences, namely preferences immutability. Ensuring the immutability property does not mean the user is forced to make choices that cannot change over time, but rather that it is verifiable that at time ***t*** the user has given consent to the privacy preference ***p(t)***, in ***t+1*** to ***p(t+1)*** and so on. This ensures that in case the user makes a change, thereby establishing ***p(t+1)***, they can still prove that at ***t*** they gave consent to ***p(t)***. Of course, this may also be used against a user attempting to falsely claim a preference that differs from the one that they originally defined: it should be impossible for the user to successfully repudiate the original preference.

3. Hyperledger Overview

Although some described features could be obtained by leveraging any type of blockchain, we chose Hyperledger Fabric for considerations more suited to the automotive context. Hyperledger Fabric [3] is an open-source blockchain framework proposed by the Linux Foundation, operating as a distributed ledger technology (DLT). Among the most obvious security properties of modern ICT systems, e.g., TLS for securing the transactions, Hyperledger stands out for the following features revolving around the concept of permissioned blockchain "where each component and actor has an identity, and policies define access control and governance". The transactions in Fabric require endorsements from designated peers to be considered valid. The blocks are delivered to the peers within an ordering service, towards their validation and commitment to the ledger.

3.1. Hyperledger Components

The main components of Hyperledger are:

- **Channels**: Fabric enables the creation of private subnetworks within the blockchain where specific participants can transact and share data confidentially.
- **Chaincode (Smart Contracts)**: they automate business processes by encoding rules and logic for transactions.
- **Peers**: they are the nodes in the network hosting the ledger, executing the chaincode, and endorsing transactions.

From a security point of view, the following features are unquestionably advantageous:

- **Membership Service Providers**: A Membership Service Provider (MSP) is a part that establishes the guidelines governing an organisation's admissible identities. Using X.509 certificates issued by a Certificate Authority (CA) as identities, the default MSP implementation in Fabric follows the conventional Public Key Infrastructure (PKI) hierarchical paradigm.

- **Identities**: In Fabric, there are several players such as administrators, client apps, peers, and more. Each one of these actors has a digital identity that is included in an X.509 digital certificate that has been granted by a Certificate Authority (CA). These identities are important because they specify the precise rights that participants in a blockchain network have over resources and information access.
- **Policies**: Policies in Fabric serve as the infrastructure management system. Members agree on whether to accept or reject modifications made to a channel, the network, or a smart contract by using policies.

3.2. Operations on Hyperledger

The conventional flow of operations in Hyperledger is as follows:

- A client submits a transaction proposal to the endorsing peers, who simulate the transaction and generate endorsements if it meets the criteria;
- The endorsed proposal is sent to the ordering service, which batches transactions into blocks and delivers them to all peers;
- Each peer validates the transactions in the block against the endorsement preference and world state. If valid, the block is committed to the ledger and the world state is updated;
- Clients can query the ledger to retrieve information about the current state of assets or the history of transactions.

The inherent security features and network-level protections that Hyperledger offers by default serve as the foundation upon which we can develop the design outlined in this study.

In the automotive context, it is useful to have a permissioned, therefore private, blockchain, because several assumptions could coexist: a) the blockchain has been proposed by specific car brands, which wish to propose the properties mentioned exclusively for their users; b) the blockchain has been created by specific privacy protection organisations, where only interested users can register their privacy preferences; c) the blockchain, due to its lower notoriety, can presumably be less exposed to external attacks, and internal ones thanks to its permissioned nature; d) last but not least, we want to avoid the use of public blockchains, where both the computational limitations are well-known due to their extension, and where there is a need to have credits to be able to operate.

4. General Design Proposed

We have illustrated that the properties we aim to ensure for the user privacy preferences support the choice of a decentralised solution. This is because a decentralised approach inherently possesses these properties. We have also explained that Hyperledger is the suitable solution for the application we envision and consequently for the automotive sector. This is due to its high-level properties related to blockchain management.

However, while among the properties it is easy to guarantee integrity and public verifiability, it is less obvious to guarantee pseudonymity. In the following subsections, we illustrate how we can obtain pseudonymity by providing designs at architectural and data levels.

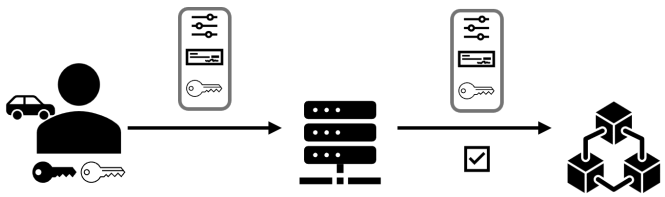

Figure 1: Design and components for pseudonymously storing the privacy preferences in the blockchain

4.1. Architectural Design

Due to the permissioned nature of Hyperledger, whoever wants to be part of the network would have to authenticate themselves, and their transactions would therefore be directly traceable due to their non-anonymous identity.

Bogatov et al. [4] have presented a cryptographic scheme tailored to Hyperledger to make certain transactions anonymous, with the possibility of revocation and auditing. The scheme however requires that the car node be part of the blockchain to operate. Conversely, in our thesis, we would like the car node from which the various preferences are set, not to be definitively part of the blockchain itself. We are examining a scenario in which each user's primary concern is that their vehicle possesses the computational and network capabilities necessary to record their preferences on the blockchain, without the desire to be an active participant in it. Furthermore, every car node might have the will to communicate not with a single blockchain peer but with a multitude of peers, implementing a loose coupling with the blockchain.

The first choice of the proposed design consists of enforcing the discussed loose coupling between the user (at the head of the car node) and the blockchain. The user (car node) remains uncoupled and in contact with the blockchain employing one or more peers as intermediate nodes between the car and the blockchain. The peer should be able not only to produce transactions towards the blockchain but also to accept, via any protocol such as SSH and FTP, the data that the user sends. The peer receives the information that the user wants to keep in the blockchain and records it in the blockchain within the appropriate transaction. The car node would therefore not be part of the blockchain.

In this way, the transactions on the blockchain, although still authenticated on the peer side, are pseudonymous on the user side. It is not possible to know who the user was who asked to upload their privacy preferences to the blockchain.

Within this design, in the basic case there would be at least three agents/components:

1. the user **U**, as not part of the blockchain-recognised identities. U (via his car) needs to know a blockchain peer for operating on the blockchain and performing all necessary operations on behalf of U.
2. the blockchain peer **P**, as responsible for intermediating between U and the blockchain, respectively for receiving the blockchain assets and submitting the assets with a proper blockchain transaction. P is a permissioned node for the blockchain therefore its identity is *not* pseudonymous.
   The peer P could be able to associate the pseudonym of the user U with the U itself. In case P was an attacker, the security property of the pseudonymity of U would be compromised since P may improperly share information about U. Thus, in our setting, we assume P to be trusted.
3. the blockchain **HF**, as responsible for making permanent the received privacy preferences and being reliable when receiving proper queries on the registered ones.

4.2. Data Design

But why pseudonymity and not just anonymity? The user would not only like his preferences to be recorded anonymously but also to be able to claim them in the future, therefore demonstrating to be the true owner, in case of disputes about the improper use of data with the various service providers. To achieve this, with our design, the user would send not only the privacy preferences but also a digest obtained by digitally signing them. In this way, the user can claim to be the author since they own the private key used to produce the digest sent. The pseudonym-public key (digest) pairing allows to bind the pseudonym (hence the user) to the preferences they have set due to the digital signature mechanism and the consequent possession of the private key.

Considering this assumption, as part of the design, the user must send a certain set of data. The data (hereafter called *blockchain asset* because it is an entire module that will then be sent to the blockchain without any modifications by the peer) contains both the privacy preferences that the user would like

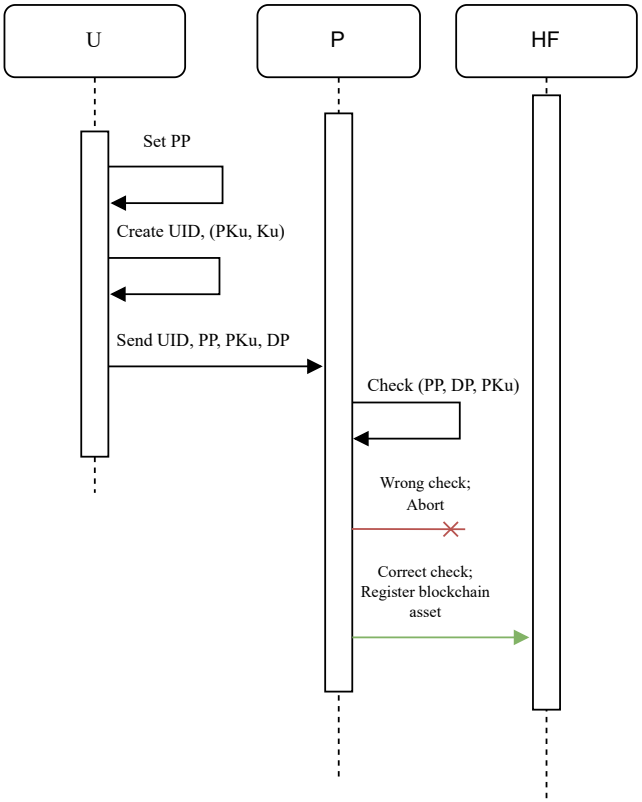

Figure 2: Sequence diagram for setting and storing the user's privacy preferences

to keep and other information necessary for the pseudonymised pairing between the user and his preferences. The blockchain asset is composed of 3 elements, as follows:

1. the wrapped privacy preferences (e.g., a json/XML file) where on top of the preferences is also included the pseudonym generated by the user;
2. the digest of the file, obtained through the digital signature of the previous file;
3. the public key, as the counterpart of the private key used in the digital signature, necessary for verifying the digest.

5. Operational Flow

Storing the privacy preferences in the blockchain, hence mutually guaranteeing the pseudonymity, public verifiability, and integrity (immutability) properties, is part of a flow where we assume a user can interact with the car services and properly set the privacy preferences for each service.

Once the privacy preferences are set, they are wrapped together to be mutually signed with the digital keys and sent to be registered.

The general operational flow includes two macro operations. The first is of setting, the second is of sending, checking and storing. The entire flow considering only the two macro operations is represented in the sequence diagram in Figure 2.

In particular, the macro operation of setting requires that:

- ***The user U sets privacy preferences for the related service providers***. U sets his privacy preferences **PP** for a single service via the user interface and dynamically generates a file where all the service providers' privacy preferences are stored. U digitally signs PP and retrieves the digest **DP**. PP and DP are part of the blockchain asset.

- ***The user U creates the pseudonym***. U sets the pseudonym **UID** within the UUID v4 format. The pseudonym is uniquely related to each user, but each user can choose to have multiple pseudonyms. The pseudonym is part of the blockchain asset.
- ***The user U creates an asymmetric key pair***. The asymmetric key pair comprising a public and a private half (**PKu**, **Ku**, respectively) are generated and kept within the car environment, i.e., a service with a reserved location, and will be used when the proper operations must be executed. The couple (UID, PKu) provides the uniqueness and proof that U is the real setter of specific privacy preferences. PKu is part of the blockchain asset.
  The user may also decide to create several pairs, i.e., to avoid direct profiling (although pseudonymised) between the public key and the pseudonym that will be registered in the blockchain (but this can lead to an attack, as in Section 6 specified).

The second macro operation requires that:

1. ***The user U sends to P the blockchain asset***. U sends the blockchain asset **BA** composed of UID, PP, PKu, DP, to P. This is not yet a transaction on the blockchain, but an asset exchange through any data exchange mechanism.
2. ***P queries HF to find other preferences bound to UID***. P checks whether a real or fraudulent attempt to upload a preference to the blockchain is underway. If a transaction with the same UID exists, jump to the next step, not otherwise.
3. ***P verifies if at least another PP among those found is signed with Ku***. P verifies DP with PKu. This single step prevents an attacker from impersonating another user, as this would require knowledge of Ku used by a certain UID for signing the PPs. Consequently, the digital signature with a private key allows everyone to verify that the privacy preference belongs to its real owner. The signature scheme also protects users who have decided to have multiple pseudonyms, since they are the real owners of the relative signature keys.
4. ***P makes a transaction on HF***. Once the previous check has passed, P aggregates UID, PP, PKu, DP in a proper blockchain asset that will be the object of the final transaction performed by P on HF.

Nextly to the the storing, the following steps might be used to retrieve and verify the assets.

1. ***A third-party TP queries the blockchain to read the privacy preferences.*** TP can simply read the blockchain without any particular requirement if the blockchain is set as searchable. Whether the blockchain is permissioned and cannot be publicly searchable, TP should authenticate themself with it. This operation would intrinsically inhibit improper and malicious readings in the blockchain
2. ***User U wants to prove to be the real creator of some privacy settings.*** U provides the pseudonym to the blockchain, the blockchain returns the privacy preferences (or multitude of) associated with that pseudonym, and the user proves their identity by unlocking the digital signature(s) of the privacy preferences with the proper keys.

6. Attack Scenario and Intrinsic Mitigations

The proposed design could be subject to different attacks; therefore, it is necessary to analyse potential attack scenarios in more depth, albeit we already anticipated the malicious situation above, along with some intrinsic countermeasures. The presented attack scenarios are not exhaustive, but serve the purpose of showing the robustness of our presented solution.

- ***An attacker attempts to create a privacy preference for another user***. An attacker would need to know the private key used to sign the previous transactions, hence, this attack would fail.

- ***An attacker tries to store a privacy preference using public information (pseudonym and public key), i.e., a replay attack***. Even if an attacker had the possibility of recovering a public pseudonym and the related public key, they would not be able to modify anything within the preference, hence, they would not be able to bypass the uniqueness check performed by the peer before the upload of the preference to the blockchain, which serves the purpose to understand if the transaction is fresh or if a replay attack is underway.
- ***An attacker creates fake privacy preferences and uploads them to the blockchain***. Such an attack could compromise the stability of the service, in the case of a DDoS, or simply add superfluous information to the blockchain. In the first case, we remark that users do not upload objects directly to the blockchain, therefore, having to go through peers, attackers could be blocked by the peers if an attack is detected; in the second case, there would not be an urgent security problem: the blockchain would simply be loaded with useless privacy preferences, with the only drawback of wasting memory space.
- ***An attacker generates a new key pair and loads a privacy preference using an existing pseudonym.*** To establish a strong security assumption, a pair of keys is associated only with a pseudonym. Hence, for a single pseudonym, thus for a single user, there is only one valid pair of keys. Therefore, an attacker would not have the possibility to propose the insertion of a privacy preference using an existing pseudonym and a different key pair to the blockchain. If a user loses their key pair (e.g., if they wipe the head unit), they will need to generate a new pseudonym and a new key pair. However, in our design, we have considered the hypothesis that a user can create more than one key pair, although we are aware that this can lead to impersonation attacks. In fact, by eliminating the hypothesis that a user is protected by a single pseudonym-private key pair, in the case of the possibility of generating multiple keys, an attacker could retrieve a pseudonym from the blockchain and link it to a new key pair. Such a choice would derive from a user's preference to take a security risk to obtain some sort of benefit in terms of profiling, in particular not to see the pseudonym profiled as permanently associated with a public key.

7. State of The Art

Several works have been presented about the use of blockchain in the automotive context, both for privacy and security reasons. The main solutions consist of providing innovative cryptographic schemes, tailor-made for the automotive context so that the desired properties can be obtained in any type of blockchain. Kumar et al. [5] introduce BDTwin, an integrated framework based on deep learning and blockchain that improves security and privacy in CT-driven V2X applications. The blockchain system specifically uses a smart contract-based enhance-Proof-of-Work (ePoW) and Zero Knowledge Proof (ZKP)-based verification mechanism to guarantee safe communication between cars, roadside units, CT-edge server, and cloud server. Smart contracts are used to automatically and non-deniably enforce rules and regulations that control how V2X entities behave. Also in the work of Huang et al. [6] a cryptographic solution is illustrated. They present a privacy-preserving personalised car insurance (PCI) scheme using a consortium blockchain to address transparency and privacy concerns in data collection. By establishing a blockchain for public verification, insurance companies (ICs) can deploy contracts that allow secure interactions with drivers through a protocol utilising partially homomorphic encryption and zero-knowledge proofs. A third-party auditor (TPA) is authorised to audit encrypted data, minimising fraud through a recursive inspection game model. Dorri et al. [7] also address the problem of user privacy in the automotive context through a blockchain. However, the architectural solution presented assumes that the vehicle is a node that is part of the blockchain, an assumption that we have discussed not wanting to implement to ensure that there is loose coupling and therefore independence. Huang et al. [6] present a privacy-preserving personalised car insurance (PCI) scheme using a consortium blockchain to address transparency and privacy concerns in data collection. By establishing a blockchain for public verification, insurance companies (ICs) can deploy contracts that allow secure interactions with drivers through a protocol utilising partially homomorphic

encryption and zero-knowledge proofs. A third-party auditor (TPA) is authorised to audit encrypted data, minimising fraud through a recursive inspection game model. These works, as anticipated, use cryptographic schemes specifically crafted. Our solution can also be distinguished by considering the problem not from a mathematical point of view but from an architectural one.

Considering the scope of regulations, notable is the work of Campanile et al. [8]. The authors of the paper propose a reference model for a blockchain-based system that manages data in the Internet of Vehicles (IoV) while ensuring compliance with the General Data Protection Regulation (GDPR). The system can track responsibility in accidents or maintenance-related issues, thereby aiding in legal processes. Although this work, similarly to ours, uses pseudonyms generated directly by the user, in this case too it is mandatory to register the vehicle as part of the blockchain.

8. Conclusion

In this work we illustrated how to guarantee the security properties of pseudonymity, public verifiability, and integrity of users' privacy preferences in the automotive context. The purpose of these properties is to offer users a robust solution for demonstrating that they have made specific privacy choices concerning a set of services, and then to verify their validity in a legal context. The latter point is crucial to demonstrate for any purpose the choices made by users, thus forcing service providers to act as requested.

We then showed how the desired properties can be achieved by considering a decentralised solution such as blockchain. Among the various blockchains, the choice fell to Hyperledger Fabric.

Even in the Hyperledger context where authentication of peers is mandatory, users can record their preferences in a pseudonymous manner. In the proposed design, the decoupling of the user from the blockchain occurs by means of a blockchain peer devoted to verifying that the information sent is provable.

The solution can be extended in every context outward the automotive domain where privacy settings should be stored in a pseudo-anonymous and publicly demonstrable way.

Acknowledgements

Gianpietro Castiglione acknowledges financial support from: PRIN 2022 MUR project FuSeCar (E53D23008220006).
Sergio Esposito acknowledges financial support from: PNRR MUR project PE0000013-FAIR.
Daniele Francesco Santamaria acknowledges financial support from: Italian PNRR 2022 SERICS Spoke 6, Task 1.2, Project "SCAI — Supply Chain Attack Avoidance'.
Giampaolo Bella acknowledges financial support from: Italian PNRR 2022 SERICS Spoke 7, BaC, Project "SCAR4SUD - SCAR's Four Security-Unravelling Dimensions", CUP C69J24000320008.
Daniele Francesco Santamaria and Giampaolo Bella acknowledge the Research Program *PIAno di inCEntivi per la Ricerca di Ateneo 2024/2026 – Linea di Intervento I "Progetti di ricerca collaborativa" - Università di Catania - Progetto "Semantic Web of EveryThing through Ontological Protocols" (SWETOP).*

Declaration on Generative AI

During the preparation of this work, the authors did not use Generative AI.

References

[1] Mozilla Foundation, It's official: Cars are terrible at privacy and security, 2023. URL: https://foundation.mozilla.org/en/privacynotincluded/articles/its-official-cars-are-the-worst-product-category-we-have-ever-reviewed-for-privacy/.

[2] G. Bella, G. Castiglione, S. Esposito, M. Raciti, S. Riccobene, Not sure your car withstands cyberwarfare, 2024. URL: https://arxiv.org/abs/2410.14320. arXiv:2410.14320.

[3] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolić, S. W. Cocco, J. Yellick, Hyperledger fabric: a distributed operating system for permissioned blockchains, in: Proceedings of the Thirteenth EuroSys Conference, EuroSys '18, Association for Computing Machinery, New York, NY, USA, 2018. URL: https://doi.org/10.1145/3190508.3190538. doi:10.1145/3190508.3190538.

[4] D. Bogatov, A. De Caro, K. Elkhiyaoui, B. Tackmann, Anonymous transactions with revocation and auditing in hyperledger fabric, in: Cryptology and Network Security: 20th International Conference, CANS 2021, Vienna, Austria, December 13-15, 2021, Proceedings, Springer-Verlag, Berlin, Heidelberg, 2021, p. 435–459. URL: https://doi.org/10.1007/978-3-030-92548-2_23. doi:10.1007/978-3-030-92548-2_23.

[5] R. Kumar, P. Kumar, R. Tripathi, G. P. Gupta, S. Garg, M. M. Hassan, Bdtwin: An integrated framework for enhancing security and privacy in cybertwin-driven automotive industrial internet of things, IEEE Internet of Things Journal 9 (2022) 17110–17119. doi:10.1109/JIOT.2021.3122021.

[6] C. Huang, W. Wang, D. Liu, R. Lu, X. Shen, Blockchain-assisted personalized car insurance with privacy preservation and fraud resistance, IEEE Transactions on Vehicular Technology 72 (2023) 3777–3792. doi:10.1109/TVT.2022.3215811.

[7] A. Dorri, M. Steger, S. S. Kanhere, R. Jurdak, Blockchain: A distributed solution to automotive security and privacy, IEEE Communications Magazine 55 (2017) 119–125. doi:10.1109/MCOM.2017.1700879.

[8] L. Campanile, M. Iacono, F. Marulli, M. Mastroianni, Designing a gdpr compliant blockchain-based iov distributed information tracking system, Inf. Process. Manage. 58 (2021). URL: https://doi.org/10.1016/j.ipm.2021.102511. doi:10.1016/j.ipm.2021.102511.