A Data Protection-Oriented System Model Enforcing Purpose Limitation for Connected Mobility

Sarah Syed-Winkler
Sebastian Pape*
Ahmad Sabouri
sarah.syed-winkler@continental.com
sebastian.pape@continental.com
ahmad.sabouri@continental.com
Continental Automotive Technologies GmbH, Software and Central Technologies
Frankfurt, Hesse, Germany

ABSTRACT

Cars are getting rapidly connected with their environment allowing all kind of mobility services based on the data from various sensors in the car. Data privacy is in many cases only ensured by legislation, i. e., the European General Data Protection Regulation (GDPR), but not technically enforced. Therefore, we present a system model for enforcing purpose limitation based on data tagging and attribute-based encryption. By encrypting sensitive data in a way only services for a certain purpose can decrypt the data, we ensure access control based on the purpose of a service. In this paper, we present and discuss our system model with the aim to improve technical enforcement of GDPR principles.

CCS CONCEPTS

• **Security and privacy** → *Human and societal aspects of security and privacy*; **Privacy protections**; **Usability in security and privacy**; • **Computer systems organization** → *Special purpose systems.*

KEYWORDS

privacy model, data protection implementation, automotive, data tagging, attribute-based encryption

ACM Reference Format:
Sarah Syed-Winkler, Sebastian Pape, and Ahmad Sabouri. 2022. A Data Protection-Oriented System Model Enforcing Purpose Limitation for Connected Mobility. In *Computer Science in Cars Symposium (CSCS '22), December 8, 2022, Ingolstadt, Germany.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3568160.3570231

1 INTRODUCTION

With the increasing connectivity of (autonomous) vehicles, the automotive industry is facing major changes. The current trend [21] of connecting vehicles with local infrastructures and cloud backends opens great potential for data-driven applications, improved user experiences, and new business models. Like mobile phones, cars may hold massive information about their drivers such as where they are driving, what speed they are driving at, and even whether they are tired. As a result, the vehicle changes from being a private space to being a part of the internet. Drivers' mobile lives are recorded and made available to various (3rd) parties. However, one major challenge is still to respect the users' privacy when providing data-driven applications. The problem is being addressed through several legislative initiatives from a legal perspective such as the European General Data Protection Regulation (GDPR) [29] or the upcoming ePrivacy Regulation [10]. In many cases key aspects such as transparency or purpose limitation are not technically enforced. If at all, many of them are only implemented within the processes of the data handler. However, this approach has several drawbacks. On the one hand, the realization is not enforced but rather implemented manually which may cause problems in terms of transparency, and consistency of the implemented guidelines. On the other hand, with an increasing number of 3rd parties and a rapidly changing environment around the Internet of Things (IoT) this approach is also error-prone and labour-intensive. This paper presents a data protection-oriented system model for connected mobility with the aim of technically enforcing privacy regulations. Since the GDPR is quite extensive, this paper cannot cover all aspects. While there are privacy principles in the GDPR, e. g., integrity and confidentiality of certain data that are technically feasible with state-of-the-art technologies, other principles are not as straightforward to realize, e. g., storage limitation. This paper has a specific focus on the technical implementation of "Purpose limitation" (Article 5(1)(b) GDPR) as well as "Data protection by design and by default" (Article 25 GDPR) with state-of-the-art privacy-preserving technologies. The contribution of this paper is a system model for connected mobility where the technical enforcement of purpose limitation is incorporated into the system design in an early design phase. To the best of our knowledge, there is no academic or industrial solution for the technical assurance of data-protection goals in vehicles, yet. Related applications of data tainting in the internet of things are focused on data flow analysis and enforcement, but not on purpose limitation (cf. Sect. 2). The remainder of this paper is structured as follows: Section 2 presents background and related work. Sections 3 and 4 describe the methodology and the underlying use case. The main Section 5 presents the proposed system model which is discussed and evaluated in Section 6. Section 7 concludes the paper.

CSCS '22, December 8, 2022, Ingolstadt, Germany

ACM ISBN 978-1-4503-9786-5/22/12…$15.00
https://doi.org/10.1145/3568160.3570231

2 BACKGROUND AND RELATED WORK

This section briefly introduces the two relevant concepts of data tagging and attribute-based encryption. In the following subsection, related work is discussed.

2.1 Data Tagging

Data tagging is a prominent concept with a multitude of use cases ranging from data classification to data leak prevention [43]. A few of the most important goals mentioned by Zamfir [43] are, for instance, to track and assure that sensitive data is not exposed to outsiders, or to meet compliance requirements for reporting who is accessing data internally and externally. The basic idea of tagging is to connect data with metadata. Metadata per definition is "data about data" - information that describes data and thus allows to interact with them to obtain the required knowledge [30]. In other words, data tagging refers to any process which amends data payloads with classification metadata in a structured manner. This way the context of a specific data type is available at any place and decisions on sensitive data handling are made based on tagged data without requiring knowledge about actual data representations. However, the method alone does not specify the technical implementation. Taint markings can be wrapped in a standardized container [36], stored adjacent to variables in memory [13], or be embedded in a file's metadata [44]. One approach for employing data tagging in privacy-centric applications is to amend privacy properties of the data, like a tag for personally identifiable information (PII). Once sensitive data is distinguished from non-sensitive data, the application can perform additional data protection measures for those particular sets of records. Supplementary procedures are, for instance, data flow tracking, anonymization, or encryption. If machine-readable policies are attached to data, in the literature the approach is referred to as sticky policies [20]. The word 'sticky' implies that the policies travel together with the data.

2.2 Attribute-Based Encryption

Attribute-based encryption (ABE) is a new type of identity-based encryption (IBE) scheme [32]. The concept of identity-based cryptosystems was first proposed by Shamir [35] with the idea of enabling public key encryption without requiring public key certificates. Instead, the user's identity is used to derive an identity-based key. In other words, if Bob wants to send a message to Alice using public key cryptography, Bob does not need to contact a trusted Key Distribution Center (KDC) to receive Alice's public key. Instead, he can use Alice's identity-based key to encrypt his message. The protocol does not require a Public Key Infrastructure (PKI) to establish a secure communication. In IBE systems identities are viewed as a string of characters. Sahai and Waters [32] introduce a new type of IBE scheme, namely attribute-based encryption, that views an identity as a set of descriptive attributes. In their scheme an entity encrypts a dataset for all users that comply with a certain set of attributes. The rules for allowing decryption of the dataset are attached to the data in form of an access policy. Access policies are boolean formulas defined on some attributes which describe the interested consumer or the encrypted data itself [37]. The encrypting party uses a public and individual encryption key whereas decrypting parties have their own private and individual decryption keys [17]. As an example, if Alice wishes to decrypt a message for multiple recipients, it is not necessary for her to encrypt the message for each individual receiver. Instead, she encrypts the message once using her public and individual encryption key and defines some attributes as access policy. As long as Bob's private decryption key conforms to the attributes listed in the access policy, he will be able to decrypt the message. Otherwise, the decryption process will fail. One major advantage of ABE is that it mathematically enforces an access control mechanism through the policy and attributes so that only decryption keys with adequate access rights can eventually decrypt. With conventional symmetric and asymmetric encryption schemes, data producers must encrypt datasets individually for each receiver, i.e., using public keys in asymmetric cryptography or shared secret keys in symmetric cryptosystems. As a result of ABE, a dataset is encrypted only once, yet be shared with multiple recipients while maintaining its confidentiality. ABE is a promising cryptographic technique that integrates data encryption with access control. However, the efficiency problem of ABE is considered a bottleneck limiting its development and application [12]. Feasibility studies [3, 4, 15] have shown the adaptation of ABE in IoT systems, but the computational overhead, in particular with underlying pairing-based cryptography [6], is excessive in practical applications, especially for devices with limited computational resources and power supply, like embedded systems in the automotive domain. Meanwhile, lightweight alternatives are investigated [12, 22, 42], promising higher efficiency. In general, ABE reduces the cost of multiple-receiver end-to-end encryption [17], as datasets only need to be encrypted once and can still be decrypted by multiple receivers. There is a wide range of literature discussing the feasibility of lightweight attribute-based encryption schemes on limited-resource devices, like smartphones [4], IoT devices [3, 15], or even smart home scenarios [37]. One paper also focuses on automotive platforms [17], where La Manna et al. test the impact of ABE schemes on a real hardware automotive platform. However, their measurements only reflect the decryption time of ABE in a vehicle as the vehicle receives an ABE encrypted software update packet from the server and thus, only performs the decryption operation. The proposed system model in this work requires the vehicle system to perform encryption operations. With ABE, the encryption method is computationally expensive, which might be a problem with resource-constrained devices [7]. For the technical details of ABE, we refer to [6].

2.3 Related Work

Research on enforcing GDPR principles technically is scarce in the automotive sector. Most research is focused on a legally compliant implementation where stakeholders are rather legally than technically restricted. Similar to our work, the few works considering technical enforcement, are restricted to individual principles of the GDPR, such as data erasure [33]. In the area of automotives, Krauß [16] describes an architectural concept for self-data protection in a connected vehicle. The paper presents an abstract data protection architecture, which aims to combine the aspects of risk assessment, user transparency and self-determined control. The architectural considerations can be used as a starting point for creating a transparent and privacy-friendly user experience in vehicles. However, there is no direct mapping to GDPR requirements. In contrast, our paper aims at the technical realization of GDPR requirements. In addition to the system model, our work reflects on implementation considerations by selecting suitable privacy-preserving technologies (PPT) that are required for a technical realization in the vehicle, which can be challenging itself [19, 28]. The concept of the Privacy Manager (cf. Sect. 5.3) in this work is inspired by the privacy policy framework from Al-Shomrani et al. [2]. The data protection system model in this work is based on formal privacy policies, which allow to convey and enforce user preferences throughout the system and beyond system boundaries. Furthermore, the policy decision point is separated from the actual technical realization of privacy enforcement. PRICON [41] is a user-centred privacy-aware control system which allows users to define self-determined privacy policies which are applied to the vehicular system. It could be connected with our system model to define the privacy policies. A study on privacy concerns and data sharing from connected cars [9] highlights the importance of control. They added drivers' feelings of possession toward their driving data to the privacy calculus and explain with it why individuals are reluctant sharing even low-sensitivity data that do not raise privacy concerns.

In the area of the internet of things, data tainting is used to allow users to control data flow patterns [14] and for information flow tracking and analysis [8], but not for purpose limitation, i. e., the aim is to deny certain parties access to data or detecting data leaks by analysing its flow. However, purpose limitation rather focuses on the reason the data is used for than allowing or denying access by a certain party. Rahulamathavan et al. [27] make use of ABE for data aggregation, a wideley used privacy pattern in the IoT architecture [23].

3 METHODOLOGY

In this section, we will briefly describe how the system model was developed and evaluated. We will first discuss the use case and the requirement elicitation and then describe the iterative process to develop the system model.

3.1 Use Case and Requirement Elicitation

Before the system model and the respective use case can be discussed, it is important to define the scope of the model which strongly depends on the considered use case. The use case was defined in collaboration with a consortium of a research project, which consists of organisations from academia and industry. The development of the data protection-oriented system model is focused on and limited to connected vehicle systems. Apart from the vehicle, entities interacting with vehicles, like backend agents or 3rd parties, must also comply with data protection. However, the privacy implications on external parties are not in the focus of our work, and thus not addressed. Furthermore, a legal evaluation which datasets in vehicles are regarded as sensitive is out of scope and presumed to be done beforehand. A security evaluation of the proposed system model is also not performed and outside of the scope of this work. The research objective focuses on technical measures for assuring the selected GDPR requirements in vehicle systems. The use case is further described in Sect. 4.

3.1.1 Assumptions. To further specify the focus of our model, we explicitly describe the following assumptions which may also need some effort to be implemented, but which are outside of the scope of this paper:

- The application is running in a secure environment with state-of-the-art security measures, i. e., integrity measures to ensure the datasets are not illegitimately modified and the secrecy of the data is ensured against external attackers.
- The data classification process to determine whether the data contains sensitive information was performed beforehand. Data classification relies on a preliminary privacy risk analysis and system modelling with legal support to define specific and business-related data classification rules, which is out of scope for this work. In this sense, we assume that it is known at this stage which data types are considered sensitive and which are not.
- The mobility services are following the privacy by default principles, i. e., a service only requests data that are necessary to provide the service functionality.

3.1.2 Requirements. The system model shall incorporate privacy-preserving technologies in connected vehicles to accomplish purpose limitation (Art. 5(1)(b), GDPR), privacy-friendly default settings, and controllability to users over their sensitive data from the early design phase (Art. 25, GDPR). The principle of purpose limitation has two components: 1) personal data shall only be collected for specified, explicit, and legitimate purposes; and 2) collected data shall not be processed in a different or incompatible manner than for the initial purpose. Technically, one possibility to achieve purpose limitation with cryptographic measures is by the means of encryption and decryption. Only functions with legitimate purposes are allowed to decrypt the personal data to guarantee purpose binding. This way the controller monitors if functions receive the decryption key, which allows access to specific data. Article 25 GDPR addresses two main aspects: 1) Privacy by Design; and 2) Privacy by Default. The first aspect puts a general obligation on data controllers to implement appropriate technical and organizational measures ensuring that principles related to the processing of personal data are met. Since purpose limitation is the only principle within the scope of this work, the implementation of appropriate measures is also limited to achieving purpose binding. The second aspect is meant to assure privacy-friendly default settings for data collection. Although Privacy by Design and by Default include both, technical and operational measures, this work focuses on the technical measures.

3.2 System Model Development and Evaluation

For the system model development, we followed an iterative approach with three cycles. Feedback was given by an expert round consisting of four post-docs and two industrial engineers. The first round consisted of the development of the Policy Decision Point (cf. Sect. 5.3.1) which defines rules for dealing with personal data in vehicles. The focus was on data tagging as a suitable candidate to achieve controllability of sensitive data in the vehicle system. The fundamentals of data tagging and the application in the automotive context were then evaluated with the help of the expert round. In the second round, the focus was on the Policy Enforcement Point (cf. Sect. 5.3.2) and the analysis which PPTs are suitable. At the end, attribute-based encryption was evaluated for automotive applicability and presented to the same expert round. In the third and final round, the combination of both technologies, which form the system model, were presented and evaluated accordingly. After each feedback round, the system model was reworked, and the feedback was incorporated into the system model design.

4 USE CASE: "CONNECTED MOBILITY"

In this section we present the considered use case. For the sake of generalizability, we chose a rather abstract use case. The user is operating a car which has access to some external service running or exchanging information in the cloud. The interaction of the user with the car is realized by a Human Machine Interface (HMI). The HMI allows the user to activate or deactivate automotive services and to permit or decline access from the services to his personal data. If the user changes settings, the Privacy Manager (PM) is responsible for implementing them by managing the privacy policy and applying and monitoring data protection measures. The PM is also responsible for reacting to changes in the privacy policy and adapting appropriate PPTs to meet pre-defined data protection requirements. The Privacy Manager is the core of our system model and will be explained in more detail in the next section.

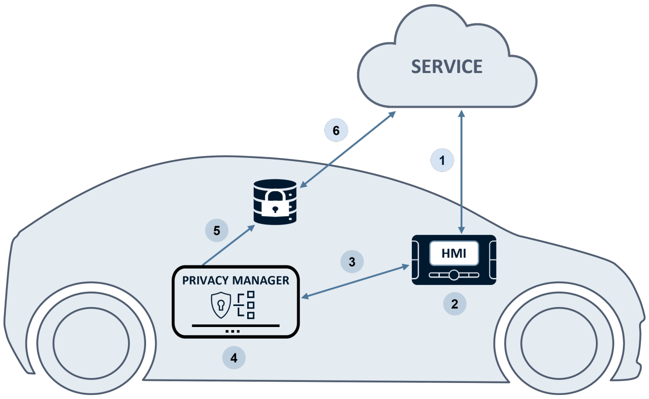

Figure 1: Use Case Overview

Figure 1 shows a high-level perspective of the proposed system model. Connected mobility refers to seamlessly connecting users, vehicles, and services over the internet [1]. For the sake of simplicity, we depict connected mobility services generically as an external cloud. The arrows in Figure 1 are labelled with numbers describing individual steps of the service usage. In the following, we are describing the steps for the typical scenario of a connected mobility service:

(1) Users want to run a connected mobility service in their vehicle, which requires access to a specific set of sensitive data. In this model, services without legitimate interest are by default denied access to personal data to ensure user-friendly privacy settings. So as a first step, the service requests access to the required sensitive data in the vehicle.
(2) The HMI offers users the possibility to interact with the vehicle system. Users enter their privacy preferences which either permits access to continue with the use of the service or declines access via the HMI display panel. If users have changed their privacy settings, e. g., from declined to allowed, the changes must be updated in the privacy policy of the vehicle system, which is managed by the PM.[1] If there was no change in the privacy settings, no further action is required.
(3) For the remainder of this use case, we assume that in step 2 the user changed the access policy for a service. Thus, the changes must be reflected in the vehicle system accordingly. The, the PM is informed about an updated privacy policy. As a reminder, the privacy policy contains information such as which services are permitted to access what type of data.
(4) Now the PM comes into play. The application provides suitable PPTs to achieve controllability over data in the vehicle system and to only permit access to sensitive data for services with legitimate interest. If the mobility service received user permission to access the personal data related to the requesting service, the PM is responsible for applying suitable PPTs to facilitate the access. When users change the policy from "enable" to "deny", the access for this service must be revoked. The PPTs ensure that only permitted services can actually read the data.
(5) After the PM applied suitable privacy protection measures, the data is stored in a database. As a note, the policy update does not affect previously stored data, but only data which is stored after the policy update. In the current model, the database is located inside the vehicle.
(6) The service accesses the database if permission is granted. If access was denied, the service is not able to read the data.

5 SYSTEM MODEL

This chapter introduces the data protection-oriented system model. The objective of the system model is the technical realization of two selected GDPR articles as specified in Section 3 and is therefore focused on implementing purpose limitation, privacy-friendly default settings, and to enable users to control their sensitive data. Therefore, the model is strongly based on the assumptions defined in Sect. 3. Several steps are required to coherently implement data protection in a system. First, rules for handling sensitive data are defined. Then, those rules are enforced in the vehicle system with suitable technologies. Figure 2 illustrates the elements of the model and their interactions. The figure depicts a component called Privacy HMI (PHMI) which allows the communication between the vehicle and the users. Further details regarding the PHMI will follow in the next section.

The numbered steps are explained in the following:

(1) Over the PHMI, the user specifies his privacy preferences.
(2) Then, the PHMI interprets the user input into a machine-readable format and updates the privacy policy accordingly.

[1] Remark: Users can change their privacy preferences independent from service requests. The HMI, as any conventional interface, offers a dedicated section to privacy settings, where users can proactively make changes to the privacy policy at any time. In this case, step 1 is skipped, and the model begins with step 2.

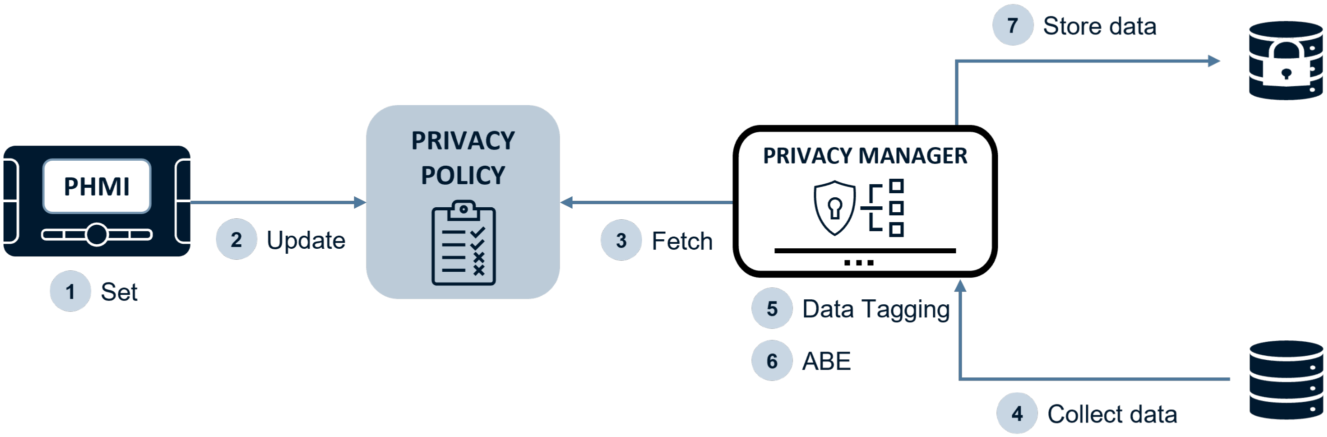

Figure 2: High-Level Design of the Model

(3) The Privacy Manager retrieves the changes from the privacy policy. From that point forward, the modified policy settings are effective. The changes will not affect previously processed datasets.
(4) The PM collects the requested data sources from a dedicated database, which is not accessible to services.
(5) During the first step of processing, the defined data tags from the privacy policy are applied to the plain data collected earlier.
(6) In this step, the privacy tag is checked. Data that contains sensitive information is encrypted using attribute-based encryption.
(7) Finally, the processed data is stored in a data storage that is accessible by mobility services.

5.1 Privacy Human Machine Interface

The Privacy HMI (see Figure 3) can be seen as an extract of a conventional HMI in the vehicle that only contains the privacy settings. The primary objective of the PHMI design is to ensure transparency to the user by conveying an honest and comprehensive representation of the system settings. Additionally, it provides explicit control over the storage, processing, and sharing of personal information. Users can enter their privacy preferences transparently, which are then reflected in the privacy policy. Afterwards, the PM can enforce any necessary changes.

5.2 Privacy Policy

Privacy Policies in our system model are referring to structured and machine-readable files containing information on all verified mobility services, data sources, and their correlations. Note that privacy policies, on the one hand, have content that is identical for all users in a vehicle, but, on the other hand, also contain preferences that differ for each vehicle user. As an example, the default settings of a privacy policy are generally identical for all users. However, differences occur when one driver allows the sharing of his sensitive data while another driver refuses to share it. If the affected vehicle is maintaining different user profiles, it is advised to also link the privacy preferences to those user profiles. Table 1 lists the structure and data fields for each data source, i. e., a sensor. The first three fields are generic information, which include a unique identifier as a reference to the data type, the name of the data type, and a concise description of the content of this data. The next three fields relate to the privacy policy of the data identifiable information (PII) is introduced. Then, the status of the data type is noted. The policy status can either be 'enabled' by the user, 'disabled' by the user, 'mandatory' due to legal reasons, or 'unspecified'. The default setting for any sensitive data is source. First, a field to identify if the dataset contains personally 'unspecified', which implies that access to the data source is not permitted (i. e., Privacy by Default). Finally, the last field lists all legitimate purposes, for which the data source is collected. The example in the last column shows how data could look like for a GPS sensor.

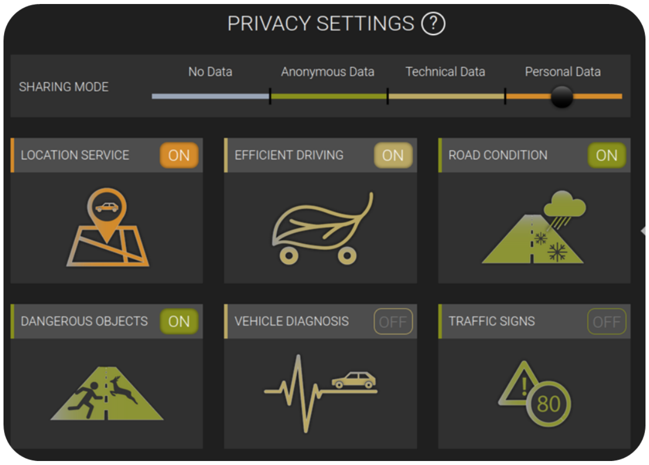
PRIVACY SETTINGS ?

SHARING MODE: No Data | Anonymous Data | Technical Data | Personal Data

LOCATION SERVICE ON
EFFICIENT DRIVING ON
ROAD CONDITION ON
DANGEROUS OBJECTS ON
VEHICLE DIAGNOSIS OFF
TRAFFIC SIGNS OFF

Figure 3: Screencapture of Privacy HMI [11]

Table 2 lists the structure and data fields for mobility services. Each service has a unique identifier, a name, and a comprehensive description. In the privacy policy, the status of each service is set to 'disabled' by default. Additional data fields are a list of purposes that are realized with this service and a list of all data types that are

Table 1: Data Source Policy Description

Field	Description	Example
Data ID	Unique identifier for the data type	DS_30
Name	Name of the data type to be displayed to the user	GPS
Description	Concise description on the information of this data type	Provides the location information at the time of measurement
PII	Determines if data type is sensitive or not	sensitive
Status	Determines if the collection of this data type is 'enabled' by the user, 'disabled' by the user, 'mandatory', or 'unspecified'	disabled
Purposes	A list of legitimate purposes for which the data type is collected	navigation, road condition, efficient driving

required for this service. Here, it is assumed that the services are following the design principles of privacy by default, meaning the services only request access to data which are necessary to provide the main functionalities. The data sources can contain sensitive as well as non-sensitive information. The example in the last column defines a location and navigation service.

Table 2: Service Policy Description

Field	Description	Example
Service ID	Unique identifier for the service MS_46	
Name	Name of the service	Location Service
Description	Comprehensive description of the service to be displayed to the user	The Location and Navigation Service traces your trips to provide reliable traffic information, alternative routes, possible points of interest, or simply navigate you to desired destinations
Status	Determines if the service is 'enabled' or 'disabled' by the user	disabled
Purposes	A list of purposes that are realized with the service	navigation
Data Sources	A list of data types that are required to realize this service	DS_03, DS_25, DS_30, DS_39

5.3 Privacy Manager

The Privacy Manager consists of multiple architectural components, which are represented in Figure 4. On the top level, the application is split into the Policy Decision Point (PDP) and the Policy Enforcement Point (PEP). The PDP makes use of the privacy policy described in the previous section and determines rules how to deal with (personal or sensitive) data. The PEP then executes the defined privacy rules by applying PPTs.

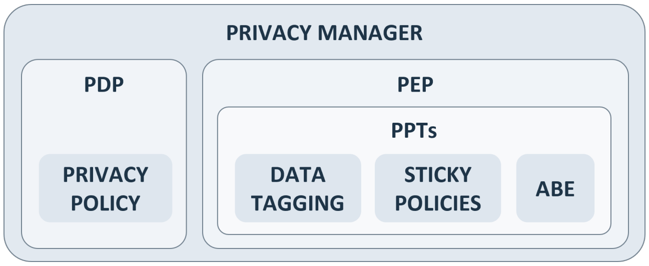

Figure 4: Architectural Components of the Privacy Manager

Purpose limitation is accomplished through sticky policies (a data tagging variant) and attribute-based encryption (ABE). Data protection by default and by design is accomplished with the help of privacy policies and data tagging. We describe the data processing which shows the interplay between privacy policies, data tagging and ABE in the next subsection. In the following subsections, we dive into the PEP, and how data tagging, and ABE is applied in our system model.

5.3.1 Data Processing. Before we describe the involved components in more detail, we start with a high-level description of the data flow. For that purpose, we describe how the dataset M is processed within the application and how the key components of the PM interact with each other as shown in Figure 5.

For a dataset M, metadata describing the dataset and in particular if M contains personal identifiable information (PII) is attached to the dataset. If M does not contain PII, it is considered to be non-sensitive data and stored in plain. If M contains PII, it is considered to be sensitive data and stored encrypted (along with its metadata). In this case, the application knows that the dataset contains sensitive information and treats it differently. For the encryption an ABE scheme is used to implement an access policy only allowing legitimate services to decrypt the dataset. The following steps refer to the numbers in Figure 5 and describe the process in more detail.

(1) The PM receives a dataset M as input.
(2) As a first step, the privacy policy for the dataset's data type is checked. As already described, the privacy policy states which data types are considered sensitive.
(3) Next, metadata is attached to the dataset. As an example, the additional data fields contain an ID as a reference to the dataset, the name of the data type and a description of the content.
(4) In addition to the generic metadata, information relating to the privacy policy of the dataset is amended (here referred to as sticky policies), e. g., a flag to identify if the dataset contains PII. Also, the current policy status is attached (cf.

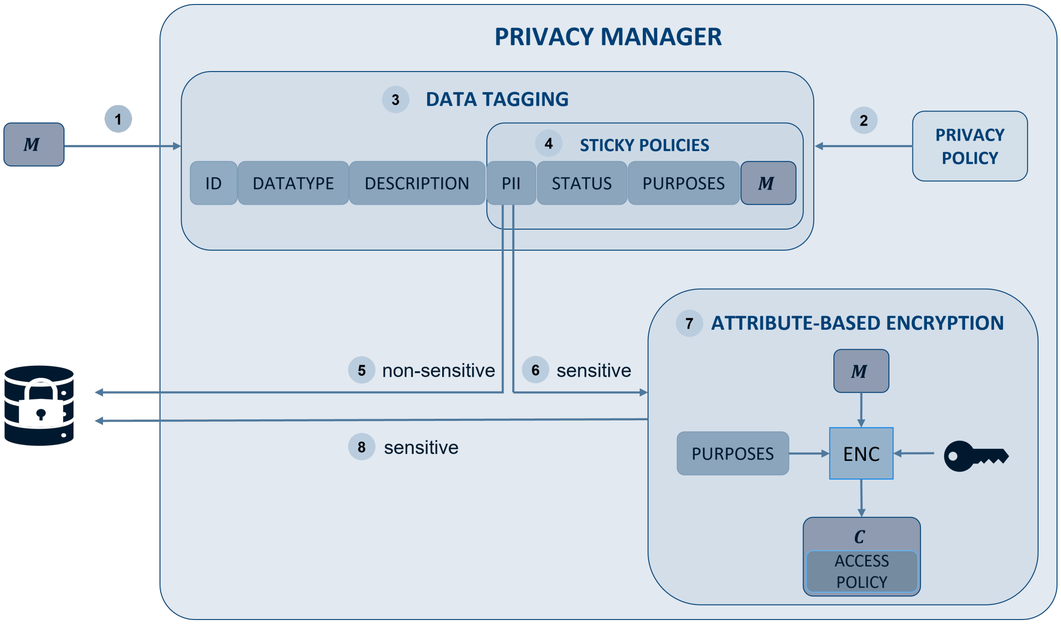

Figure 5: Data Processing within the Privacy Manager

Table 1). The data tag 'purposes' lists all legitimate purposes for which the dataset is collected. Technically, there is no difference in the implementation of sticky policies compared to any other metadata. All data tags are added to the payload the same way.

(5) The non-sensitive dataset is stored in a database without any additional privacy measures. Note that the tagged datasets are stored in the database, which are of the data structure: `(ID|DATATYPE|DESCRIPTION|PII|STATUS|PURPOSES|M)`

(6) Only sensitive data will undergo additional privacy preserving technologies. The distinction rule aims to limit the performance footprint of the approach on automotive architectures. In this model, the next step is to encrypt dataset M using attribute-based encryption.

(7) Attribute-based encryption, other than any conventional encryption scheme, adds an access policy to the encrypted data. The encryption algorithm takes the dataset M, attributes describing services that are allowed to access the dataset, and an encryption key as input. The result is a Cipher C with an embedded access policy.

(8) Instead of adding dataset M in plaintext to the payload, for sensitive data the corresponding cipher C is attached to the metadata, meaning: `(ID|DATATYPE|DESCRIPTION|PII|STATUS|PURPOSES|C)`

5.3.2 Policy Enforcement Point. In the last section, the rules for handling sensitive data and mobility services were defined. Now, these rules need to be translated and enforced in practice by a PEP. The logic of enforcing the defined privacy rules is solely realized inside the PEP. One core aspect of the system model is the different handling of sensitive and non-sensitive data. Because of the different treatment, data is tagged in a way that the sensitivity is visible without disclosing the actual information. This way, the system can distinguish between sensitive and non-sensitive data and additional measures can be enforced for sensitive information. This also increases the controllability of the data within the vehicle system. For this purpose, data tagging is used as a technique. Since the PEP is able to selectively differentiate sensitive from non-sensitive data, as a next step, the access to sensitive data must be secured for legitimate purposes only. In the privacy policy, the prerequisite for purposes to be listed under the data field Purposes has already been defined as a criterion. As a next step, the PEP needs a suitable component to technologically facilitate the criteria. We make use of attribute-based encryption for this purpose which offers the possibility to cryptographically secure access to sensitive data with purpose limitation.

5.3.3 Data Tagging. This section discusses considerations necessary for applying privacy-centric data tagging in the automotive environment and implementing it according to the described system model. Although, on a conceptual level, the processing of data according to a policy seems straightforward, on a technical level, automatically differentiating various data types may prove to be difficult. A serialization format is required to harmonize between different data types providing structured inputs for the policy enforcement. Generally, any serialization format can be used or even an own standard could be defined. However, using standardized containers is the preferred solution in this work, as they simplify the abstraction of different formats and provide a unified way to enforce sticky policies. Once a generic data container is defined, all datasets are standardized with the same data structure. In the following, we describe the sequence of data tagging within our model. First, the PM obtains the plain data from the dataset requested by a service. The member values are taken from the privacy policy for the respective data type. When the PII flag is set to "false", all further privacy-related processing is skipped. As a last step, the plain data is appended as a payload to the data structure and the object is then stored in a database. With the help of sticky policies, users can directly control how their data should be processed, handled, and shared by explicitly expressing their preferences and data handling policies [25]. To function properly, sticky policies must be compatible with generally applicable privacy policies, and vice versa. Based on the dynamic taint policies proposed by Schwartz et al. [34], the following rules are derived:

(1) Tag Introduction: Tag introduction specifies how, when, and which type of data tags are introduced into a system. This model introduces the data fields outlined in Table 1 as tags to a dataset. Only datasets required by services (i. e., data sources) are tagged, which reduces performance overhead. Each data source (e. g., velocity, location, time) has its own policy definition (cf. Example in Table 1). The Privacy Manager analyses the requested datasets, retrieves their definition from the privacy policy, and adds the data fields as tags to the plain data. Depending on the sensitivity, either additional measures are applied to the tagged dataset, or it is directly stored.

(2) Tag Propagation: Often privacy implications arise from combining different data types, for example, when data are coupled with vehicle identifiers. These datasets are referred to as aggregated data. As a general rule, the aggregation of data should be avoided if not necessary for service functionalities. In the exceptional case where aggregation is needed (e. g., GPS requires aggregation of location and time), tag propagation defines the rules for inheriting data tags, specifically if sensitive data is involved. If two data sources 'DS1' and 'DS2' are combined to create a new data source, the resulting data source 'DS3' is set to sensitive, if at least one of them is sensitive. The remaining data fields must be specified in the privacy policy.

(3) Tag Checking: Sensitive data awaits further data processing, while non-sensitive data is made available to services without any additional measures. Since the PM is able to selectively differentiate sensitive from non-sensitive data, the access to sensitive data must be secured for legitimate purposes only. As a first step, the data field PII must be checked. Furthermore, the items listed under Purposes must be respected for earmarked data processing. A technical solution to integrate purpose limitation is ABE and described the next section.

If the purpose of data tagging is to track the data flow of sensitive information inside a system, a minimal implementation such as setting a single bit for reflecting the PII status is sufficient. As an example, Taintdroid [13] provides data tracking at multiple granularities with minimal taint tags which results in a performance overhead of 32% on a CPU-bound microbenchmark. It is an extension to the Android mobile-phone platform that tracks the flow of privacy-sensitive data through third-party applications. Its primary goal is to detect when sensitive data leaves the system. TaintDroid automatically labels data from privacy-sensitive sources and when tainted data leaves the network, e. g., transmitted over the network, the data's label, the application responsible for transmitting the data, and the data's destination are logged. Such real-time feedback is intended for users to give insight into what applications are doing with sensitive data and to potentially identify misbehaviour. The same system design approach is adjustable to vehicle environments. Thus, the data flow of sensitive information becomes trackable inside the vehicle to identify leakage points and to react accordingly.

5.3.4 *Attribute-Based Encryption.*

One way to achieve purpose limitation is to control access to personal data. Any conventional encryption mechanism can be applied to satisfy this requirement. A major advantage of attribute-based encryption is that it reduces the cost of multiple-receiver end-to-end encryption [17] without compromising security, since datasets only need to be encrypted once for multiple receivers. There are two possibilities with ABE to obtain data access control: 1) CP-ABE; and 2) KP-ABE. With CP-ABE the access policy travels with the data while KP-ABE defines policies for the decryption keys. CP-ABE offers better control to the data producer on his data than KP-ABE [37]. Also, the CP-ABE approach resembles the sticky policy paradigm presented in Sect. 5.3.3. Since the scope of this work is to ensure purpose limitation for data processing and to grant user controllability by reflecting their data sharing preferences, we chose CP-ABE. Next, we explain the initial steps that are required to set up a CP-ABE environment. Afterwards, the primitives for implementing a CP-ABE scheme are explained in more detail. At the end, the approach for realizing purpose limitation with ABE and the help of previously declared data tags is described. In general, mechanisms based on ABE need a trusted authority (TA) to generate and distribute secure keys for authorized entities. In the infrastructure of the automotive industry, the OEM backend can take over the role of a trusted authority. The OEM has to be considered as a trusted party anyway since the OEM is able to control software and hardware used in the vehicle. Users can decide which OEM they want to trust by purchasing their vehicle from their prefered OEM. The OEM acts as Key Distribution Centre (KDC) and needs to setup and manage the following two functions for each vehicle system [37]:

(1) generate and distribute the key used for encryption called encryption key EK, which is unique for each vehicle system
(2) generate and assign each data consumer a decryption key DK with an embedded set of attributes γ

The EK is used by the data producer, the car, to encrypt its data and the decryption key is used by the data consumer, the service, to decrypt the encrypted data. Once the encryption and decryption keys are generated and assigned to the correct entities, the ABE mechanism can commence between data producers and data consumers without further involvement from a third party. The generation and distribution of an EK is a one-time task for each vehicle system. Thus, the required encryption key shall be generated and provisioned into the vehicle at a production facility. Each data consumer participating to the scheme is considered an authorized entity and is securely provisioned with a DK and an embedded set of attributes γ. However, if a DK is compromised, a key revocation procedure shall be executed. Here, the procedure by Sicari et al. can be used. The interested reader can refer to [37] for more details. The TA is responsible for detaining a list of all services and their attribute sets and to verify that declared attributes describe the consumer. Formally, all CP-ABE schemes are modelled by at least the following four primitives [15, 37]:

(1) Setup: $\mathsf{Setup}(\kappa) \rightarrow (MK, EK)$
This algorithm initializes the CP-ABE scheme. The setup algorithm is executed by the TA and takes a security parameter κ as input to generate a master key MK and an associated encryption key EK. The master key is kept secret by the TA.

(2) Key Generation:: $\mathsf{KeyGen}(MK, \gamma) \rightarrow DK$
In this the TA generates a decryption key DK for a data consumer. KeyGen embeds its input, the master key MK and a set of attributes γ into the decryption key DK.

(3) Encryption: $\mathsf{Encrypt}(M, T, EK) \rightarrow C$
For the encryption, one enters a plaintext M with the access policy T and the encryption key EK. The algorithm outputs the encrypted data C, which embeds the access policy T. The encryption algorithm can be executed by any component of the vehicle network. This step is computationally demanding, which makes its execution challenging on resource-constrained devices [3].

(4) Decryption: $\mathsf{Decrypt}(C, DK) \rightarrow M$
For decryption, the encrypted data C and a decryption key DK is taken as input. The output is the content of the plaintext message M if the consumer satisfies the embedded access policy T. Otherwise, decryption is unsuccessful and the algorithm outputs nothing. Mathematically, the decryption is successful only if the attribute set γ embedded in DK satisfies the access policy T embedded in C. As a remark, the access policy is a Boolean formula composed by certain attributes. If an attribute inside the access policy belongs to γ, it is considered true for the policy evaluation. The decryption algorithm is executed by a data consumer holding the appropriate decryption key DK. This step is also computationally demanding for resource-constrained devices but proven to be manageable by modern IoT devices (e. g., tablets, smartphones), as verified in [4].

As previously declared, the main motivation for applying ABE in the system model is to achieve purpose limitation. So far, the rightful purposes, for which the access to data sources is allowed, were defined as a data tag in the data source policy (cf. Table 4). Now those purposes need to be reflected in the access policy of the sensitive dataset. Figure 6 shows how data tagging and ABE interact with each other. The left box depicts the tagged dataset. The right box shows the building blocks of CP-ABE. All legitimate purposes are listed in the data tag Purposes, which are taken as the access policy for encryption with CP-ABE. As a reminder, only sensitive data is handled this way.

As for the services, all purposes that are realized through a service are outlined in the privacy policy (cf. Table 2). Those purposes must be incorporated as attributes into their individual decryption keys. It is only then that services complying with the access policy can decrypt the data. Regarding the automotive trends [24], modern vehicles are well equipped for processing vast amounts of data, unlike typical IoT devices. Therefore, it is assumed that lightweight implementations of ABE developed for resource constrained IoT devices are feasible for vehicle infrastructures as well the additional overhead for the encryption can be handled.

6 DISCUSSION AND EVALUATION

The presented system model technically enforces purpose limitation by combining privacy policies with data tagging and attribute-based encryption (ABE). To ensure purpose limitation, only services with legitimate purposes are able to decrypt personal data. An implementation with ABE is advantageous in systems with multiple recipients since the data only needs to be encrypted once. This allows the user to control the data with reasonable effort. As a result, the data controller can control which services may receive the key for decryption, which provides access to the protected data. Besides enforcing the users' preferences, this also raises ethical questions [40] since it enables the data controller to allow or deny 3rd party services in a similar manner than app stores for mobile phones are guarding the users' mobile phones with the difference that due to the encryption this central guard may not be by passed, e. g. with the help of a 3rd party app store. While this is desirable for the prevention of undesired data flows, it also has consequences for the underlying business models and might result in law suits of 3rd parties in trying to get access to the ecosystem which was already seen for app stores on mobile phones. The described model showed how privacy policies are defined for one vehicle system. If users want to have the same privacy settings for a specific dataset in another system, they would have to adjust the preferences at each of the target systems again. Going one step further, machine-readable policies can stick to data defining allowed usage and obligations when travelling across multiple parties. In other words, if data is transferred to a target system, the corresponding policy is transmitted with it. When policies are attached to datasets, the receiving party can adopt the privacy preferences in the corresponding system. This has the benefit that users only need to set their preferences once instead of managing them individually for every new system (e.g., vehicle, smartphone, smart home, etc.). However, the user must still have the possibility to change his preferences at any access point. The EnCoRe project [25] has already developed a technical solution for privacy management enabling users to improve control over their personal information, which is suitable for use in a broad range of domains. The presented system model is also prepared for future developments, i.e., due to the use of encryption to technically enforce the purpose limitation, there is no disadvantage should data be stored outside of the car, e.g., in a cloud environment. Without a correct decryption key, the data stays encrypted, and can not be used for other purposes. Depending on the implemented granularity of fine-grained access control either each service could have its own label or services are grouped within several categories. The former requires more computational overhead since the access tree must contain each permitted service ID individually. The higher the number of attributes included in the access policy, the larger the computational overhead. The latter requires the change of keys each time a certain service becomes undesired and should lose access to the data.

6.1 Limitations

The proposed system model has several limitations. First of all, the assumptions defined in Sect. 3.1.1 are varyingly difficult to address. While state of the art hardware already aims to ensure integrity and confidentiality measures (Assumption 1), the data classification

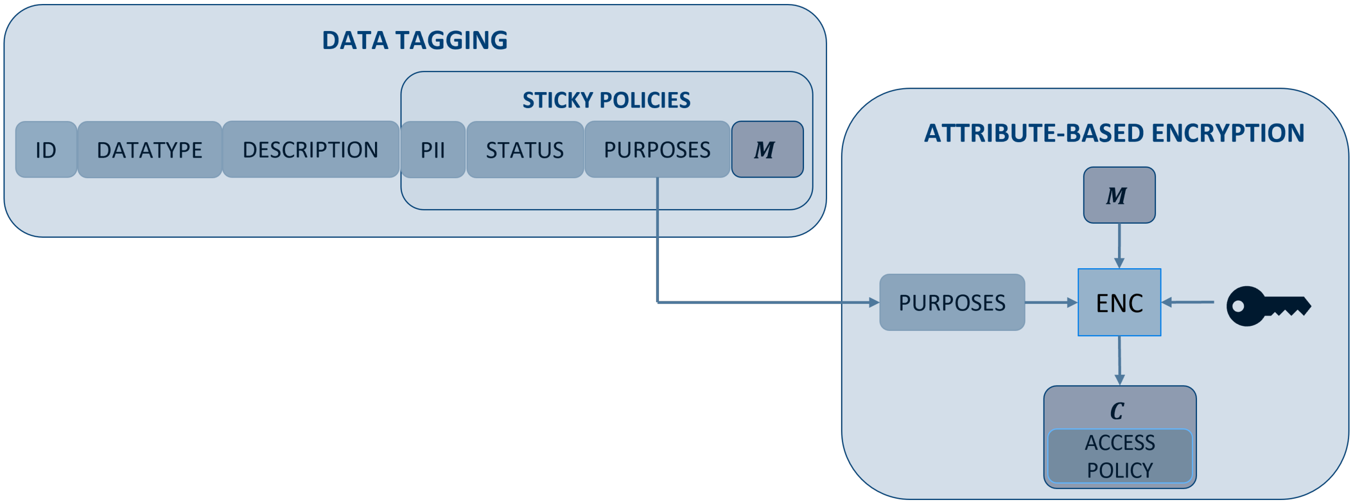

Figure 6: Purpose Limitation with ABE

process is harder (Assumption 2). On the one hand, care has to be taken that data linking doesn't allow inference attacks [24]. With inference attacks seemingly non-sensitive data could be combined to possibly allow a (re-)identification of persons resulting in sensitive data [38]. This has already been demonstrated for data extracted from logs of the CAN bus [18] or for sensors embedded in a car's seat [31]. If the (re-)identification is desired, i.e. for personalizing the car and assistance system, or a threat to the users' privacy, i.e. using the data to the driver's disadvantage by claiming drug usage after an accident, strongly depends on the use case. The problem to distinguish between sensitive and insensitive data becomes even more difficult in case the data is unstructured [39]. The compliance of the involved 3rd party services is also hard to ensure (Assumption 3). Once the data has left the system and the service is granted access to the data, because it fulfils the conditions for providing a legitimate service, the data is unencrypted and uncontrolled usage or data sharing can not be prevented anymore. As a consequence, a strict monitoring or auditing of 3rd parties would be desirable to avoid the loss of control.

6.2 Evaluation

As already discussed, the system model followed an iterative development with three cycles. The intermediate result after each round was presented to an expert group and after each feedback round, the system model was revised, and the relevant suggestions were incorporated into the system model design. The first version of the system model did not take taint propagation into account. Thus, the second version defined rules for taint propagation. After expanding the system model with ABE, the impact of key distribution using ABE in the automotive industry were not reflected. Those aspects were introduced in the final version of the system model. With the assumptions and scope described in Sect. 3.1 and the limitations discussed in the previous subsection, the expert group gave a positive evaluation of the final system model.

7 CONCLUSION AND FUTURE WORK

The data protection-oriented system model presented in this work illustrates the realization of a technical enforcement for purpose limitation as required by Art. 5(1)(b) of the GDPR. In this model, purpose limitation is accomplished through sticky policies and attribute-based encryption. Furthermore, data protection by default and by design is accomplished with the help of privacy policies and data tagging. Data flow analysis showed to be a good starting point to provide users real-time feedback on their data handling inside the vehicle. Smart cities are the future [5], where vehicles, smart devices, smart homes, and so on are all connected. Ideally, there will be one central point for identity management, where users can also set their privacy preferences (cf. Self-Sovereign Identity [26]). Once the preferences are set, they can easily be adopted to connected systems. In order to realize this, every connected system must be able to read, apply, and exchange privacy policies attached to datasets. It will be a future challenge. A further challenge is to technically enforce other privacy rights defined in the GDPR. Once it is known how to technically enforce each of them, the next challenge is to enforce them altogether as an integrated mechanism. In this work, a theoretical concept of the data protection-oriented system model was presented. Future work should set the motivation for the simulation part, which will provide a proof of concept for the presented system model. Following this, the performance of the simulation program should be evaluated since vehicles come with computational limitations. As a final step, architectural considerations that are fundamental to applying the model to an automotive environment shall be discussed.

ACKNOWLEDGMENTS

This work was supported by the German Federal Ministry of Education and Research funded project AUTOPSY (grant no. 16KIS1382). We are grateful to Margareta Hirsekorn for her support with the graphics.

REFERENCES

[1] 2022. Bosch Mobility Solutions. https://www.bosch-mobility-solutions.com/en/mobility-topics/connected-mobility/.

[2] Abdullah Al-Shomrani, Fathy Fathy, and Kamal Jambi. 2017. Policy enforcement for big data security. In *2017 2nd international conference on anti-cyber crimes (icacc)*. IEEE, 70–74.

[3] Moreno Ambrosin, Arman Anzanpour, Mauro Conti, Tooska Dargahi, Sanaz Rahimi Moosavi, Amir M Rahmani, and Pasi Liljeberg. 2016. On the feasibility of attribute-based encryption on internet of things devices. *IEEE Micro* 36, 6 (2016), 25–35.

[4] Moreno Ambrosin, Mauro Conti, and Tooska Dargahi. 2015. On the feasibility of attribute-based encryption on smartphone devices. In *Proceedings of the 2015 Workshop on IoT challenges in Mobile and Industrial Systems*. 49–54.

[5] Michael Batty, Kay W Axhausen, Fosca Giannotti, Alexei Pozdnoukhov, Armando Bazzani, Monica Wachowicz, Georgios Ouzounis, and Yuval Portugali. 2012. Smart cities of the future. *The European Physical Journal Special Topics* 214, 1 (2012), 481–518.

[6] John Bethencourt, Amit Sahai, and Brent Waters. 2007. Ciphertext-policy attribute-based encryption. In *2007 IEEE symposium on security and privacy (SP'07)*. IEEE, 321–334.

[7] Joakim Borgh. 2016. Attribute-based encryption in systems with resource constrained devices in an information centric networking context.

[8] Z Berkay Celik, Leonardo Babun, Amit Kumar Sikder, Hidayet Aksu, Gang Tan, Patrick McDaniel, and A Selcuk Uluagac. 2018. Sensitive information tracking in commodity {IoT}. In *27th USENIX Security Symposium (USENIX Security 18)*. 1687–1704.

[9] Patrick Cichy, Torsten Oliver Salge, and Rajiv Kohli. 2021. Privacy Concerns and Data Sharing in the Internet of Things: Mixed Methods Evidence from Connected Cars. *MIS Quarterly* 45, 4 (2021).

[10] European Commission. 2017. Proposal for an ePrivacy Regulation. https://digital-strategy.ec.europa.eu/en/policies/eprivacy-regulation.

[11] Continental, Cybersecurity Lab. 2021. Privacy Human Machine Interface. Internal Document.

[12] Sheng Ding, Chen Li, and Hui Li. 2018. A novel efficient pairing-free CP-ABE based on elliptic curve cryptography for IoT. *IEEE Access* 6 (2018), 27336–27345.

[13] William Enck, Peter Gilbert, Seungyeop Han, Vasant Tendulkar, Byung-Gon Chun, Landon P Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N Sheth. 2014. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. *ACM Transactions on Computer Systems (TOCS)* 32, 2 (2014), 1–29.

[14] Earlence Fernandes, Justin Paupore, Amir Rahmati, Daniel Simionato, Mauro Conti, and Atul Prakash. 2016. {FlowFence}: Practical Data Protection for Emerging {IoT} Application Frameworks. In *25th USENIX security symposium (USENIX Security 16)*. 531–548.

[15] Benedetto Girgenti, Pericle Perazzo, Carlo Vallati, Francesca Righetti, Gianluca Dini, and Giuseppe Anastasi. 2019. On the feasibility of attribute-based encryption on constrained IoT devices for smart systems. In *2019 IEEE International Conference on Smart Computing (SMARTCOMP)*. IEEE, 225–232.

[16] Christoph Krauß. 2019. Selbstdatenschutz im vernetzten Fahrzeug und dessen technische Umsetzung. In *Grundrechtsschutz im Smart Car*. Springer, 227–244.

[17] Michele La Manna, Luigi Treccozzi, Pericle Perazzo, Sergio Saponara, and Gianluca Dini. 2021. Performance evaluation of attribute-based encryption in automotive embedded platform for secure software over-the-air update. *Sensors* 21, 2 (2021), 515.

[18] Szilvia Lestyan, Gergely Acs, Gergely Biczók, and Zsolt Szalay. 2019. Extracting vehicle sensor signals from CAN logs for driver re-identification. *arXiv preprint arXiv:1902.08956* (2019).

[19] Sascha Löbner, Frédéric Tronnier, Sebastian Pape, and Kai Rannenberg. 2021. Comparison of de-identification techniques for privacy preserving data analysis in vehicular data sharing. In *Computer Science in Cars Symposium*. 1–11.

[20] Daniele Miorandi, Alessandra Rizzardi, Sabrina Sicari, and Alberto Coen-Porisini. 2019. Sticky policies: A survey. *IEEE Transactions on Knowledge and Data Engineering* 32, 12 (2019), 2481–2499.

[21] Trevor Neumann. 2021. Seven Automotive Connectivity Trends Fueling the Future. https://www.jabil.com/blog/automotive-connectivity-trends-fueling-the-future.html.

[22] Nouha Oualha and Kim Thuat Nguyen. 2016. Lightweight attribute-based encryption for the internet of things. In *2016 25th International Conference on Computer Communication and Networks (ICCCN)*. IEEE, 1–6.

[23] Sebastian Pape and Kai Rannenberg. 2019. Applying privacy patterns to the internet of things'(iot) architecture. *Mobile Networks and Applications* 24, 3 (2019), 925–933.

[24] Sebastian Pape, Jetzabel Serna-Olvera, and Welderufael B Tesfay. 2015. Why open data may threaten your privacy. In *Workshop on Privacy and Inference*.

[25] Siani Pearson and Marco Casassa-Mont. 2011. Sticky policies: An approach for managing privacy across multiple parties. *Computer* 44, 9 (2011), 60–68.

[26] Norbert Pohlmann. 2022. Self-Sovereign Identity (SSI). In *Cyber-Sicherheit*. Springer, 645–671.

[27] Yogachandran Rahulamathavan, Raphael C-W Phan, Muttukrishnan Rajarajan, Sudip Misra, and Ahmet Kondoz. 2017. Privacy-preserving blockchain based IoT ecosystem using attribute-based encryption. In *2017 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*. IEEE, 1–6.

[28] Kai Rannenberg, Sebastian Pape, Frederic Tronnier, and Sascha Löbner. 2021. *Study on the technical evaluation of de-identification procedures for personal data in the automotive sector*. Technical Report. Technical Report. Goethe University Frankfurt. https://doi.org/10.21248 ....

[29] Protection Regulation. 2018. General data protection regulation. *Intouch* 25 (2018).

[30] Jenn Riley. 2017. Understanding metadata. *Washington DC, United States: National Information Standards Organization (http://www. niso. org/publications/press/UnderstandingMetadata. pdf)* 23 (2017).

[31] Silvia Rus, Moritz Nottebaum, and Arjan Kuijper. 2021. Person Re-Identification in a Car Seat: Comparison of Cosine Similarity and Triplet Loss based approaches on Capacitive Proximity Sensing data. In *The 14th PErvasive Technologies Related to Assistive Environments Conference*. 97–104.

[32] Amit Sahai and Brent Waters. 2005. Fuzzy identity-based encryption. In *Annual international conference on the theory and applications of cryptographic techniques*. Springer, 457–473.

[33] Subhadeep Sarkar, Jean-Pierre Banatre, Louis Rilling, and Christine Morin. 2018. Towards enforcement of the EU GDPR: enabling data erasure. In *2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*. IEEE, 222–229.

[34] Edward J Schwartz, Thanassis Avgerinos, and David Brumley. 2010. All you ever wanted to know about dynamic taint analysis and forward symbolic execution (but might have been afraid to ask). In *2010 IEEE symposium on Security and privacy*. IEEE, 317–331.

[35] Adi Shamir. 1984. Identity-based cryptosystems and signature schemes. In *Workshop on the theory and application of cryptographic techniques*. Springer, 47–53.

[36] Kaiyu Shi, Xiangzhan Yu, and Yue Zhao. 2020. Fuzzing Improving Techniques Applied and Evaluated on a Network Traffic Analysis System. In *Proceedings of the 2020 International Conference on Cyberspace Innovation of Advanced Technologies*. 543–548.

[37] Sabrina Sicari, Alessandra Rizzardi, Gianluca Dini, Pericle Perazzo, Michele La Manna, and Alberto Coen-Porisini. 2021. Attribute-based encryption and sticky policies for data access control in a smart home scenario: a comparison on networked smart object middleware. *International Journal of Information Security* 20, 5 (2021), 695–713.

[38] Latanya Sweeney. 2002. k-anonymity: A model for protecting privacy. *International journal of uncertainty, fuzziness and knowledge-based systems* 10, 05 (2002), 557–570.

[39] Welderufael B Tesfay, Jetzabel M Serna, and Sebastian Pape. 2016. Challenges in Detecting Privacy Revealing Information in Unstructured Text.. In *PrivOn@ ISWC*.

[40] Frédéric Tronnier, Sebastian Pape, Sascha Löbner, and Kai Rannenberg. 2022. A Discussion on Ethical Cybersecurity Issues in Digital Service Chains. In *Cybersecurity of Digital Service Chains*. Springer, Cham, 222–256.

[41] Jonas Walter, Bettina Abendroth, Thilo Von Pape, Christian Plappert, Daniel Zelle, Christoph Krauß, G Gagzow, and Hendrik Decke. 2018. The user-centered privacy-aware control system PRICON: An interdisciplinary evaluation. In *Proceedings of the 13th International Conference on Availability, Reliability and Security*. 1–10.

[42] Xuanxia Yao, Zhi Chen, and Ye Tian. 2015. A lightweight attribute-based encryption scheme for the Internet of Things. *Future Generation Computer Systems* 49 (2015), 104–112.

[43] Cristian Zamfir. 2020. Data tracing vs. data tagging. https://www.cyberhaven.com/blog/data-tracing-vs-data-tagging/.

[44] David Zhu, Jaeyeon Jung, Dawn Song, Tadayoshi Kohno, and David Wetherall. 2011. TaintEraser: Protecting sensitive data leaks using application-level taint tracking. *ACM SIGOPS Operating Systems Review* 45, 1 (2011), 142–154.

All URLs have been last accessed on Sep 12th 2022.