$\mathcal{V}$-tokens for Conditional Pseudonymity in VANETs

Florian Schaub, Frank Kargl, Zhendong Ma, and Michael Weber
Institute of Media Informatics, Ulm University, Germany
firstname.lastname@uni-ulm.de

***Abstract*—Privacy is an important issue in future vehicle networks, because vehicles broadcast detailed information. Also of importance is accountability due to safety critical applications. Conditional pseudonymity, i.e., the usage of resolvable pseudonyms, is a common approach to address both. Often, resolvability of pseudonyms is achieved by authorities maintaining pseudonym-identity mappings. However, these mappings are privacy sensitive and require strong protection to prevent abuse or involuntary leakage. We present a new approach that does not rely on pseudonym-identity mappings to be stored by any party. Resolution information is directly embedded in pseudonyms and can only be accessed when multiple authorities cooperate for identity resolution. Our privacy-preserving pseudonym issuance protocol ensures that pseudonyms contain valid resolution information but prevents issuing authorities from creating pseudonym-identity mappings.**

I. Introduction

In future inter-vehicular networks, also known as VANETs, wireless communication between vehicles will facilitate cooperative applications enhancing road safety, traffic efficiency, and driving convenience. Collision avoidance, real-time traffic information, lane merge assistance, and warnings of approaching emergency vehicles are some of the envisioned applications. It is generally agreed that security and privacy are mandatory requirements for the deployment of VANETs. As a practical security approach, the management of vehicle IDs and authentication by digital signatures and public key certificates is proposed by research projects [1] and standardization efforts [2]. Privacy issues arise from the frequent dissemination of beacon messages that contain detailed vehicle-related information (e.g., position, speed, heading), which can be abused for tracking and profiling of individuals. However, solving it is a challenging task because privacy approaches for VANETs are constrained by special network characteristics and security requirements [3].

One solution often proposed in related work is the use of frequently changing pseudonyms [1]. Here, a pseudonym is a public key certificate for an arbitrary key pair which does not contain information linking it to a vehicle, driver, or any other pseudonym. However, in order to achieve accountability, some authorities have to be able to resolve a pseudonym to a vehicle identity in certain situations, e.g., for law enforcement after hit-and-run accidents. Thus, only conditional pseudonymity can be provided in VANETs.

Several steps are involved in the pseudonym lifecycle: During *pseudonym issuance*, a vehicle obtains a set of pseudonyms from a CA. In the process, the vehicle has to be authenticated and information for potential pseudonym-identity resolution has to be retained. *Pseudonym usage* is a vehicle's use of pseudonyms to authenticate messages when communicating with other vehicles and infrastructure nodes. If required (e.g., for law enforcement), a restricted number of authorities may perform *identity resolution* to trace a pseudonym back to an identity by utilizing some information retained during pseudonym issuance. *Revocation* is an optional step, in which a vehicle's identity or pseudonyms can be revoked to exclude it from the network. To mitigate scalability issues, a vehicle can also be passively revoked by denying the issuance of further pseudonyms.

Such pseudonym solutions provide privacy for vehicles and also fulfill the accountability requirement. But they also require vehicles to trust pseudonym issuing authorities to store and manage resolution information securely, and to only provide this information to other authorities when justified. If resolution information would leak or be openly available, the privacy protection provided by pseudonyms would be undermined.

We propose to reconsider the common assumption that authorities can be fully trusted with managing information that can render privacy protection mechanisms ineffective. Instead, authorities should follow the principles of *minimum disclosure* and *separation of concerns*, i.e. only the entities responsible for identity resolution should be able to access the relevant information while other entities, e.g., involved in pseudonym issuance, should neither store nor have access to it. This raises several questions in the VANET context: How can accountability be achieved without trusting pseudonym issuing authorities with resolution information? How can it be ensured that resolution information can only be used in legitimate situations by a limited number of authorities? And to what extend should linking information be released then?

In this work, we propose a new approach for conditional pseudonymity in VANETs that addresses these questions. Our approach achieves accountability without requiring authorities to store resolution information and prevents them from keeping it. As a benefit, drivers have to place less trust in authorities. The scheme also benefits authorities by helping them comply with privacy regulations and reducing the amount of sensitive information to be managed. Further, we enforce the cooperation of several authorities for pseudonym-identity resolution, to ensure that a certain number of authorities agree that it is justified. The resolution protocol also provides *perfect forward privacy* [3], i.e., only linking information for the current pseudonym is made available while other pseudonyms and messages of that user remain unlinkable. Before presenting and analyzing our solution, we first review related work.

II. Related Work

Privacy and pseudonymity have been discussed in many research projects like PRIME and there are resulting frame-

IEEE Wireless Communications and Networking Conference (WCNC 2010), Sydney, Australia, April 2010.
Author version. Definitive version at http://dx.doi.org/10.1109/WCNC.2010.5506126

works like Idemix[1]. However, they are focusing mainly on Internet-like scenarios that are very different from the VANET scenarios we are considering herein.Privacy protection is generally considered mandatory for successful VANET deployment. Most approaches are based on pseudonyms with identity resolution as proposed by major research projects like SeVeCom[2] or PREDRIVE-C2X[3] and standardization efforts, e.g., carried out by ETSI TC ITS WG5. We conclude that pseudonym-based solutions are considered the most practical and promising privacy protection mechanisms in VANETs and focus on them in our work.

In [3], we analyze the specifics of VANETs and what requirements this creates for privacy solutions. We also carry out a broad review of current proposals in the light of those requirements. While basic schemes work as described in the introduction, more advanced schemes try to reduce the created overhead, e.g., with self-signed certificates [4]. Self-signed certificates create a Sybil attack problem as one cannot limit the amount of pseudonyms a vehicle controls. A recent approach tries to contain this problem [5]. [6] proposes a scheme that enforces collaborative identity resolution. However, resolution authorities also have to participate in pseudonym issuance which is not desirable. Ideally, we would like to have a strict separation of concerns so that each entity in our system model has a clear task and can be implemented independently of other functionality. Furthermore, actual pseudonym resolution has a high computational complexity. The basic motivation for our approach was to enhance privacy protection even under the assumption of (partly) malicious authorities while still providing a clean and efficient solution, which we will present after our system model in the next Section.

III. System Model

Our system model is similar to the system model of SeVeCom [1]. A vehicle V is identifiable by a unique long-term identifier id_V, e.g., an identity certificate and the corresponding key pair. V is registered with an authority CA_h, its home CA, which issued id_V and manages V's virtual identity. CA_h is identified by id_{CA_h}. In practice, a regional vehicle registration authority could take on this role, thus consolidating authority over V's virtual and physical license plates.

V can obtain pseudonyms P_i from pseudonym providers PP_k. Pseudonym providers are independent from CA_h so that V can request P_i from arbitrary PP_k. Before new P_i are issued, V needs to be authenticated and it should be verified that V has not been revoked. A pseudonym P_i is a public key certificate for a key pair (PK_{P_i},SK_{P_i}) that contains no information linking P_i to V or any P_j, $j \neq i$. When communicating, V signs messages with secret key SK_{P_i} of the current pseudonym P_i. The signature and P_i are attached to the message for verification by receivers.

Only a restricted number of resolution authorities RA_l can take part in pseudonym-identity resolution. A subset of them has to cooperate in the process. RA_l should be independent from authorities involved in the issuance of a pseudonym, i.e., CA_h or PP_k.

[1]http://www.prime-project.eu/

[2]http://www.sevecom.org/

[3]http://www.pre-drive-c2x.eu/

IV. Embedding Identities in Pseudonyms

Our approach is based on the idea of embedding resolution information directly in pseudonym certificates rather than having authorities store them. id_V, id_{CA_h}, and a unique randomization factor r are encrypted with PK_{RA}, the commonly known public key of the resolution authorities. Resulting ciphertexts, we call them $\mathcal{V}$-tokens, are unlinkable. For randomized encryption schemes, like ElGamal, r is implicitly part of the encryption scheme, while r must be explicitly included for deterministic encryption schemes, like RSA.

Pseudonyms with embedded $\mathcal{V}$-tokens are issued in a two phase protocol, which ensures that $\mathcal{V}$-token content is valid but prevents issuing authorities form linking pseudonyms to vehicles (see Sec. IV-A). V uses the resulting pseudonym P_i for normal message authentication. V signs messages with SK_{P_i} and attaches P_i to a message. Receivers verify P_i and the signature. Embedding the $\mathcal{V}$-token in P_i, does not affect how P_i is used in communications.

Pseudonym-identity resolution has to be performed collaboratively by a minimum number of authorities. They need to jointly decrypt the $\mathcal{V}$-token embedded in a pseudonym to retrieve the linking information. Sec. IV-B details the resolution protocol.

A. Privacy-preserving Pseudonym Issuance

The privacy-preserving issuance protocol employs a blind signature scheme to prevent issuing authorities from learning linking information. In the *authentication phase*, a vehicle V first obtains blindly signed $\mathcal{V}$-tokens from CA_h. Subsequently, $\mathcal{V}$-tokens are used in the *acquisition phase* to obtain pseudonyms from a pseudonym provider PP. The full protocol is given in Fig. 1 and will be discussed in detail in the following.

1) Authentication phase: The authentication phase between V and CA_h results in one or more $\mathcal{V}$-tokens blindly signed by CA_h. The protocol description has been generalized to remain independent from a specific signature scheme. We only assume that a blind signature extension exists for the signing algorithm, e.g., for RSA [7] or EC-ElGamal [8]. A simplified abstract notation is used for blinding operations, i.e., $(m)^b$ indicates a message m blinded with a blinding factor b, and unblinding is represented by $((m)^b)^{b^{-1}} = m$ with b^{-1} being the corresponding unblinding factor. Actual blinding and unblinding operations depend on the employed blind signature scheme and may consist of multiple steps.

We step through the protocol in the following. In (1) V sends a $\mathcal{V}$-token request req to CA_h signed with SK_V to prove identity id_V. The structure of req depends on the chosen authentication scheme and may entail further message exchange. (2) CA_h verifies the signature $\sigma_V(req)$ with V's public key PK_V and checks internally that V has not been revoked. CA_h then returns to V the composed identifier $id = id_{CA_h} \parallel id_V$ to be included in the $\mathcal{V}$-token, the public key PK_{RA} of the resolution authorities, id_{CA_h}, expiration date exp, and a request for n commitments. The expiration date exp is set to a discrete value, i.e., last day of the week or midnight, to prevent linking based on individualized exp.

V verifies that id is correct. Then, (3) V creates n $\mathcal{V}$-tokens $\mathcal{V}_i$ by choosing a unique random r_i that is appended to id,

Authentication phase:

$$V \longrightarrow CA_h \quad : \quad (id_V, req, \sigma_V(req)) \tag{1}$$
$$V \longleftarrow CA_h \quad : \quad (id, PK_{RA}, id_{CA_h}, exp, n) \tag{2}$$
$$V \quad : \quad \mathcal{V}_i = E_{PK_{RA}}(id \parallel r_i) \tag{3}$$
$$V \quad : \quad C_i = (m_i)^{b_i} = (\mathcal{V}_i \parallel exp \parallel id_{CA_h})^{b_i} \tag{4}$$
$$V \longrightarrow CA_h \quad : \quad (C_1, \ldots, C_n) \tag{5}$$
$$V \longleftarrow CA_h \quad : \quad \mathcal{I} \tag{6}$$
$$V \longrightarrow CA_h \quad : \quad \left\{ \left(b_i^{-1}, r_i\right) \mid i \in \mathcal{I} \right\} \tag{7}$$
$$CA_h \quad : \quad (C_i)^{b_i^{-1}} = (m_i)^{b_i b_i^{-1}} = m_i \tag{8}$$
$$CA_h \quad : \quad m_i \stackrel{?}{=} (E_{PK_{RA}}(id \parallel r_i) \parallel exp \parallel id_{CA}) \tag{9}$$
$$V \longleftarrow CA_h \quad : \quad \{\sigma_{CA_h}(C_j) \mid j \notin \mathcal{I}\} \tag{10}$$
$$V \quad : \quad (\sigma_{CA_h}(C_j))^{b_j^{-1}} = \sigma_{CA_h}(m_j) = \sigma_{CA_h}(\mathcal{V}_j \parallel exp \parallel id_{CA_h}) \tag{11}$$

Acquisition phase:

$$V \stackrel{*}{\longrightarrow} PP \quad : \quad E_{PK_{PP}}(\mathcal{V}_i, exp, id_{CA_h}, \sigma_{CA_h}(\mathcal{V}_i \parallel exp \parallel id_{CA_h}), PK_{P_i}, \sigma_{P_i}(\circ)) \tag{12}$$
$$PP \quad : \quad P_i = (PK_{P_i}, \mathcal{V}_i, exp_{P_i}, id_{PP}; \sigma_{PP}(\circ)) \tag{13}$$
$$V \stackrel{*}{\longleftarrow} PP \quad : \quad P_i \tag{14}$$

Fig. 1. Pseudonym issuance protocol.

before encrypting it with PK_{RA}. exp and id_{CA_h} are appended to each $\mathcal{V}_i$. The expiration date limits the lifetime of a $\mathcal{V}$-token. id_{CA_h} indicates the issuing authority for verification purposes in the later acquisition phase. (4) V then chooses n random distinct blinding factors b_i with inverse b_i^{-1}. Each m_i is blinded, resulting in commitments $C_i = (\mathcal{V}_i)^{b_i}$. (5) V sends $C_1, \ldots, C_n$ to CA_h, and stores the corresponding b_i^{-1} and r_i.

To make sure that V is committed to the content encoded in all C_i in the sense that it cannot manipulate or change the content anymore, V has to prove probabilistically to CA_h that the encoded content contains id as provided by CA_h in (2). This is done by a commitment scheme, in which CA_h randomly asks V to reveal the content of some of the C_i. For this purpose, (6) CA_h randomly chooses $h \geq n/2$ commitments C_i and requests the corresponding b_i^{-1} and r_i. The selected indices i are organized in the indices set $\mathcal{I}$ which is send to V. (7) V sends b_i^{-1} and r_i, $i \in \mathcal{I}$, to CA_h. CA_h can now verify the content of $\mathcal{V}_i$ by first (8) unblinding the commitments $\{C_i \mid i \in \mathcal{I}\}$ with b_i^{-1} to obtain m_i. Then, (9) CA_h computes the corresponding $\mathcal{V}$-token with r_i. The result has to be compared to m_i. If all unblinded m_i are correct, the remaining $n - h$ commitments C_j ($j \notin \mathcal{I}$) are also correct except for an exponentially small probability, i.e., the probability that V managed to cheat is negligible. This is due to V not knowing which C_i will be unblinded later when it creates the commitments, and not being able to change them when CA_h selects the commitments to be opened. See [9] for a formal analysis of the security of commitment schemes. By adjusting the ratio of $h : n$, CA_h can control the cheating probability in trade-off with required overhead.

(10) CA_h signs the remaining commitments C_j with its secret key SK_{CA_h}, yielding $n-h$ blind signatures $\sigma_{CA_h}(C_j)$ which are sent to V. In the last step, (11) V unblinds each $\sigma_{CA_h}(C_j)$ by applying the corresponding b_j^{-1} ($j \notin \mathcal{I}$). This way, V obtains $n-h$ $\mathcal{V}$-tokens $\mathcal{V}_j$, each encrypted with PK_{RA} and signed by CA_h.

2) Acquisition phase: Once in possession of signed $\mathcal{V}$-tokens, V can interact with one or more pseudonym providers PP_k to obtain a pseudonym P_i for each signed $\mathcal{V}$-token $\mathcal{V}_i$. The signed $\mathcal{V}$-token is used as an anonymous authentication credential. It implicitly certifies that its owner has been authenticated successfully by CA_h, identified by id_{CA_h}. To ensure the anonymity of V when interacting with PP and to ensure unlinkability between resulting pseudonyms and V, an anonymous communication channel is required between the two parties (denoted by $\stackrel{*}{\rightarrow}$). Either V uses a previously issued pseudonym to communicate anonymously or an anonymization mechanism like onion routing [10] can be used otherwise.

The acquisition phase starts with (12) V generating a new key pair (PK_{P_i}, SK_{P_i}) as a pseudonym key pair. Here, the key generator function of the signature scheme for authentication in vehicular communication is used. V stores SK_{P_i} securely. V sends a pseudonym certification request to PP containing PK_{P_i} and a signed $\mathcal{V}$-token $\mathcal{V}_i$ (including exp and id_{CA_h}). V signs the request with SK_{P_i} to prove its ownership ($\sigma_{P_i}(\circ)$ indicates a signature over the whole message). The request is also encrypted with PP's public key PK_{PP} to protect confidentiality of the signed $\mathcal{V}$-token.

PP decrypts the request and verifies $\sigma_{P_i}(\circ)$. PP checks the validity of the presented $\mathcal{V}$-token by verifying signature $\sigma_{CA_h}(\ldots)$ with CA_h's well-known public key PK_{CA_h}, identified by id_{CA_h}. If valid, PP proceeds by checking that $\mathcal{V}_i$ has not expired and that it has not been used before (see Sec. IV-A3). If all checks succeed, (13) PP includes the plain $\mathcal{V}$-token $\mathcal{V}_i$ (without σ_{CA_h}, exp and id_{CA_h}) in a pseudonym certificate P_i for PK_{P_i}. P_i also contains an expiration date exp_{P_i} and id_{PP}. (14) PP sends P_i to V. V can now use P_i for authentication in vehicular communication.

In the above, we only show the acquisition of one pseudonym. V can repeat the acquisition phase for each $\mathcal{V}$-token $\mathcal{V}_i$ obtained in the authentication phase. As stated before, V can acquire pseudonyms from different pseudonym providers PP by engaging with multiple PPs in the acquisition phase. This can be advantageous in a region where a specific pseudonym provider is dominant, i.e., one provider issued the majority of pseudonyms used by vehicles in that region. While V may usually use pseudonyms of its preferred pseudonym provider PP_a it can obtain pseudonyms from PP_b to prevent *sticking out* when travelling through a region dominated by PP_b. In theory, this issue could be avoided by only allowing one pseudonym provider in the system. However, in practical systems it can be expected that several pseudonym providers will exist, e.g., in different countries or that this services will be provided by different entities.

3) Double spending prevention: The issuance protocol enables V to obtain pseudonyms anonymously from different pseudonym providers, but V could also present one signed $\mathcal{V}$-token $\mathcal{V}_i$ to multiple PP_k in order to obtain more pseudonyms

than the number of signed $\mathcal{V}$-tokens issued by CA_h. Pseudonyms containing the same $\mathcal{V}_i$ would be trivially linkable, but by using them at inherently different spatiotemporal positions linking could be rendered unlikely. Double spending, i.e., multiple use of tokens, is a well-known problem of electronic cash and credential systems [11].

Double spending of $\mathcal{V}$-tokens can be prevented by extending the functionality of pseudonym providers. Pseudonym providers can operate a distributed $\mathcal{V}$-token clearing house CH in which hash values of used $\mathcal{V}$-tokens are stored. When a pseudonym provider receives a pseudonym request with signed $\mathcal{V}$-token $\mathcal{V}_i$ in (13), it computes $H(\mathcal{V}_i)$ and queries CH for it. Let $H(x)$ be a collision-resistant hash function known by all PPs. $H(\mathcal{V}_i)$ is rejected if it is in CH and added to it otherwise. Optionally, exp could be stored with $H(\mathcal{V}_i)$ to enable automated deletion of expired entries. Only using hash values of $\mathcal{V}$-tokens $H(\mathcal{V}_i)$ in CH instead of actual $\mathcal{V}$-tokens $\mathcal{V}_i$ reduces storage size and ensures that CH does not contain any (encrypted) linking information. CH can be realized as a distributed hash table (DHT) to provide scalable lookups (i.e., $\mathcal{O}(log(n))$ for n nodes [12]).

B. Collaborative Identity Resolution

While identity resolution is a part of conditional pseudonymity to prevent misuse and abuse of a system, it also exposes users to potential privacy infringement. Therefore, the information required for identity resolution needs to be protected properly, so that it is only available to certain authorities in very specific situations. Separation of duties is a common principle to prevent intentional or unintentional misuse of certain information or processes. We apply separation of duties to the protection of identity resolution information. For this purpose, we distribute the ability to perform identity resolution between a number of authorities and enforce their collaboration to perform identity resolution successfully with a threshold encryption scheme.

In our system, identity resolution corresponds to the decryption of a $\mathcal{V}$-token $\mathcal{V}_i$ embedded in a pseudonym P_i to obtain id_V that links P_i to vehicle V. The secret key of the resolution authorities SK_{RA} is split among n resolution authorities, so that each holds only a share of SK_{RA}. Cooperation of a subset of k of n RAs is required to decrypt a $\mathcal{V}$-token, which has been encrypted with PK_{RA}.

For the following outline of the protocol, we assume three resolution authorities: a law enforcement agency L, a judge or juridical institution J, and a data protection agency DP. L is interested in identifying the message sender with pseudonym P_i, J decides if evidence provided by L is sufficient to justify identity resolution, and DP surveys privacy breaches. We will discuss later how the protocol can be extended for more complex scenarios. It is assumed that a common public key PK_{RA} has been published and that the corresponding secret key SK_{RA} has been divided into three shares SK_{RA}^{L}, SK_{RA}^{J}, and SK_{RA}^{DP}. We use a $(3,3)$-threshold scheme, i.e., all three shares need to be applied to successfully decrypt a $\mathcal{V}$-token $\mathcal{V}_i = E_{PK_{RA}}(id \parallel r_i)$. The use of secret sharing homomorphisms [13] and a homomorphic encryption scheme (e.g., ElGamal [14]) enable homomorphic threshold decryption that prevents SK_{RA} or its shares from being disclosed in the decryption process. Each party applies its secret share to $\mathcal{V}_i$, and only when the k-th entity applies its secret share, $E_{PK}(m)$ is decrypted.

$$
\begin{aligned}
L \longrightarrow J &: (\mathcal{V}_i, \mathfrak{E}_i) && (1)\\
J &: \mathcal{V}_i^J = D_{SK_{RA}^J}(\mathcal{V}_i) && (2)\\
L \longleftarrow J &: \left(\mathcal{V}_i^J, \sigma_J(\mathfrak{E}_i)\right) && (3)\\
L \longrightarrow DP &: \left(\mathcal{V}_i^J, \mathfrak{E}_i, \sigma_J(\mathfrak{E}_i)\right) && (4)\\
DP &: \mathcal{V}_i^{J,DP} = D_{SK_{RA}^{DP}}\left(\mathcal{V}_i^J\right) && (5)\\
L \longleftarrow DP &: \left(\mathcal{V}_i^{J,DP}\right) && (6)\\
L &: \mathcal{V}_i^{J,DP,L} = D_{SK_{RA}^{L}}\left(\mathcal{V}_i^{J,DP}\right) && (7)\\
&\quad = D_{SK_{RA}^{L}}\left(D_{SK_{RA}^{DP}}\left(D_{SK_{RA}^{J}}(\mathcal{V}_i)\right)\right)\\
&\quad = D_{SK_{RA}}\left(E_{PK_{RA}}(id)\right)\\
&\quad = id = id_{CA_h} \parallel id_V\\
L \longrightarrow CA_h &: (id) && (8)\\
L \longleftarrow CA_h &: info_V && (9)
\end{aligned}
$$

Fig. 2. Collaborative identity resolution protocol with 3 authorities.

First, (1) L extracts $\mathcal{V}_i$ from P_i and gathers evidence $\mathfrak{E}_i$. L forwards $\mathcal{V}_i$ and $\mathfrak{E}_i$ to J with a request for identity resolution. (2) J assesses $\mathfrak{E}_i$ and either supports or declines identity resolution on basis of the provided evidence. If J supports resolution, it decrypts $\mathcal{V}_i$ with partial secret SK_{RA}^J. J also signs $\mathfrak{E}_i$ to certify its approval for identity resolution. This is optional but can serve for audit purposes. (3) $\mathcal{V}_i^J$ and $\sigma_J(\mathfrak{E}_i)$ are returned to L. Note that as long as $\mathcal{V}_i$ has been *decrypted* by less than $k-1$ RAs, no information about the plaintext is revealed.

Next, (4) L forwards $\mathcal{V}_i^J$ and the evidence signed by J to DP. DP verifies $\sigma_J(\mathfrak{E}_i)$ with J's well-known public key PK_J. If the signature is valid, DP can decide to trust J's assessment of $\mathfrak{E}_i$ or perform its own assessment of the evidence. (5) If DP decides to support identity resolution, it decrypts $\mathcal{V}_i^J$ with its partial secret SK_{RA}^{DP}. (6) DP returns $\mathcal{V}_i^{J,DP}$ to L.

Now, (7) L can apply its own secret share SK_{RA}^L to $\mathcal{V}_i^{J,DP}$ yielding $\mathcal{V}_i^{J,DP,L}$. The threshold $k = 3$ is reached, thus, $\mathcal{V}_i^{J,DP,L}$ equals the decrypted plaintext identifier id. Note, that only L learns the linking information id because it applies its secret share last.

(8) Based on id, L can contact the regional CA (CA_h) responsible for the long-term identity id_V to request further information about vehicle V. CA_h looks up id_V in its database and returns information about V to L. If required, CA_h can revoke V's long-term identity to prevent V from obtaining new $\mathcal{V}$-tokens in an additional step.

L has successfully linked pseudonym P_i to vehicle V and has sufficient information to hold V accountable. The protocol provides a straightforward approach for identity resolution with enforced distribution of resolution authority. It is also extensible and flexible. For example, the order in which entities apply their secret share is irrelevant as long as the k-th entity is the one that should learn the plaintext. We used a simplified scenario with only three RAs to outline the protocol, but hierarchical secret sharing schemes exist [15] that can model multilevel hierarchies with different threshold values for different subtrees. Such a scheme can be instantiated to reflect the external and internal organizational structure of RAs and how secret shares are distributed and divided further.

Another aspect to consider is the initial computation of the key pair (PK_{RA}, SK_{RA}) and splitting of SK_{RA}, which should not rely on a trusted party. Instead, a secure multi-party computation (MPC) protocol, such as in [16], should be used that allows participating RAs to jointly compute (PK_{RA}, SK_{RA}) and individual secret shares, without revealing SK_{RA} in the process. The setup of an MPC scheme for key initialization is out of scope of this work.

V. Analysis

Our analysis focuses on the protocols' ability to resist security and privacy attacks. We have identified two general categories of potential attacks: repudiation attacks and linking attacks. In a repudiation attack V tries to cheat the issuance protocol in order to evade accountability. In a linking attack other entities aim to link pseudonyms or $\mathcal{V}$-tokens to V or each other. We assume that an adversary actively participating in the issuance or resolution protocol behaves semi-honest, i.e., the adversary aims to fulfill its attack goal but still adheres to defined protocol steps. Thus, denial of service attacks are excluded in the following. For linking attacks, we additionally assume that the adversary does not have access to sensitive key material of V. This also includes $\mathcal{V}$-tokens signed by CA_h. This assumption can be realized in practical systems by storing such data in a tamper-resistant hardware security module in the vehicle [1].

A. Repudiation Attacks

A vehicle V could try to mount a repudiation attack with the aim of evading non-repudiation. Thus, the attack goal is to prevent that correct identity information is embedded in pseudonyms in the issuance protocol (see Fig. 1).

In the authentication phase, V could try to include a wrong identifier in $\mathcal{V}_i$ in step (3). This is prevented by the commitment scheme [9], which ensures that CA_h would detect a wrong identifier with exponentially large probability in step (9). At the same time, it is not possible for CA_h to include a wrong identifier because V generates the $\mathcal{V}$-token itself.

In the acquisition phase, V could try to submit an arbitrary bitstring instead of a $\mathcal{V}$-token to PP, or a real $\mathcal{V}$-token extracted from a pseudonym of another vehicle. Both attacks would not be successful, because PP requires a valid signature by a CA, i.e., CA_h, on a $\mathcal{V}$-token to accept it. $\mathcal{V}$-tokens that have already been embedded in a pseudonym do not carry a CA signature any more and would also be detected by querying the distributed clearinghouse in (13) (see Fig. 1).

B. Linking Attacks

In a linking attack, an adversary tries to link pseudonyms or $\mathcal{V}$-tokens to their respective holder, i.e., vehicle V. Adversaries in a linking attack can either be entities actively participating in the issuance or resolution protocols or external entities, that are not involved in the protocols. Note, that linking attacks based on vehicle tracking are out of scope of this work.

An external adversary may perform a linking attack in order to infer vehicle movement patterns, which afterwards could be combined with further external information that allows inference of the vehicle identity. By definition, pseudonym certificates contain no linkable information. Encoded public keys and certificate identifiers are generated randomly. Pseudonyms can also not be linked based on $\mathcal{V}$-tokens embedded in them, due to the randomization factor r that ensures that $\mathcal{V}$-token ciphertexts are randomized and unlinkable. However, id_{PP} could facilitate linking of pseudonyms if V successively uses pseudonyms issued by one PP, in a region where most vehicles use pseudonyms issued by another PP. As discussed before, this can be thwarted by obtaining pseudonyms from multiple providers or from the PP most dominant in a specific region. Thus, vehicle V can control the success likelihood of such a linking attack by its choice of PP for a given context.

Potential linking attacks that involve protocol participants are discussed separately per protocol.

1) During pseudonym issuance: In the pseudonym issuance protocol, CA_h, PP, or both could act as adversaries. We can analyze what information each party learns during protocol execution by defining their respective knowledge sets $K(CA_h)$ and $K(PP)$. CA_h knows id_V because it maintains V's information. It learns the opened commitments, which however do not contain new information. The blind signature scheme in steps (4)-(11) prevents CA_h from learning which $\mathcal{V}$-tokens it signed. So at the end of the acquisition phase the knowledge set of CA_h is

$$K(CA_h) = \{id_{CA_h}, id_V, req, id, exp, C_1, \ldots, C_n, m_i (i \in \mathcal{I})\}.$$

PP learns the presented $\mathcal{V}$-token $\mathcal{V}_i$ and the pseudonym P_i it issues, but not id_V:

$$K(PP) = \{id_{PP}, \mathcal{V}_i, exp, id_{CA_h}, exp_{P_i}, P_i\}.$$

Further, we define the identity set $I(V) = \{id_V\}$ and the anonymity set $A(V) = \{\mathcal{V}_i, P_i\}$ for vehicle V. An adversary can only link a pseudonym to V if it knows at least one item from $I(V)$ and one from $A(V)$ after protocol execution. Thus, to prevent linking the following condition must be fulfilled:

$$K(X) \cap I(V) = \emptyset \vee K(X) \cap A(V) = \emptyset.$$

This holds true for CA_h and also for PP:

$$K(CA_h) \cap I(V) = I(V),\ K(CA_h) \cap A(V) = \emptyset$$
$$K(PP) \cap I(V) = \emptyset,\ K(PP) \cap A(V) = A(V).$$

Therefore, neither CA_h nor PP can link P_i and id_V on their own. We can further show that linking is not possible even if CA_h and PP collude. Because authentication and acquisition phase are decoupled, a shared information set between CA_h and PP would be required for linking:

$$K(CA_h) \cap K(PP) = \{id_{CA_h}, exp\}.$$

Thus, CA_h and PP could only encode linking information in id_{CA_h} and exp. Although CA_h originally specifies id_{CA_h} and exp in the authentication phase, V can ultimately verify them in step (4). V can prevent CA_h from issuing traceable $\mathcal{V}$-tokens by requiring a fixed identifier id_{CA_h} and that exp adheres to a fixed expiration scheme, e.g., noon, midnight, or end of the week. Therefore, the pseudonym issuance protocol is robust against linking attacks by any of the involved parties.

2) During identity resolution: The identity resolution protocol is flexible in terms of definition and structure of secret sharing schemes and thresholds in order to be adjusted to organizational requirements. Participants of the secret sharing scheme should be selected in a way that reduces incentives for collusion, e.g., because of inherently divergent interests. We assume that participants have been chosen in a way that results in a negligible probability of a collusion of $\geq k$ parties, for decryption threshold k.

Returning to our example from Sec. IV-B with authorities L, J, and DP and $k = 3$, it is apparent that no information about the content of $\mathcal{V}$-token $\mathcal{V}_i$ is revealed until all parties applied their secret shares and the threshold is reached. By analyzing the knowledge sets after protocol execution of each party, we see that J and DP do not gain information about V through execution of the protocol:

$$K(L) = \{P_i, \mathcal{V}_i, \mathfrak{E}_i, id_V, info_V\}\,,\ K(J) = K(DP) = \{\mathcal{V}_i, \mathfrak{E}_i\}\,.$$

Thus, J and DP can participate in the protocol without learning id. Only L learns the content of $\mathcal{V}_i$. However, L is supposed to have access to this information after execution of the identity resolution protocol. The protocol cannot prevent L from sharing id with other parties after resolution. But this is an inherent problem of any protocol in which sensitive information needs to be revealed to a party, e.g., credit card transactions.

When L and CA_h exchange information in steps (8) and (9) (see Fig. 2), P_i and $\mathcal{V}_i$ have already been linked to V, as is the purpose of the protocol. However, neither L nor CA_h gain direct information about any other P_j or $\mathcal{V}_k$ $(j, k \neq i)$ belonging to V. Therefore, perfect forward privacy [3] is achieved, i.e., the resolution of one pseudonym to an identity does not facilitate linking of other pseudonyms of that user.

What is left to analyze is if it is feasible for an entity that knows id_V and PK_{RA}, e.g., L or CA_h, to compute all possible $\mathcal{V}$-tokens for vehicle V with an exhaustive search over r. The purpose would be tracking of a single vehicle V by linking the P_i and $\mathcal{V}_i$ to V. In the case that id_V and PK_{RA} are known to the adversary, the security of the $\mathcal{V}$-token depends on the bitsize of the randomization factor r. By choosing r sufficiently large, such an attack is rendered infeasible. But larger r entail larger $\mathcal{V}$-tokens and pseudonyms and, thus, a tradeoff between security and communication costs is required. Due to space limitations, we will provide an analysis of this attack and tradeoff in future work.

VI. Conclusion

The outlined approach for conditional pseudonymity in vehicular networks does not require pseudonym-identity mappings to achieve accountability. Instead, resolution information is embedded as encrypted unlinkable $\mathcal{V}$-tokens in pseudonym certificates. As a result, the privacy of vehicles is enhanced in multiple ways. No authorities need to be trusted to protect privacy sensitive resolution information, identity resolution requires the cooperation of several authorities in order to be successful, and perfect forward privacy is provided. At the same time, authorities can still determine the identity of a pseudonym holder when necessary, but without the need to manage large amounts of critical information requiring secure storage and protection. With the $\mathcal{V}$-token approach, each vehicle carries its own resolution information thus also providing a scalability advantage.

We have also shown that the issuance and resolution protocols are resistant against repudiation and linking attacks. The security of $\mathcal{V}$-tokens can be controlled but entails a tradeoff with communication costs. In future work, we will provide an extended analysis of this tradeoff. We are also currently evaluating with simulations how the additional overhead of embedded $\mathcal{V}$-tokens in pseudonyms affects inter-vehicular communications in scenarios with varying traffic density. As a future extension, we also plan to include pseudonym revocation in our scheme.

References

[1] P. Papadimitratos, L. Buttyan, T. Holczer, E. Schoch, J. Freudiger, M. Raya, Z. Ma, F. Kargl, A. Kung, and J.-P. Hubaux, "Secure vehicular communication systems: Design and architecture," *IEEE Communications*, Nov. 2008.

[2] IEEE P1609.2 working group, "IEEE trial-use standard for wireless access in vehicular environments - security services for applications and management messages," *IEEE Std 1609.2-2006*, 2006.

[3] F. Schaub, Z. Ma, and F. Kargl, "Privacy requirements in vehicular communication systems," in *Intl. Symposium on Secure Computing (SecureCom09), IEEE PASSAT09*, Vancouver, Canada, August 2009.

[4] F. Armknecht, A. Festag, D. Westhoff, and K. Zeng, "Cross-layer privacy enhancement and non-repudiation in vehicular communication," in *4th Workshop on Mobile Ad-Hoc Networks (WMAN07)*, March 2007.

[5] L. A. Martucci, M. Kohlweiss, C. Andersson, and A. Panchenko, "Self-certified sybil-free pseudonyms," in *Proc. 1st ACM Conf. on Wireless Network Security (WISEC 2008)*, USA, March 2008, pp. 154–159.

[6] L. Fischer, A. Aijaz, C. Eckert, and D. Vogt, "Revocable anonymous authenticated inter-vehicle communication (SRAAC)," in *Embedded Security in Cars (ESCAR 2006)*, Berlin, Germany, 2006.

[7] D. Chaum, "Blind signature systems," US Patent 4759063, July 1988.

[8] D. Jena, S. K. Jena, and B. Majhi, "A novel untraceable blind signature based on elliptic curve discrete logarithm problem," *IJCSNS*, vol. 7, no. 6, pp. 269–275, 2007.

[9] I. Damgard, "Commitment schemes and zero-knowledge protocols," *LNCS*, vol. 1561, pp. 63–86, 1999.

[10] D. Goldschlag, M. Reed, and P. Syverson, "Onion routing for anonymous and private internet connections," *Comm. of the ACM*, vol. 42, no. 2, pp. 39–41, 1999.

[11] D. Chaum, A. Fiat, and M. Naor, "Untraceable electronic cash," in *Proc. on Advances in Cryptology*. New York: Springer, 1990, pp. 319–327.

[12] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," in *Proc. of Conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM 01)*. New York, NY, USA: ACM, 2001, pp. 149–160.

[13] J. C. Benaloh, "Secret sharing homomorphisms: Keeping shares of a secret secret (extended abstract)," in *CRYPTO '86*, ser. LNCS, vol. 263. Springer, August 1986, pp. 251–260.

[14] T. ElGamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," *IEEE Trans. on Inf. Theory*, vol. 31, no. 4, pp. 469–472, 1985.

[15] G. J. Simmons, "An introduction to shared secret and/or shared control schemes and their application," in *Contemporary Cryptology*. IEEE Press, 1992, pp. 441–498.

[16] R. Cramer, I. Damgård, and J. B. Nielsen, "Multiparty computation from threshold homomorphic encryption," in *EUROCRYPT*, ser. LNCS, vol. 2045. Springer, 2001, pp. 280–299.