Privacy Requirements in Vehicular Communication Systems

Florian Schaub, Zhendong Ma, and Frank Kargl
Institute of Media Informatics, Ulm University, Germany
{ *florian.schaub* | *zhendong.ma* | *frank.kargl* }*@uni-ulm.de*

***Abstract*—A primary goal of vehicular communication systems is the enhancement of traffic safety by equipping vehicles with wireless communication units to facilitate cooperative awareness. Privacy issues arise from the frequent broadcasting of real-time positioning information. Thus privacy protection becomes a key factor for enabling widespread deployment. At the same time, stakeholders demand accountability due to the safety-critical nature of many applications. Earlier works on privacy requirements for vehicular networks often discussed them as a part of security. Therefore many aspects of privacy requirements have been overlooked. In this paper, we identify a structured and comprehensive set of privacy-related requirements for vehicular communication systems, and analyze the complex inter-relations among them. Our results enable system designers to better understand privacy issues in vehicular networks and properly address privacy requirements during the system design process. We further show that our requirements set facilitates the comparison and evaluation of different privacy approaches for vehicular communication systems.**

I. Introduction

Equipping vehicles with Dedicated Short Range Communication (DSRC) radios will enable a multitude of new cooperative applications on the road. It is envisioned that in Vehicular Communication Systems (VCS), vehicles periodically exchange information related to their movements such as current positions and headings. The goals are to enhance road safety, traffic efficiency, and driver convenience. Envisioned applications in VCS include collision avoidance, real-time traffic data, lane merge assistance, and infotainment. In recent years, numerous research projects in Europe (CVIS, SeVeCom), USA (VSC, VII), and other countries have worked on all fronts to let this vision become widespread reality in the near future.

VCS research and development is now entering a second phase, in which standardization and field operational trials will play a key role to push VCS further towards deployment. At the same time, refinement of previously developed network architectures and protocols is equally important. Two major issues that have gained attention in recent years are security and privacy of VCS. Security is important in vehicle networks to safeguard safety-critical communication. Efforts like SeVeCom [1] or the IEEE 1609.2 WG [2] have addressed security and provide adequate solutions.

However, privacy remains an issue in VCS. Due to frequently broadcasting their positions, vehicles are susceptible to tracking and profiling. The potential misuse of VCS as a large-scale road surveillance tool remains a concern of the public[1]. Thus, privacy is becoming one of the key issues that could delay large-scale deployment of vehicular communication networks, despite the safety advantages. At the same time, privacy enhancing mechanisms have to be tailored for vehicular communications due to unique network characteristics, as well as liability issues in VCS. So far, privacy has often been treated as one of the many requirements but not as the main focus. Only a few recent projects like PRECIOSA[2] started focusing on privacy preserving mechanisms for VCS. Some requirements for privacy in VCS have been addressed [1], [3]–[7], but the discussion is mainly focused on security. A comprehensive set of dedicated privacy requirements for VCS is missing so far.

In this paper, we dissect the single and abstract requirement *privacy* first to distinguish its details and identify the different forces and orthogonal requirements that influence privacy in vehicular communication systems. First, a VCS system model is given (Sec. II). Following a systematic approach we derive a structured and comprehensive set of requirements for VCS, and also identify new privacy requirements not discussed before (Sec. III). Next, we show that our requirements set not only supports privacy-friendly system design but can also be utilized for comparison of existing VCS privacy approaches (Sec. IV). We evaluate our approach in a state of the art analysis by assessing which of the privacy requirements have been addressed by certain approaches. A discussion of the current level of privacy in VCS and open issues concludes the paper (Sec. V).

II. System Model

Vehicular communication systems contain many entities. Vehicles are the majority with potentially millions of participants after the technology has been deployed. They can communicate with each other or with infrastructure entities, e.g., roadside units (RSU), which provide access to service providers and applications in the backend. Many safety applications will publicly disseminate information about vehicles (e.g. their position, heading, and speed) using (geographically restricted) broadcast communication to enable cooperative awareness on the road.

Only registered vehicles should be able to take part in the vehicular network to prevent abuse by unauthorized devices, e.g., modified laptops. Thus, an authority infrastructure is

[1]The Guardian, *Big Brother is watching*, 31 March 2009, http://www.guardian.co.uk/uk/2009/mar/31/surveillance-transport-communication-box

[2]PRECIOSA project website: http://www.preciosa-project.org/

IEEE International Conference on Privacy, Security, Risk, and Trust (PASSAT 2009), Symposium on Secure Computing, Vancouver, BC, 2009
Author version. Definitive version at http://dx.doi.org/10.1109/CSE.2009.135

required to manage vehicle registration, and grant and revoke network membership. In addition, a privacy infrastructure should be in place to protect privacy of drivers and vehicles. It also has to provide mechanisms for accountability due to the safety critical nature of some applications, so that legal authorities can assign liability to individual entities in case of inapproriate or rogue behavior.

An adversary can infringe a driver's or vehicle's privacy in several ways. Eavesdropping on communication may reveal potentially sensitive information. An adversary with extended scope may be able to track vehicles based on communication messages to infer movement patterns. If tracking information can be linked with identities, high resolution profiling of individuals becomes possible. Privacy infringement is not only possible in the vehicle domain but may also result from attacks against entities of the communication, trust, or privacy infrastructure, as well as against service providers.

Privacy requirements for VCS have to take all these different stakeholders and potential points of attack into consideration.

III. Privacy-related Requirements

In VCS, privacy requirements have strong dependencies with other requirements. Therefore, in order to identify a comprehensive and correct set of privacy requirements, we first start our requirement analysis with basic system and security requirements that constrain and influence privacy. Then, we identify and derive privacy requirements that take the earlier requirements into consideration.

A. Basic system requirements

VCS have many unique characteristics. Requirements in this category mostly arise from the communication environment and have to be fulfilled in order to guarantee system functionality. Thus they are of relevance for privacy as well.

1) Real-time constraints: Many VC applications, e.g., safety applications, require that the latency of communication and message processing is kept to the minimum.

Due to the high vehicle mobility and the short communication range, the network topology of vehicular networks can change quickly. A very short communication window among vehicles is the result. Utilization of available communication time should be maximized, which means that bandwidth must be used efficiently and communication overhead must be kept as low as possible. Safety-critical communications are extremely time-sensitive. For example, a vehicle receiving a collision warning message must process it as quickly as possible to give the driver enough time to respond. Therefore, communication and processing efficiency pose strong real-time constraints on VCS.

2) Robustness: VCS should remain operational despite high mobility and frequent network topology changes, and also exhibit high resilience to disruptions of the network.

Robustness means that VCS can provide high availability of communication services during most of system run-time. Safety-critical applications require robust communication mechanisms they can rely on. This also includes routing, security, and privacy mechanisms. The system should be fault tolerant and availability should not be easily disturbed.

3) Scalability: Applications and communication mechanisms have to scale to a network with many nodes.

In the long run, it is envisioned that most vehicles will be VC enabled, implying a potentially large number of network nodes. Applications and mechanisms have to scale with the number of vehicles to achieve adequate availability and provide sufficient performance.

4) VC support: VCS have to support special communication patterns often based on broadcast communication and function with sporadic infrastructure access.

To support a multitude of safety and non-safety applications, VCS need to support a set of unique communication patterns [8], e.g., beaconing and geobroadcast. Some applications require accurate location information, which therefore is reflected in these communication patterns. Because access to infrastructure services may be limited and only available in certain areas and situations, VCS should function autonomously without relying on centralized infrastructure services.

B. Security requirements

The wireless and open nature of VCS make it a target of various security attacks. Protecting the system against attacks is crucial for road safety and normal functioning of the system. Thus, security requirements express the need for securing communications in the hostile environment of vehicular networks. On the other hand, security requirements constrain the achievable privacy.

1) Authentication: Authentication is the process of verifying the authenticity of certain claims, e.g., sender identity, certain roles or privileges, or other message properties.

Authentication facilitates the establishment of trust in received information. This is required to enable cooperation in loosely coupled but safety-critical systems such as VCS. Verifying authenticity of received information is crucial in order to utilize it for potentially precarious decision making. Through authentication of certain claims a degree of trust can be established between entities, even if they have had no previous contact. Message authentication includes sender authentication and message integrity. Message integrity ensures that the content of a message has not been altered in transit. Usually, sender authentication enables receivers to corroborate the identity of a network member. In some cases, however, the knowledge of a sender's identity may not be required, it may be sufficient to verify that the sender is a legitimate network member, e.g., by authenticating an anonymous credential issued by a trusted third party, and that the message originated from it. Specific roles or properties of a sender could also be authenticated in VCS, e.g., a vehicle may have properties that define its details (long vehicle) or describe privileges (emergency vehicle).

2) Accountability: A VCS has to provide sufficient information to hold individuals accountable and assign liability.

Accountability is an essential requirement in VCS due to the safety-critical runtime environment. For example, an accident caused by a forged warning message may have lethal consequences. In such a case, law enforcement agencies must be able to reliably identify the perpetrator in order to hold them accountable. Thus, accountability is required for assigning liability. Accountability also implies non-repudiation, i.e., a sender of a message cannot deny having sent it. Thus a receiver can prove who sent a message by authenticating the sender. If anonymous credentials are employed for authentication, the receiver may not learn the real identity of a sender but only its ephemeral credential. However, such an ephemeral credential must be resolvable to the sender's real identity to achieve accountability. Only trusted resolution authorities should be able to perform identity resolution, i.e., map an ephemeral credential to its holder, in well-defined situations.

3) Restricted credential usage: Usage of authentication credentials should be restricted in time and parallel use.

A credential is a cryptographic token that is used to achieve authentication and accountability. Free and uncontrolled usage of credentials would encourage misuse. A malicious vehicle could mount a Sybil attack [9] by obtaining a set of anonymous credentials and using them in parallel to impersonate a number of vehicles. Thus, the number of credentials that a vehicle can use in parallel should be restricted. A credential's validity period must also be limited to prevent an adversary from accumulating credentials for Sybil attacks. On the other hand, short life-times of ephemeral credentials may increase the frequency with which new credentials have to be obtained. So there is a trade-off between life-time of ephemeral credentials and the frequency of acquiring new sets of fresh credentials.

4) Credential revocation: It should be possible to exclude a vehicle from the VCS by invalidating its credentials.

Malfunctioning or misbehaving nodes should be isolated from VCS and denied further network access [3]. The biggest challenge of credential revocation in VCS is scalability. Thus, either efficient means for distribution of revocation information have to be provided, or credentials have to be valid only for a very short time to eliminate the need for revocation. Credential revocation can refer to the revocation of ephemeral or long-term identity credentials.

C. Privacy requirements

Privacy requirements are essential to achieve privacy protection in VCS. The aim is to provide an adequate level of privacy protection for users under various constraints.

1) Minimum disclosure: The amount of information that a user reveals in communication should be kept to the minimum, i.e., no more information than required for normal functionality of VC applications.

The exposure of personal information on a *need-to-know* basis is principal in most privacy protection methods. As the exchange of information is fundamental to realize the cooperation among nodes in VCS, a user is required to reveal some information about itself, but should only do so in a controlled way that keeps the disclosure of personal information to a minimum. There are several aspects of minimum disclosure: (1) the disclosure of information should be adaptive to specific application requirements; (2) disclosed information should be as coarse as possible and as fine-grained as necessary; (3) conflicting requirements on information granularity by multiple applications have to be resolved on the user side; (4) from a user-centric perspective minimum disclosure also means that the exposure of information to any authorities should be kept minimal.

2) Anonymity: A sender of a message should be anonymous within a set of potential senders, the *anonymity set* of the message. But due to the requirement of accountability, only conditional anonymity is possible in VCS.

To achieve privacy in VCS it is important to provide anonymity for senders. If messages are not sent anonymously, message content can be trivially linked to distinct vehicles. Sender anonymity requires that it should not be possible to link a message to a sender based on message content. An anonymity set of a message m contains all vehicles that are equally likely to have sent m. It has to be large enough to generate sufficient uncertainty for any adversary that tries to determine m's origin. At the same time, identity resolution has to be supported, i.e., authorities have to be able to link an anonymous credential to an individual in order to fulfill the accountability requirement.

3) Unlinkability: Unlinkability requires that the relations between two or more items of interest cannot be linked.

Among the various items of interest, identifiable persons are the primary ones. Only when it can be linked to an identifiable person, the information in VCS becomes privacy-relevant. With this line of thought, we can derive that unlinkability of items such as a credential, a message, or a vehicle, etc. to an identifiable person is the primary requirement of unlinkability. Unlinkability is also a requirement on the transitive relations of items. For example, if an attacker can link a message to a particular vehicle, and the vehicle can be linked to a particular person, the attacker can link the message to that person. Depending on the items of interest in the unlinkability relation, unlinkability can also be used to describe other privacy concepts. For example, unlinkability of a sender to the messages she sent is equivalent to sender anonymity. Unlinkability of a message to its originator is equivalent to untraceability. Furthermore, unlinkability of consecutive messages from the same vehicle is equivalent to the ability to avoid tracking.

4) Distributed resolution authority: The capability of identity resolution should be distributed between authorities so that cooperation of a number of distinct authorities is required to link an anonymous credential to an individual.

Distribution of resolution authority is essential to providing conditional anonymity while still maintaining a high

level of user privacy. By ensuring that the power of identity resolution does not reside with a single party, its impact on privacy can be limited. If the resolution authority is split between several entities, no single authority can abuse or misuse resolution information, thus a *multi-eye principle* [10] is enforced. This also holds true in case an authority gets hijacked or corrupted. To prevent collusion, the entities that have to cooperate for identity resolution should have different intrinsic interests. Distribution of resolution authority and cooperation of different entities can be achieved with regulations and policies. But enforcing it technically is desirable to ensure adherence to privacy policies.

5) Perfect forward privacy: Resolution of one credential to an identity should not reveal any information that decreases unlinkability of other credentials of the same user.

Perfect forward privacy is inspired by *perfect forward secrecy* which states that compromised long-term keys should not compromise previously used session keys.Adapted to VCS, this means that identity resolution of one anonymous credential only enables linking of messages sent under this credential, but no information is provided about other credentials held by the same user. Perfect forward privacy may seem similar to the requirement of unlinkability but they have very different notions. Unlinkability is an inherent requirement to ensure that anonymous credentials cannot be trivially linked to each other or an individual. However, in case of identity resolution additional information may be available that breaks unlinkability. Perfect forward privacy ensures that even then different credentials of the same user remain unlinkable. Thus, minimal disclosure in face of identity resolution is provided by preventing linking of items not relevant to the current identity resolution process. If other messages are suspected to be relevant, they can also undergo the formal resolution process, but previously sent benign messages remain private.

D. Inter-relations of requirements

It follows from the above definition and discussion of requirements that some requirements exhibit strong inter-relations. Figure 1 illustrates these inter-relations on two levels: on a general level among the three categories and on a detailed level among individual requirements. It can be observed that constraining relations (red-solid) only occur between requirements of different categories, while supporting relations (green-dotted) only occur between requirements of the same category. Intuitively, this makes sense because requirements of the same category share a specific goal: basic system requirements enable VCS, security requirements secure VCS, and privacy requirements enhance privacy of VCS users. At the same time, these goals conflict with each other. One of the biggest challenges in the design and development of VCS is the incorporation of all these requirements w.r.t. their inter-relations. We facilitate this by discussing the significant inter-relations in detail.

The basic system requirements reflect inherent characteristics of VCS and have to be fulfilled by any system in order

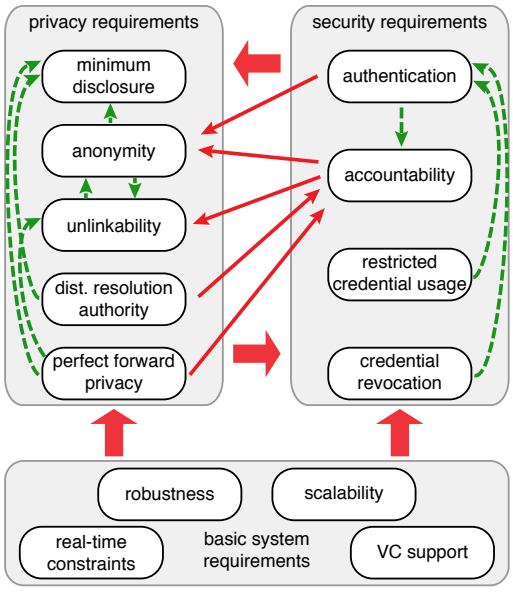

Figure 1. Constraining (red-solid) and supporting (green-dotted) inter-relations of basic system, security, and privacy requirements.

to function properly. They influence all other requirements indirectly and place constraints on potential mechanisms. For example, the real-time requirement constrains mechanisms for anonymity because multi-hop anonymization mechanisms, e.g., onion routing, are not feasable.

Security requirements place strong constraints on some privacy requirements. Accountability and authentication limit the level of anonymity that can be provided by requiring mechanisms for identity resolution, so that only conditional anonymity can be achieved in VCS. In the same way, accountability constrains unlinkability. On the other hand, the different security requirements support each other. Authentication enables accountability. Restricted credential usage and credential revocation are requirements derived from authentication and accountability, in the sense that they strengthen authentication by restricting how authentication credentials can be used and by whom.

The main privacy requirements anonymity, unlinkability, and minimum disclosure all support each other in some way. Due to the constraints placed on privacy by security requirements additional privacy requirements have been derived to ameliorate the effect of these constraints. Distributed resolution authority and perfect forward privacy do not prevent accountability but they constrain its extend to an appropriate level. Both support minimum disclosure by ensuring that identity resolution reveals no more information than required for accountability. Perfect forward privacy also supports unlinkability by restricting the extend of linking information that can be gained from identity resolution.

We conclude that adequate privacy in VCS can be achieved even when accountability is required if its impact is restricted through the distribution resolution authority and implementation of perfect forward privacy. However, meet-

ing these conflicting requirements at once is a challenging endeavor.

IV. State of the Art Analysis

Some approaches have been proposed that claim to effectively provide privacy protection in VCS. However, privacy requirements are often only implicitly stated. The explicit set of privacy requirements identified above allows us to assess the actual level of privacy protection achieved by an approach. VCS privacy approaches can be coarsely divided into five general categories. In the following, we select representative approaches from these categories and discuss how they fulfill the requirements. Due to the page limit, we can only give a short overview of the approaches and refer interested readers to the original papers for further details.

A. Basic pseudonym approaches

In the context of VCS, pseudonyms refer to pseudonymous public key certificates. These certificates do not contain any identifiable information and cannot be used to link to a particular user or to another pseudonymous certificate. Vehicles are equipped with pseudonyms and their corresponding secret keys. When sending a message, a vehicle signs it with its secret key and attaches the signature and the pseudonym certificate to the message so that receivers can verify the signature. Vehicles also have to change pseudonyms often to make it hard for an attacker to link different messages from the same sender.

The *SeVeCom approach* [1] employs a hierarchical CA structure, in which CAs manage and issue long-term identities to vehicles. Pseudonyms are issued by pseudonym providers (PP) and are only valid for a short period of time. When issuing pseudonyms, a PP authenticates a vehicle by its long-term identity and keeps the pseudonyms-to-identity mapping in case of liability investigation. The secret keys of the pseudonyms are stored and managed by a Hardware Security Module (HSM), which is tamper-resistant to restrict the parallel usage of pseudonyms. Provided with a pseudonym, pseudonym resolution authorities can resolve an identity by accessing the pseudonyms-to-identity mappings at a PP. Due to the short lifetime of pseudonyms, the need for credential revocation is minimized. Basically, only a vehicle's long-term identity is revoked to prevent it from acquiring new pseudonyms from a PP. Consequently, CAs only need to distribute certificate revocation lists (CRLs) to PPs, which are part of the infrastructure network.

B. Extended pseudonym approaches

Approaches in this category aim to either improve or enhance specific aspects of the basic pseudonym approaches.

The *PKI+ approach* [11] is based on bilinear mappings on elliptic curves. A user obtains a master key and certificate from a CA after it proves its identity and knowledge of a user secret x to the CA. The user can then self-generate pseudonyms by computing a public key from the master certificate, the secret x, and a random value. A certificate is computed as a signature of knowledge proof s over the public key and the master public key. The certificate also includes the version number Ver of the CA public key for revocation purposes. The user signs a message m by computing the signature of knowledge proof m_s on m. A receiver of m can verify the message with the public key in the pseudonym. When revoking a user, the CA publishes a new version information Ver', which has to be used by all users to update their keys. Ver' is chosen so that it is incompatible with the master key and master certificate of the revoked user. As an advantage, vehicles do not need to contact a CA or pseudonym provider to obtain new pseudonyms. Drawbacks are that Sybil attacks based on unlinkable pseudonyms are hard to detect and that the CA has no means to control the amount of self-generated pseudonyms. The proposed revocation mechanism also does not scale well with a large user base.

The *blind signature approach* [10] uses blind signatures and secret sharing in the pseudonym issuance protocol to enforce distributed pseudonym resolution. In the pseudonym issuance process, a user blinds the public key to be signed and presents shares of it to a number of CAs. Each CA holds a partial secret of the secret key shared by all CAs in a secret sharing scheme. Each CA signs the presented blinded key part with its partial secret key, returns it to the user, and stores a corresponding partial resolution tag in its database. The user can unblind and combine the received results, yielding a certificate which can be verified with a public key common to all CAs. The certificate is only valid if k of n CAs participated in the issuance process, because otherwise the threshold of the secret sharing scheme is not reached, thus resulting in an incomplete signature. To resolve a pseudonym, more than t CAs have to cooperate in a second secret sharing scheme to compute a joint resolution tag for the presented pseudonym and compare it to all tags in the database. The scheme effectively prevents misuse of resolution authority, but incurs considerable overhead by requiring a number of authorities to take part in the certification of a single pseudonym. Furthermore, pseudonym resolution requires comparisons with all tags stored in the revocation database, and therefore, does not scale well with the number of users.

C. Symmetric key approaches

Symmetric cryptography schemes require less computational effort than asymmetric operations, thus they are more efficient for time-critical applications. However, symmetric encryption has to somehow emulate asymmetric properties to achieve authentication.

The *TESLA approach* [12] is based on the TESLA broadcast authentication protocol. TESLA utilizes time to create asymmetric properties similar to public key cryptography, assuming that network nodes are loosely synchronized. A user computes a key chain and releases keys subsequently in fixed time intervals. Each message is authenticated with a key that has not been released yet, and receivers have

to buffer messages until the corresponding key is released and the message can be verified. The authenticity of a message can be verified with any key higher up in the chain. The advantage is that TESLA keys are much shorter in length than public keys and are thus more efficient. To enhance trust, each vehicle also has a set of pseudonyms signed by a CA. Pseudonyms are only used to sign anchors of the key chains. When two vehicles enter each others reception range, they first exchange certificates to obtain each others TESLA anchors. Subsequently, they only use symmetric TESLA keys to authenticate messages. Keys belonging to the same key chain as the presented anchor can be traced back to it and thus verified. The TESLA approach provides efficient authentication while reducing certificate exchanges to a minimum. However, the delay in authentication can create problems for time-critical safety applications. Additionally, TESLA keys are unsuitable for multi-hop forwarding because the keys expire too quickly and actual receivers might not receive disclosed keys.

D. Group signature approaches

Group signatures are a signature scheme to provide conditional anonymity to members of a group. Each group member can create signatures which can be verified with a common group public key. Only the group manager is able to determine the identity of a signer, because it assigns either individual secret keys or membership certificates to the group members.

The *hybrid approach* [13] uses group signatures to reduce the overhead of key and pseudonym management. Vehicles are members of a group and possess individual secret keys. Each vehicle generates random public/secret key pairs to be used for pseudonymous communications. The public keys are signed with the group secret key, yielding a pseudonym certificate that can be verified with the group public key. When communicating, vehicles sign the outgoing messages with the secret key of the pseudonym and attach the pseudonym to the message. A receiver of such a message can verify with the group public key that the pseudonym was created by a legitimate group member. The group manager, however, is able to open group signatures and retrieve the signer's identity, if necessary. The scheme obviates the need to acquire new pseudonyms periodically, but revocation of group membership is a scalability issue nevertheless.

The *GSIS approach* [14] is based on short group signatures. In the approach, a CA acts as the group manager. The CA computes a group public key and group secret keys for each vehicle in the group from their unique identifiers. With the identifier and a part of the secret key, a CA is able to determine the identity of a group member. Thus accountability can be achieved while at the same time impersonation attacks are prevented. A vehicle signs messages with its own secret key and receivers can verify them with the group public key. Revocation is achieved by distributing revocation lists. One difference to other schemes is that revocation lists are only allowed to grow to a threshold t to avoid increasing verification times. When t vehicles have been revoked, the group key and individual secret keys are updated. However, the issue of distributing revocation information persists.

Table I
REQUIREMENTS FULFILLMENT OF VCS PRIVACY APPROACHES

Approach	real-time constr.	robustness	scalability	VC support	authentication	accountability	rest. cred. usage	cred. revocation	min. disclosure	anonymity	unlinkability	dist. res. auth.	perf. forw. priv.
SeVeCom	●	●	◐	●	●	●	●	●	◐	●	●	○	○
PKI+	◐	●	◐	●	●	●	○	●	◐	●	●	○	◐
Blind sig.	◐	●	◐	●	●	●	○	●	●	●	●	●	●
TESLA	◐	◐	◐	◐	●	●	○	◐	◐	●	●	○	○
Hybrid	●	●	◐	●	●	●	○	◐	◐	●	●	○	◐
GSIS	◐	●	◐	●	●	●	○	●	◐	●	●	○	◐
ECPP	◐	●	◐	●	●	●	○	●	◐	●	●	○	◐

not fulfilled (○), partially fulfilled (◐), fulfilled (●)

E. IBC approaches

Identity-based cryptography (IBC) derives public keys from the identity of a user. Presented with a signature, a verifier can check its validity merely by knowing the sender's identity.

The *ECPP approach* [15] utilizes both IBC and group signatures. A trusted authority TA sets up an IBC scheme and publishes its system parameters. Vehicles have a unique identity, which is used to authenticate with the TA to obtain a pseudonym. TA generates a pseudonym, or pseudo-identifier, by encrypting the vehicle identifier with its public key and extracting a corresponding private key from it. The vehicle can use the resulting key pair as a pseudonym in authentication processes with RSUs under control of TA. These transactions are anonymous but linkable, because only one pseudonym is issued per vehicle. When a vehicle enters the vicinity of a RSU, it requests a short-time anonymous key certificate to take part in a local group signature scheme. The group identifier is thereby also used as the group public key. The RSU checks that the presented pseudonym is not listed on a CRL, and issues a group membership certificate, valid only for a short period of time. The RSU further retains a mapping between group membership certificate and pseudonym. Afterwards, the vehicle can perform group signatures on messages by proving possession of a membership certificate, and therefore communicate anonymously with other vehicles. Identity resolution is realized by the TA opening the group signature of a message, and retrieving the identifier of the RSU that issued the group membership certificate. The RSU can then be contacted and returns the pseudonym corresponding to the presented membership certificate. In the last step, the TA decrypts the pseudonym with the symmetric key, used for encryption in the beginning, yielding the real vehicle identifier.

F. Requirement fulfillment

Although not exhaustive, the selected approaches are representative in their respective categories, and thus provide

a good overview on the trend and extent of privacy enhanced designs in VCS. Table I summarizes these approaches w.r.t. all requirements we identified in Section III. Thanks to the structured and comprehensive requirements set, we are able to evaluate and compare the proposed privacy approaches in a systematic way. We define three outcomes for each requirement fulfillment: a (○) denotes the requirement is not fulfilled, a (◐) denotes the requirement is partially fulfilled to some extent, and a (●) denotes the requirement is fulfilled. It should be noted that some of the requirements are not explicitly addressed in the approaches. Therefore, in our analysis we map some aspects of the approaches to our requirements.

From the table we can see that except the SeVeCom approach, most approaches do not fulfill the requirement on restricted credential usage. However, this can be solved by storing key materials in a tamper-resistant device in the vehicles. It becomes apparent that current solutions mostly protect privacy against other vehicles but privacy protection against authorities is inadequate or even not considered by many approaches. Especially distributed resolution authority and perfect forward privacy are not enforced by most approaches. Although some solutions show promising approaches towards integrating security and privacy requirements in their system, the balance is often tipped in favor of security. The reason why security requirements are better fulfilled is probably the fact that security in the context of VCS has been extensively studied and is well understood. Privacy on the other hand has been recognized as important but not been treated accordingly. Future research is required to bridge this gap and provide holistic approaches to fully address all identified privacy requirements.

V. Conclusion

In this paper, we derive a comprehensive set of privacy-related requirements for vehicular communication systems. Requirements are structured in three categories: basic system requirements addressing characteristics of the VCS environment, security requirements enabling secure communications, and privacy requirements preserving privacy of vehicles and their drivers. We identified the constraining and supporting inter-relations among requirements and discussed them in detail. Understanding these inter-relations will help system designers to both address security and privacy requirements at design and development time. Based on the identified requirements, we analyzed existing VCS privacy approaches. Our analysis reveals that the issue of privacy protection against authorities, e.g., accountability with sufficient user privacy, has not been considered in most approaches, and has to be addressed by future systems.

Such issues are addressed by the PRECIOSA project which aims at developing a privacy-friendly architecture and design process for intelligent transportation systems (ITS). In this context, we are currently designing a pseudonym-based privacy approach for VCS that achieves conditional pseudonymity while addressing all privacy requirements presented in this paper.

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 communications: Design and architecture," *IEEE Communications Magazine*, vol. 46, no. 11, pp. 100–109, Nov. 2008.

[2] IEEE P1609.2 WG, "Wave security services for applications and management messages," *IEEE Std 1609.2-2006*, 2006.

[3] F. Kargl, Z. Ma, and E. Schoch, "Security engineering for VANETs," in *ESCAR*, 2006.

[4] P. Papadimitratos, V. Gligor, and J.-P. Hubaux, "Securing vehicular communications: Assumptions, requirements, and principles," in *ESCAR*, 2006.

[5] M. Raya and J.-P. Hubaux, "Securing vehicular ad hoc networks," *Journal of Computer Security (JCS), special issue on Security of Ad Hoc and Sensor Networks*, vol. 15, no. 1, pp. 39–68, January 2007.

[6] F. Dötzer, "Privacy issues in vehicular ad hoc networks," in *Privacy Enhancing Technologies*, vol. 3856. Springer, 2006, pp. 197–209.

[7] A. Aijaz, B. Bochow, F. Dötzer, A. Festag, M. Gerlach, R. Kroh, and T. Leinmüller, "Attacks on inter-vehicle communication systems - an analysis," in *3rd Intl. Workshop on Intelligent Transportation*, 2006.

[8] E. Schoch, F. Kargl, T. Leinmüller, and M. Weber, "Communication patterns in vanets," *IEEE Communications Magazine*, vol. 46, no. 11, p. 2–8, Nov. 2008.

[9] J. R. Douceur, "The sybil attack," in *IPTPS '01: First Intl Workshop on Peer-to-Peer Systems*. Springer, 2002, pp. 251–260.

[10] L. Fischer, A. Aiijaz, C. Eckert, and D. Vogt, "Secure revocable anonymous authenticated inter-vehicle communication (SRAAC)," in *Proc. ESCAR '06*, 2006.

[11] F. Armknecht, A. Festag, D. Westhoff, and K. Zeng, "Cross-layer privacy enhancement and non-repudiation in vehicular communication," in *WMAN '07: 4th Workshop on Mobile Ad-Hoc Networks*, March 2007.

[12] Y.-C. Hu and K. P. Laberteaux, "Strong VANET security on a budget," in *4th Workshop on Embedded Security in Cars (ESCAR)*, 2006.

[13] G. Calandriello, P. Papadimitratos, J.-P. Hubaux, and A. Lioy, "Efficient and robust pseudonymous authentication in VANET," in *VANET '07: Proc. 4th ACM Intl workshop on Vehicular ad hoc networks*. ACM, 2007, pp. 19–28.

[14] X. Lin, X. Sun, P.-H. Ho, and X. Shen, "Gsis: A secure and privacy-preserving protocol for vehicular communications," *IEEE Trans. Vehicular Technology*, vol. 56, no. 6, pp. 3442–3456, Nov. 2007.

[15] R. Lu, X. Lin, H. Zhu, P.-H. Ho, and X. Shen, "Ecpp: Efficient conditional privacy preservation protocol for secure vehicular communications," *INFOCOM 2008: 27th Conference on Computer Communications*, pp. 1229–1237, 2008.