PREparing SEcuRe VEhicle-to-X Communication Systems

Deliverable 5.2

Deployment Issues Report V2

Project:	**PRESERVE**
Project Number:	**IST-269994**
Deliverable:	**D5.2**
Title:	**Deployment Issues Report V2**
Version:	**1.0**
Confidentiality:	**Public**
Editor:	**Jonathan Petit**
Cont. Authors:	**N. Bissmeyer, M. Feiri, F. Kargl, S. Gisdakis, M. Moser, P. Papadimitratos**
Date:	**2013-01-30**

Part of the Seventh Framework Program
Funded by the EC-DG INFSO

Document History

Version	Date	Main author	Summary of changes
v0.1	2012-09-18	J. Petit (UT)	Initial version
v0.2	2012-10-05	N. Bissmeyer (SIT)	Misbehavior detection, PKI structure
v0.3	2012-10-16	S. Gisdakis (KTH), P. Papadimitratos (KTH)	Questionnaire
v0.4	2012-11-12	M. Feiri (UT)	Certificate omission
v0.5	2012-12-02	J. Petit (UT)	Pseudonym management
v0.6	2012-12-02	F. Kargl (UT)	Differential privacy
v0.7	2013-01-23	M. Moser (ESCRYPT)	ASIC cost model
v0.9	2013-01-24	J. Petit (UT)	Introduction and conclusion, additions and modifications
v1.0	2013-01-30	J. Petit (UT), F. Kargl (UT) S. Gisdakis (KTH), P. Papadimitratos (KTH)	Integration of partners comments and final version

Approval		
	Name	Date
Prepared	Jonathan Petit	2013-01-25
Reviewed	All Project Partners	2013-01-28
Authorized	Frank Kargl	2013-02-02

Circulation	
Recipient	Date of submission
Project Partners	2013-02-05
European Commission	2013-02-05

Contents

List of Figures

List of Tables

1 Glossary

Abbrev	Synonyms	Description	Details
API		Application Programming Interface	An API is a particular set of specifications that software programs can follow to communicate with each other.
AU		Application Unit	Hardware unit in an ITS station running the ITS applications
ASN.1		Abstract Syntax Notation One	ASN.1 is a standard and flexible notation that describes data structures for representing, encoding, transmitting, and decoding data.
CA		Certificate Authority	A CA is an entity that issues digital certificates.
CAM		Cooperative Awareness Message	CAMs are sent by vehicles multiple times a second (typically up to 10 Hz), they are broadcasted unencrypted over a single-hop and thus receivable by any receiver within range. They contain the vehicle's current position and speed, along with information such as steering wheel orientation, brake state, and vehicle length and width.
CAN		Controller Area Network	A CAN is a vehicle bus standard designed to allow microcontrollers and on-board devices to communicate with each other.
CCM		Communication Control Module	Module responsible for protecting on-board communication. Originates from the EVITA project.
CCU		Communication & Control Unit	Hardware unit in an ITS station running the communication stack
CE		Consumer Electronics	Electronic devices like smartphone or MP3 player of the vehicle driver or a passenger

Abbrev	Synonyms	Description	Details
CL		Convergence Layer	Module that connects the external on-board entities (e.g. communication stack or applications) to the PRESERVE Vehicle Security Subsystem (VSS)
CPU		Central Processing Unit	
CRC		Cyclic Redundancy Code	Is used to produce a checksum in order to detect errors in data storage or transmission.
CRS		Cryptographic Services	Module acting as proxy for accessing different cryptographic algorithm implementations. Originates from the EVITA project
DoS		Denial of Service	A DoS is a form of attack on a computer system or networks.
DENM	DNM	Decentralized Environmental Notification Message	A DENM transmission is triggered by a cooperative road hazard warning application, providing information to other ITS stations about a specific driving environment event or traffic event. The ITS station that receives the DENM is able to provide appropriate HMI information to the end user, who makes use of these information or takes actions in its driving and traveling. Fehler: Referenz nicht gefunden
EAM		Entity Authentication Module	Module responsible for ensuring entity authentication of in-vehicle components. Originates from the EVITA project
ECC		Elliptic Curve Cryptography	ECC is an approach to public-key cryptography based on the algebraic structure of elliptic curves over finite fields.
ECU		Electronic Control Unit	
ECR		ECU configuration register	Register used for secure boot and authenticated boot inside the HSM (similar to platform configuration register inside a TPM)
FOT		Field Operational Test	

Abbrev	Synonyms	Description	Details
G5A		ITS road safety communication (802.11p)	Frequency band between 5.875 GHz and 5.905 GHz - reserved for ITS road safety communication
G5B		ITS non-safety communication (802.11p)	Frequency band between 5.855 GHz and 5.875 GHz - reserved for ITS road non-safety communication
G5C	C-WLAN	5GHz WLAN communication (802.11a)	
GNSS	GPS	Global Navigation Satellite System	Generic term for an Global navigation satellite system (GPS, GLONAS, Galileo)
HMI		Human-Machine Interface	
HSM		Hardware Security Module	
HU		Head-Unit	
I2V	I2C	Infrastructure-to-Vehicle	Communication between infrastructure components like roadside units and vehicles
I2I		Infrastructure-to-Infrastructure	Communication between multiple infrastructure components like roadside units
ICS		ITS Central Station	ITS station in a central ITS subsystem
ILP		Inter Layer Proxy	Component introduced by the SeVeCom project, that captures and allows modification of messages between different layers of a communication stack
IDK	Module Authentication Key	Device Identity Key	The Device Identity Key is introduced by EVITA and is used for HSM identification. The IDK can also be certified by a manufacturer authentication key.
IMT	GSM, GPRS, UMTS	Public cellular services (2G, 3G, ...)	
IPR		Intellectual Property Right	

Abbrev	Synonyms	Description	Details
ITS		Intelligent Transportation Systems	Intelligent Transport Systems (ITS) are systems to support transportation of goods and humans with information and communication technologies in order to efficiently and safely use the transport infrastructure and transport means (cars, trains, planes, ships).
ITS-S		ITS Station	Generic term for any ITS station like vehicle station, roadside unit, ...
IDM		ID & Trust Management Module	Module responsible for ID management originating from SeVeCom project.
IVC	ITSC, ITS Communications	Inter-Vehicle Communication	Combination of V2V and V2I
IVS	OBU	ITS Vehicle Station	The term "vehicle" can also be used within PRESERVE
LDM	Environment Table	Local Dynamic Map	Local geo-referenced database containing a V2X-relevant image of the real world
LTC		Long-Term Certificate	PRESERVE realization of an ETSI Enrolment Credential. The long-term certificate authenticates a stations within the PKI, e.g., for PC refill and may contain identification data and properties.
LTCA		Long-Term Certificate Authority	PRESERVE realization of an ETSI Enrollment Credential Authority that is part of the PKI and responsible for issuing long-term certificates.
MAC		Media Access Control	The MAC data communication protocol sub-layer is a sublayer of the Data Link Layer specified in the seven-layer OSI model.
OBD		On-Board Diagnosis	OBD is a generic term referring to a vehicle's self-diagnostic and reporting capability that can be used by a repair technician to access the vehicles sub-systems.

Abbrev	Synonyms	Description	Details
OEM		Original Equipment Manufacturer	Refers to an generic car manufacturer
OBU	IVS	On-Board Unit	An OBU is part of the V2X communication system at an ITS station. In different implementations different devices are used (e.g. CCU and AU)
PAP		Policy Administration Point	Module related to the PDM originating from EVITA project
PC	Short Term Certificate	Pseudonym Certificate	A short term certificate authenticates stations in G5A communication and contains data reduced to a minimum.
PCA		Pseudonym Certificate Authority	Certificate authority entity in the PKI that issues pseudonym certificates
PDM		Policy Decision Module	Module responsible for enforcing the use of policies originating from EVITA project
PDP		Policy Decision Point	Module related to the Policy Decision Module originating from EVITA project
PeRA		Privacy-enforcing Runtime Architecture	Module responsible for enforcing privacy protection policies originating from PRECIOSA project
PEP		Policy Enforcement Point	Module related to the Policy Decision Module originating from EVITA project
PIM		Platform Integrity Module	Module responsible for ensuring in-vehicle component integrity originating from EVITA project
PKI		Public Key Infrastructure	A PKI is a set of hardware, software, policies, and procedures needed to create, manage, distribute, use, store, and revoke digital certificates.
PMM		Pseudonym Management Module	Module responsible for management of the station's pseudonym certificates originating from SeVeCom project
RSU	IRS, ITS Roadside Station	Roadside Unit	A RSU is a stationary or mobile ITS station at the roadside acting as access point to the infrastructure.

Abbrev	Synonyms	Description	Details
SAP		Service Access Point	Informative functional specification that enables the interconnection of different component implementations.
SM		Security Manager	Module responsible for securing the V2X communication with external ITS stations originating from SeVeCom project
SCM		Secure Communication Module	A generic name for the complete secure communication stack
SEP		Security Event Processor	Module responsible for security event management (e.g. checking message plausibility, station reputation calculation)
TPM		Trusted Platform Module	A TPM is both, the name of a published specification detailing a secure crypto-processor that can store cryptographic keys, as well as the general name of implementations of that specification, often called the "TPM chip" or "TPM Security Device".
UML		Unified Modeling Language	UML is an object modeling and specification language used in software engineering.
UTC		Coordinated Universal Time	UTC is the primary time standard by which the world regulates clocks and time.
V2I	C2I	Vehicle-to-Infrastructure	Direct vehicle to roadside infrastructure communication using a wireless local area network
V2V	C2C	Vehicle-to-Vehicle	Direct vehicle(s) to vehicle(s) communication using a wireless local area network
V2X	C2X	Vehicle-to-Vehicle (V2V) and/or Vehicle-to-Infrastructure (V2I)	Direct vehicle(s) to vehicle(s) or vehicle(s) to infrastructure communication using a wireless local area network
VIN		Vehicle Identification Number	Unique serial number of a vehicle
VSA		Vehicle Security Architecture	General outcome of PRESERVE work package 1

Abbrev	Synonyms	Description	Details
VSS		V2X Security Subsystem	Close-to-market implementation of the PRESERVE VSA that is the outcome of PRESERVE work package 2
WLAN		Wireless Local Area Network	
XML		Extensible Markup Language	XML is a set of rules for encoding documents in machine-readable form.

2 Introduction

The Work Package 5 (WP5) is concerned with security and privacy related challenges in ITS, especially those related to key matters of later deployment and standardization of secure and privacy protecting V2X communication systems.

The focus of this deliverable is on scalability of secure communication and the future PKI structure. Before that, we investigate the deployment issues of PRESERVE, more especially, how the PRESERVE platform is seen and received by industries (see Section 3). We introduce the PRESERVE questionnaire that was created and disseminated to stakeholders in automotive industry and beyond to analyze their awareness of security and privacy. While we introduce this questionnaire already in this document, its final evaluation will only be presented in D5.3.

Secondly, Section 4 describes the development of life-cycle management components, namely the Public Key Infrastructure (PKI).

Section 5 gives an initial analysis of the cost model for ASIC development. A cost model is created to relate functionality of an ASIC to its costs which is very useful during early phases of the chip development. A refinement of this model will be provided after ASIC development has concluded.

Section 6 reports the research work relating to various security and privacy aspects performed throughout 2012 by the PRESERVE partners. We also shortly introduce an extensive research survey on V2X pseudonym schemes that we submitted to ACM Surveys. We hope that this will serve as a standard reference for standardization and development to access the broad body of literature on this topic.

The results presented herein correspond to the subtasks 5210 and 5110 of the PRESERVE Description of Work.

The focus in this report is on forward looking issues, beyond the PRESERVE architecture and security subsystem. In addition, the WP5 reports provide a track record of all related research output. Accordingly, the V2X Security Subsystem (VSS) does not integrate all schemes presented in this deliverable: the details regarding VSS, notably its first version, are available in deliverables of WP2 and WP4, and the field trial related material in deliverables of WP3. It is expected that the second version of the VSS will integrate some schemes and elements that are results of the ongoing WP5 work.

3 Broadening awareness on the PRESERVE platform

There is a consensus being formed, in terms of basic technological aspects for security and privacy in ITS. Nonetheless, many questions concerning the actual deployment of these systems are not addressed yet. In addition, issues such as product life-cycles and costs for ITS products and services have to be defined, so that vehicular communication solutions can be brought to market. These are in fact important factors for the PRESERVE project and more generally for the ITS community.

In order to gauge the perception of the broader ITS community regarding the security and privacy needs for ITS and the PRESERVE architecture, we have designed and disseminated a questionnaire that seeks answers to the above and serves as an extension of the investigations related to this work package of PRESERVE. This section provides an overview of the structure of our survey along with the methodology for its design. The questionnaire can be found on the PRESERVE website[1].

We are currently in the process of collecting responses. Rather than including here limited results and thus providing a limited analysis, we shall update this report in the Deliverable D5.3.

3.1 Overview of the Survey

Our survey is designed in a way that no prior knowledge of the responders is presumed. We begin by asking the responder to provide us with input on her background. This helps us to better analyse the responses and weight them accordingly. Moreover, we treat each individual response as anonymous and strictly confidential. We emphasize that answers reflect the opinions of the individual responder alone and not of the institutions they represent. Similarly, once the responses are collected the resulting analysis will present aggregate responses. We use three types of questions; *multiple choice*, *free text* and *matrix questions*. For the latter ones, we utilize a scale from 0 (low) to 4 (high). The questionnaire comprises six sections:

- **Introductory Questions**: This section contains general questions concerning the background of the responder in terms of security and privacy for ITS. In addition, we try to capture the understanding of the responder on the PRESERVE architecture.

[1] http://www.preserve-project.eu/node/43

- **Questions on Safety Applications**: These questions focus on security and privacy requirements for specific safety applications as defined in the survey. We also inquire on the suitability of the PRESERVE architecture for protecting these applications.
- **Questions on Infotainment and Miscellaneous Applications**: These two sections focus on infotainment and miscellaneous applications. Similarly to the previous section, we are interested in the security and privacy requirements of these applications and in the applicability of PRESERVE's VSA for these application types.
- **Questions Regarding Financial Aspects**: These questions target responders whose role in the institutions they represent is of managerial nature.
- **Questions Regarding Technical Aspects**: This category contains questions of technical nature that target the part of the audience/responders with technical security and privacy expertise.

3.1.1 Introductory Questions

This section includes seven (7) questions. The first questions (*Q1* and *Q2*) ask for the responders' personal information. Given we treat the answers of individuals anonymously, fields such as the responders' name, email and phone are only *optional*. The only pieces of information we require are the organization position and the organization type for the responder. Based on this question we can have an understanding of her background.

In *Q3*, we ask the responders about their familiarity with the PRESERVE project and various standardization bodies active in the area of ITS (IEEE 1609.2-WG[2], ETSI-WG5[3] and C2C-CC[4]). If the respondent is familiar with the above, she is considered to be a specialist when it comes to technical aspects for ITS and her answers will be analyzed accordingly.

Q4 asks the responders how important they consider security and privacy requirements to be. These requirements are extracted from the state-of-the-art research and technical literature.

Q5 tries to identify whether the broader ITS community considers applications built on top of collaborative, ad hoc communication (IEEE 802.11p [1]) warrant stronger and more involved security protection scheme compared to the ones that rely on cellular networks (e.g., 2G/3G/LTE).

Q6 and *Q7* are specific to the PRESERVE architecture. Based on the approach for *Q5*, in *Q6* we ask the responders if they consider the results of PRESERVE applicable for ITS applications that built on cellular networks.

Q7 asks the opinion of the responders regarding the applicability of PRESERVE to applications specific to various different domains.

[2] http://vii.path.berkeley.edu/1609_wave/
[3] http://www.etsi.org/website/technologies/intelligenttransportsystems.aspx
[4] http://www.car-to-car.org/

3.1.2 Safety Applications Questions

In this section, the survey focuses on safety applications. We consider the following list of safety applications [2]:

- **Road Hazard Warning**: Sudden slow-down warning, vehicle safety function out of normal condition warning
- **Cooperative Awareness**: Emergency vehicles notification, slow vehicle notification, motorcycle notification
- **Cooperative Collision Avoidance**: Vulnerable user warning
- **Traffic Hazard Warning**: Wrong way driving notification, stationary vehicle notification, traffic jam notification, signal violation notification

This section contains four (4) questions whose purpose is to help us understand *if* and *how* PRESERVE's VSS can be utilized to guarantee the security and privacy requirements of the four safety applications presented above.

Q8 and *Q9* probe the familiarity of the respondent concerning security and privacy requirements for safety applications. As the core focus of PRESERVE is on safety applications, it is critical to understand the opinion of the ITS community concerning the suitability of PRESERVE for these applications. This is exactly the purpose of *Q10*.

In case the respondent answered with "Strongly Disagree" or "Disagree" in *Q10*, a contingent question follows (*Q11*) where it is requested that the respondent explain briefly why and provide, if known, approach(es) that would be more applicable or more effective in satisfying the aforementioned security and privacy requirements of safety applications.

The purpose of the last question in this section is to get the opinion of the responder regarding the possible overhead that security and privacy mechanisms introduce for safety applications.

3.1.3 Traffic Efficiency and Infotainment Applications Questions

These two sections follow the same structure and mentality with the previous one. Their difference is that they focus on traffic efficiency applications (Sec. 3 of the questionnaire) and infotainment applications (Sec. 4). To facilitate the answering of the questions in these sections we provide the responders with lists of traffic efficiency and infotainment applications according to [2].

3.1.4 Financial Aspects Questions

This section of the survey targets responders whose role in the company or the institution they represent is of managerial/business (non-technical) nature. This is the purpose of *Q21*, which gauges the understanding of the respondents on the business aspects of ITS systems.

Q22 asks the respondent's opinion on the potential commercial value of the PRESERVE ITS solution.

Q23 tries to understand the motives that drive organizations and institutions to introduce security and privacy solutions into their ITS related products and services.

In *Q24* we ask for the level of security that the organization of the respondent is applying and subsequently in *Q25* the reader should provide an estimation of the budget that their company can afford for introducing security and privacy protection to their products and services.

Similarly, in *Q26* we request from the respondents to estimate the product life cycle cost of ITS related products.

In *Q27* the responders are asked to elaborate on the cost distribution among the different components of the product life cycle of ITS products and services. *Q28* and *Q29* gauge the value that the responder attributes to the security and privacy for safety and traffic efficiency application respectively.

Q30 works towards the direction of specifying business models for security and privacy in ITS. More specifically, we ask the responder to answer on who she thinks that should be the responsible for paying the cost of security and privacy. A similar question regarding the cost of deploying the needed RSUs and infrastructure is presented in *Q31*.

3.1.5 Technical Aspects Questions

The final section of our survey is concerned with technical aspects of security and privacy for ITS. We begin by asking the responders in *Q32* about their technical background and understanding of technical aspects of security and privacy of ITS.

In *Q33* we require the responders to provide their input on the impact of a set of security and privacy threats. In questions *Q34*, *Q35* and *Q36* we ask the responder on the suitability of different cryptographic schemes in the context of safety, traffic efficiency and infotainment applications. We conclude this section with *Q37* which asks the responders to provide their input on the technical challenges towards the deployment of secure and privacy protecting ITS.

3.2 Survey Dissemination

We have created an on-line version of our survey which allows enhanced dissemination and analysis capabilities. The link to the survey has been uploaded on the web-page of the PRESERVE project [5] The survey disseminated to various responders such as standardization bodies and experts in the are of ITS. We gathered responding volunteers during the ITS World Congress held in Vienna from 22 to 26 of October 2012. In addition, we advertised our survey during the proceedings of C2C-CC Forum held in Göteborg (Sweden) on 13 and 14 of November 2012 and in the EIT-ICT Safe Mobility chapter [6]. We have continued with collaborating FOT projects, with a US-EU Harmonization Working Group, and select researchers in the broader ITS area.

3.3 Analysis of the Responses

Currently we are in the process of collecting the responses from the responders. We already have a number of responses but for best results we refrain from presenting any conclusions with a limited set. We rather extend our effort and wait for responses to come in. Once this phase is over, with satisfactory number of responses, we will analyse the responses.

We shall present all responses as aggregates, identify the correlations among them and the trends that emerge. We will then include the corresponding technical discussion on the exact possibilities and ways of integration of the PRESERVE solution notably towards the possibilities that are identified as more relevant or plausible. Based on those, and the upcoming steps within the project, we should be able to outline a roadmap towards further integration of PRESERVE; this relates to future possibilities that materialize towards further testing of the PRESERVE system.

[5] http://preserve-project.eu/
[6] http://www.eitictlabs.eu/action-lines/intelligent-mobility-and-transportation-systems/

4 Development of Life-Cycle Management Components

4.1 Introduction

4.1.1 Purpose of this chapter

This section aims to document and discuss implementation related topics of a C2X PKI. The PKI architecture used in this document is based on the C2C-CC PKI Memo [3]. The PKI developed in the PRESERVE project should be flexible and scalable so that the application in different Field Operational Tests (FOTs) is possible.

This document is describing current work in progress and must be handled confidential. Distribution without permission is not allowed.

In section 4.2, different operational PKI issues are discussed that have to be specified and commonly accepted inside a PKI domain. Section 4.2.6 presents the different certificate types in detail. This description may help to integrate and implement the IEEE 1609.2 v2 [4] certificates. As existing formats in IEEE 1609.2 v2 seems to be not sufficient for the PRESERVE PKI, we propose additions of the standard in Section 4.1.3. The IEEE 1609.2 v2 formats extended by own formats are used in Section 4.2.6 in order to describe the API between the ITS-S and the PKI.

4.1.2 Overview of PKI integration into the V2X communication

In order to protect the integrity and authenticity of messages in V2X communications, all participating ITS-S have to be equipped with digital certificates. As a precondition for being able to request pseudonym certificates, a vehicle must be registered at a Long-Term CA and own a valid long-term certificate. Pseudonym certificates are issued by a trusted third party (i.e. PKI) as described in the following.

The vehicle sends a request to a predefined Pseudonym CA (e.g. Home PCA). If the vehicle's preferred PCA is not available, the vehicle may send the request to another available PCA. The pseudonym certificate request includes the signer ID of the long-term certificate, only one public key or a list of public keys, the current position, and an ID or address of the vehicle's Long-Term CA (i.e., the Long-Term CA that the vehicle is registered at). If a Pseudonym CA receives a request, it checks if it is able to issue pseudonyms for the requested region based on the stated region in the certificate request.

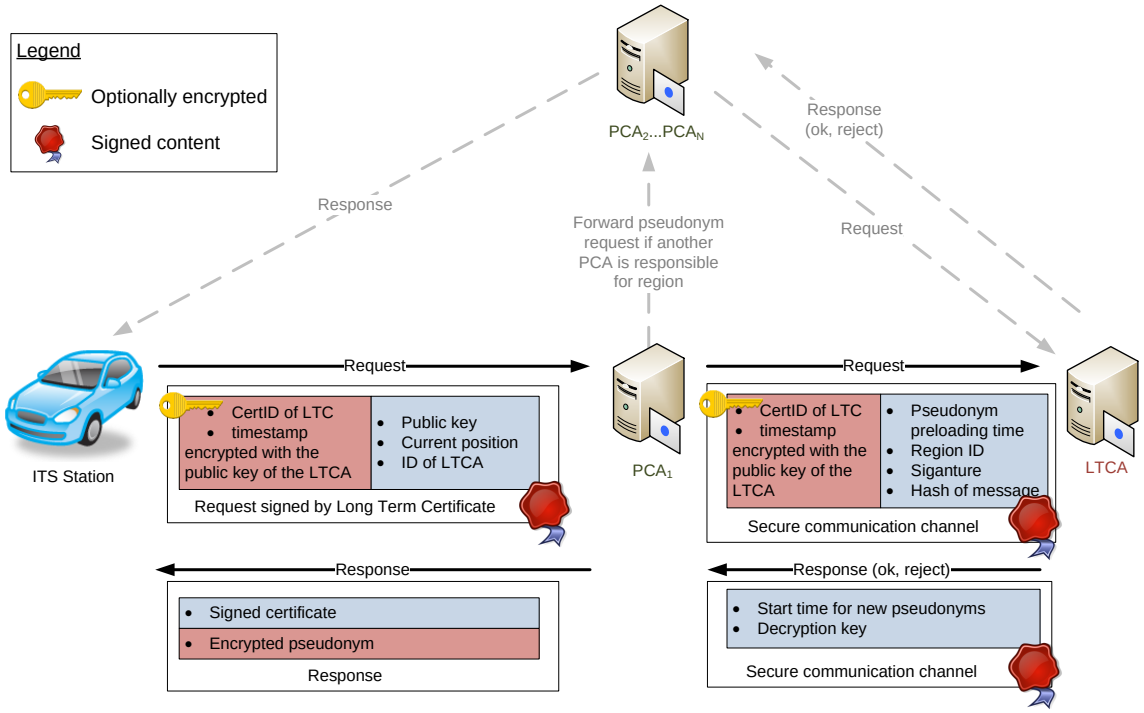

Figure 4.1: Pseudonym certificate request

If another Pseudonym CA is responsible for the region, the request is forwarded to an appropriate Pseudonym CA such as displayed in the Fig. 4.1. For privacy reasons the signer ID of the sender can be encrypted with a freshly generated symmetric AES key which is subsequently encrypted with the public key of the Long-Term CA using ECIES. In this case, the Pseudonym CA is not able to create a link between the pseudonyms and the long-term ID of the vehicle. Neither can the Long-Term CA create such a link because the Pseudonym CA does not forward public keys or pseudonym certificates. Based on legislation, the signer ID can also be transmitted unencrypted such that the Pseudonym CA is able to operate a database with links between a long-term ID and corresponding pseudonym certificates. If the signer ID is encrypted, the appropriate Pseudonym CA sends a request with the (encrypted) signer ID of the requester, a calculated preloading time and the region ID to the Long-Term CA for verification. The Long-Term CA maintains a database that stores the timestamp until a vehicle has valid pseudonyms for a distinct region. Only if the following checks are successful the Pseudonym CA will get a positive response from the Long-Term CA.

- Signer ID of LTC can be found at the LTCA and the certificate is not deactivated.
- The issuance of the requested pseudonyms would not result in exceeding the number of allowed parallel pseudonyms for the given region ID until the given preloading time (See section 4.2.5 for details about the number of allowed pseudonyms).

The Long-Term CA sends a positive or negative response with a start timestamp for new pseudonyms back to the Pseudonym CA. With this procedure, the PKI can prevent vehi-

cles from requesting more pseudonyms than allowed for the same time interval from various Pseudonym CAs. Please note, that the process described in this section is based on constantly available PCAs and LTCAs. An offline generation of Pseudonyms on the PCA without contact to the responsible LTCA is not possible in this concept as the permissions and the timestamps for pseudonym certificate validity have to be requested individually for every pseudonym public key from the LTCA.

4.1.3 Standards

The PKI described and discussed in this document bases in general on concepts described in the C2C-CC PKI Memo [3], ETSI TS 102 731 [5], ETSI TS 102 940 [6] and IEEE 1609.2 v2 D9 [4]. Furthermore, the content of this document is related to content of ETSI 102 941 [7] and makes possible concrete proposal for open issues.

The PKI should be aligned with applicable ETSI standards, especially ETSI TS 102 867 on mapping of IEEE 1609.2, regarding reference points between ITS-S stations and authorities for enrolment and authorizations.

4.2 Operational PKI Issues

The deployment, operation and usage of a PKI depends primarily on the trustworthiness of the involved CAs. Thus, it is necessary to develop not only a secure implementation of the CA software, but also to define operational rules and policies and to follow them. This is even more important, if different CA operators exist. In this case, it is especially important to provide a specification of common interfaces that all participants can rely on.

4.2.1 Common interfaces

This section lists the network interfaces that are provide by CAs of the PKI used in PRESERVE.

4.2.1.1 Interfaces for ITS stations

Long-Term CAs and Pseudonym CAs provide a UDP interface for sending requests to the CA. Over this interface, a serialized message in the format of IEEE 1609.2 v2 D9 [4] or in a format of the IEEE 1609.2 v2 D9 [4] extensions defined in section 4.2.7 can be sent. To overcome the limitations of UDP, it is reasonable to also provide a web service (e.g. SOAP) that could be considered in future work.

4.2.1.2 Interfaces for Inter-CA communication

During processing of a pseudonym certificate request, the processing Pseudonym-CA can communicate with the appropriate Long-Term CA using a web service (SOAP). The data transmitted over this web service is represented as a serialized message defined in the extensions of IEEE 1609.2 v2 D9 [4] in section 4.2.7.

4.2.2 Revocation of CA certificates

Revocation of CA certificates is a subtle problem in PKI scenarios as the validity of every issued certificate depends also on the certificate of the signing CA. As an illustration of the problem, assume a CA with a certificate that is valid at time t_1 and expires at time t_3 and assume further that this certificate is revoked at time t_2 ($t_1 < t_2 < t_3$). An important question in this case is what happens to the validity of certificates issued by the CA before time t_2 when the CA certificate is revoked at time t_2. In literature, mainly two validity models for CA revocation are discussed: Shell- and Chain-model [8].

- In the shell model, the validity of the issued certificates is directly related to the validity of the issuer certificate. If the issuer certificate is revoked, also the issued certificate is immediately invalid. In the verification process of the issued certificate the validity of the issuer's certificate is checked using the current timestamp.

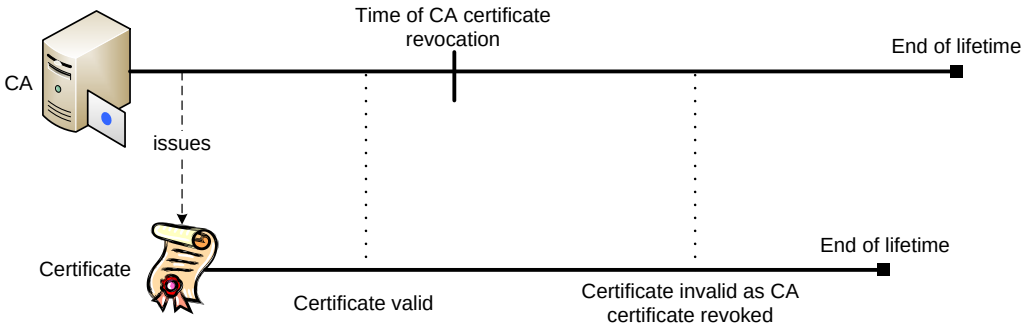

Figure 4.2: Verification of certificate validity following the shell model

Example: A long-term CA issues long-term certificates for many ITS stations with validity for 2 years. If a specific permission is then withdrawn at the LTCA after one year, then all long-term certificates of affected stations must be updated.

- In the chain model the validity of the issued certificate is not directly related to the issuer certificate. If the issuer certificate is revoked, the issued certificate is further valid until the expiry of the issued certificate. In the verification process of the issued certificate, the validity of the issuer certificate is checked, using the timestamp of issuance of the certificate. In this model revocation information of the past has to be available at the ITS stations which increases the complexity and storage requirements.

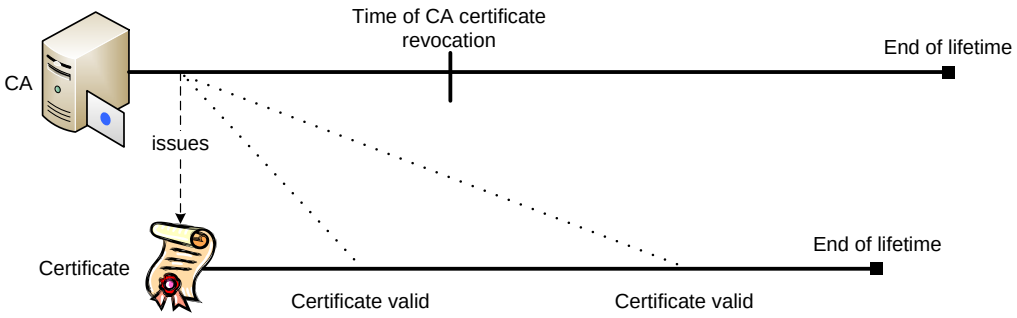

Figure 4.3: Verification of certificate validity following the chain model

Example: A long-term CA issues long-term certificates for many ITS stations with validity for 2 years. If a specific permission is then withdrawn at the LTCA after one year, then the LTCA certificate is revoked but the issued long-term certificate is still valid as the timestamp of issuance is used to verify the long-term certificate. This means, that the issued certificates are valid until their lifetime is exceed.

4.2.3 Permission Management of CAs and ITS-Stations

As defined in IEEE 1609.2 v2 [4], the concept of Provider Service Identifier (PSID) with related Service Specific Permissions (SSP) is used to define permissions. In general, every certificate may contain a list of PSIDs that determines the rights of the certificate owner. The PSID encodes the general right to provide general information such as CAMs or DENMs. The specific permission of the sender is encoded with the SSP. For example, only emergency vehicles should be able to send CAMs with the emergency profile that indicates beside others the use of blue lights.

An important aspect is the addition and removal of PSIDs and SSPs of existing CA certificates.

- **Addition of permission:** Assuming a CA should get a new permission in form of a PSID in its certificate then the root CA can issue a new certificate for the CA with the new permissions and new lifetime. The old certificate can be further used in parallel with the new certificate as both certificates have different certIDs. One the one hand, certificates that are issued with the old certificate is valid and can be verified as their issuer certificate is not revoked and the PSIDs are matching. One the other hand, new certificates with the added PSID are issued by the new CA certificate that contains also the required PSID. The parallel usage of several CA certificates with the same private / public key pair is a well-known concept in PKIs.
- **Remove of permission:** The deletion of permissions from CA certificates is more complex and is related to the issues discussed in section 4.2.2. If a specific permission is withdrawn at a CA then the issuer certificate is revoked and a new certificate is created using the same private / public key pair. The exchange of issued certificates (i.e. long-term certificates) can be done in the process of requesting new

pseudonym certificates from the PKI. The revocation of PCAs is not critical as issued pseudonym certificates have a relatively short lifetime and regular request of pseudonym certificates are part of the concept. Nevertheless, if a PCA certificate is revoked then all issued pseudonyms are immediately invalid and new pseudonyms must be requested.

4.2.3.1 Permissions of Root CA

The RCA certificate has specific significance in the PKI concept. Inside the ITS station the root certificate has to be protected against modification and arbitrary exchange. Furthermore, the root certificate has in general a longer lifetime and can only be substituted by another root certificate if this is update is permitted by the old root certificate. If permissions are part of the root certificate in form of PSIDs, then newly issued certificates can only contain a subset of the issuer's certificate. This is a problem if new functions with new PSID should be introduced. Then the new root certificate cannot contain more PSIDs than the old root certificate. In order to solve this dilemma, it is proposed that the root certificate do not contain PSIDs and the RCA is able to issue other CA certificates with any permission that is necessary. This reduces the complexity at the ITS station for updating the root certificate and enables a more flexible lifetime definition of root certificates. In IEEE 1609.2 v2 [4], it is also defined that no PSID is provided in the root certificate.

4.2.3.2 Permissions of Long-Term CA

The LTCA is able to provide PSIDs listed in Table 4.1.

Service Identifier (PSID)	Description
108	CAM
109	DENM

Table 4.1: Possible service identifiers of the LTCA

4.2.3.3 Permissions of Pseudonym CA

The PCA is able to provide PSIDs listed in Table 4.2.

Service Identifier (PSID)	Description
108	CAM
109	DENM

Table 4.2: Possible service identifiers of the PCA

4.2.3.4 Permissions of ITS Stations

The ITS station is able to provide PSIDs listed in Table 4.3.

Service Identifier (PSID)	Specific Permission (SSP)	Description
108		CAM with default profile
108	109	CAM with public transportation profile
108	110	CAM with emergency vehicle profile
109		DENM
109	109	DENM with message type SPAT

Table 4.3: Possible service identifiers and service specific permissions of the ITS station

4.2.4 Lifetime of CA certificates

The lifetime of the CA certificates is restricted in general. Therefore, an update of CA certificates has to be considered. As the Root CA certificate is used as trust anchor all CA certificate updates can be protected by the root CA. In order to have valid CA certificates also if ITS stations are not constantly connected to the PKI, the CA certificates shall have overlapping validity time periods as shown in Figure 4.4.

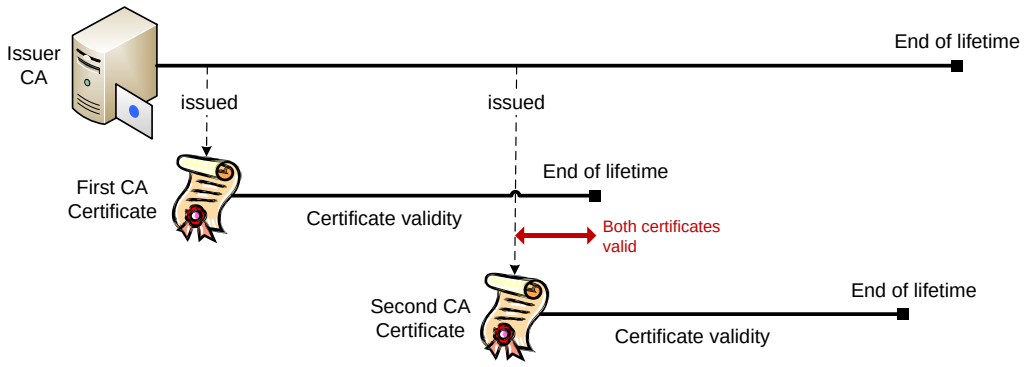

Figure 4.4: Overlapping lifetime of CA certificates

4.2.4.1 Lifetime and Update of Root CA certificate

The RCA certificates contain only an expiration timestamp and no validity start time as shown in section 4.2.6.1.

The predefined expiration time of the two RCA certificates within PRESERVE certificate is **2013-12-31**. Before reaching the expiration timestamp of the old certificate, the root CA

certificate has to be updated and exchanged by a new certificate considering following process:

1. Parallel to the old RCA credentials, that are still valid, a second ECDSA-256 key pair is generated and a self-signed root certificate is created and stored. After this point in time, two root certificates are valid and can be used equally.
2. ITS stations that request the CA certificate will get the new root CA certificate. If the root CA has no older certificate that is valid in parallel, then the root certificate is encapsulated in an unsecured 1609Dot2Message. Otherwise, the root certificate is signed by the old root certificate that is still valid.
3. The ITS station shall store the new root certificate in parallel if the signature of the downloaded certificate can be correctly verified with the old root CA certificate. If the old root certificate on the ITS station is expired before the new root certificate is downloaded then an update of the VSS software is necessary in order to get the new root certificate which acts as trust anchor.
4. ECIES encrypted messages that are sent to the RCA will be decrypted with the keys belonging to the new certificate.

4.2.4.2 Lifetime and Update of LTCA certificate

The LTCA certificates contain only an expiration timestamp and no validity start time as shown in section 4.2.6.2. The CertificateContentFlags are not set.

The predefined expiration time of LTCA certificates within PRESERVE is **2013-12-31**. Before reaching the expiration timestamp of the old certificate, the LTCA certificate has to be updated and exchanged by a new certificate considering following process:

1. Parallel to the old LTCA credentials, that are still valid, a second ECDSA-256 key pair is generated for the new LTCA certificate.
2. The new public key is sent to the RCA in order to get the LTCA certificate signed by the root CA. In this process, the certificate request used is described in [4] section 5.3.33.
3. The LTCA stores the old certificate with related keys as long as it is valid whereas every requester gets the new LTCA certificate. The new credentials are also used for signing long-term certificate requests.
4. Long-term certificate requests that are encrypted with ECIES and sent to the LTCA will be decrypted with the keys belonging to the latest LTCA certificate.

4.2.4.3 Lifetime and Update of PCA certificate

The PCA certificates contain only an expiration timestamp and no validity start time as shown in section 4.2.6.3. The CertificateContentFlags are not set.

The predefined expiration time of the PCA certificates within PRESERVE is **2013-12-31**. Before reaching the expiration timestamp of the old certificate, the PCA certificate has to be updated and exchanged by a new certificate considering following process:

1. Parallel to the old PCA credentials, that are still valid, a second ECDSA-256 key pair is generated for the new PCA certificate.
2. The new public key is sent to the RCA in order to get the PCA certificate signed by the root CA. In this process the certificate request is used that is described in [4] section 5.3.33.
3. The PCA stores the old certificate with related keys as long as it is valid whereas every requester gets the new PCA certificate. The new credentials are also used for signing pseudonym requests.

Pseudonym certificate requests that are encrypted with ECIES and sent to the PCA will be decrypted with the keys belonging to the latest PCA certificate.

4.2.5 Roaming

The certificate format defined by IEEE 1609.2 [4] allows restricting CA certificates and pseudonyms validity to a geographical area. In this PKI implementation, geographic region IDs are used to restrict certificate validity to nations, countries or states as proposed by C2C-CC [3] and ETSI [7]. As extension to the certificate formats described in IEEE 1609.2 [4], a numeric region ID can be added to the certificates in order to encode specific geographic areas in a flexible way but limit the certificate size of pseudonyms. A region list is used that contains polygons describing the borders of a geographical region that is linked to the region ID. This list is signed by the root CA and should be available on all ITS stations of the PKI domain. Table 4.4 describes the region IDs used in the PRESERVE project. It is proposed to use country calling codes in order to define unique IDs for the regions. More detailed regions can be defined by adding phone prefix numbers of counties or cities. Using an unsigned Integer (UInt32) that is encoded with flexible length, numbers can be assigned between 0 and $2^{28} - 1$. Details for the region ID format can be found in Section 4.2.7.

Region ID	Description	Polygon Points
49	Country calling code of Germany	
33	Country calling code of France	

Table 4.4: List of regions used in certificates

4.2.6 Certificate Formats

In the following section, example certificates are described for the three CA entities RCA, LTCA and PCA. Furthermore, a long-term certificate is described as well as an example of a pseudonym certificate.

4.2.6.1 Root CA Certificate Format

Figure 4.5: General structure and size of a root CA certificate

- UInt8 – versionAndType = 2 (1 octet)
- ToBeSignedCertificate (97 octets)
  - SubjectType → Root_CA (1 octet)
  - CertificateContentFlags – encryption_key (1 octet)
  - CertSpecificData – scope of certificate (21 octets)
    - RootCAScope (21 octets)
      - UInt8[] – name of scope ”PRESERVE_Root CA” (1 + 16 octets)
      - SubjectTypeFlags – Permission to issue other certificates having subject type: ”MessageCA” (1 octets)
      - PsidArray (2 octets)
    - GeographicRegion → NONE (1 octet)
  - Time32 – expiration (expires on 2013/12/31) (4 octets)
  - CrlSeries - 1 (4 octets)
  - PublicKey for verification (34 octets)
    - PKAlgorithm – algorithm to verify message, ECDSA256 (1 octet)
    - EccPublicKey – PublicKey to verify message (33 octets)
      - EccPublicKeyType – COMPRESSED (1 octet)
      - Opaque[32] x – key data (32 octets)
  - PublicKey for encryption (35 octets)
    - PKAlgorithm – algorithm to encrypt message, ECIES_NISTP256 (1 octet)
    - SymmAlgorithm - aes_128_ccm (1 octet)
    - EccPublicKey – PublicKey to verify message (33 octets)
      - EccPublicKeyType – COMPRESSED (1 octet)
      - Opaque[32] x – key data (32 octets)
- Signature (65 octets)
  - EcdsaSignature (65 octets)
    - EccPublicKey signature1 (33 octets)
      - EccPublicKeyType – X_ONLY (1 octet)
      - Opaque[32] x
    - Opaque[32] signature2 (32 octets)

4.2.6.2 Long-Term CA Certificate Format

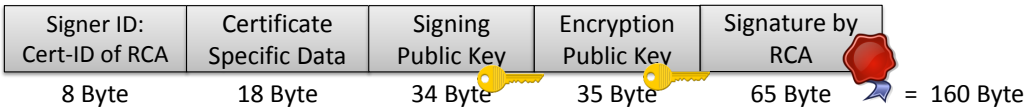

Figure 4.6: General structure and size of a Long-term CA certificate

- UInt8 – versionAndType = 2 (1 octet)
- ToBeSignedCertificate (115 octets)
  - SubjectType → Message_CA (1 octet)
  - CertificateContentFlags – encryption_key (1 octet)
  - CertId8 – signer_id (RootCA) (8 octets)
  - PKAlgorithm – signature_alg (1 octets)
  - CertSpecificData – scope of certificate (27 octets)
    - MessageCAScope (27 octets)
      - UInt8[] – name of scope ”PRESERVE_Long-Term_CA” (1 + 21 octets)
      - SubjectTypeFlags – Permission to issue other certificates having subject type: ”MessageCSR” (1 octets)
      - PsidArray (3 octets)
    - GeographicRegion → NONE (1 octet)
  - Time32 – expiration (expires on 2013/12/31) (4 octets)
  - CrlSeries - 1 (4 octets)
  - PublicKey for verification (34 octets)
    - PKAlgorithm – algorithm to verify message, ECDSA256 (1 octet)
    - EccPublicKey – PublicKey to verify message (33 octets)
      - EccPublicKeyType – COMPRESSED (1 octet)
      - Opaque[32] x – key data (32 octets)
  - PublicKey for encryption (35 octets)
    - PKAlgorithm – algorithm to encrypt message, ECIES_NISTP256 (1 octet)
    - SymmAlgorithm - aes_128_ccm (1 octet)
    - EccPublicKey – PublicKey to verify message (33 octets)
      - EccPublicKeyType – COMPRESSED (1 octet)
      - Opaque[32] x – key data (32 octets)
- Signature (65 octets)
  - EcdsaSignature (65 octets)
    - EccPublicKey signature1 (33 octets)
      - EccPublicKeyType – X_ONLY (1 octet)
      - Opaque[32] x
    - Opaque[32] signature2 (32 octets)

4.2.6.3 Pseudonym CA Certificate Format

- UInt8 – versionAndType = 2 (1 octet)

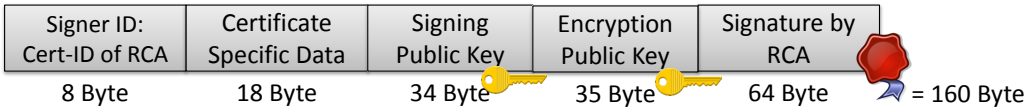

Figure 4.7: General structure and size of a Pseudonym CA certificate

- ToBeSignedCertificate (115 octets)
  - SubjectType → Message_CA (1 octet)
  - CertificateContentFlags – encryption_key (1 octet)
  - CertId8 – signer_id (RootCA) (8 octets)
  - PKAlgorithm – signature_alg (1 octets)
  - CertSpecificData – scope of certificate (27 octets)
    - MessageCAScope (27 octets)
      - UInt8[] – name of scope ”PRESERVE_Pseudonym_CA” (1 + 21 octets)
      - SubjectTypeFlags – Permission to issue other certificates having subject type: ”MessageIdentifiedLocalized” and ”MessageIdentifiedNotLocalized” (1 octets)
      - PsidArray (3 octets)
    - GeographicRegion → NONE (1 octet)
  - Time32 – expiration (expires on 2013/12/31) (4 octets)
  - CrlSeries - 1 (4 octets)
  - PublicKey for verification (34 octets)
    - PKAlgorithm – algorithm to verify message, ECDSA256 (1 octet)
    - EccPublicKey – PublicKey to verify message (33 octets)
      - EccPublicKeyType – COMPRESSED (1 octet)
      - Opaque[32] x – key data (32 octets)
  - PublicKey for encryption (35 octets)
    - PKAlgorithm – algorithm to encrypt message, ECIES_NISTP256 (1 octet)
    - SymmAlgorithm - aes_128_ccm (1 octet)
    - EccPublicKey – PublicKey to verify message (33 octets)
      - EccPublicKeyType – COMPRESSED (1 octet)
      - Opaque[32] x – key data (32 octets)
- Signature (65 octets)
  - EcdsaSignature (65 octets)
    - EccPublicKey signature1 (33 octets)
      - EccPublicKeyType – X_ONLY (1 octet)
      - Opaque[32] x
    - Opaque[32] signature2 (32 octets)

4.2.6.4 Long-Term Certificate Format

- UInt8 – versionAndType = 2 (1 octet)
- ToBeSignedCertificate (88 octets)
  - SubjectType → Message_CSR (1 octet)

Figure 4.8: General structure and size of a Long-term certificate

  - CertificateContentFlags – none (1 octet)
  - CertId8 – signer_id (LTCA) (8 octets)
  - PKAlgorithm – signature_alg (1 octets)
  - CertSpecificData – scope of certificate (27 octets)
    - MessageCsrScope (35 octets)
      - UInt8[] – name of scope "Example ITS-Station 4711" (1 + 24 octets)
      - SubjectTypeFlags – Permission to issue other certificates having subject type: "MessageIdentifiedLocalized" and "MessageIdentifiedNotLocalized" (1 octets)
      - PsidSspArray (8 octets)
    - GeographicRegion → NONE (1 octet)
  - Time32 – expiration (expires on 2013/12/31) (4 octets)
  - CrlSeries - 1 (4 octets)
  - PublicKey for verification (34 octets)
    - PKAlgorithm – algorithm to verify message, ECDSA256 (1 octet)
    - EccPublicKey – PublicKey to verify message (33 octets)
      - EccPublicKeyType – COMPRESSED (1 octet)
      - Opaque[32] x – key data (32 octets)
- Signature (65 octets)
  - EcdsaSignature (65 octets)
    - EccPublicKey signature1 (33 octets)
      - EccPublicKeyType – X_ONLY (1 octet)
      - Opaque[32] x
    - Opaque[32] signature2 (32 octets)

4.2.6.5 Pseudonym Certificate Format

Figure 4.9: General structure and size of a Pseudonym certificate

- UInt8 – versionAndType = 2 (1 octet)
- ToBeSignedCertificate (56 octets)
  - SubjectType → Message_identified_not_localized (1 octet)
  - CertificateContentFlags – use_start_validity, lifetime_is_duration (1 octet)

  - CertId8 – signer_id (PCA) (8 octets)
  - PKAlgorithm – signature_alg (1 octets)
  - CertSpecificData – scope of certificate (5 octets)
    - IdentifiedNotLocalizedScope (5 octets)
      - UInt8[] – none (1 octets)
      - PsidSspArray (4 octets)
  - Time32 – expiration (expires on 2013/12/31) (4 octets)
  - CertificateDuration – hours, 2880 = 4 month (2 octets)
  - CrlSeries - 1 (4 octets)
  - PublicKey for verification (30 octets)
    - PKAlgorithm – algorithm to verify message, ECDSA224 (1 octet)
    - EccPublicKey – PublicKey to verify message (29 octets)
      - EccPublicKeyType – COMPRESSED (1 octet)
      - Opaque[28] x – key data (28 octets)
- Signature (65 octets)
  - EcdsaSignature (65 octets)
    - EccPublicKey signature1 (33 octets)
      - EccPublicKeyType – X_ONLY (1 octet)
      - Opaque[32] x
    - Opaque[32] signature2 (32 octets)

4.2.7 Extensions to Existing 1609.2 v2 Messages

In this section, new formats for the request of certificates are proposed as well as additions to existing formats of [4].

4.2.7.1 Variable-length unsigned Integer

`uint32<var>`

This proposed structure is based on the variable length PSID that is encoded as follows according to [4]: *A field of type psid contains a Provider Service Identifier as defined in IEEE Std 1609.3. This is a variable-length type that encodes the length inside the data rather than in an external length field.* The `uint32<var>` is encoded in (length, value):

- if value < 27, length is the bit " 0 " and the value allocates the remaining 7 bits (e.g. bits `0xxx xxxx` with `x` as a bit of the value)
- else if value < 214, length is the bits "10" and the value allocates the remaining 14 bits (e.g. bits `10xx xxxx xxxx xxxx` with `x` as a bit of the value)
- else if value < 221, length is the bits "110" and the value allocates the remaining 21 bits (e.g. bits `110x xxxx xxxx xxxx xxxx xxxx` with `x` as a bit of the value)
- else if value < 228, length is the bits "1110" and the value allocates the remaining 28 bits (e.g. bits `1110 xxxx xxxx xxxx xxxx xxxx` with `x` as a bit of the value)

For example:

- value 0 is encoded as `0x00`
- value 1 is encoded as `0x01`
- 3 is encoded as `0x03`
- 127 is encoded as `0x7f` (bits `0111 1111`)
- 128 is encoded as `0x80 0x80` (bits `1000 0000 1000 0000`)

4.2.7.2 RegionType

```
enum {from_issuer(0), circle(1), rectangle(2), polygon(3), none (4),
// new material
    region_id(5),
// end new material
(2^8-1)} RegionType;
```

The new type `region_id` determines that a list of identifiers is added to the `GeographicRegion`. A motivation for this new type and details about its usage can be found in section 4.2.5.

4.2.7.3 GeographicRegion

```
struct {
    RegionType region_type;
    select(region_type){
    case from_issuer: ;
    case circle:
        CircularRegion  circular_region;
    case rectangle:
        RectangularRegion rectangular_region<var>;
    case polygon:
        PolygonalRegion     polygonal_region;
    // new material
        case region_id:
        uint32<var>     region_id<var>;
    // end new material
    case none: ;
    unknown:
        opaque      other_region<var>;
    }
} GeographicRegion;
```

A list of region identifiers can be added in the `GeographicRegion` structure in order to restrict the validity of a certificate to an area that is mapped to the given `region_id`. Every ITS station has to maintain a list that maps a `region_id` to a specific polygon that reflects the geographical region as detailed in section 4.2.5. Using a list of region identifiers may produce smaller certificate sizes instead of adding large lists of polygon to certificates (i.e. pseudonym certificates).

4.2.7.4 RegionIdList

```
struct {
    SignerIdentifier         signer;
    ToBeSignedRegionIdList   unsigned_list;
    signature              signature;
} RegionIdList;
```

The usage of region ID lists is described in section 4.2.5.

4.2.7.5 ToBeSignedRegionIdList

```
struct {
    RegionIdListEntry        entries<var>;
} ToBeSignedRegionIdList;
```

The usage of region ID lists is described in section 4.2.5.

4.2.7.6 RegionIdListEntry

```
struct {
    uint32<var>           region_id;
    PolygonalRegion           polygonal_region;
} RegionIdListEntry;
```

The usage of region ID lists is described in section 4.2.5.

4.2.7.7 CertificateRequestFlags

```
flags { use_resolution_id (0), use_integrity_data (1)
} CertificateRequestFlags;
```

The `CertificateRequestFlags` are used to determine whether additional information is added to certificate requests (e.g. integrity data in a `ToBeSignedLongTermCertificateRequest` or resolution information in a `ToBeSignedAuthorizationValidationRequest`. Both flags are optional but the introduction of CertificateRequestFlags enables future extensibility.

4.2.7.8 MessageCsrScope

```
struct {
    uint8       name<var>;
    SubjectTypeFlags    permitted_subject_types;
    PsidSspArray    permissions;
    GeographicRegion    region;
} MessageCsrScope;
```

The new `MessageCsrScope` is required for long-term certificates. In contrast to the other scopes defined in [4] (i.e. `MessageCaScope`), the LTC has to specify a `PsidSspArray` and `permitted_subject_types` in order to permit the request of PCs with respective permissions. As discussed in section 4.2.3, the LTC should contain the permissions that are used to set the permissions of requested PCs.

4.2.7.9 CertSpecificData

```
struct {
    extern SubjectType      subject_type;
    select(subject_type){
    case root_ca:
        RootCaScope     root_cascope;
// new material
    case message_ca:
        MessageCaScope  message_ca_scope;
    case message_csr:
        MessageCsrScope     message_csr_scope;
// end new material
    case wsa_ca, wsa_csr:
        WsaCaScope      wsa_ca_scope;
    case crl_signer:
        CrlSeries       responsible_series<var>;
```

```
    case message_identified_not_localized:
        IdentifiedNotLocalizedScope idnonloc_scope;
    case message_identified_localized,
        IdentifiedScope     id_scope;
    case message_anonymous:
        AnonymousScope  anonymous_scope;
    case wsa:
        WsaScope        wsa_scope;
    unknown:
        opaque      other_scope<var>;
    }
} CertSpecificData;
```

We propose the distinction between `MessageCaScope` as defined in [4] and a `MessageCsrScope` for the LTC.

4.2.7.10 SignerIdentifierType

```
enum {
    self (0) ,
    certificate_digest_with_ecdsap224 (1),
    certificate_digest_with_ecdsap256 (2),
    certificate(3),
    certificate_chain (4),
    certificate_digest_with_other_algorithm (5),
// new material
    encrypted (6),
    id_with_ecdsap256 (7),
// end new material
    reserved (240..255),
(2^8-1)} SignerIdentifierType;
```

Two new signer identifier types are introduced.

- The type `encrypted` is used when an ITS station requests new pseudonyms from a PCA. In this case, the request is signed using the private long-term signature key of the ITS-S (the private key corresponding to the signature key of the LTC), but the PCA should not be able to read this long-term identifier. As result, the `SignerIdentifier` is encrypted with the LTCA's public encryption key.
- The second type `id_with_ecdsap256` is used in the `LongTermCertificateRequest` when the signer of the request is to be identified by a canonical identifier that has been used to register the respective ITS-S. During registration of the ITS-S at the LTCA, this canonical identifier is stored together with the public key that has to be used for verifying signatures in case the `SignerIdentifier` is of type `id_with_ecdsap256`.

4.2.7.11 SignerIdentifier

```
struct {
    SignerIdentifierType type;
    select (type) {
        case self: ;
        case certificate_digest_with_ecdsap224 :
        case certificate_digest_with_ecdsap256 :
            CertId8 digest;
        case certificate:
            Certificate certificate;
        case certificate_chain:
            Certificate certificates<var>;
        case certificate_digest_with_other_algorithm :
            PKAlgorithm algorithm;
            CertId8 digest;
// new material
        case encrypted:
            PKAlgorithm algorithm;
            EncryptedMessage encrypted_signer_identifier;
        case id_with_ecdsap256:
// end new material
        unknown:
            opaque id<var>;
    }
} SignerIdentifier;
```

Two new signer identifier types are introduced.

- The encrypted signer identifier is used to hide the signer ID against the PCA in case of requesting a pseudonym certificate. Only the LTCA should be able to identify the requester by decrypting the SignerIdentifier. A `PKAlgorithm` is used to state which verification operation has to be executed at the receiver. If `ecdsa_nistp224_without_hash` or `ecdsa_nistp256_without_hash` is given as `PKAlgorithm` then the receiving station must omit the hashing process in verification process. In the specific case of requesting a PC, the PCA has to create a hash value over the `ToBeSignedPseudonymCertificateRequest` and provide this to the LTCA. Only the LTCA is able to verify the signature by using the verification key of the LTC and the hash that is provided by the PCA.
- When `id_with_ecdsap256` is given as `SignerIdentifierType`, an identifier is used with variable length. This type is used in the `LongTermCertificateRequest` to identify the canonical identifier of the ITS-S that was used in the bootstrapping process.

4.2.7.12 PKAlgorithm

```
enum {  ecdsa_nistp224_with_sha224 (0),
        ecdsa_nistp256_with_sha_256 (1),
        ecies_nistp256 (2),
// new material
        ecdsa_nistp224_without_hash (3),
        ecdsa_nistp256_without_hash (4),
// end new material
        reserved (240..255), (2^8-1)
} PKAlgorithm;
```

The new types `ecdsa_nistp224_without_hash` and `ecdsa_nistp256_without_hash` are used in a pseudonym certificate request. When the signer is encrypted then the PCA is not able to verify the request. In this case, the PCA creates the hash of the request and sends it to the LTCA. The LTCA can identify the signer and verify the request but without hashing the given data, which is indicated by the new types `ecdsa_nistp224_without_hash` and `ecdsa_nistp256_without_hash` contained in the `SignerIdentifier`.

4.2.7.13 ContentType

```
enum {  unsecured (0), signed(1), encrypted (2),
        certificate_request(3), certificate_response(4),
        anonymous_certificate_response(5),
        certificate_request_error(6), crl_request(7),
        crl(8),
        signed_partial_payload(9),
        signed_external_payload(10),
        signed_wsa(11),
        certificate_response_acknowledgment (12),
// new material
        certificate_retrieval_request (225),
        certificate_retrieval_response (226),
        long_term_certificate_request (227),
        long_term_certificate_response (228),
        pseudonym_certificate_request (229),
        pseudonym_certificate_response (230),
        decryption_key_request (231),
        decryption_key_response (232),
        authorization_validation_request (233),
        authorization_validation_response (234),
        certificate_request_signer (235),
        crl_req (236),
        crl_req_error (237),
```

```
        misbehavior_report_req (238),
        misbehavior_report_ack (239),
// end new material
        reserved (240...255), (2^8-1)
} ContentType;
```

We propose the following new message types:

- The `certificate_retrieval_request` and `certificate_retrieval_response` is used to request certificates from the ITS stations. For example, vehicles could request the current CA certificate from the PKI entities. Additionally, a CA can request a specific certificate from another CA of the PKI or a vehicle requests the pseudonym certificate from another vehicle.
- The `long_term_certificate_request` and `long_term_certificate_response` is used by vehicles or RSUs to request new long-term certificates from the LTCA.
- The `pseudonym_certificate_request` and `pseudonym_certificate_response` is used by vehicles or RSUs to request new pseudonym certificates from the PCA.
- The `decryption_key_request` and `decryption_key_response` is used by vehicles or RSUs to request a decryption key for a bunch of PCs that are encrypted by the LTCA. This feature is currently not implemted.
- The `authorization_validation_request` and `authorization_validation_response` is used by the PCA to request authorization from the LTCA to issue new pseudonym certificates.
- The `certificate_request_signer` is used in a `ToBeEncrypted` structure when the `SignerIdentifier` is encrypted. This is used by a ITS-S when a new PC is requested and only the LTCA should be able to read the `certID8` of the LTC.
- The types `crl_req` and `crl_req_error` are used to request revocation lists from the CA. Only CA certificates can be part of a revocation list. Revocation of long-term certificates and pseudonym certificates is not planned in PRESERVE. These types are not implemented.
- The `misbehavior_report_req` and `misbehavior_report_ack` is used by vehicles or RSUs to send reports to the infrastructure indicating possible misbehavior. This feature is currently not used.

4.2.7.14 PublicKey

```
struct {
    PKAlgorithm algorithm;
    select(algorithm){
```

```
        case ecdsa_nistp224_with_sha224:
        case ecdsa_nistp256_with_sha256:
// new material
        case ecdsa_nistp224_without_hash:
        case ecdsa_nistp256_without_hash:
// end new material
            EccPublicKey public_key;
        case ecies_nistp256:
            SymmAlgorithm supported_symm_alg;
            EccPublicKey public_key;
        unknown:
            opaque other_key<var>;
    }
} PublicKey;
```

See description of new public key algorithms in section 4.2.7.13.

4.2.7.15 ToBeEncrypted

```
struct {
    ContentType type;
    select(type) {
        case unsecured:
            opaqueExtLength plaintext;
        case signed, signed_external_payload,
            signed_partial_payload :
            SignedMessage signed_message;
        case certificate_request :
CertificateRequest request;
// new material
        case certificate_response, long_term_certificate_response :
// end new material
            ToBeEncryptedCertificateResponse response;
        case anonymous_certificate_response :
            ToBeEncryptedAnonymousCertResponse anon_response;
        case certificate_request_error:
            ToBeEncryptedCertificateRequestError request_error;
        case crl_request :
            CrlRequest crl_request;
        case crl :
            Crl crl;
        case certificate_response_acknowledgment:
            ToBeEncryptedCertificateResponseAcknowledgment ack;
// new material
```

```
        case long_term_certificate_request :
            LongTermCertificateRequest long_term_cert_request;
        case pseudonym_certificate_request:
            PseudonymCertificateRequest pseudonym_cert_request;
        case pseudonym_certificate_response:
        ToBeEncryptedPseudonymCertificateResponse
          pseudonym_cert_response;
        case certificate_request_signer :
            SignerIdentifier request_signer;
        case authorization_validation_request :
        AuthorizationValidationRequest
          authorization_validation_request;
        case authorization_validation_response :
        AuthorizationValidationResponse
          authorization_validation_response;
// end new material
        unknown:
            opaque message<var>;
    }
} ToBeEncrypted;
```

See description of new types in section 4.2.7.13.

4.2.7.16 LongTermCertificateRequest

```
struct {
    SignerIdentifier            request_signer;
    ToBeSignedCertificateRequest    unsigned_csr;
    signature                request_signature;
} LongTermCertificateRequest;
```

This new structure is used by an ITS-S to request a new long-term certificate from the LTCA. This structure is based on the `CertificateRequest` as described in [4].

4.2.7.17 ToBeSignedLongTermCertificateRequest

```
struct {
    uint8       version_and_type;
    Time32      request_time;
    SubjectType     subject_type;
    CertificateContentFlags cf;
    CertificateRequestFlags crf;
    CertSpecificData    type_specific_data;
```

```
    Time32      expiration;
    if_set(cf, use_start_validity) {
        if_set(cf, lifetime_is_duration) {
            CertificateDuration lifetime;
        }
        if_not_set(cf, lifetime_is_duration) {
            Time32 start_validity;
        }
    }
    PublicKey       verification_key;
    if_set (cf, encryption_key) {
        PublicKey   encryption_key;
    }
    if_set (crf, use_integrity_data) {
        opaque  integrity_data<var>;
    }
    if_value_not_in (cf,
        use_start_validity, encryption_key) {
        opaque  other_cert_content<var>;
    };
    PublicKey       response_encryption_key;
} ToBeSignedLongTermCertificateRequest;
```

This new structure is part of a `LongTermCertificateRequest` that is used to request a new long-term certificate from the LTCA. This structure is based on the `ToBeSignedCertificateRequest` structure as described in [4]. New elements and differences to the ToBeSignedCertificateRequest are the following:

- `CertificateRequestFlags` are added to indicate whether integrity data of the requesting system are added. This information should be optional. According to [6], a not further specified kind of integrity data could be used in the process of requesting long-term credentials from the PKI.
- The `integrity_data<var>` can be optionally used to provide integrity information of ITS-S system software or hardware.

4.2.7.18 ToBeSignedPseudonymCertificateRequest

```
struct {
    uint8                   version_and_type;
    Time32                  request_time;
    SubjectType             subject_type;
    CertificateContentFlags cf;
    CertificateRequestFlags crf;
    CertSpecificData    type_specific_data;
```

```
    Time32              batch_expiration;
    if_set(cf, use_start_validity) {
        if_set(cf, lifetime_is_duration) {
            CertificateDuration batch_lifetime;
        }
        if_not_set(cf, lifetime_is_duration) {
            Time32  batch_start_validity;
        }
    }
    PublicKey           verification_keys<var>;
    if_set(cf, encryption_key) {
        PublicKey       encryption_keys<var>;
    }
    if_set (crf, use_integrity_data) {
        opaque  integrity_data<var>;
    }
    if_value_not_in (cf,
        use_start_validity, encryption_key) {
        opaque      other_cert_content<var>;
    }
    PublicKey               response_encryption_key;
} ToBeSignedPseudonymCertificateRequest;
```

This new structure is part of a `PseudonymCertificateRequest` that is used to request a new pseudonym certificate from the PCA. This structure is based on the `ToBeSignedCertificateRequest` structure as described in [4]. New elements and differences to the `ToBeSignedCertificateRequest` are the following:

- `CertificateRequestFlags` are added to indicate whether integrity data of the requesting system are added. This information should be optional.
- A bunch of PCs can be requested with this structure by using the following fields:
  - The `batch_expiration`, `batch_lifetime` and `batch_start_validity` entries are used for all public keys that should be used in the new PCs. Only one start / duration and one expiration time for all requested PCs. The lifetimes of a single pseudonym certificate will be set by the LTCA verifying the request. If special validity times for single PCs are desired, separate PseudonymCertificateRequests have to be created and sent to the PCA.
  - In the field `verification_keys<var>` a variable number of public keys can be added.
  - In the field `encryption_keys<var>` a variable number of public keys can be added.
  - In the field `other_cert_content<var>` a variable number of other certificate content can be added.

- The integrity_data<var> can be optionally used to provide integrity information of ITS-S system software or hardware.

4.2.7.19 PseudonymCertificateRequest

```
struct {
    SignerIdentifier        info;
    ToBeSignedPseudonymCertificateRequest unsigned_request;
    Signature           signature;
} PseudonymCertificateRequest;
```

This new structure is used by an ITS-S to request a new pseudonym certificate from the PCA. This structure is based on the CertificateRequest as described in [4].

4.2.7.20 ToBeEncryptedPseudonymCertificateResponse

```
struct {
    extern uint8    version_and_type;
    flags           f;
    Certificate     pseudonym_certificates<var>;
    Certificate     certificate_chain<var>;
} ToBeEncryptedPseudonymCertificateResponse;
```

This new structure is part of an EncryptedMessage that is used by the PCA in order to transmit the issued pseudonym certificates. This structure is based on the ToBeEncryptedCertificateResponse structure as described in [4]. New elements and differences to the ToBeEncryptedCertificateResponse are the following:

- A bunch of PCs can be responded with this structure by using the following fields:
  - The certificate_chain<var> a variable number of pseudonym certificates can be added.
- The field pseudonym_certificates<var> contains the certificate path of the new PCs. This path is in order, with the most local certificate (the newly issued one) being last and each preceding certificate signing the one before it. The path should be complete with the first certificate being a trust anchor.

4.2.7.21 ToBeSignedAuthorizationValidationRequest

```
struct {
    opaque      certificate_request_hash_value[32];
    SignerIdentifier    certificate_request_signer;
    Signature       certificate_request_signature;
    SubjectType     subject_type;
    CertificateContentFlags cf;
    CertificateRequestFlags crf;
    CertSpecificData    type_specific_data;
    uint32<var>     number_requested_certificates;
    Time32      expiration;
    if_set(cf, use_start_validity) {
        if_set(cf, lifetime_is_duration) {
            CertificateDuration lifetime;
        }
        if_not_set(cf, lifetime_is_duration) {
            Time32 start_validity;
        }
    };
    if_set(crf, use_resolution_id) {
        CertId8     resolution_ids<var>;
    }
    if_set (crf, use_integrity_data) {
        opaque  integrity_data<var>;
    }
    if_value_not_in (cf,
        use_start_validity, encryption_key) {
        opaque  other_cert_content<var>;
    };
} ToBeSignedAuthorizationValidationRequest;
```

This new structure is part of an `AuthorizationValidationRequest` that is used by the PCA to request the authorization and validation from a LTCA to issue new PCs. This structure is based on the `ToBeSignedCertificateRequest` as described in [4].

- The `certificate_request_hash_value` contains the 32 bytes of the hash (SHA256) over the `ToBeSignedPseudonymCertificateRequest` from a `PseudonymCertificateRequest` that is received by a PCA. If the `SignerIdentifier` of the PC request is encrypted, the PCA has to request the authorization and validation from the LTCA but the LTCA should not be able to read the public keys of the requested PCs. Therefore, the PCA creates that hash over the request and send this to the LTCA for verification.
- The `certificate_request_signer` contains the (encrypted) signer information of the ITS-S that requests new pseudonym certificates.

- The `certificate_request_signature` contains the signature over the `ToBeSignedPseudonymCertificateRequest` from a `PseudonymCertificateRequest` that is received by a PCA.
- The `subject_type` shall be copied from the `ToBeSignedPseudonymCertificateRequest`.
- The `cf` shall be copied from the `ToBeSignedPseudonymCertificateRequest`.
- The `crf` shall be copied from the `ToBeSignedPseudonymCertificateRequest`.
- The `type_specific_data` shall be copied from the `ToBeSignedPseudonymCertificateRequest`.
- The `number_requested_certificates` is the number of public keys contained in the `ToBeSignedPseudonymCertificateRequest`.
- The `expiration` shall be copied from the `ToBeSignedPseudonymCertificateRequest`.
- The `lifetime` shall be copied from the `ToBeSignedPseudonymCertificateRequest`.
- The `start_validity` shall be copied from the `ToBeSignedPseudonymCertificateRequest`.
- The `resolution_ids<var>` is optional and can be used to provide resolution.
- The `integrity_data<var>` is optional and can be used to provide integrity information of ITS-S system software or hardware.

4.2.7.22 AuthorizationValidationRequest

```
struct {
    SignerIdentifier info;
    ToBeSignedAuthorizationValidationRequest
      unsigned_auth_valid_request;
    Signature       signature;
} AuthorizationValidationRequest;
```

This new structure is used by the PCA to request the authorization and validation from a LTCA to issue new PCs. This structure is based on the `CertificateRequest` as described in [4].

4.2.7.23 ApprovedCertificateTime

```
struct {
    extern CertificateContentFlags cf;
    uint32<var> number_approved_certificates;
    Time32      expiration;
    if_set(cf, use_start_validity) {
        if_set(cf, lifetime_is_duration) {
            CertificateDuration lifetime;
        }
        if_not_set(cf, lifetime_is_duration) {
            Time32  start_validity;
        }
    }
} ApprovedCertificateTime;
```

This new structure defines an approval to issue a certain number of certificates with certain content flags in the specified time. This structure is to be used as a list item within the structure `ToBeSignedAuthorizationValidationResponse`.

- The content flags `cf` describe which content flags the PCA may include in the issued certificate.
- The field `number_approved_certificates` is the number of certificates that may be issued by the PCA with the validity period defined by `expiration` and `lifetime` or `start_validity`.
- In the current specification, the flag `use_start_validity` must be set. For possible changes in the future, this structure, also allows specifying a validity period without a start time.

4.2.7.24 ToBeSignedAuthorizationValidationResponse

```
struct {
    opaque          certificate_request_hash_value[32];
    SubjectType         subject_type;
    CertificateContentFlags     cf;
    CertSpecificData    type_specific_data;
    ApprovedCertificateTime     approved_certificates<var>;
} ToBeSignedAuthorizationValidationResponse;
```

This new structure is used by the LTCA to respond to an `AuthorizationValidationRequest` by the PCA. This structure is based on the `CertificateRequest` as described in [4].

- • The `certificate_request_hash_value[32]` is used to map the received response to the `AuthorizationValidationRequest` at the PCA.

- • The `subject_type`, `cf` and `type_specific_data` are used to tell the PCA which content may be written into the issued pseudonym certificates.
- • The `list approved_certificates<var>` contains all possible validity periods that the PCA may issue certificates for. The PCA shall issue certificates for any of the given periods as the LTCA assumes the successful issuance for all given validity periods.

4.2.7.25 AuthorizationValidationResponse

```
struct {
    SignerIdentifier    info;
    ToBeSignedAuthorizationValidationResponse
      unsigned_auth_valid_response;
    Signature       signature;
} AuthorizationValidationResponse;
```

This new structure is used by the LTCA to respond to an authorization and validation request from a PCA during processing of a pseudonym certificate request. This structure is based on the `CertificateRequest` as described in [4].

4.2.7.26 1609Dot2Message

```
struct {
    uint8   protocol_version;
    ContentType type;
    select (type) {
    case unsecured :
        opaque data<var>;
    case signed, signed_partial_payload, signed_external_payload:
        SignedMessage signed_data;
    case signed_wsa:
        SignedWsa signed_wsa;
    case encrypted :
        EncryptedMessage encrypted_data;
    case crl_request :
        CrlRequest crl_request;
    case crl :
        Crl crl
// new material
    case certificate_retrieval_request:
        CertificateRetrievalRequest cert_retrieval_req;
    case certificate_retrieval_response:
        CertificateRetrievalResponse cert_retrieval_res;
```

```
// end new material
    unknown:
        opaque data<var>;
    }
} 1609Dot2Message;
```

See the message structures above for descriptions of the new message types `certificate_retrieval_request` and `certificate_retrieval_response`.

4.2.7.27 CertificateRetrievalFlags

```
flags { use_cert_id (0) use_signature (1)
} CertificateRetrievalFlags;
```

The formats for certificate retrieval are new and not present in [4]. The formats can be used to get certificates from ITS stations on demand. The `CertificateRetrievalFlags` are used to indicate whether a specific certificate is requested in a `CertificateRetrievalRequest` or if a `CertificateRetrievalResponse` is to be signed.

4.2.7.28 CertificateRetrievalRequest

```
struct {
    CertificateRetrievalFlags crf;
    if_set(crf, use_cert_id) {
        CertId8 requesting_cert;
    }
} CertificateRetrievalRequest;
```

The formats for certificate retrieval are new and not present in [4]. The formats can be used to get certificates from ITS stations on demand. The `CertificateRetrievalRequest` can be used to request a certificate from an ITS station. If the no `CertId8` is set, indicated by the `CertificateRetrievalFlags`, then the main certificate shall be responded. For the PKI entities, the following main certificates shall be responded if no is CertId8 given:

- The RCA provides its own currently valid root certificate
- The LTCA provides its own currently valid LTCA certificate
- The PCA provides its own currently valid PCA certificate
- If the CertId8 is given, the ITS station shall provide the requested certificate as `CertificateRetrievalResponse` as long as local policies are not violated. For example, the LTCA or PCA may not be allowed to provide issued long-term certificates or pseudonym certificates.

4.2.7.29 CertificateRetrievalResponse

```
struct {
    CertificateRetrievalFlags crf;
    if_set(crf, use_signature) {
        SignerIdentifier    info;
        Signature       signature;
    }
    Certificate         certificate_chain<var>;
    Crl                 crl_path<var>;
} CertificateRetrievalResponse;
```

The formats for certificate retrieval are new and not present in [4]. The formats can be used to get certificates from ITS stations on demand. The `CertificateRetrievalResponse` format is based on `ToBeEncryptedCertificateResponse` type of 1609.2.

- The `CertificateRetrievalFlags` are used to define whether the `CertificateRetrievalResponse` contains a `SignerIdentifier` and a `Signature`.
- `SignerIdentifier` and `Signature`: Normally, it is not needed to sign a certificate response as the certificate itself is signed by the issuer. However, in the specific case that a root certificates expires, a new root certificate is created that overlaps the validity time of the old root certificate. In this case, the old root certificate can be used to sign the request of the new root certificate.
- A chain of certificates can be added to the `CertificateRetrievalResponse` in order to provide all certificates up to the root certificate. This path is in order, with the newly issued certificate being last and each preceding certificate signing the one before it. The path should be complete with the first certificate being a trust anchor
- If one of the provided certificates in the certificate chain could appear on a CRL then the most recent version of the CRL series on which the issued certificate would appear shall be added. Please notice that only CA certificates may appear on a CRL.

4.2.7.30 CertificateRequestErrorCode

```
enum {  verification_failure(0), csr_cert_expired(1),
csr_cert_revoked(2), csr_cert_unauthorized(3),
request_denied(4), csr_cert_unknown (5),
canonical_identity_unknown (6),
// new material
        ca_not_available(7), message_processing_error(8),
        request_in_process(9), message_parsing_error(10),
// end new material
(255)
```

```
} CertificateRequestErrorCode;
```

The following additional error codes are introduced in PRESERVE to provide more detailed information to the ITS-S in case of a failed certificate request:

- `ca_not_available`: This error code is returned by a PCA if another CA that is required for processing the request is not available. For example, a PCA requires communication to an LTCA if the signer identifier of a pseudonym certificate request is encrypted. If the required LTCA is not available, the PCA shall return `ca_not_available`.
- `message_processing_error`: An internal error happened during processing the message at the CA.
- `request_in_process`: For long-term certificate requests, the LTCA shall return this error code if a new long-term certificate request is received while another long-term certificate request for the same ITS-S is still being processed. For pseudonym certificate requests, the LTCA shall return this error code to the PCA if a new authorization validation request from a PCA is received while another authorization validation request for the same ITS-S is still being processed. The PCA shall send an error message with this error code to the requesting ITS-S if it receives an authorization validation response from the LTCA with this error code.
- `message_parsing_error`: The received message could not be parsed.

5 Cost model for ASIC development

A cost model is created to relate functionality of an ASIC to its costs which is very useful during early phases of the chip development. However, creating such a cost model for the production of an ASIC-based C2C-HSM is a difficult task, as it is based on two parameters which are in principle unknown:

- **Performance**: the performance of an ASIC chip can only be estimated until an ASIC is actually produced. Such estimations of the performance depend on many different factors, but may be given based on previous experience with similar technologies. As the number of verifications per second is the key performance factor for ASICs in a C2C environment, we will use the verification speed as an indicator for the overall performance of the ASIC.
- **Absolute costs**: absolute costs of ASIC production depend on many factors such as produced quantities, design size, supported features, technologies, customer-supplier relationship and many more. In short, an OEM ordering an ASIC highly specialized for a certain use-case in large quantities (millions) for series production will get a totally different price than a smaller organization producing only small quantities of research ASICs. Hence, a cost model including absolute costs is not really meaningful and we will concentrate on relative costs instead.

Based on these difficult preconditions, we try to give numbers and estimations for the given parameters to the best of our knowledge in the following.

5.1 Performance

Assuming that only one ECC core is implemented, the key factor for the verification speed is the technology (node size) in use. Generally speaking, a smaller gate size yields higher clock rates of the chip and thus better performance in terms of verification speed.

The verification speed can also be improved by implementing more than one ECC core in the chip design which can be used in parallel. However, the number of verifications does not scale linearly with the number of ECC cores, as there are several other limiting factors, e.g.:

- Busload on the AHB bus
- AHB bus frequency
- System software complexity

The more ECC cores are running in parallel, the more weight will fall to these limiting factors. If, e.g., the busload is already fully exceeded, adding additional ECC cores will not add any additional verification performance. The number of ECC cores that can be implemented is also limited by the number of gates available on the chip. Using a smaller technology will result in a higher number of gates on a chip of the same size. For example, on a chip of size 4mm x 4mm we can estimate the following numbers:

- ASIC 180nm: approx. 1.4 million gates
- ASIC 90nm: approx. 3 million gates
- ASIC 55nm: approx. 8 million gates

Using the 180nm technology will only yield enough space for one ECC core, whereas 90nm will allow for up to ten ECC cores and 55nm will allow for even more. Of course, this also depends heavily on the remaining components on the chip (e.g. RAM, interfaces, other cores) and how much chip space they require. Furthermore, we assume that more than 10 ECC cores are not reasonable with respect to the limiting factors.

Based on these numbers, we estimated the maximal numbers of verifications per second that can be achieved with a highly specialized and optimized chip design. As mentioned, these are only estimations and concrete numbers can only be given once an ASIC is produced and tested. The results can be seen in Table 5.1.

Technology	**Max clock rate**	**Verifications per second with**		
		1 ECC	**5 ECC**	**10 ECC**
ASIC 180nm	200 MHz	100	-	-
ASIC 90nm	400 MHz	200	750	1100
ASIC 55nm	700 MHz	320	1200	1760

Table 5.1: ASIC performance estimation

5.2 Relative costs

The costs stated in this section are relative costs based on evaluations done within the PRESERVE project. They are useful to compare different options/technologies and show, how different performance requirements on the one hand are reflected in the costs/prices on the other hand.

At the center of the cost estimation is the slowest option (option 1), i.e. the 180nm technology with only one ECC. The costs of the other options are then given as additional costs relative to option 1. We distinguish the following categories of costs for the ASIC production:

- Fixed costs: only applicable once in the production process which are mostly given by the following two items

  - Design costs (frontend and backend design)
  - Prototyping costs (production of prototype/silicon mask and first shuttle)
- Costs per item: costs for each additional unit that is being produced

Altogether, the considerations result in the following cost model described in Table 5.2.

Opt.	Verifications/s	Technology	ECCs	Cost relative to option 1		
				Design	Prototype	Item Costs
1	100	180nm	1	0	0	0
2	200	90nm	1	+ 9 %	+ 175 %	+ 83 %
3	320	55nm	1	+ 22 %	+ 175 %	+ 116 %
4	750	90nm	5	+ 30 %	+ 175 %	+ 83 %
5	1100	90nm	10	+ 51 %	+ 175 %	+ 83 %
6	1200	55nm	5	+ 43 %	+ 175 %	+ 116 %
7	1760	55nm	10	+ 64 %	+ 175 %	+ 116 %

Table 5.2: ASIC cost model

Analyzing the above cost model, one will find many interesting aspects. Of course, option 1 is the cheapest, but offers also the weakest performance. This is only an option for validation purposes, but not for applications in realistic C2C environments. The other options offer more possibilities in these terms. However, moving to a smaller technology will increase all costs items. While prototype costs will be equal for 90nm and 55nm, design costs and costs per item will increase significantly for a smaller gate size. The number of ECCs does only influence the design costs, as more ECCs result in a bigger design and thus in higher design efforts.

An interesting aspect can also be found by comparing options 5 and 6, since both result in a similar performance, but different prices. While option 5 uses a bigger technology and a bigger design, option 6 makes use of a higher clock rate. With the lower design costs and slightly better performance, option 6 is a good choice for a research environment. Yet on the other hand, it also comes with higher costs per item and thus option 5 is more suitable for a mass production environment.

6 Research

In this section, we detail the research work done during the year 2 of PRESERVE. First, data plausibility checks are investigated in Sections 6.1-6.2. Indeed, plausibility check is a security mechanism that is generally used for every message received, and therefore, the scalability of such mechanism has to be analyzed. Then, a risk analysis of ITS is proposed in Sections 6.3-6.4. Also related to scalable security mechanism, a certificate omission scheme is proposed in Section 6.5.1. As a balance between security and privacy has to be found, we compare the pseudonym schemes and draw conclusions and future challenges in Section 6.6-6.7.

6.1 Assessment of Node Trustworthiness in VANETs Using Data Plausibility Checks with Particle Filters

In Vehicular Ad-Hoc Networks (VANETs), the exchange of location data (i.e. absolute position, heading, time) for traffic safety applications plays an important role. The trustworthiness of this information is crucial as false data affects applications heavily and might endanger human lives. Beside cryptographic solutions that ensure sender authenticity and message integrity, the data plausibility check is an important mechanism to ensure positional reliability. In this paper, we show that a particle filter is an appropriate instrument to perform plausibility checks in order to assess the trustworthiness of neighbor nodes. Our approach allows the aggregation of information from different data sources directly in one particle filter per neighbor. Thus, dependencies and relationships between individual sources can be fully accounted for and the framework is easily extensible and scales well. The concept is implemented as a Java-OSGi bundle for a field operational test framework and evaluated using both manually generated traces and recorded data from real vehicle trips. We show that the detection of several types of location-based attacks is possible under consideration of errors and system inherent deviations in sensor data.

In V2X communication cryptographic mechanisms allow basic authentication and message integrity verification, but beyond that, the detection of faulty nodes and malicious attackers is still a huge security challenge. Internal attackers that are in possession of valid cryptographic credentials are able to distribute bogus V2X messages in order to exploit the future Intelligent Transportation System (ITS) or simply to disturb its operation.

In V2X, mobility information is broadcasted periodically to all VANET nodes in communication range at a rate of up to 10 Hz using, e.g., Cooperative Awareness Messages (CAMs) [9]. To address the risks caused by nodes sending forged information a misbehavior detection framework is proposed based on data plausibility checks. It verifies all

incoming mobility information (i.e. absolute GPS position, heading, velocity, timestamp) received with V2X messages. In order to increase the quality of plausibility checking, the framework leverages on different independent information sources that confirm or disprove a specific situation. The position of a neighboring vehicle is for example verified using received mobility data from CAMs sent by the target node and other neighbors, as well as data from vehicle-local sensors (e.g. digital road map, Radar, Lidar, directional antennas). Earlier approaches deployed separate rating modules to process the different information sources as proposed, e.g., in [10–13]. In this paper we show that the probabilistic particle filter [14,15] is an appropriate instrument to implement data plausibility checks in VANETs. On the one hand, we are able to combine all available different location information from a broad variety of input sources using only one instance of a particle filter per single-hop neighbor node. On the other hand, our approach benefits from the possibility to directly assesses the trustworthiness of these nodes as shown in Fig. 6.1.

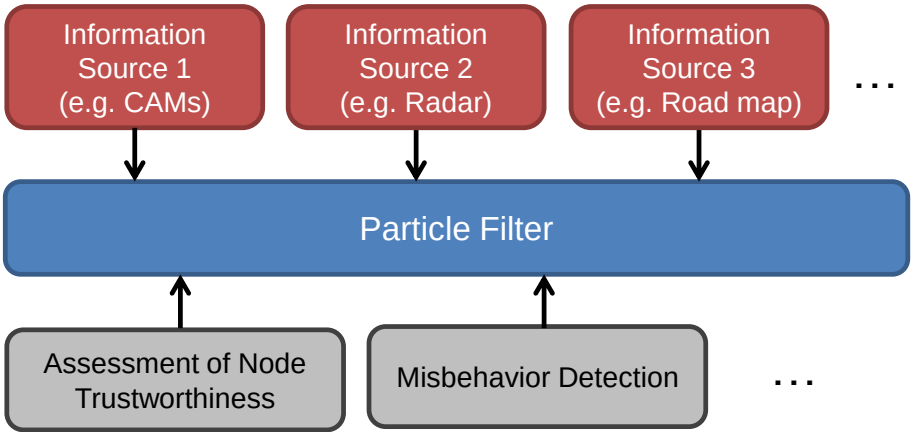

Figure 6.1: Data source aggregation and node trustworthiness assessment for data plausibility checking in VANETs using a particle filter

Our strategy allows the weighting of data sources and avoids duplicating information by having only one particle filter per neighbor node. Additionally, we avoid a complex aggregation of information sources and intermediary results that may lead to mutual interference. Nonetheless, the particle filter based plausibility check can be flexibly extended in order to detect inconsistencies introduced by faulty nodes or malicious attackers.

Our approach provides a local consistency check of location-based data in each vehicle. This contrasts with other proposals [16] where vehicles only report information to a central entity that applies then a global consistency checking. There are two main reasons why we advocate the application of data plausibility checks in all vehicles:

1. local applications are directly able to decide whether information from untrustworthy neighbors should be handled with caution without referring to a backend service that might not be reachable at that moment and
2. detected inconsistencies can be filtered much better before being reported to a global misbehavior evaluation authority that would then identify and revoke faulty nodes and attackers.

6.1.1 Adversary Model

We consider an insider attacker that is able to tamper with a vehicle's Communication and Control Unit (CCU) or Application Unit (AU) to make it send V2X messages (e.g. CAMs) that have a valid signature and certificate but where the attacker can still forge the content of the message. An attacker may also be able to extract valid credentials from vehicle CCUs or AUs and then use other devices to send forged messages. By just modifying the mobility data (i.e. absolute position, heading, time) contained in V2X messages, the attacker is able to perform a wide range of location based attacks.

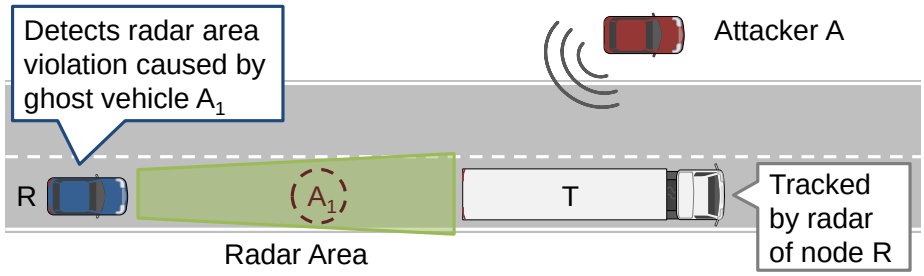

Figure 6.2: Attacker model used as basis for mobility data plausibility checks

As shown in Fig. 6.2, attacker A can misuse the vehicular communication system by sending messages with false mobility data in order to create a non-existent "ghost" vehicle A_1 on the road. We assume that an attacker may even be able to use different identifiers in parallel to create the illusion of several ghost vehicles. As a result, other vehicles (e.g. node R) would assume that a real vehicle A_1 is present at the specified location and time which could trigger false application warnings or other undesired and potentially dangerous system behavior. Besides such roadside attackers with static sending positions, we also assume mobile attackers. However, in both cases, there is a chance that ghost vehicles cause inconsistencies with other information received either in messages from other vehicles or via local sensors. Fig. 6.2 shows an example where A_1 creates a radar area violation. Vehicle R receives a radar echo from truck T driving in front and is able to determine a radar detection area. Other inconsistencies could be caused by A_1 due to position overlaps with real vehicles on the road, sudden node appearances, map violations or positional jumps. Our approach aims to detect all such inconsistencies based on individual sensors.

6.1.2 Position Tracking with Particle Filters

Particle filters belong to the family of Bayesian filters and consist in general of predict / update cycles that are performed repeatedly to estimate the state of a dynamic system from sensor measurements [15]. The first step is the prediction, where a new believe state is calculated based on the prior believe state and a control/input induced transition. The second step is the so called measurement update. Here, the predicted estimate is corrected using sensor observations. The basic idea of particle filters is that any Probability Density

Function (PDF) can be approximated by a set of samples. With a sufficient amount of samples, the density of samples in a given area represents the probability of that area. With particle filters, each sample is represented by a particle, containing a whole set of state variables. This enables the sampling of arbitrary density functions and therefore of several complex models.

For the proposed mobility data consistency check a particle filter algorithm using the common Sequential Importance Resampling (SIR) approach is chosen [15]. Each particle $x_t^{[m]}$ is a concrete instantiation of the system state at a time t and represents a sample of the posterior distribution. χ_t is the particle set at time t containing all particles $x_t^{[m]}$ (with $1 \leq m \leq M$) of that time step where M denotes the total number of particles. The belief $bel(x_t)$ reflects the internal knowledge about the state of the environment or the system. In particle filters, the belief is represented by the posterior distribution which is approximated by the set of samples χ_t. For a transition from one belief state to another, it is required that a new control information u_t is available. The transition itself is described by a state transition distribution $p(x_t|u_t, x_{t-1})$. The likelihood for a state hypothesis x_t to be included in the particle set χ_t should be proportional to this distribution.

The algorithm depicted in Fig. 6.3 takes a set of particles χ_{t-1} together with the most recent control information u_t and the most recent sensor measurement z_t as an input. If

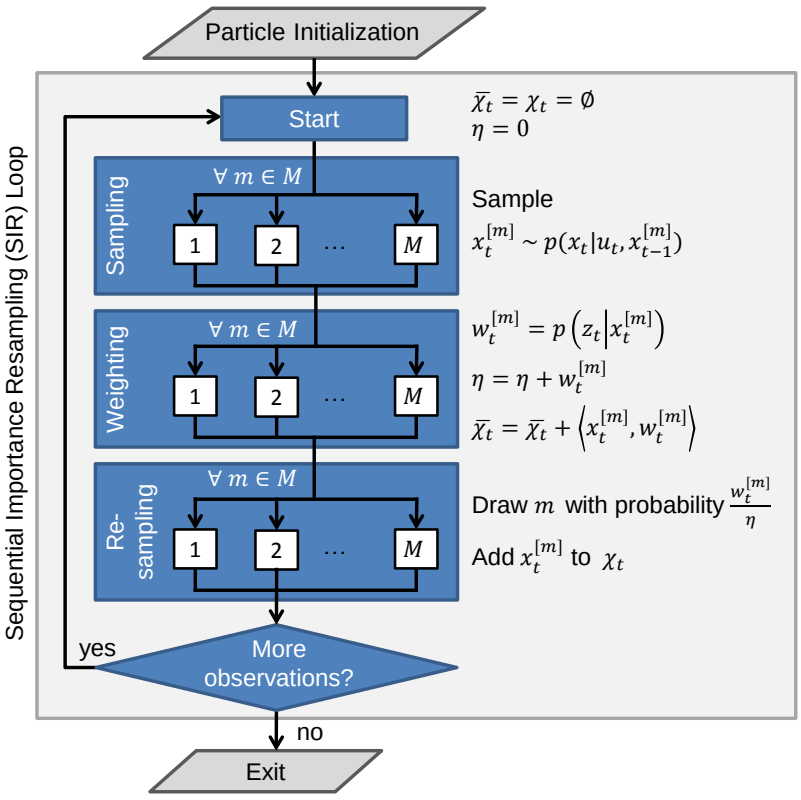

Figure 6.3: The particle filter algorithm using sequential importance resampling

no particle set exists yet, a new set with uniformly distributed particles needs to be created

first in the initialization phase. Then, two new empty particle sets $\overline{\chi_t}$ and χ_t are created. For normalization purposes, a running counter η is used which sums up all particle weights in the process. The first two steps of the algorithm are performed for each particle in the given particle set χ_{t-1} as shown in Fig. 6.3: In the sampling step, a new particle $x_t^{[m]}$ is created, based on the knowledge about the control input u_t and the respective particle $x_{t-1}^{[m]}$ of the particle set χ_{t-1}. In order to calculate the required state shift, the state transition distribution $p(x_t|u_t, x_{t-1})$ is sampled.

In the second step, the new particle is weighted. The goal is to correct estimation errors of the prediction using sensor measurements z_t. The weight of a particle $x_t^{[m]}$ is calculated using the conditional probability $w_t^{[m]} = p(z_t|x_t^{[m]})$. This is the probability of the measurement under the condition that the state is according to the given particle. After the weighting is done, the particle is added to a new temporary particle set $\overline{\chi_t}$.

The most important step of the particle filter algorithm is the *resampling* as shown in Fig. 6.3. The resampling algorithm draws with replacement M particles from the temporary particle set $\overline{\chi_t}$. The probability of drawing a particle corresponds to its normalized particle weight $w_t^{[m]}/\eta$. Finally, the drawn particles are added to the output particle set χ_t. Regarding the necessity of the resampling, the following explanation is given in [15]: "The resampling step is a probabilistic implementation of the Darwinian idea of *survival of the fittest*: It refocuses the particle set to regions in state space with high posterior probability. By doing so, it focuses the computational resources of the filter algorithm to regions in the state space where they matter the most." The resulting particle set is used again when a new observation occurs. For further processing each intermediary result could be used as an output of the particle filter.

6.1.3 Trust Based Node Assessment using Particle Filters

In this work we use a particle filter algorithm to perform a data fusion of several location-related data sources in order to check mobility data plausibility of single-hop neighbor nodes. In this approach, a separate particle filter is used for each tracked vehicle. Particle filters show a high efficiency with respect to tracking purposes and allow the inclusion of negative and positive weighting factors. However, the VANET scenario differs from typical usage scenarios of particle filters where a hypothesis is corrected by fully trusted sensor data values. On the one hand, the incoming position values of the tracked vehicles, which represent an essential part of the "sensor data" used to correct the sampling, can be forged or flawed. On the other hand, the goal of the tracking is not to identify the most likely position itself but to determine the plausibility of the claimed position.

With our proposed particle filter based scheme for plausibility checking of mobility data, we are able to integrate all location verification methods proposed in earlier work:

- Tracking of adjacent nodes to detect position jumps
- Integration of sensor information to confirm or disprove a claimed neighbor node position (e.g. Radar, Lidar, cameras, directional antennas)

- Integration of knowledge to confirm or disprove a claimed neighbor node position (e.g. digital maps, sudden appearance areas [11], maximum communication ranges [10])
- Support of specific checks using the particle filters (e.g. node overlap detection [17], minimum distance moved observation [18], pseudonym change detection [19], tracking of the own position)

6.1.4 Misbehavior Detection with Particle Filters

Besides the usage of the particle filter to assess the trust values, there are other possibilities to make use of the particle filter.

On the one hand, it is possible to check if an object at a given location is matching with the particle cloud of one of the tracked vehicles. This mechanism can be used to test if any of the tracked vehicles is the one detected by the radar or if a tracked vehicle has performed a pseudonym change and appears subsequently with a new identifier. On the other hand, the particle cloud could be used to check whether the claimed position of neighboring vehicles overlap. As two or more vehicles should not be at the same location at the same time, this could be used to identify suspicious vehicles.

Finally, an additional particle filter could be used to track the own vehicle position. It does not matter if imprecise map data, winding roads, or an inaccurate own position is the cause, the own vehicle should always be able to serve as a reference with respect to plausibility. If the own vehicle is not able to achieve a good vehicle trust value, the whole plausibility check should be suspended.

6.1.5 Quality of Node Trustworthiness Assessment

The practicability of our scheme is tested under real condition, where three vehicles were used to drive various maneuvers on a testing area. For the following evaluations we used one particle filter per neighboring vehicle and thresholds are used to indicate suspicious behavior.

The tests are based on real traces recorded on a dedicated test area where several simple maneuvers, like sudden braking and evasion of obstacles were performed. The test results shown in Fig. 6.4 address the impact of radar object detection similar to the attack discussed in Section 6.1.1: A tracked ghost vehicle A_1 drives along with the vehicle R that runs the plausibility checker. At the beginning A_1 moves with a constant speed, identical to the speed of R, and keeps a constant distance. Then the tracked ghost vehicle A_1 starts a sudden overtaking maneuver and cuts in the gap between the vehicle R and a heading vehicle T at time t_1. Afterwards, at time t_2 the ghost vehicle leaves the radar area but stays in communication range and performs some further driving maneuvers. Fig. 6.4 shows the decrease of vehicle trust and message trust below the threshold at time t_1 which indicates non plausibility behavior of the tracked vehicle A_1. Similar to the

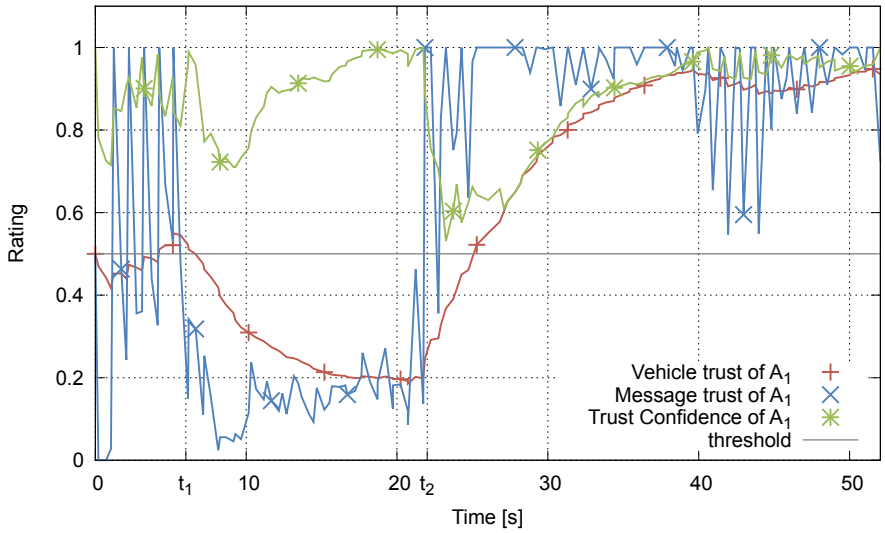

Figure 6.4: Trust and confidence values of a real overtaking maneuver with radar area violation

test results under laboratory conditions, the vehicle trust increases as soon as the ghost vehicle leaves the radar area at time t_2. In spite of jumpy message trust values caused by abrupt driving, the expectations are fulfilled.

Therefore, we can show that the particle filter algorithm is able to handle real vehicle data without giving wrong warnings under typical driving behavior. At the same time, attacks according to our attacker model can be detected reliably.

6.1.6 Accuracy and Performance of the Particle Filter

As the quality of the plausibility check is directly related to the number of particles in the filters, we analyzed the accuracy and performance using our vehicle traces. The optimal number of particles is always dependent on the use case and its requirements. An increase of particles leads to a higher accuracy but needs more processing power. Fig. 6.5(a) presents vehicle trust graphs for multiple different numbers of particles, starting from 10 particles up to 1000 particles. For these performance evaluations, we reused the recorded real vehicle traces presented in Section 6.1.5. As reference for the accuracy evaluation we use the graph for vehicle trust of A_1 shown in Fig. 6.4. All graphs depicted in Fig. 6.5(a) that exhibit small deviations from the reference vehicle trust graph can be assumed to have appropriate particle numbers.

The best results can be achieved with particle numbers between 500 and 1000. The graphs of those two particle numbers are nearly identical and centered within the other graphs. In theory, a particle filter using more particles should produce better results as it should converge to the optimal solution. However, in our tests we have observed that filters with more than 2000 particles cannot be processed fast enough due to limited processing

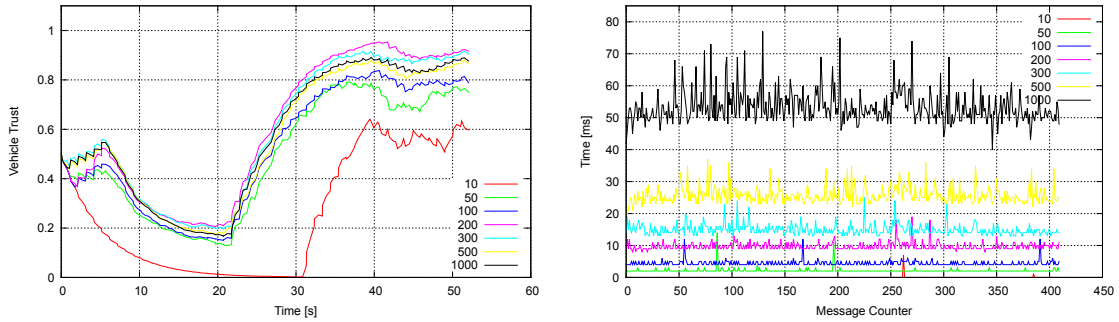

(a) Accuracy of measurements with different numbers of particles

(b) Runtimes of the particle filer algorithm in dependence of the number of particles

Figure 6.5: Accuracy and Performance of the Particle Filter

power on automotive systems. With less than 300 particles, the results are still usable but get less and less accurate. For reliable results, 100 particles should be the lower bound.

Fig. 6.5(b) shows the performance measurements of the particle filter with varying numbers of particles similar to the accuracy evaluation shown in Fig. 6.5(a). As the complexity of particle filters is $O(M)$, an increase of particles causes a linear increase of computational effort, which might be a problem in resource restricted environments. Using only 100 particles per particle filter, it would be possible to handle up to 200 incoming messages per second. Using between 500 and 1000 particles per filter, approximately 40 messages can be processed. Consequently, the particle filter algorithm may be adapted to incoming message rates and only relevant neighbors may be tracked.

6.2 Central Misbehavior Evaluation for VANETs based on Mobility Data Plausibility

Trustworthy communication in vehicular ad-hoc networks is essential to provide functional and reliable traffic safety and efficiency applications. A Sybil attacker that is simulating "ghost vehicles" on the road, by sending messages with faked position statements, must be detected and excluded permanently from the network. Based on misbehavior detection systems, running on vehicles and roadside units, a central evaluation scheme is proposed that aims to identify and exclude attackers from the network. The proposed algorithms of the central scheme are using trust and reputation information provided in misbehavior reports in order to guarantee long-term functionality of the network. A main aspect, the scalability, is given as misbehavior reports are created only if an incident is detected in the VANET. Therefore, the load of the proposed central system is not related to the total number of network nodes. A simulation study is conducted to show the effective and reliable detection of attacker nodes, assuming a majority of benign misbehavior reporters. Extensive simulations show that a few benign nodes (at least three witnesses) are enough to significantly decrease the fake node reputation and thus identify the cause of misbehavior.

In case of colluding attackers, simulations show that if 37% of neighbor nodes cooperate, then an attack could be obfuscated.

Vehicles and Roadside Units (RSU) are broadcasting messages with traffic related data. Due to the wireless property of the communication channel, the transmitted messages, used by ITS functions, have to be protected by cryptographic security solutions. As proposed by the IEEE 1609.2 draft standard [4], digital certificates are used to ensure the authentication and authorization of message senders. A Public Key Infrastructure (PKI) issues certificates, but only for authenticated nodes of a Vehicular Ad-hoc Network (VANET). Nevertheless, it is assumed that attacks on the VANET are possible as cryptographic keys could be misused or malicious software could be installed on some nodes. Also faulty nodes could disturb the functionality of ITS communication unintentionally but permanently. As a result, misbehavior detection and evaluation is necessary to keep overall functionality of ITS communications.

A local misbehavior detection system on the network nodes is able to detect inconsistencies in mobility data (i.e. absolute position, heading, speed) by applying plausibility checks. However, this detection is restricted to nodes that are in communication range of the attacker at current moment in time. Furthermore, the local misbehavior evaluation suffers from pseudonym changes of attackers and therefore only a reduced set of information may be available. Hence, we propose the transmission of locally created misbehavior reports to a central evaluation entity. This permits to detect attacks, based on a larger set of information, and identifies attackers with higher probability. Our main concept is based on the computation of trust information regarding neighboring VANET nodes. Successful detections are used subsequently to exclude disturbing nodes from the VANET until they have proven their benignity. As great attention is given to scalability, flexibility and practicability, the proposed scheme aims to provide a basis for automated misbehavior detection in ITS.

6.2.1 Adversary Model

We focus on internal attacker that can forge messages to generate ghost vehicles. This malicious behavior is also known as Sybil attack.

An attacker, that places for example a non-existing broken down vehicle on a road segment, would be able to reroute other vehicles if their navigation systems process the faked messages, or worse, affect safety applications of neighboring vehicles. Based on this kind of malicious mobility data modification in broadcasted messages, a wide range of different attacks is imaginable. An exemplary situation is depicted in Figure 6.6. The claimed position of a faked vehicle $a \in V$ is overlapping a real benign vehicle $b \in V$. Other witness nodes $w_1, w_2, ..., w_5 \in V$ in the communication range are able to detect this inconsistency autonomously [17]. In the remaining sections, we also assume a and b as suspected nodes.

As this adversary model is relatively generic, we aim to detect a wide area of location based attacks by applying the proposed misbehavior detection and evaluation framework.

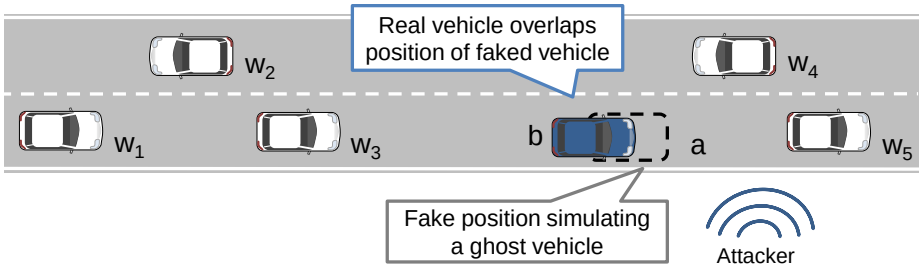

Figure 6.6: Attacker simulates a faked ghost vehicle on the road that is overlapped by real vehicle positions

6.2.2 Misbehavior Report

A Misbehavior Report (MR) is used to send information regarding possible misbehavior from network nodes to a Misbehavior Evaluation Authority (MEA). A report stores the type of detected misbehavior (e.g. node overlap, unexpected position jump), the pseudonymous ID of the reporter node, a list of overlapping nodes including their trust statements and a list of neighbor nodes surrounding the reporter. Figure 6.7 gives on overview of the misbehavior report structure and its content.

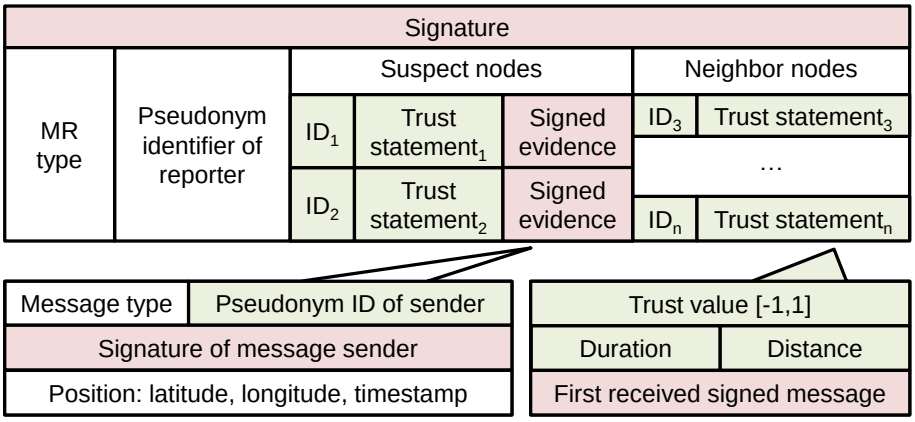

Figure 6.7: Structure of misbehavior report

In every report an evidence of the observed event is added. In case of a position overlap, two signed CAMs are added that attest to the overlap of node polygons as detailed in [17]. The suspect nodes and relevant one-hop neighbor nodes are reported by providing their pseudonymous ID and a trust statement. The latter contains a trust value of the target that is calculated by the local misbehavior detection system of the reporter, as well as the contact duration and the traveled distance of this node. In order to attest to the values for distance and duration, an appropriate Cooperative Awareness Messages (CAM) is appended, and thus, can be verified by the MEA. Finally, the complete report is signed and encrypted before it is sent to the central MEA. In the case of dense traffic, the size of a misbehavior report could be limited by adding only relevant neighbors which has

observed the misbehavior autonomously. Only selected one-hop neighbor nodes shall be added, prioritized on the distance between the node and the location of the detected inconsistency. The probability that nearby neighbors have also detected the inconsistency autonomously is higher than for distant neighbors.

Due to the signed evidence proving the overlap, an attacker cannot accuse arbitrary nodes in the network. Therefore, cooperative attacks with several malicious reporters are spatially and temporarily limited. Furthermore, the lightweight MR format allows the transmission to the infrastructure via roadside stations using G5A. Our implementation has shown that the size of a misbehavior report is approximately 1 KB and will be increased by 200 Bytes for every additional neighbor node. It has to be considered that the temporary MR storage on the network node should be persistent but has no requirements regarding security or tamper protection.

6.2.3 Certification of Misbehavior Reports

When receiving a new MR at the central MEA, the contained signatures are first verified by using the public keys of the related pseudonym certificates. As the MEA has a connection to the PKI that has issued the certificates, it is sufficient to store only short certificate IDs in the MR instead of the complete certificate structure. With the certificate ID, the MEA is able to request the appropriate certificate. In a second step, the evidence of overlaps is checked by verifying the signature of CAMs. The overlap scenario can be verified by comparing the position vectors of the appended messages as shown in Figure 6.7. Subsequently, all information of neighbor nodes that are appended to the MRs is verified by comparing the position vector of the signed CAM with given duration and distance values. If the MEA detects a noteworthy differences between claimed values for duration and distance and the CAM position, the report is discarded and will not be used in the further evaluation process. Furthermore, duplicated reports from the same node, using different pseudonymous IDs, are discarded. A possible infringement of privacy due to this required resolution of pseudonym ID to their related long-term ID is not further discussed in this paper.

The verified MRs are stored in order to amass enough reports from nodes that are involved, or have observed an inconsistency, caused by node overlaps. Having enough reports for an evaluation, a *session* is created for every misbehavior scenario. This session contains a list of all involved nodes. The list of neighbors from the MRs is used to identify possible witnesses. Based on a policy, the number of needed witnesses can be defined before starting the further evaluation. It is therefore a requirement that every involved node should be able to detect an overlap autonomously, create a report and send it to the MEA.

6.2.4 Node Assessment Concept

The goal of the central misbehavior evaluation is the identification of an attacker inside a set of nodes that are actively or passively involved in an overlap scenario. The evaluation process is divided into five steps, as depicted in Figure 6.8.

6.2.4.1 Generation of Cooperative Trust-Confidence

As soon as enough reports are collected by the MEA, the assessment process is started for the respective session. In order to use the reported information for detecting and identifying an attacker, the trustworthiness of all involved nodes is calculated. Figure 6.8 provides an overview of the process to calculate trust-confidence pairs that are reported by network nodes. In the first step shown in this figure, the confidence of reported trust is calculated.

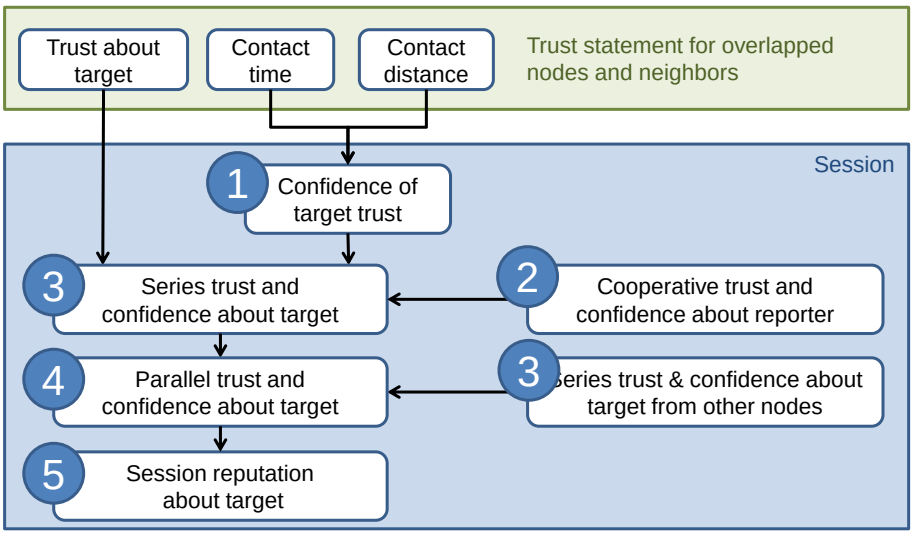

Figure 6.8: Overview of central processing of reputation information

The given trust and calculated confidence from Misbehavior Reports are used to express the trustworthiness regarding reported misbehaving nodes or witnesses. After extraction of the trust-confidence pairs for all session nodes, a cooperative trust and confidence value is calculated. This tuple determines the trustworthiness of every involved node as shown in step 2 of Figure 6.8. In the following process description, the set V contains all session nodes. In step two, the cooperative trust and the cooperative confidence for node $e \in V$ at time k are calculated where node e is evaluated by all other session nodes $a \in V$ and $a \neq e$.

Having this aggregated cooperative trust-confidence pair of every node in a session, the evaluation of suspected nodes is started in the third step of Figure 6.8. All nodes that are reported to be overlapped by another node, are suspects. The trust regarding a suspect,

provided by a witness node, is weighted with the cooperative trust-confidence pair of this witness node. The computation of trust-confidence data is denoted as *series* combination as the trustworthiness of a reporter node a is used to rate a target node e. By applying this approach, a cooperatively less trusted node has a lower impact in the assessment of suspects than a node with high cooperative trust. This assessment task is repeated for every suspected node, using the cooperative trust and confidence of every node involved in a session. In step 4, all series trust-confidence pairs are combined in parallel in order to get the final values for trust and confidence of a suspect.

In the last step, it is checked whether the suspected nodes are rated positively or negatively. Therefore, the results of the parallel combination of trust and confidence are used to combine them in a reputation value. The higher the parallel confidence value the more the trust value is considered. If a reported trust of a suspicious node has low confidence or the reporter itself has a low confidence then a neutral reputation for the suspicious node $e \in V$ is given.

Based on the concept, presented in Figure 6.8, the reputation for suspicious nodes is combined in the fifth step so that a single reputation value is generated. This reputation value contains aggregated information from all other involved reporters in a session. Based on a benign majority and local misbehavior detection mechanisms on the reporter nodes, a ghost vehicle is rated with a negative trust value and a real vehicle is rated with a positive trust value.

6.2.5 Quality of Malicious Node Detection

Based on a simulation, different quality measurements of the central MEA concept are presented in the following. In order to work with realistic simulation data, the evaluations presented in this Section are working with incomplete sets of Misbehavior Reports. It is assumed that 30% of all nodes are not able to transmit a recorded Misbehavior Report to the infrastructure. Therefore, this fraction of neighbors is not listed in Misbehavior Reports used in the simulation. In Figure 6.9, a situation is evaluated with one faked node that is overlapped by one benign node. This attack is also illustrated and described in Section 6.2.1. In order to get the information on “how many witness nodes are needed for reliable detection of an attacker", the number of benign witnesses is increasing on the x-axis. The three graphs, shown in Figure 6.9, describe the reputation of the benign and the faked node as well as the mean reputation of all witness nodes. As shown in this graph, the decrease of faked nodes’ reputation attenuates with 7 witness nodes.

In Figure 6.10, a similar situation is simulated. A fixed number of 10 benign nodes are generating misbehavior reports where 5 nodes are overlapping actively the one faked ghost vehicle. The other 5 benign nodes are acting as witnesses. In this simulation the number of malicious reporters is increased in order to measure the impact of several cooperating attackers. According to our simulations, it is sufficient if 37% of the participants are cooperating attackers in order to hide a real attack on the road as shown in Figure 6.10.

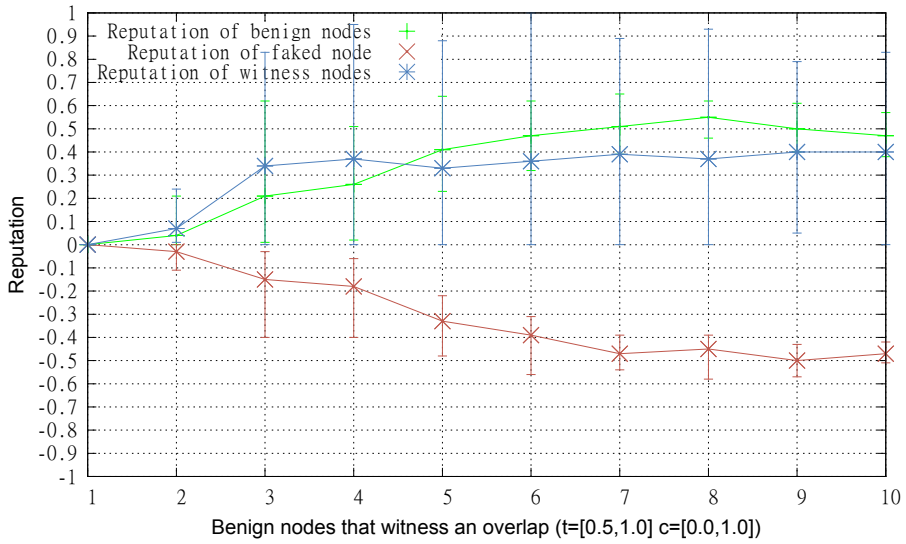

Figure 6.9: Ghost vehicle attack with increasing number of benign witnesses

However, the effort for an attacker is very high to mount an attack where several manipulated vehicles are at the same location at a given time. Furthermore, the MEA is able to link different pseudonyms that are used by the same node.

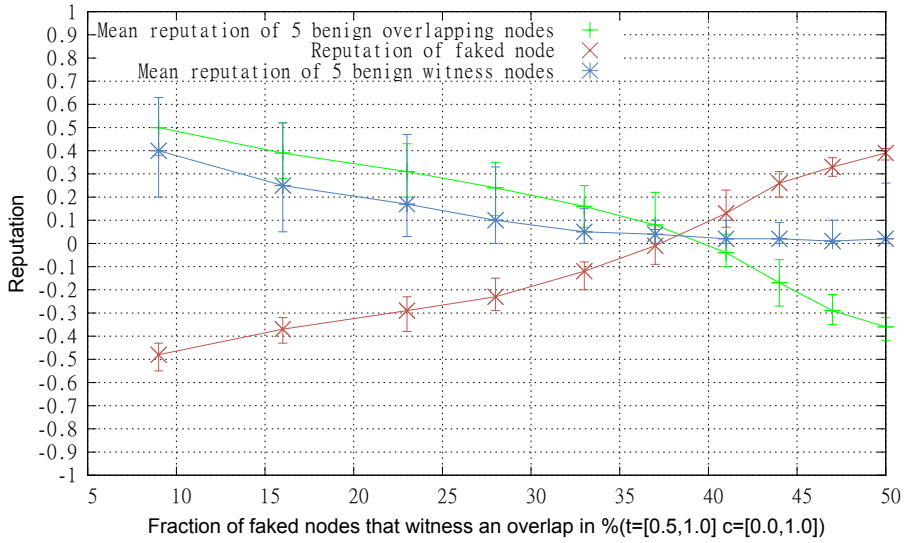

Figure 6.10: Ghost vehicle attack with increasing number of faked witnesses

6.2.6 Consolidated Findings

Based on the results, shown in Figure 6.9 and Figure 6.10, the number of needed MR reporters is dependent on their provided trust and confidence regarding a ghost node.

If both involved nodes of an overlap scenario has transmitted misbehavior reports, then the number of needed witnesses may be lower. Otherwise, a higher number of benign witnesses is needed which would automatically increases the effort for a cooperative attack (see Figure 6.10). The proposed concept for central Misbehavior Report evaluation is used to decide which node is probably a faked ghost vehicle created by an attacker and which is a benign node. It has to be considered, that false positive detections where two benign nodes overlap each other virtually on the road, are not detected by this entity. Therefore, it is proposed to collect the final reputation of suspicious nodes in a global reputation table. Only if the same node has several overlaps with different other nodes, in a specific time frame, a reaction may be reasonable to protect the network against malicious or defective nodes.

6.3 How to Secure ITS Applications?

In this paper, based on recent standardization activities, we give an overview of ITS applications and we detail a classification of road safety applications. We investigate the security issues of cooperative ITS applications and we present their security profiles. Taking into account the communication architecture of an ITS Station, we advance a new application oriented security approach. In order to have an optimal security solution, we propose to have specific security solution to each application (message) type/classes depending on security requirements. So we propose to use security within higher layers in order to respect specificity and security requirements of each ITS application. Like this, we define a security service oriented approach in which security depends on the nature of an ITS service (application) and not on the communication type. The advantage of this approach is to have an extensible and flexible ITS system that can be easily updated with new ITS applications.

6.4 Risk Analysis of ITS Communication Architecture

In this paper, we present the ITS framework and communication architecture in Europe. Then, we identify inherent threats in this architecture. And using ETSI threat analysis methodology TVRA, we propose a risk analysis of ITS. Our risk analysis proves that threats to security of ITS applications are as critical as threats to security of vehicular communications. For example, Sybil and illusion attacks, occurred mostly in applications level, cause serious disruptions to safety applications operation. the associated motivation is high because it consists on changing ITS applications environment and giving false information to neighbors. The main motivation of these attacks is to gain advantage on road. Moreover, in these attacks, a lot of time is needed to identify and revoked the attacker. This analysis supports our advanced application security approach.

6.5 Evaluation of Congestion-based Certificate Omission in VANETs

Telematic awareness of nearby vehicles is a basic foundation of electronic safety applications in Vehicular Ad hoc Networks (VANETs). This awareness is achieved by frequently broadcasting beacon messages to nearby vehicles that announce a vehicle's location and other data like heading and speed. Such safety-related beacons require strong integrity protection and high reliability, two properties that are hard to combine because the communication and computation overhead introduced by security mechanisms affects reliability. This applies especially to the signatures and certificates needed for authentication.
We propose a mechanism to reduce the communication overhead of secure safety beacons by adaptively omitting the inclusion of certificates in messages. In contrast to similar earlier proposals, we control the omission rate based on estimated channel congestion. A simulation study underlines the advantages of the congestion-based certificate omission scheme compared to earlier approaches. Moreover, we show that the benefits of certificate omission outweigh the negative effect of cryptographically unverifiable beacons.

In vehicular ad hoc networks (VANETs), vehicles are broadcasting beacon periodically with a frequency of 10 Hz [4]. While the upcoming European standards [20] foresee also adaptive beaconing rates between 1 and 10 Hz, we stick to the 10 Hz rate in this paper. These beacon messages are either processed directly by applications trigger certain effects, e.g., warning the driver of a potentially imminent collision. Or vehicles use the information to build a so-called Local Dynamic Map that different applications use for purposes like traffic advise or collision warnings.

If attackers succeed to inject spoofed information into the system, this might have severe consequences, e.g. drivers misreacting due to wrong warnings. Therefore, most proposals foresee a basic integrity protection and authentication of messages based on ECDSA signatures and certificates issued by a Public Key Infrastructure (PKI) [4,21,22]. Thus, authorized vehicle have a private/public key pair and receive a certificate from a Certification Authority (CA) that declares the vehicle a valid participant in the VANET[1]. The vehicle then signs every beacons with its private key and appends the certificate to the message. Any receiver then has to verify the certificate and the signature of the beacon before further processing of the message. Therefore, security creates a communication overhead (i.e., packet size increases) and a computational overhead (i.e., time to process the packet). As was already investigated in [23–27], these two overheads introduce a scalability problem that can affect reliable communication and thus traffic safety in high density scenario. For example, a vehicle surrounded by 100 vehicles will receive approximately $100 \times 10 = 1000$ messages per second and has to perform 1,000 signature verifications per second plus at maximum another 1000 certificate verifications. Beyond, the signature and certificate enlarge the beacon message by roughly 200 bytes, increasing the channel load and chance for collisions.

[1] For simplicity, we are not addressing pseudonym schemes here

One proposal to deal with computational overhead is to include a dedicated cryptographic accelerator in the On-Board Unit (OBU) that can handle the approximately required 1,000 verifications per second.

So [23, 24] investigated approaches to selectively skip certain steps during the communication process, e.g., by attaching certificates only to specific packets, or by skipping verification of some signatures. For instance, one could skip attaching a certificate to every packet as vehicles may cache certificates received in earlier messages. This certificate omission is the focus of this paper and was also investigated in [23–27].

These strategies all come at the risk that a vehicle A may receive a beacon from vehicle B without attached certificate before having cached the missing certificate of B from an earlier beacon with attached certificate. To prevent potential attackers from injecting spoofed packets, A would have to discard the beacon leading to what we term as *cryptographic packet loss*. If we, on the other hand, reduce the size of a lot of beacons by omitting certificates, the overall channel load, and thus also packet loss caused by collisions, will be reduced. With this paper, we want to investigate the trade-off between cryptographic- and communication-channel-induced packet loss by analyzing and comparing two of the existing schemes, namely Neighbor-based Certificate Omission (NbCO) and Periodic certificate omission (POoC) schemes. Based on this, we propose a new Congestion-based Certificate Omission scheme (CbCO) that combines the advantages of NbCO and POoC.

Our main goal is to reduce overall packet loss even when the density of vehicles changes. In a simulation-based comparison to four other schemes, namely no omission, full omission, NbCO, and POoC, we show the advantages and greater flexibility of CbCO.

6.5.1 Certificate Omission

Our protocol is based on the Periodic Omission of Certificates (POoC) [25–27] and the Neighbor-based Certificate Omission (NbCO) [23, 24]. We present those protocols in this section and discuss their advantages and disadvantages.

6.5.1.1 Periodic Omission of Certificates

The idea of the POoC [27] is to add the certificate every n beacons.[2] Figure 6.11 gives an example with $n = 3$. The overhead reduction depends linearly on the certificate period.

But omitting certificates on a periodic schedule, creates a window of $n-1$ beacon periods where a vehicle that has not yet cached the certificate of a sender may receive a beacon that it cannot verify and has to drop it. Assuming a beacon interval of Δb, the period until a node can verify packets from a new node in its neighborhood is $(n-1) \times \Delta b$ in the worst case, and $\frac{(n-1)\times\Delta b}{2}$ on average. One major drawback of the POoC is that the scheme is independent of vehicle context, thus, n is static. This might jeopardize safety applications.

[2] called *certificate period* in the original paper

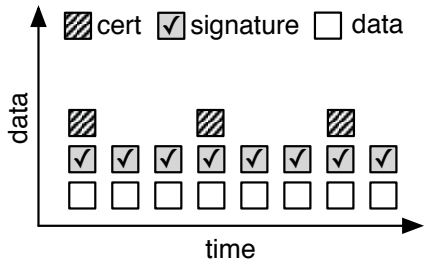

Figure 6.11: Example of POoC

Figure 6.12: Example of NbCO

For instance, with $\Delta b = 0.1s$ and $n = 10$, a vehicle has to wait for $(10 - 1) \times 0.1 = 0.9s$ in the worst case before being capable of verifying incoming messages of vehicles that newly enter communication range (no certificate cached). At a speed of 120 km/h, this corresponds to 30 meters, too much for many safety applications. On a side note, this problem also arises whenever pseudonyms are changed.

6.5.1.2 Neighbor-based Certificate Omission

Schoch et al. [24] propose a different certificate omission scheme that considers the context of a vehicle in the omission decision. The idea of NbCO is to only attach the certificate to beacons if there is a change in the neighbor table. Figure 6.12 shows an example when a change in the neighbor table appears in the fourth beacon. A node can monitor changes to the neighbor table caused by incoming packets from unknown nodes and attach a certificate only when the neighbor table has changed since the last beacon with certificate was sent. Note that the reception of an unverifiable beacon with missing certificate also needs to trigger a neighbor table change, even if the information is unverifiable. When node A is about to send a new beacon, A determines if new neighbors were added to its neighbor table since the last beacon with certificate was sent. If so, a certificate is attached to the new beacon, else it is sent without certificate.

As we have found out, the main problem with the NbCO approach is that it does not scale in high density scenarios as such situations will expose a vehicle to a constantly

high change of neighborhood so that almost all packets carry certificiates, leading to high channel load and increased collissions.

To deal with lossy channels where the packets containing certificates for a newly arrived neighbor get lost, the authors of [24] propose that nodes could solicit for certificates if a certificate is not available within Δb or that certificates should be attached at least every n beacons.

6.5.1.3 Problem Statement

As discussed in the previous sections, omission of certificates in authenticated one-hop broadcast beacons is an effective way to reduce load on a communication channel. However, this improvement requires a trade-off against the immediate verifiability of messages. Some beacons may become unverifiable due to a missing certificate at the recipient, and have to be discarded. We call it *cryptographic packet loss* (CPL). The more certificate omissions, the higher this cryptographic packet loss will be. Other factors that influence CPL are beacon rates and vehicle mobility.

To avoid CPL we can attach certificates to all packets, thus, going back to the basic scheme. This, however, will create larger packets, increasing channel load, and effectively leading to more packet drops because of longer packet queuing or collisions. We call this *network packet loss* (NPL).

The ultimate goal is to increase information awareness, i.e., the actuality of information that a vehicle has about its neighborhood. Packet loss, no matter whether caused CPL or NPL, creates additional latency until updates are received, thus decreasing information awareness. When introducing our omission scheme, we have therefore to investigate whether the induced CPL is out-weighted by the reduced NPL due to shorter messages. Then, and only then, it is reasonable to apply these strategies.

With respect to this goal, both schemes have their advantages or disadvantages. In case of a stable environment, a PoOC scheme might add too many certificates to packets, especially if n is chosen lower than necessary. On the other hand, NbCO has its limits in case of high vehicle densities and high volatility in neighborhood, as it then degenerates to the no-omission case and adds to channel congestion. While the idea to track neighborhood for omission decisions is intuitively valid, we note that in practice the behavior of this scheme is not scalable. So we argue that we need a new scheme that in addition also considers channel load as an additional factor.

Therefore, we combine the advantages of both strategies. We call our resulting approach Congestion-based Certificate Omission scheme (CbCO). Our claim is that this scheme can better address the trade-off between CPL and NPL and thus achieves better information awareness of vehicles.

Figure 6.13: Example of CbCO

6.5.2 Congestion-based Certificate Omission Scheme

In Congestion-based Certificate Omission (CbCO), we propose to optimize omission of certificates not towards maximizing the number of omissions but instead towards minimizing the overall packet loss and thus optimizing the trade-offs between communication load and CPL. To achieve this, CbCO considers the load of the communication channel as the guiding metric. If the communication channel is free, there is no need to trade in CPL for less load on the channel. And if the communication channel is congested we want to reduce the communication load by aggressively omitting certificates. While aggressive omission increases the CPL, our evaluation shows that it will likewise decrease the overall NPL due to the reduced size of messages at an even larger rate, yielding an overall positive effect on packet loss. Figure 6.13 shows an example of CbCO where a congestion is detected on the third beacon. Then, the beacons 4 through 6 are transmitted without certificate.

CbCO is based on POoC and omits certificates on a periodic schedule. However, the certificate rate n at which certificates are added is flexible and triggered by the number of vehicles in communication range (as measured by the size of our neighbor table). The larger the size of the neighbor table, the larger we choose n. If N is the size of the neighbor table, then $n = \lfloor \Omega(N) \rfloor$, where Ω is a weight function. This weight function defines the maximum number of omission in function of the channel occupancy. As Ω is a key parameter of the CbCO scheme, we analyze three different trends to determine the optimal strategy. We consider n_{max} the size of the neighbor table that should trigger maximum omission and o_{max} the maximum omission rate. The selection of appropriate values for o_{max} and n_{max} is discussed in Section 6.5.3.2. We evaluate the following functions for Ω:

$$\Omega_{linear} : y = \min\left(\frac{x}{n_{max}} \cdot o_{max}, o_{max}\right) \tag{6.1}$$

$$\Omega_{quad} : y = \min\left(\left(\frac{x}{n_{max}}\right)^2 \cdot o_{max}, o_{max}\right) \tag{6.2}$$

$$\Omega_{trig} : y = \begin{cases} -\cos\left(\frac{\pi}{n_{max}} \cdot x\right) \cdot \frac{o_{max}}{2} + \frac{o_{max}}{2}, & x < n_{max} \\ o_{max}, & x \geq n_{max} \end{cases} \tag{6.3}$$

Table 6.1: Simulation parameters

Parameter	Value
Field size	3 km × 3 km
MAC	802.11p, 3 MBit/s
Fading	Rayleigh
Pathloss	Two-ray ground
Noise	Additive
Simulation time	60 s
Simulation runs per configuration	10
Transmit power	10.9 dB
Beaconing frequency	10 Hz
Payload Size	50 Bytes
Number of nodes	100, . . . , 1300

6.5.3 Evaluation

To evaluate our omission scheme we focus on a city scenario with a varying number of vehicles that allow us to investigate the effects of the omission schemes especially under high communication load. While omission might not be critical at low vehicle densities, as the channel is free and can easily cope with larger packets, we expect significant effects in medium to high densities.

6.5.3.1 Simulation Setup

We use a simulation package based on JiST/SWANS [28] with extensions by Ulm University.[3] The simulation environment provides 802.11p radio simulation and a realistic vehicular mobility model called STRAW [29], which uses map data from the U.S. Census Bureau. This simulation package allows us to efficiently simulate scenarios with a high density of vehicles [30], which is our main interest for the evaluation of congestion-based certificate omission. We use a 3 km by 3 km urban city map in Suffolk County, U.S.A., which is the same scenario as used in previous research in omission scheme [24].

In our simulation we consider only the transfer of one-hop beacon messages. While one-hop beacon messages will not be the only safety messages in the CCH, we assume that these messages will dominate the load. The configuration of the 802.11p communication channel is set to 3 MBit/s with a fixed transmission power of 10.9 dB for robust delivery of one hop safety messages [31].

The basic parameters for our simulation are in line with previous works by Schoch et.al [24] and the current draft version of IEEE 1609.2 [4]. A summary of relevant parameters is given in Table 6.1. For the format of beacon messages we closely follow the Basic Safety

[3] Website: http://www.vanet.info

Table 6.2: Cryptographic settings

Parameter	Value
PKAlgorithm	nistp256
ECC Key Type	compressed
Cert Size	140 Bytes
Signature Size	65 Bytes

Message (BSM) format as specified in SAE J2735 [32], delivered as a 45 byte DER encoded payload in a IEEE 1609.2 data message [4]. We do not consider any optional Part II attributes of the BSM format or optional parts of the 1609.2 message format. The security services specified in IEEE 1609.2 offer different options for the cryptographic additions to messages. From these options we selected the compressed representation of nistp256 keys and signatures. We do not consider certificates chains in this study. But we note that certificates chains would increase the benefit of certificate omission as the crypto payload would get even larger. A summarized description of the cryptographic additions to our simulated messages is included in Table 6.2. Adding the 45 byte BSM and 5 byte for headers in the payload to the cryptographic material, the total size of one beacon message is 255 bytes with certificate and 115 bytes when omitting the certificate.

The beaconing rate in our simulations is fixed at 10 Hz, as recommended by SAE J2735. A full simulation run simulates 60 seconds of traffic. During this time we do not simulate pseudonym changes. We expect the rate of pseudonym changes to be low enough to not be a relevant factor for the bandwidth optimization of beaconing services.

To test the efficiency of omission schemes under high loads, we scale the number of vehicles in the simulation scenario between 100 and 1300 vehicles on a 3km x 3km road network. On our map, this leads to an average of 5 to 68 vehicles in communication range and 18 to 252 vehicles in sensing range.

6.5.3.2 Analysis

For the analysis of our scheme we first investigate the settings for congestion-based certificate omission. The guiding metric we use as the foundation for congestion-based certificate omission is the number of neighbors in communication range. This metric and a basic model of a 802.11p CCH with a 10Hz BSM application on top enable us to estimate the congestion on the channel. In our simulations we identified an approximate limit of 1000 BSMs per second to saturate one communication channel in 802.11p wireless communications. We derive that 100 vehicles in communication range sending BSMs at 10Hz represent a natural limit of the communication channel.

The authenticated delivery of BSMs is a cornerstone of various safety applications. To achieve a robust delivery of verifiable BSMs, it is reasonable to consider an upper bound on the maximum number of omissions our scheme allows. As a guideline we use a recommendation in Annex B2.2 of IEEE 1609.2 v2 D12 [4] and in Annex B of ETSI TS 102637-2

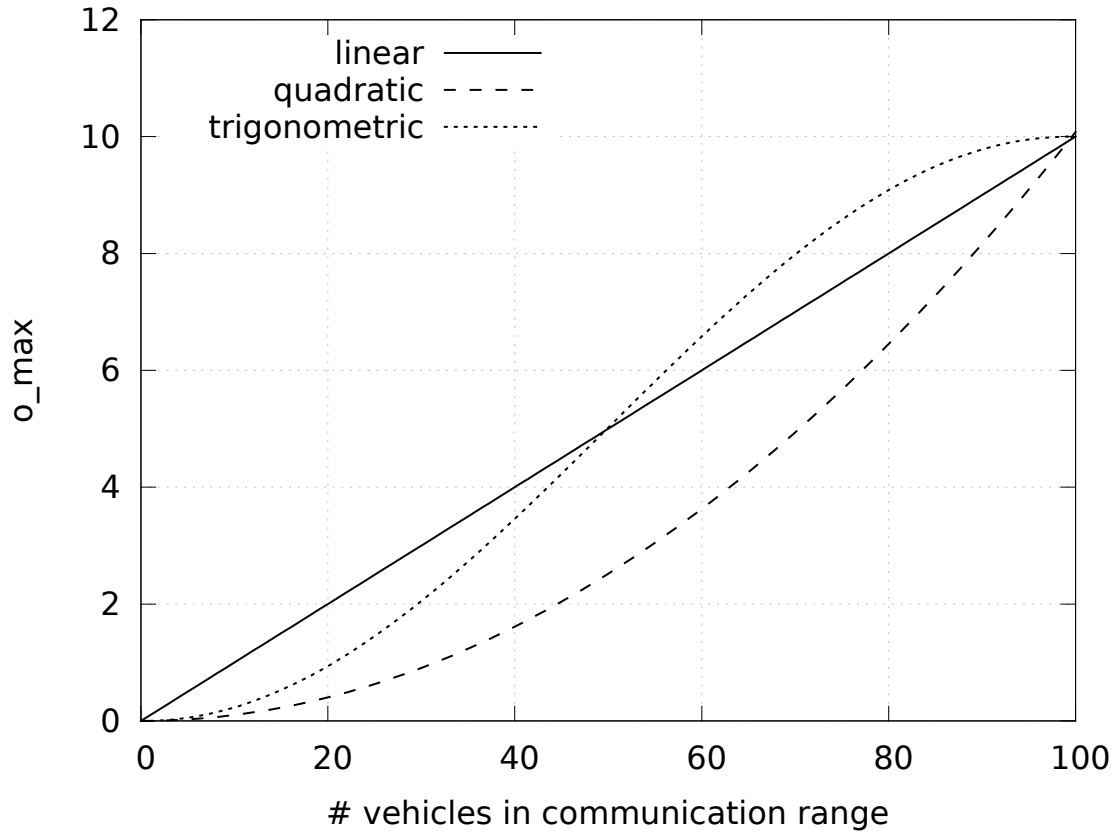

Figure 6.14: Omission rates strategies for congestion-based certificate omission

draft [20]. The IEEE 1609.2 recommends to include a full chain of certificates instead of just a single certificate at least once per second. In ETSI TS 102637-2 there is a description of Cooperative Awareness Messages (CAM), which are the equivalent to BSMs in the European standardization process. There, we find a set of informative rules for context adaptive beaconing rates, which specifies a maximum time between beacon generation of one second. From this, we deduce that an interval of one second between the inclusion of a full set of authentication material should be considered as an upper bound or $n_{max} = 10$.

With the bound on the communication channel and the bound on the maximum number of omissions we have a framework to define specific values for our CbCO scheme. Figure 6.14 shows the resulting Ω functions for $o_{max} = 10$ and $n_{max} = 100$.

While the linear function is a simple baseline to scale the number of omissions directly related to the number of neighboring vehicles, the other functions reduce the omission rate at lower vehicle densities to prevent CPL when there is no direct benefit in reduced NPL. We generally want to keep the number of omission low until the channel needs to counter increasing NPL. For this reason we propose two additional ways to calculate the number of omissions. A quadratic function lets the number of omissions grow slower in less

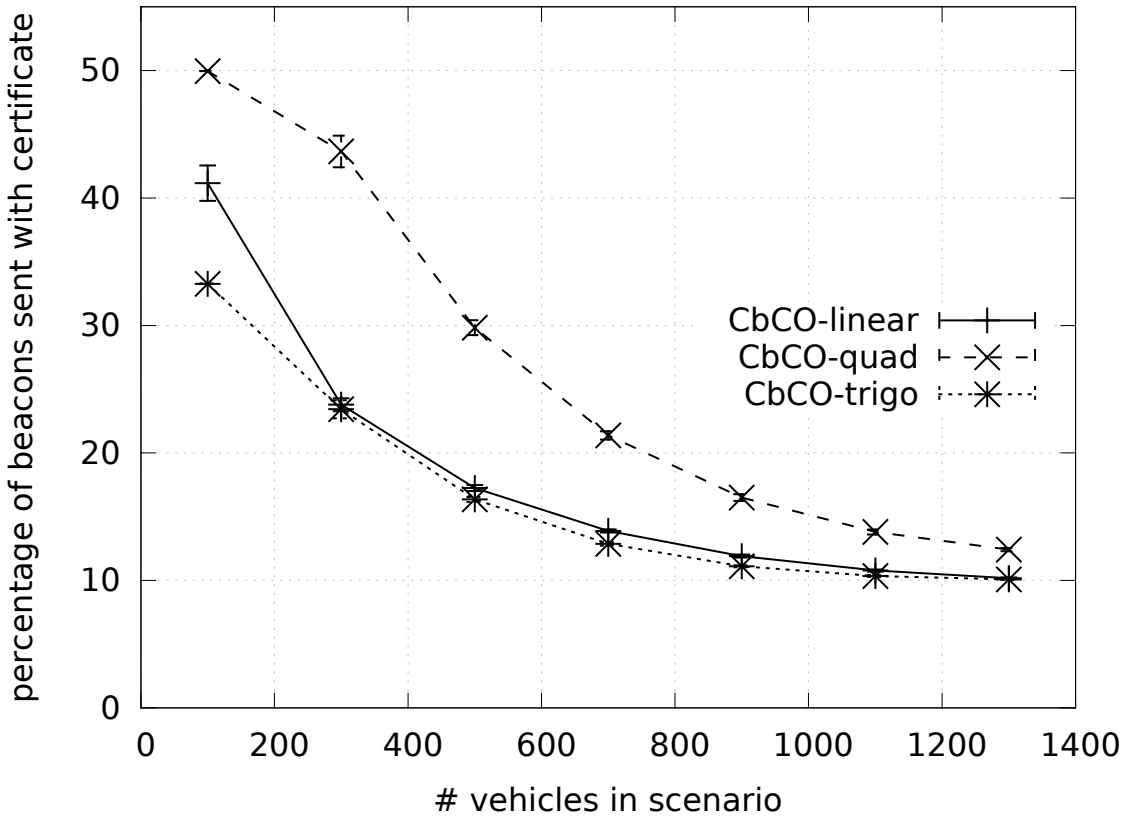

Figure 6.15: Average percentage of certificate omissions in CbCO

densely connected environments. And a trigonometric function produces similarly slow growth of omissions on sparsely connected environment while accelerating the increase of omissions more aggressively in densely populated environments.

To assess the quality of our CbCO, we analyze the number of omissions and the amount of collision based on CPL. Figure 6.15 shows the average percentage of beacons sent with a certificate attached to it. This is the inverse of the average number of omissions. We see the linear and trigonometric curves closely matching each other, while the quadratic calculation of omissions results in less omissions.

Next, we want to investigate the consequences of these different functions in terms of omissions. We measure this as cryptographic packet loss, i.e. the relative number of unverifiable messages that are dropped, and then the receiver misses a certificate to verify them. This is shown in Figure 6.16. We see that again the linear and trigonometric approaches match quite closely, while the quadratic method results in fewer unverifiable messages.

In practice we have to consider a secondary effect of omitting certificates. The goal of certificate omissions is to reduce the load on the network in order to have fewer colli-

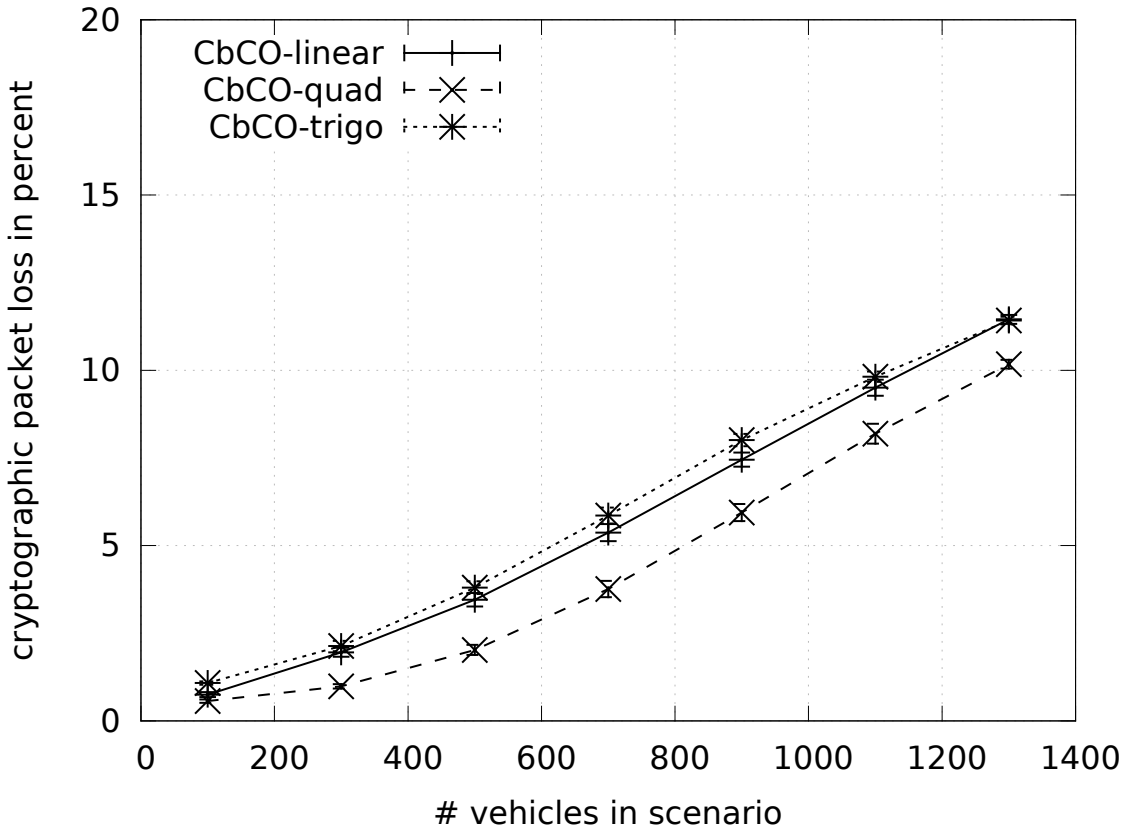

Figure 6.16: Average percent of unverifiable messages among received messages

sions and thus fewer packet loss. To see this effect we calculate a baseline of successful message delivery without any inclusion of any certificates. Using this baseline we can calculate the added packet loss due to the inclusion of certificates. Since different omissions schemes result in different numbers of omissions we see different characteristics for each scheme.

The graph in Figure 6.17 shows the average increase in network packet loss (NPL) relative to packets without certificates. As a reference, we also show the additional packet loss for no omissions (NoOm). As one can see, CbCO achieves a significantly reduced packet loss due to reduced message size compared to the NoOm scheme. One can also see that the quadratic Ω function performs a little worse than the two other.

Our goal is to decrease overall packet loss, considering NPL *and* CPL. This is shown in Figure 6.18. First of all, we can again observe the benefits of the omission schemes compared to attaching certificates to all packets. There is also a slight advantage of the linear and trigonometric Ω functions. Additionally, we note that above 1000 vehicles, we see the effect of the bounding of omissions, as the different Ω functions converge. Figure 6.19 illustrates the composition of network packet loss and cryptographic packet loss using CbCO-linear as an example. So while CbCO introduces additional CPL, it is

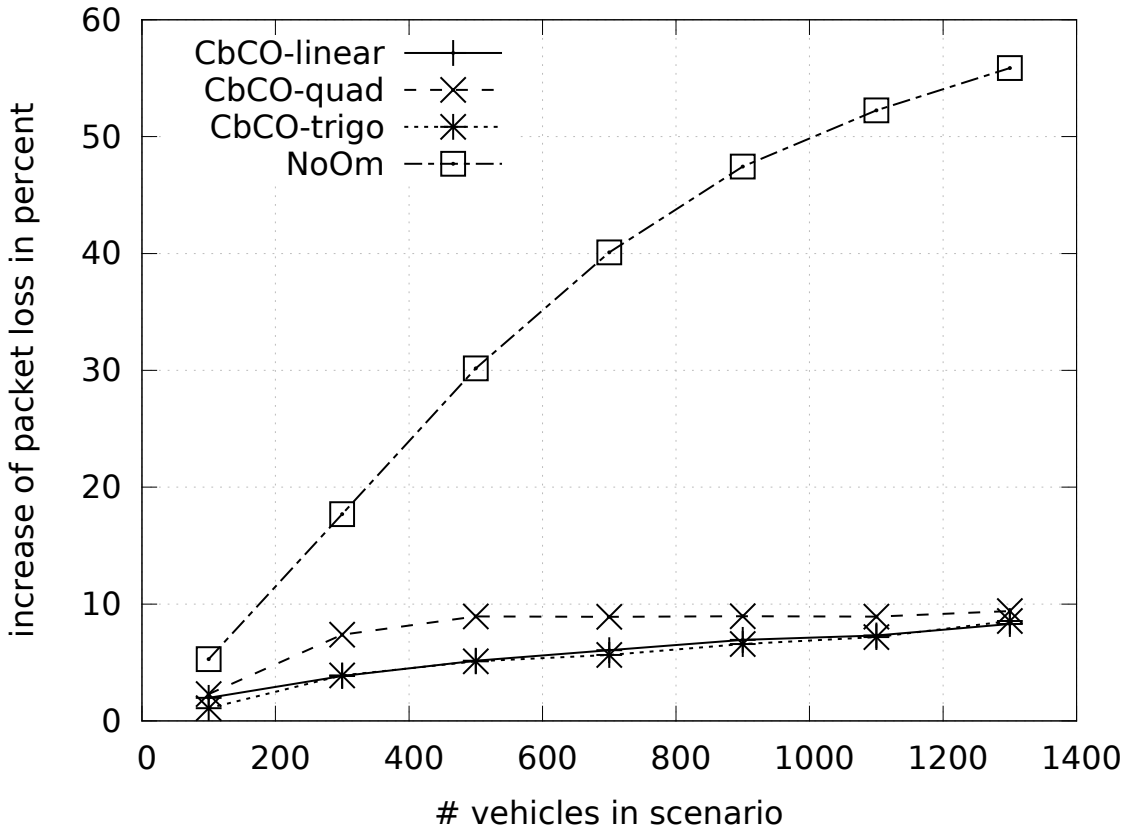

Figure 6.17: Increase of packet loss due to inclusion of certificates for different variants of CbCO (NPL only)

evident that the saved NPL outweighs this by far. However, we still need to investigate whether this comes at the expense of increased latency until a communication partner receives the certificate required to start authenticating messages.

Figures 6.20 and 6.21 illustrate the average and maximum number of unverifiable beacons until arrival of the certificate. As for the number of omissions, we notice that the quadratic method has a lower latency until messages become verifiable. We note that the linear way to calculate the congestion-based omission seems to provide a slightly improved latency characteristic compared to the trigonometric function.

We conclude that the linear and trigonometric approaches perform very similar, with slight advantages for the linear approach. We consider the advantage of the linear approach to be rooted in the faster increase of omissions in situations with high connectivity. These situations generally have more impact on the simulation results and keeping the number of omissions down until the channel is overloaded is an effective approach. The quadratic function shows the limit of following this line of thought. We see the advantage for the quadratic approach in the latency until a message becomes verifiable. Finally we remark

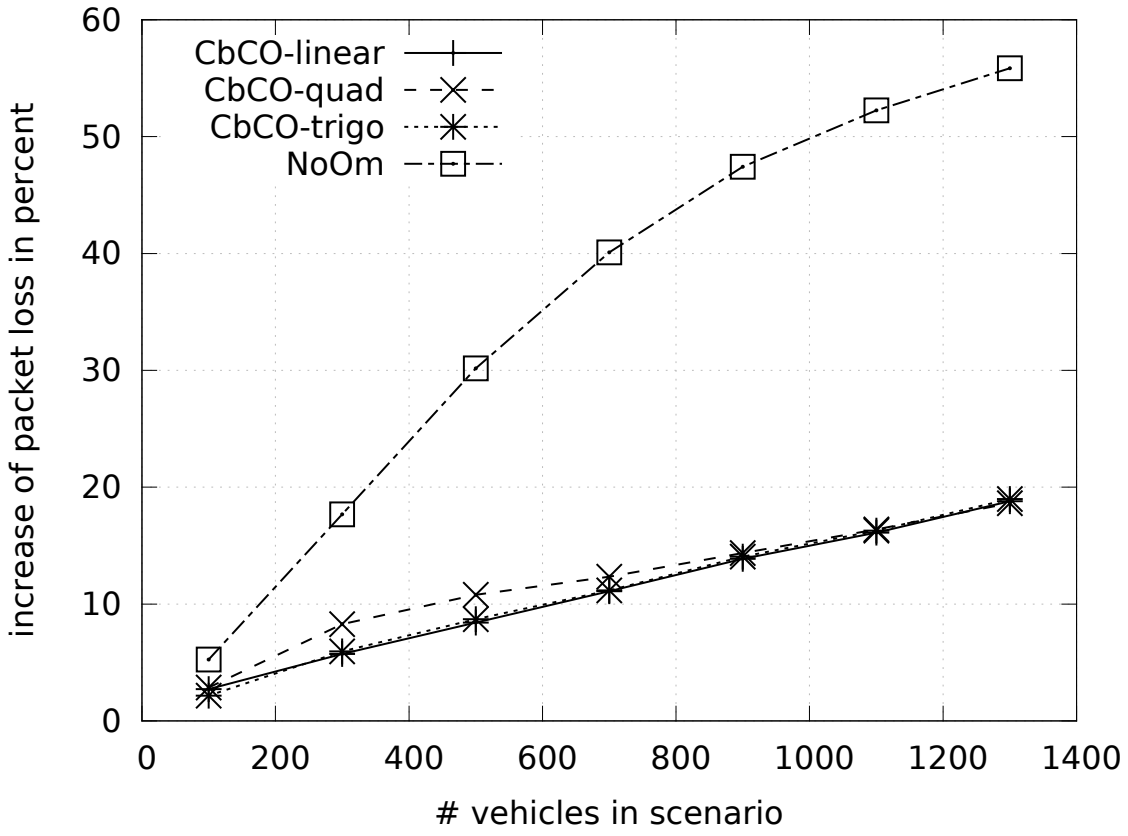

Figure 6.18: Increase of packet loss due to inclusion of certificates for different variants of CbCO, counting NPL + CPL

that the quadratic scheme showed a slightly worse overall increase of packet loss compared to the other approaches.

6.5.3.3 Comparison

To asses the utility of congestion based certificate omission we conduct comparisons to the previously proposed omissions schemes. We select Periodic Omission of Certificates (POoC) as described in [27], using the parameter $\alpha = 10$, and Neighor-based Certificate Omission proposed in [24]. Where applicable we also compare the schemes against a baseline of having no certificate omissions at all. An overview of the schemes is given in Table 6.3.

The basic percentage of certificates included in messages is an indicator of the performance of each scheme. In Figure 6.22 we remark that the congestion based omission scheme is converging to the same 90% omission rate as the POcC-10 scheme. On the other hand, the neighbor-based certificate omission scheme reduces omissions in densely

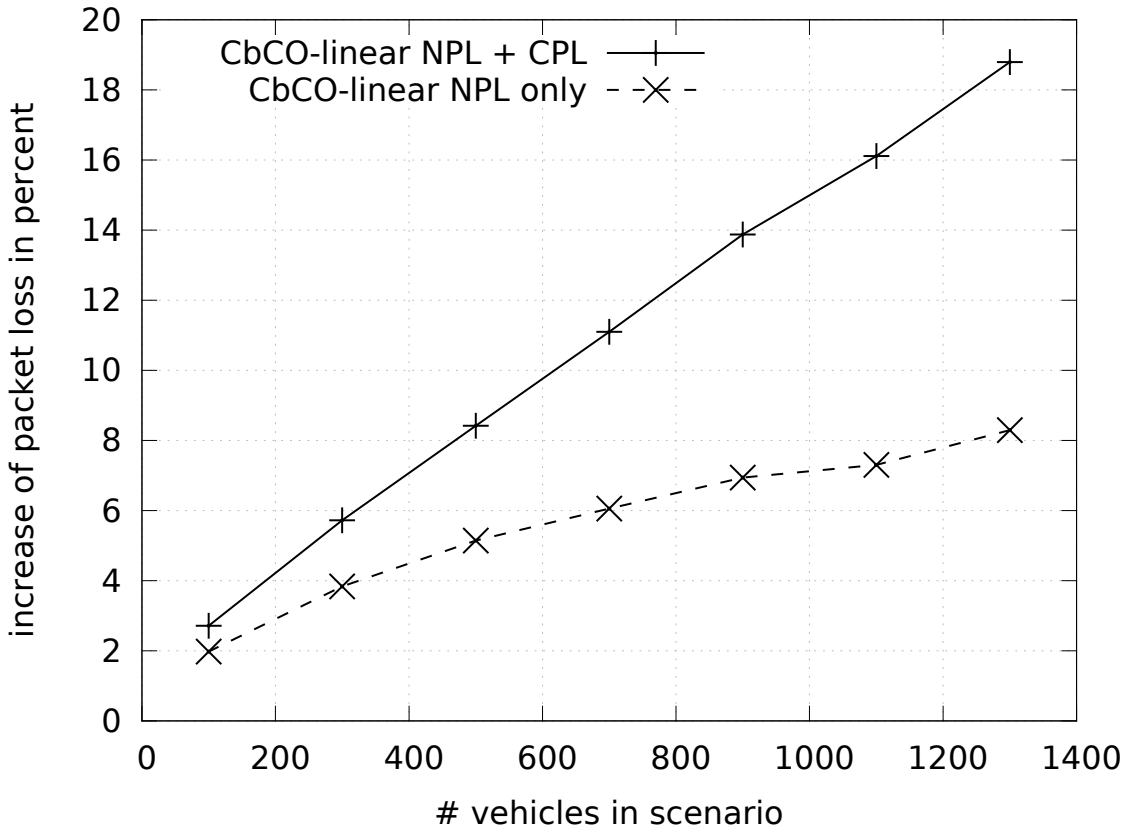

Figure 6.19: Illustration of the effect of counting cryptographic packet loss as regular packet loss

Table 6.3: Omission Schemes

Name	Options	Abbreviaion
Periodic Omission [27]	$\alpha = 10$	POoC-10
Periodic Omission [27]	$\alpha = 3$ [24]	POoC-3
Neighbor-based [24]	-	NbCO
Congestion-based	Linear	CbCO-linear
Congestion-based	Quadratic	CbCO-quad
Congestion-based	Trigonometric	CbCO-trig
No omissions	-	NoOm

populated scenarios due to the increased amount of neighbor changes in the network. This of course helps to keep down the CPL for the NbCO scheme, as can be seen in Figure 6.23. But the price for this low amount of CPL is a much higher amount of regular CPL due to collisions in the communication channel as can be seen in Figure 6.24.

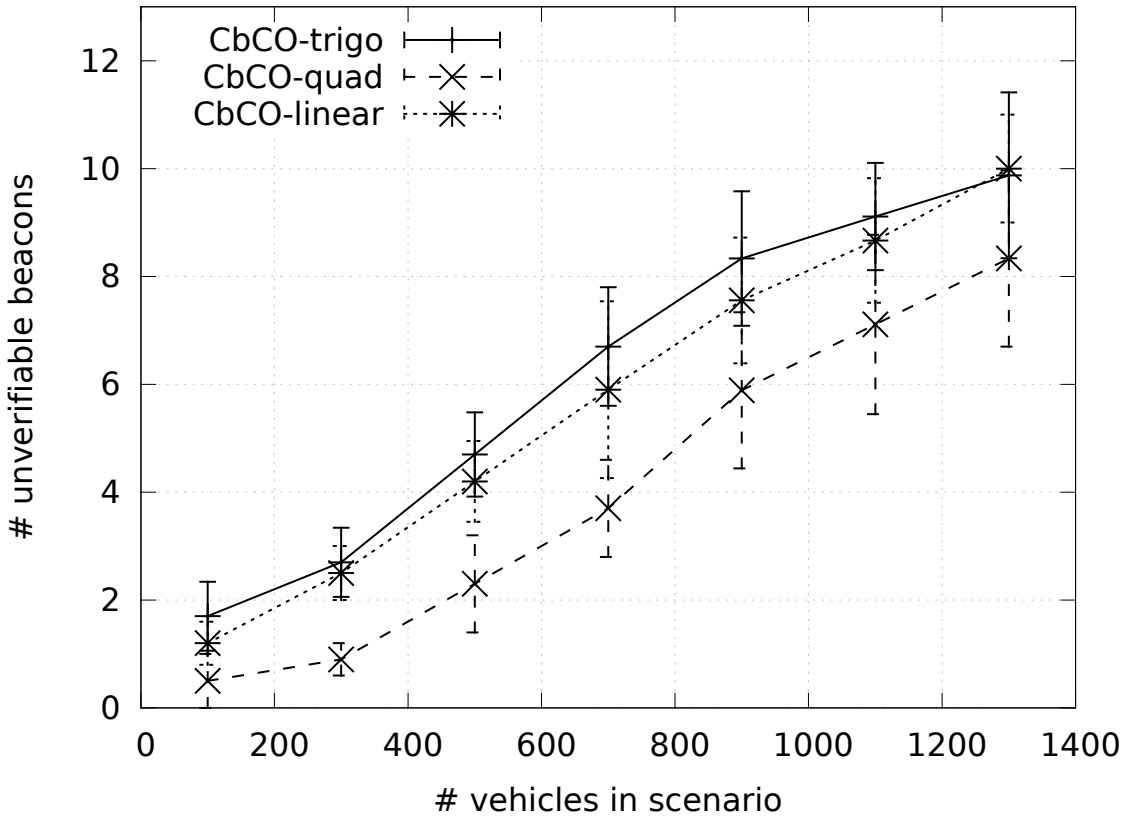

Figure 6.20: Average number of unverifiable beacons until arrival of certificate

All values are again relative to a baseline where we do not attach certificates at all (for NPL) or where every packet is assumed to be verifiable (for CPL). We note that the POoC scheme on the other hand performs well in terms of minimizing network packet loss but shows problems with regard to cryptographic packet loss.

Finally, Figure 6.25 presents an amortized total results for packet loss induced by certificate inclusion. In this graph we consider unverifiable packets to be cryptographic packet loss and see that the congestion based omission schemes deliver the best scalability in this overall view on the communication performance.

6.5.4 Conclusion and Future Work

We investigated the problem of scalability of security mechanisms in VANETs, especially with respect to communication overhead created by attaching certificates to all messages. Following earlier proposals, we suggest to adaptively omit certificates when sending beacons to reduce the channel load based on a Congestion-based Certificate Omission scheme (CbCO). This scheme uses an estimate of the channel congestion to decide

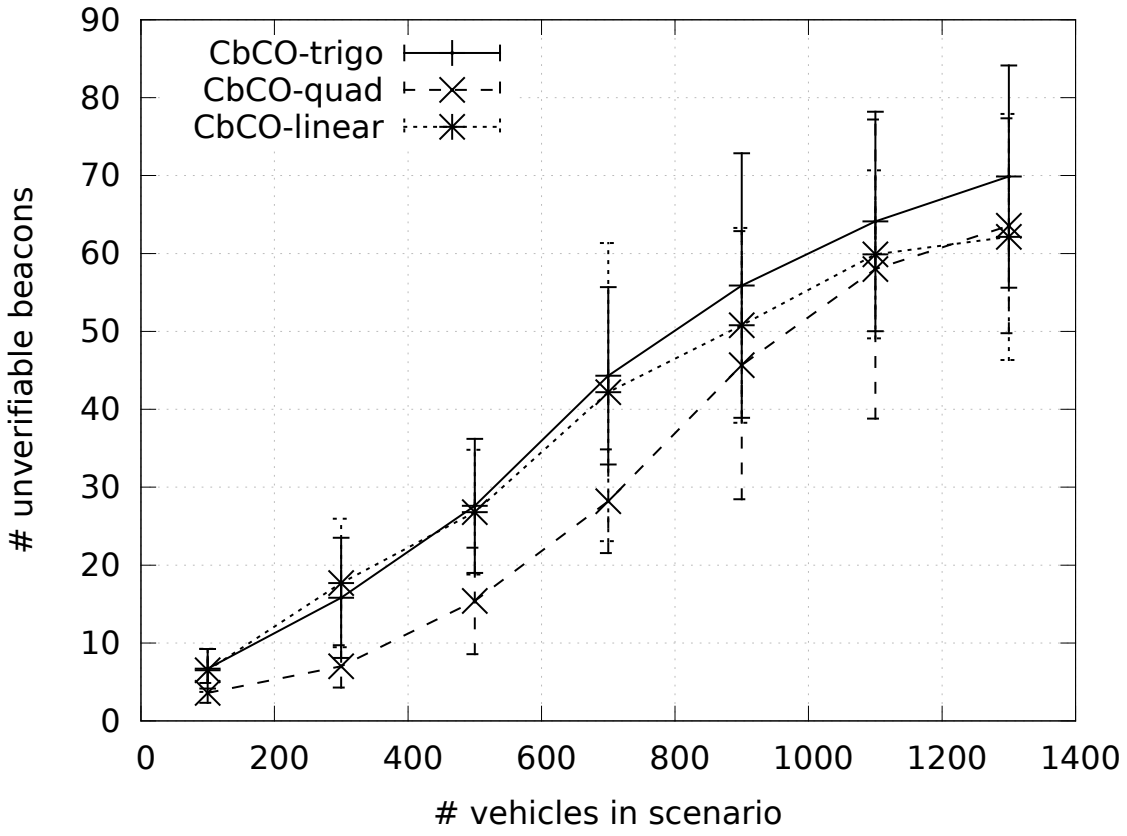

Figure 6.21: Maximum number of unverifiable beacons until arrival of certificate

whether to omit certificates. Using a simulation study, we investigate if the number of neighbors can be used to control the omission rate. The use of omission schemes leads to cryptographic latency due to intermittently missing certificates or even cryptographic packet loss if we consider unverifiable packets to be useless. Simulation results show that CbCO achieves a good balance between this effect and overall packet loss due to large messages. This shows that our scheme reduces the overall packet loss compared to the standard security mechanism that does not use certificate omission. Furthermore, we have shown that our schemes adapts better to varying vehicle densities than previous proposals.

As future work, we envision a cross-layer scheme in order to use more direct information about congestion in communication channels. This could be part of a larger effort to improve the overall quality of service in secure communication systems. Security components in communication systems can and should use cross layer information to make better decisions about security trade-offs while preserving a general separation of concerns. In this context we also propose to analyze the impact of higher bandwidths, adaptive beacon sizes and adaptive beaconing rates on the behavior of CbCO. Adaptive beaconing rates in particular represent a higher level omission scheme for entire beacons and it is

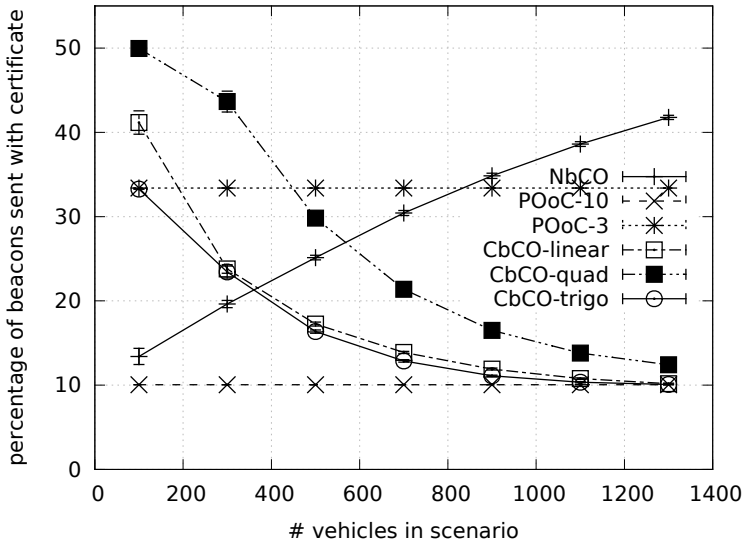

Figure 6.22: Average percentage of certificate omission in other protocols

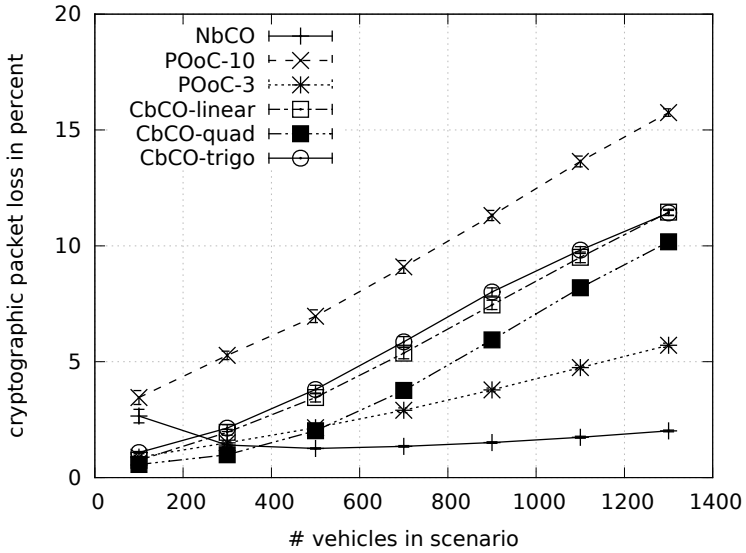

Figure 6.23: Average percent of unverifiable packets for various proposed omission schemes

necessary to investigate the effects of using omission schemes concurrently on multiple layers. While we see still some room for improvement, our results strongly suggest the consideration and adoption of certificate omission in IEEE and ETSI standards.

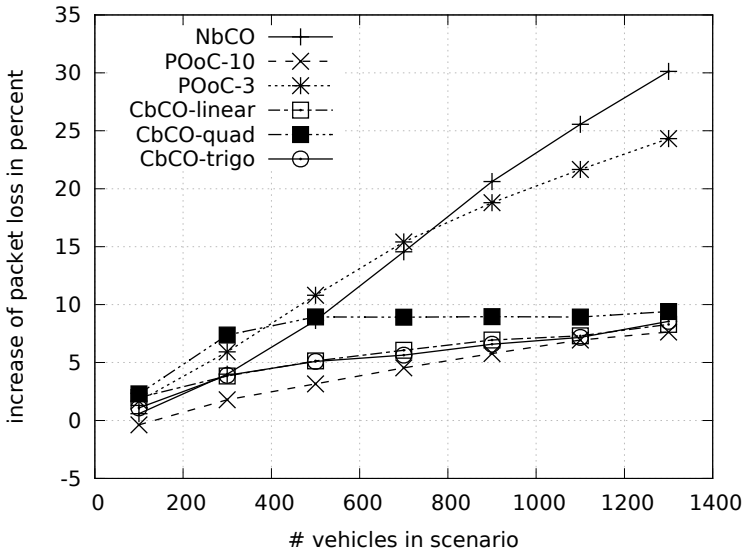

Figure 6.24: Increase of packet loss due to inclusion of certificates for different omission schemes (NPL only)

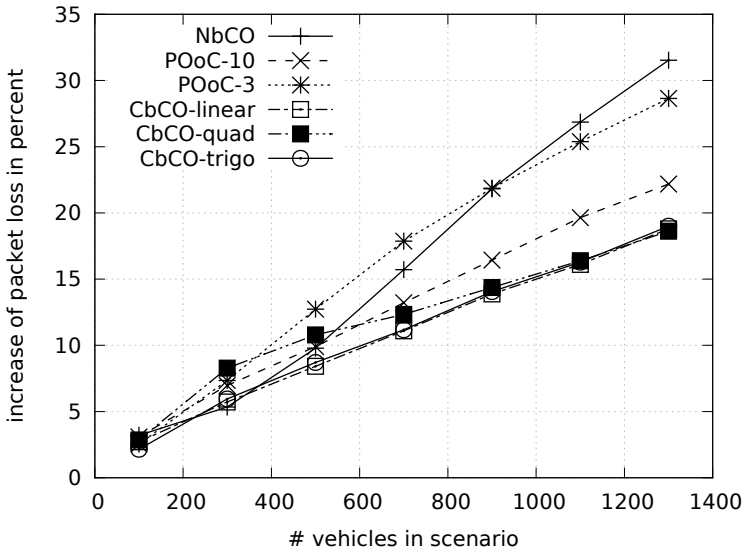

Figure 6.25: Increase of packet loss due to inclusion of certificates for different omission schemes, counting NPL + CPL

6.6 Pseudonym Management: a comparison and future challenges

In this section we complete the work done in D5.1 on pseudonym management. We provide a comparison and discussion of the four different pseudonym schemes before identifying future challenges.

6.6.1 Comparison and discussion

After providing a comprehensive overview of different pseudonym approaches, we now compare them with the help of the pseudonym lifecycle. Table 6.4 summarizes the key characteristics of each category.

The four presented categories all use pseudonyms, but the employed type of pseudonym differs. In asymmetric cryptographic schemes, a pseudonym is composed by a signature and a certificate. Therefore, the sender must have a valid public key certificate to be authenticated properly by receivers. The pseudonym used in IBC is shorter as no certificate is attached. Without certificate, the sender is then authenticated based on its identity. The group signature schemes use the same type of pseudonym as the asymmetric schemes, but the pseudonym's scope is broader as it affects every members of the group. The group certificate authenticates the group and that the sender is a valid member of the group. The symmetric cryptographic schemes use an even shorter pseudonym than IBC schemes since only a MAC is used. We now compare how the pseudonym is used in the four approaches.

Pseudonym issuance The asymmetric and identity-based approaches require backend connectivity for pseudonym issuance. Indeed, they require contact to PP or TA for pseudonym issuance. To overcome the issue of permanent infrastructure connectivity, pseudonym pre-loading or self-generation have been proposed. Nevertheless, preloading and self-generation have to be controlled to prevent Sybil attacks. On the opposite, group signature and symmetric approaches rely on vehicle collaboration or RSU contact for pseudonym issuance. At first glance, group signature and symmetric approaches appear as a cost-effective solution as they do not need infrastructure. However, vehicle collaboration raises new issues, such as group manager election. Moreover, the reliance on backend connectivity of the first two approaches is shifted to the reliance and stability of the group manager. So the core of the problem remains unchanged.

Pseudonym use Regarding the pseudonym use phase, the asymmetric approach requires a public key certificate to be attached to messages for sender authentication. Therefore, the communication overhead is higher than for identity-based or group signature approaches, which only require a small identifier and a signature. Symmetric approaches require only a MAC. To cope with the overhead issues of asymmetric schemes, certificate omission [25,33,34] and adaptive beaconing rates [35] have been proposed. By

omitting the certificate, the packet size is reduced and scalability is improved. The adaptive beaconing rate is even more drastic as some messages may be skipped completely. The group signature approach also aims for scalability as the pseudonym management is limited to the group size. Unfortunately, vehicular networks are highly dynamic and group management creates a significant computational overhead. Another issue of group signature schemes is the lack of linkability. As group members are sharing the public key, a receiver cannot distinguish the exact number of neighboring vehicles. This can jeopardize safety applications such as collision avoidance warnings or all applications that rely on knowledge of exact vehicle density. Regarding the symmetric cryptography schemes, protocols like TESLA delay key disclosure to enable use of symmetric keys, but require two successful packet receptions instead of one, which is problematic in delay-sensitive applications.

Pseudonym change Pseudonym change is mainly relevant for privacy in asymmetric and identity-based approaches to prevent vehicle tracking based on constant identifiers. Symmetric approaches also require change of symmetric keys but rather to limit the validity of keys and prevent impersonation attacks once a symmetric key has been released. Group signature approaches do not require a pseudonym change on the authentication level. Yet, static network identifiers could also allow tracking. Therefore, all approaches must change not only pseudonyms but also the vehicle's MAC address and other identifiers. In order to avoid tracking due to radio fingerprinting, switching between different radio modules has also been proposed, but given the state of fingerprinting in real outdoor scenarios, this seems currently a rather academic proposal of low practical relevancy [36]. Thus, one advantage of group signature schemes is removed by the requirement to prevent tracking attacks on lower layers. A lot of pseudonym change strategies are using asymmetric schemes, however, such strategies could also be applied to the other categories as well as change of network identifiers. Unfortunately those strategies can not fully prevent tracking. For example, eavesdroppers can analyze where users spend most of their time to discover their home address [37]. Selfish vehicles can also refuse participation in a cooperative pseudonym change and brake the mix-zone strategy [38]. But even if an OBU changes the entire communication stack identifiers (MAC address, IP address, etc.) in addition to the pseudonym, there might still be non-volatile data, such as tire pressure sensor IDs, that can serve as an attack surface [36]. An open challenge is to investigate which pseudonym change strategy is the most appropriate. A major step forward would be to reach consensus on the metric used to assess those strategies. We discuss this challenge in Section 6.6.3.

Pseudonym resolution The resolution phase either involves a single authority or multiple authorities that need to cooperate. By default, pseudonym resolution could be realized in all categories with simple identity escrow with only one authority. In order to protect against rogue authorities, new architectures are proposed to split responsibility. Regardless of the cacitetegory, these mechanisms usually involve multiple authorities in secret sharing and threshold cryptography schemes. Besides the technical aspect of resolution, there are potential legal issues. Indeed, pseudonym resolution strategies have to be in

line with legal regulations. For example in Europe, data in vehicular communication falls under the European data protection directive 95/46/EC [39], which restricts pseudonym resolution procedures by law enforcement accordingly.

Pseudonym revocation Revoking vehicles in the asymmetric approach implies management of a certificate revocation list. However, including all individual pseudonyms in such CRLs would significantly increase the CRL size. Therefore, revocation is typically limited to the VID, which is verified when obtaining new pseudonyms. Identity-based and symmetric schemes follow similar approaches of rather revoking the vehicle identifier than individual pseudonyms. In group signature approaches, revoking a vehicle provokes changes in the whole group as the group public key has to be updated. Another approach followed by the V-token [40] and CAMP [41] approaches is to insert a linkage value into each certificate. The linkage value is basically an encrypted identifier that remains secret until the revocation of the certificate. Hence, if the encryption key is included in the CRL, all future certificates used by the revoked vehicle can be recognized. This technique permits the revocation of all future messages while preserving privacy of past messages.

To conclude the discussion, symmetric schemes have substantial drawbacks and are not practical. One could notice that the establishment of shared secret keys for safety messages would add a non negligible delay. Group signature schemes provide interesting properties, but the high computation overhead is problematic. Asymmetric and IBC schemes are similar and according to our analysis the most viable approaches for realizing pseudonymity in vehicular networks.

6.6.2 Standardization

While specific privacy requirements and data protection legislation differ significantly between nations, the need to protect privacy in ITS is generally acknowledged. In Standards Development Organizations (SDOs), such as ETSI, ISO, or IEEE, approaches to protect privacy in vehicular networks are actively being discussed. However, privacy approaches are currently considered more at a preliminary stage rather than as part of drafts or final standards already. Yet, there is a clear trend to standardize privacy protection mechanisms for ITS based on pseudonyms.

In Europe, the ETSI Technical Committee on ITS Working Group 5 is responsible for designing a privacy solution. Technical Specification 102 941 [42] specifies a functional split between an enrollment authority and an authorization authority. This corresponds to a simple pseudonym scheme with a CA and PP where the enrollment authority manages vehicle identities and issues long-term certificates, while the authorization authority is responsible for verifying the long-term enrollment of vehicles and issuing short-term pseudonymous certificates that vehicles then use for message authentication. So far, this specification misses important aspects of the pseudonym lifecycle discussed throughout this paper like pseudonym resolution, protection from misuse by authorities, and even

pseudonym change. These issues need to be refined and worked out in future specifications. To support this refinement and extensions, the security working group of the C2C-CC[4] has created a "Public Key Infrastructure Memo" [43, 44] that discusses details about PKI operation and how to separate concerns between the CA (here long-term CA) and the pseudonym provider (here pseudonym CA) in order to prevent the pseudonym provider from learning the identities of vehicles it issues pseudonyms for, and the CA from learning the pseudonyms issued for those vehicles. It is expected that this solution will strongly influence the final approach of the ETSI standards.

In Japan, the standardization of vehicular communications is led by the ITS Forum. The standard ARIB STD-T109 specifies 700 MHz band ITS [45]. Unfortunately, the current version (July 2012) does not consider security and privacy.

In the U.S., vehicular networking security is defined in the standard IEEE 1609.2v2: Section D.2.6.4 [46], which does not include privacy mechanisms as the authors argue that privacy requirements are not clear yet. Nevertheless, the 1609 working group considers anonymity—the ability of private drivers to maintain a certain amount of privacy—as one goal of the system, but notes that "this goal conflicts to some extent with other goals, such as revocation of misbehaving nodes or supporting law enforcement access under appropriate circumstances. The exact tradeoff between these goals is seen as a policy matter, to be decided by stakeholders such as the U.S. vehicle OEMs and federal and state governments. As these stakeholders have not yet communicated their specific requirements to the 1609 WG, the 1609 WG decided not to include an anonymity mechanism that might fail to meet the eventual set of requirements. Anonymity will be addressed in a future version of or amendment to this standard." [46]

Similar to the C2C-CC efforts in Europe, the Crash Avoidance Metric Partnership (CAMP) consortium has provided a detailed specification of V2X trust management that also foresees pseudonym-based privacy protection. Their design specifies a detailed solution that separates concerns between Certification Authority (CA), Registration Authority (RA), and two Linkage Authorities (LA). It is to be expected that this proposal will significantly influence work on privacy protection in future versions of IEEE 1609.2.

In 2012, a joint harmonization task force has been set up by the US Department of Transportation and the European Commission. As part of this endeavor, a dedicated working group investigates how to harmonize the EU and US security solutions for vehicular networks. Privacy protection using a pseudonym scheme has been identified as one of the major areas requiring harmonization.

6.6.3 Research challenges

The previous sections show that protecting privacy in vehicular networks with pseudonyms still poses significant research challenges. Many challenges arise for specific categories, because they are shaped by the characteristics of the employed pseudonym type and

[4] Car-2Car Communications Consortium, www.car-2-car.org

underlying cryptographic primitives. However, a number of general open research challenges can be identified that require attention.

6.6.3.1 Considering pseudonym impact on communication stack and services

The purpose of pseudonym change is to (1) mask the change even from nearby vehicles (to prevent tracking) and to (2) prevent long-term tracking. The first goal creates issues for neighborhood-based mechanisms, like cooperative collision warning. Indeed, changing pseudonyms requires to flush the communication stack to change all layer identifiers, and to avoid sending messages with inconsistent sets of identifiers. Therefore, messages may get lost and routing tables will have inconsistent entries as a result of pseudonym changes [47]. Hence, pseudonym change strategies impact the communication stack and applications, and thus, require a tradeoff between application quality and privacy level, which should be adequately reflected in respective privacy metrics.

In the simTD project[5], this tradeoff is addressed by hiding pseudonym changes from the application layer. A translation table between lower layers and the application layer makes pseudonym changes transparent for applications [48]. Another option is to block pseudonym change for Decentralized Environmental Notification Messages (DENM). Indeed, if a vehicle has stopped at the roadside and sends DENMs to warn approaching vehicles and then changes its pseudonym, receivers will conclude that there are two broken-down vehicles. Plausibility checks could help to prevent such situations and need to be investigated, because blocking pseudonym change decreases the privacy level. So, an open challenge is to investigate how often and for how long vehicles can afford to block pseudonym change without too negative effects on their privacy protection [49]. A privacy metric that captures this tradeoff could help answering this question.

6.6.3.2 Enhancing scalability and reducing computation and communication overhead

When dealing with thousands of vehicles, scalability becomes an issue. The real-world performance of security mechanisms for vehicular networks have been analyzed by [50], [51], [52], and [53]. The results of these performance studies are summarized in Deliverable D1.1 [54] of the PRESERVE project[6]. A main result is the identification of an upper bound of about 1,000 verifications per second for an asymmetric cryptography scheme. Therefore, vehicles have to be capable of supporting such load to ensure a secured service. To help solving the scalability issue, a strategy is to reduce computation and communication overhead of security and privacy mechanisms, which is directly related to the use of pseudonym schemes. For instance, [55] analyze the effects of short-term pseudonyms on certificate revocation list size to highlight the relationship between privacy and security mechanisms. Indeed, depending on the policies for the number of pseudonyms carried by vehicles and the triggers for revoking certificates, the size of the CRL may grow very

[5] http://www.simtd.org
[6] http://www.preserve-project.eu

quickly. For example, when a CA issues pseudonyms for two hours of daily driving with a one year lifetime an hourly CRL would reach a size of over 2,200 MB. When a "valid after" field is added to the pseudonym to limit the lifetime of pseudonyms, a reduction down to 42 MB is achievable. Therefore, a cost model that assesses the impact of a pseudonym scheme on computation and communication overhead would be a key metric. Similarly, [56] verify the effectiveness and overhead of group signature schemes and conclude that the delay is significant and that mitigation techniques have to be studied.

6.6.3.3 Privacy metrics

Pseudonym change needs to be considered holistically in order to effectively provide privacy. To guide the selection of appropriate strategy, privacy metrics are proposed to assess the effectiveness of different pseudonym change strategies [57, 58]. For example, entropy [59], anonymity set size [60], or degree of location privacy [61] are privacy metrics used in the context of vehicular networks. Entropy assesses the level of usefulness of information and is often used to measure privacy. However, entropy is not an intuitive metric as it uses a logarithmic scale and is unbounded. Thus drawing conclusions from entropy values is difficult, because it is hard to relate to practical privacy implications. The anonymity set size (k-anonymity) is more intuitive as it represents the number of entities that are indistinguishable (e.g., due to using the same group key). A larger anonymity set (larger k) signifies better privacy. We refer the reader to [62] for details on k-anonymity, and its optimizations l-diversity and t-closeness. The degree of location privacy indicates how long an attacker could successfully track a vehicle. In the database community, differential privacy has emerged as a major privacy preservation technique [63]. Laplacian noise is added to a database in order to perturb information. A sensitivity parameter adjusts the level of added noise, e.g., depending on the number of entries in the database and the number of possible queries. While highly effective, the application to vehicular communications is not clear as differential privacy produces noisy data while vehicular communication requires accurate information. Moreover, there is no database or centralized query processor in distributed vehicular networks. Nevertheless, differential privacy could be applied to Floating Car Data (FCD) [64]. FCD is a valuable source of up-to-date traffic information, and therefore, has to be anonymized.

In general, there is a lack of consensus on suitable privacy metrics for vehicular networks, fostered by the fact that most metrics have only been validated in limited simulations and rarely in the wild. In particular, a comprehensive assessment of proposed pseudonym change strategies with consistent metrics is missing. ultimately, standard metrics and evaluation methods need to be identified and agreed upon, which can then be used to effectively evaluate pseudonym proposals in a comparable manner.

6.6.3.4 Fundamental relationship between pseudonym change strategies and privacy level

Pseudonym change is a critical phase of the lifecycle as it directly impacts the privacy level. If not properly set, the frequency of pseudonym changes can increase the linkability of a vehicle. Moreover, a pseudonym change strategy is defined by the rate of change and the context of the vehicle (location, density of neighbors, infrastructure availability and deployment). All these parameters are linked and their interdependencies have to be formalized. [65] analyze PKI approaches and their results indicate that privacy-preserving solutions should be based on one-time pseudonyms (i.e. one pseudonym per message sent), as the reuse of certificates is the key feature that enables their tracking attacks. Furthermore, one-time pseudonyms would provide forward (and backward) security properties: even if a vehicle is tracked in a trip (e.g., because it is traveling alone in the road and does not cross any other vehicles), one-time pseudonyms would not provide any useful information for tracking the past or future trips of that vehicle. Anonymous credentials are one way of implementing one-time pseudonyms with optional anonymity revocation. Nevertheless, one-time pseudonym will create "ghost vehicles" [66] inside the Local Dynamic Map (LDM)[7], which jeopardizes safety applications like cooperative collision warning. The fundamental relationship between pseudonym change strategies and the privacy level needs to be formalized. A set of standardized but diverse simulation/experiment parameters would help the comparison of pseudonym change strategies and facilitate the exposure of strengths and weaknesses in proposed strategies.

6.6.4 Conclusions

Safety-critical applications in cooperative vehicular networks require authentication of nodes and messages. Yet, privacy of individual vehicles and drivers must be maintained. Pseudonymity can combine security and privacy requirements. Thus, a large body of work emerged in recent years, proposing pseudonym solutions tailored to vehicular networks. In this paper, we provided a comprehensive survey on the complex topic of pseudonymity in vehicular networks. The proposed abstract pseudonym lifecycle is applicable to the majority of pseudonym approaches for vehicular networks and facilitates comparison and discussion of those approaches. We identified four major categories of pseudonym approaches that overlap with the dominant research directions: pseudonym schemes based on asymmetric cryptography and PKIs, identity-based cryptography schemes, group signature schemes, and schemes based on symmetric cryptography. We discussed each category by introducing its general concepts in relation to the pseudonym lifecycle, followed by a more detailed discussion of issues and optimizations for this category. The categorization and integrative discussion of contributions provides the opportunity to establish deeper insights into the pseudonym approaches in vehicular networks, their requirements, and challenges. Our discussion and the provided comparison table in Section 6.6.1 contrast

[7] The concept of Local Dynamic Map is specified in the Deliverable D3.3.3 of the SAFESPOT integrated project. [Available] http://www.safespot-eu.org/documents/D3.3.3_local-dynamic-map-spec.pdf

the four categories of pseudonym approaches and highlight their advantages and disadvantages. To foster further research in this area, we identified a number of challenges for future research, such as pseudonym change strategies, and reduction of computation and communication overhead. This survey also highlights the fact that current standardization efforts lack behind the research results regarding pseudonym solutions. Most notably, approaches beyond public key based schemes are hardly considered standardization efforts at the moment. We hope that our work is also recognized and considered helpful in standardization bodies and contributes to their work, eventually leading to secure and privacy preserving V2X systems. Therefore, an additional challenge to the research community is to demonstrate the feasibility of proposed pseudonym mechanisms in realistic settings to convincingly communicate advantages of specific contributions. Hence, suitable metrics need to be developed that capture the required utility-privacy tradeoff and can be used to compare the suitability of different proposals.

6.7 Differential Privacy for ITS

In 2012, researchers from UT and NICTA, Sydney, started to investigate how the concept of differential privacy can be applied to the field of Intelligent Transportation Systems (ITS), focusing on the protection of Floating Car Data that is stored and processed in central Traffic Data Centers. A first publication entitled “Differential Privacy in Intelligent Transportation Systems” [67] was submitted to ACM WiSec 2013 and accepted as short paper (acceptance rate of 35 %).

The PRECIOSA project has investigated privacy policy languages and policy enforcement as means to enforce privacy throughout a distributed, cooperative ITS. Personal identifiable information, e.g., contained in Floating Car Data records (FCD) could be combined with policies that would, e.g., ensure that k-anonymity was guaranteed whenever a specific FCD was processed.

However, this will not prevent a clever attacker from deriving sensitive information through a series of queries. As we show in [67], one may, e.g., infer speeds of single vehicles from a series of queries where each single query maintains k-anonymity. The concept of differential privacy /refDMNS06 promises to solve this issue by providing provable privacy guarantees. However, the concept is a rather abstract one and has not been applied to many practical application domains yet, including ITS.

In our work, we show how event-level differential privacy can be supported in ITS by extending privacy policy languages and policy-enforcement frameworks like the ones that were proposed by the PRECIOSA project. Data records are assigned a privacy budget ϵ that is spent through queries. Depending on the nature of the query and the extent of information it reveals to the querying party, this budget may be depleted earlier or later. Once it is depleted, the specific data record needs to be deleted and cannot be used anymore. If application requirements are known, one can extend the useable duration by reducing data accuracy by adding Laplace noise or by using only a subset of the FCD records to answer queries.

As an example, an FCD database may be queried to find out about the average speed driven on a specific stretch of road. Uncontrolled answering of multiple such queries can allow an attacker to derive precise information about speed driven by individual vehicles. All attackers would have to do is clever interleaving of the road segments that are queried. Even mandating k-anonymity for query responses will not prevent this. In our case, differential privacy would limit the number of queries that a specific FCD record is involved in and thus the accuracy by which an attacker can learn the (speed) value contained in this record. If applications can tolerate some reduced accuracy, one may simply add noise to the result or use only a subset of reported FCDs in the query road segment. Both measures will reduce the amount of information about this FCD record that an attacker can derive from queries and thus allows FCDs to be used for a larger number of queries before the privacy budget is depleted.

While this so-called event-level differential privacy can be ensured rather easily, we also discuss the challenges that support of user-level differential privacy poses and outline a potential solution. User-level differential privacy is concerned with the potential leakage of PII that comes from multiple FCD records of the same vehicle contained in the same database.

Furthermore, we also identify the differential privacy mechanisms which should be integrated within the policy-enforcement framework and provide guidelines for the calibration of parameters to provide specific privacy guarantees, while still providing the required accuracy of ITS applications. We illustrate the feasibility of our approach by discussing the steps needed to extend the PRECIOSA PeRA architecture to support differential privacy. As a result, we show that differential privacy could be put to practical use in ITS and will allow a stronger protection of personal data.

Our work has foundational character and should be seen as a first step in the direction of provable privacy protection for ITS. The work on this topic will continue into 2013 and may lead to a dedicated branch of research that will investigate the combination of differential privacy and policy enforcement.

Table 6.4: Overview of each approach

	Asymmetric	**Identity-based**	**Group sign.**	**Symmetric**
Pseudonym type	Asymmetric key pair, anonymous PKI certificate.	Pseudonymous node identifier as public key.	Group-wide public key.	Short-term symmetric keys
Authentication type	Sender has valid public key certificate	Sender can perform signature for pseudonym ID	Sender can perform signature for group public key	Symmetric key known to RSU.
Pseudonym issuance	Relies on PKI. Vehicles registered to CAs to obtain long-term identity. Pseudonyms are issued by PP. Frequent communication with PP required for pseudonym refill.	Pseudonym identifiers and corresponding private keys issued by TA (PP). TA can be authority or RSU. Frequent communication with TA required for pseudonym refill.	Group public key and individual private keys generated by GM. GM can be a vehicle, RSU or authority.	Vehicle registered with OM. RSU issues individual short-term symmetric keys in its region.
Pseudonym use	Sender generates asymmetric message signature and appends pseudonym certificate. Receiver verifies signature with pseudonym certificate.	Sender generates asymmetric message signature. Receiver verifies signature with sender's pseudonym identifier.	Sender generates asymmetric message signature. Receiver verifies signature with known group public key. Batch verification possible.	Sender generates MAC with individual symmetric key. Receiver waits for RSU verification or computes MAC after delayed key release.
Pseudonym change	Pseudonym change required to avoid tracking based on public key certificate. Different change strategies exist.	Pseudonym change required to avoid tracking based on identifier. Different change strategies exist.	No obvious need of pseudonym change as group signature ensures anonymity.	Symmetric key change needed to restrict key validity in space and time.
Pseudonym resolution	PP stores identity pseudonym mapping. Resolution can require cooperation of RAs.	TA stores identity pseudonym mapping. No cooperation required.	GM can determine individual signer key. No cooperation required.	RSU and OM cooperate in identity escrow.
Pseudonym revocation	Revocation of VID to prevent pseudonym refill. Possibly CRL to revoke individual pseudonyms	Revocation of VID to prevent pseudonym refill.	Change of group parameters by GM to evict node from group. Requires update of group public key.	Revocation of VID to prevent pseudonym refill.

Bibliography

[1] I. P1609.2/D12, "Draft Standard for Wireless Access in Vehicular Environments (WAVE)," Jan. 2012.

[2] ETSI TR 102 638, "Intelligent transport systems (ITS); vehicular communications; basic set of applications; definitions," Jun. 2009.

[3] C2C-CC, "C2C-CC public key infrastructure memo," CAR 2 CAR Communication Consortium, Tech. Rep., February 2011.

[4] IEEE, "Draft standard for wireless access in vehicular environments - security services for applications and management messages," Institute of Electrical and Electronics Engineers, Tech. Rep. 1609.2 - 20011 (D9), May 2011.

[5] ETSI - European Telecommunications Standards Institute, "Intelligent transport systems (ITS); security; security services and architecture," ETSI, Technical Standard TS 102 731, September 2010.

[6] ——, "Intelligent transport systems (ITS); security; its communications security architecture and security management," ETSI, Technical Report TR 102 940, June 2012.

[7] ——, "Intelligent transport systems (ITS); security; trust and privacy management," ETSI, Technical Report TR 102 941, June 2012.

[8] H. Baier and V. Karatsiolis, "Validity models of electronic signatures and their enforcement in practice," in *EuroPKI2009 - Sixth European Workshop on Public Key Services, Applications and Infrastructures*, September 2009.

[9] ETSI - European Telecommunications Standards Institute, "Intelligent transport systems (ITS); vehicular communications; basic set of applications; part 2: Specification of cooperative awareness basic service," ETSI, Technical Standard TS 102 637-2, April 2010.

[10] H. Stübing, A. Jaeger, N. Bißmeyer, C. Schmidt, and S. A. Huss, "Verifying mobility data under privacy considerations in Car-To-X communication," in *ITS World Congress*, vol. 17th ITS World Congress, Busan, October 2010.

[11] R. K. Schmidt, T. Leinmueller, E. Schoch, A. Held, and G. Schaefer, "Vehicle behavior analysis to enhance security in VANETs," in *Proceedings of the 4th IEEE Vehicle-to-Vehicle Communications Workshop (V2VCOM2008)*, 2008.

[12] T. Leinmüller, E. Schoch, and F. Kargl, "Position verficiation approaches for vehicular ad hoc networks," *Wireless Communications, IEEE*, vol. 13, no. 5, pp. 16 –21, october 2006.

[13] G. Yan, S. Olariu, and M. C. Weigle, "Providing VANET security through active position detection," *Computer Communications*, vol. 31, no. 12, pp. 2883 – 2897, 2008, mobility Protocols for ITS/VANET.

[14] K. Hsiao, J. Miller, and H. de Plinval-Salgues, "Particle filters and their applications," *Cognitive Robotics, April*, 2005.

[15] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge: MIT Press, 2005.

[16] U. D. of Transportation Research and I. T. Administration, "Security credential management system design security system design for cooperative vehicle-to-vehicle crash avoidance applications using 5.9 ghz dedicated short range communications (dsrc) wireless communications," CAMP, VSC3, www.its.dot.gov, Tech. Rep., February 2012.

[17] N. Bißmeyer, C. Stresing, and K. Bayarou, "Intrusion detection in vanets through verification of vehicle movement data," in *Second IEEE Vehicular Networking Conference*, vol. Second IEEE Vehicular Networking Conference, December 2010.

[18] R. Schmidt, T. Leinmüller, and A. Held, "Defending against roadside attackers," in *In proceedings of 16th World Congress on Intelligent Transport Systems*, 2009.

[19] B. Wiedersheim, Z. Ma, F. Kargl, and P. Papadimitratos, "Privacy in inter-vehicular networks: Why simple pseudonym change is not enough," in *Wireless On-demand Network Systems and Services (WONS), 2010 Seventh International Conference on*, feb. 2010, pp. 176 –183.

[20] ETSI, "Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service," 2010.

[21] ——, "ETSI TS 102 731 V1.1.1; Intelligent Transport Systems (ITS); Security; Security Services and Architecture," September 2010.

[22] 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," *Communications Magazine, IEEE*, vol. 46, no. 11, pp. 100–109, november 2008.

[23] F. Kargl, E. Schoch, B. Wiedersheim, and T. Leinmüller, "Secure and Efficient Beaconing for Vehicular Networks (Short Paper)," in *5th ACM International Workshop on Vehicular Ad Hoc Networks (VANET 2008)*. San Francisco, USA: ACM, September 2008. [Online]. Available: http://doi.acm.org/10.1145/1410043.1410060

[24] E. Schoch and F. Kargl, "On the efficiency of secure beaconing in vanets," in *Proceedings of the third ACM conference on Wireless network security*, ser. WiSec '10. New York, NY, USA: ACM, 2010, pp. 111–116. [Online]. Available: http://doi.acm.org/10.1145/1741866.1741885

[25] G. Calandriello, P. Papadimitratos, J.-P. Hubaux, and A. Lioy, “Efficient and robust pseudonymous authentication in vanet,” in *Proceedings of the fourth ACM international workshop on Vehicular ad hoc networks*, ser. VANET ’07. New York, NY, USA: ACM, 2007, pp. 19–28. [Online]. Available: http://doi.acm.org/10.1145/1287748.1287752

[26] P. Papadimitratos, “ "On the Road" - reflections on the security of vehicular communication systems,” *Vehicular Electronics and Safety, 2008. ICVES 2008. IEEE International Conference on*, pp. 359–363, sept. 2008.

[27] G. Calandriello, P. Papadimitratos, J.-P. Hubaux, and A. Lioy, “On the performance of secure vehicular communication systems,” *Dependable and Secure Computing, IEEE Transactions on*, vol. 8, no. 6, pp. 898 –912, nov.-dec. 2011.

[28] R. Barr, Z. J. Haas, and R. van Renesse, *Scalable Wireless Ad hoc Network Simulation*. CRC Press, Aug. 2005, ch. 19, pp. 297–311. [Online]. Available: http://www.amazon.com/Handbook-Theoretical-Algorithmic-Wireless-Networks/dp/0849328322

[29] D. R. Choffnes and F. E. Bustamante, “An integrated mobility and traffic model for vehicular wireless networks,” in *Proceedings of the 2nd ACM international workshop on Vehicular ad hoc networks*, ser. VANET ’05. New York, NY, USA: ACM, 2005, pp. 69–78. [Online]. Available: http://doi.acm.org/10.1145/1080754.1080765

[30] E. Schoch, M. Feiri, F. Kargl, and M. Weber, “Simulation of ad hoc networks: ns-2 compared to jist/swans,” in *First International Conference on Simulation Tools and Techniques for Communications, Networks and Systems (SimuTools 2008)*, Marseilles, France, Mar. 2008.

[31] F. Martelli, M. E. Renda, and P. Santi, “Measuring IEEE 802.11 p Performance for Active Safety Applications in Cooperative Vehicular Systems,” *iitcnrit*, pp. 2–6, 2011. [Online]. Available: http://www.iit.cnr.it/staff/paolo.santi/papers/VTC2011.pdf

[32] SAE International, “DSRC Implementation Guide - A guide to users of SAE J2735 message sets over DSRC,” Tech. Rep. v20, February 2010. [Online]. Available: http://www.sae.org/standardsdev/dsrc/DSRCImplementationGuide.pdf

[33] E. Schoch and F. Kargl, “On the efficiency of secure beaconing in VANETs,” in *Proceedings of the third ACM conference on Wireless network security (WiSec ’10)*. Hoboken, New Jersey, USA: ACM, 2010, pp. 111–116. [Online]. Available: http://doi.acm.org/10.1145/1741866.1741885

[34] M. Feiri, J. Petit, and F. Kargl, “Congestion-based certificate omission in vanets,” in *Proceedings of the ninth ACM international workshop on Vehicular inter-networking, systems, and applications (VANET ’12)*. Low Wood Bay, Lake District, UK: ACM, 2012, pp. 135–138. [Online]. Available: http://doi.acm.org/10.1145/2307888.2307915

[35] R. Schmidt, T. Leinmuller, E. Schoch, F. Kargl, and G. Schafer, “Exploration of adaptive beaconing for efficient intervehicle safety communication,” *IEEE Network*, vol. 24, no. 1, pp. 14–19, 2010.

[36] I. Rouf, R. Miller, H. Mustafa, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe, and I. Seskar, "Security and privacy vulnerabilities of in-car wireless networks: a tire pressure monitoring system case study," in *19th USENIX conference on Security (USENIX Security '10)*. Berkeley, CA, USA: USENIX Association, 2010, pp. 21–21. [Online]. Available: http://dl.acm.org/citation.cfm?id=1929820.1929848

[37] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady, "Enhancing security and privacy in traffic-monitoring systems," *Pervasive Computing, IEEE*, vol. 5, no. 4, pp. 38–46, oct.-dec. 2006.

[38] J. Freudiger, M. H. Manshaei, J.-P. Hubaux, and D. C. Parkes, "On non-cooperative location privacy: a game-theoretic analysis," in *Proceedings of the 16th ACM conference on Computer and communications security*, ser. CCS '09. Chicago, Illinois, USA: ACM, 2009, pp. 324–337. [Online]. Available: http://doi.acm.org/10.1145/1653662.1653702

[39] European Commission, "Directive 95/46/EC of the european parliament and of the council of 24 october 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data," *Official Journal of the European Union*, vol. L 281, pp. 31–50, October 1995. [Online]. Available: http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:EN:HTML

[40] F. Schaub, F. Kargl, Z. Ma, and M. Weber, "V-tokens for Conditional Pseudonymity in VANETs," in *IEEE Wireless Communications & Networking Conference (WCNC '10)*. Sydney, Australia: IEEE, 2010. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5506126

[41] US DOT, "Security Credential Management System Design: Security system design for cooperative vehicle-to-vehicle crash avoidance applications using 5.9 GHz Dedicated Short Range Communications (DSRC) wireless communications," US Department of Transportation, Draft, 2012.

[42] ETSI TC ITS, "ETSI TS 102 941 v1.1.1 - intelligent transport systems (ITS); security; trust and privacy management," http://www.etsi.org/deliver/etsi_ts/102900_102999/102941/01.01.01_60/ts_102941v010101p.pdf, TC ITS, june 2012. [Online]. Available: http://www.etsi.org/deliver/etsi_ts/102900_102999/102941/01.01.01_60/ts_102941v010101p.pdf

[43] C2C-CC, "Public key infrastructure memo," Car 2 Car Communication Consortium, Tech. Rep., 2010.

[44] N. Bissmeyer, H. Stübing, E. Schoch, S. Götz, and B. Lonc, "A generic public key infrastructure for securing car-to-x communication," in *18th World Congress on Intelligent Transport Systems*. Orlando, USA: ITS America, 2011. [Online]. Available: http://itswc.confex.com/itswc/WC2011/webprogram/Paper2472.html

[45] ARIB, "T109: 700 mhz band intelligent transport systems; v1.0," http://www.arib.or.jp/english/html/overview/doc/1-STD-T109v1_0.pdf, ARIB, feb 2012. [Online]. Available: http://www.arib.or.jp/english/html/overview/doc/1-STD-T109v1_0.pdf

[46] IEEE, “Trial-use standard for wireless access in vehicular environments - security services for applications and management messages,” Institute of Electrical and Electronics Engineers, Tech. Rep. 1609.2 - 2006, July 2006.

[47] E. Schoch, F. Kargl, T. Leinmüller, S. Schlott, and P. Papadimitratos, “Impact of pseudonym changes on geographic routing in vanets,” in *Third European Workshop on Security and Privacy in Ad hoc and Sensor Networks (ESAS ’06)*, Hamburg, Germany, 2006. [Online]. Available: http://infoscience.epfl.ch/record/94376

[48] A. Jaeger, N. Bißmeyer, H. Stübing, and S. Huss, “A novel framework for efficient mobility data verification in vehicular ad-hoc networks,” *International Journal of Intelligent Transportation Systems Research*, vol. 10, pp. 11–21, 2012. [Online]. Available: http://dx.doi.org/10.1007/s13177-011-0038-9

[49] Z. Ma, F. Kargl, and M. Weber, “Pseudonym-on-demand: a new pseudonym refill strategy for vehicular communication,” in *2nd IEEE International Symposium on Wireless Vehicular Communications (WiVec ’08)*. Calgary, Canada: IEEE VTS, September 2008, pp. 1–5.

[50] A. Iyer, A. Kherani, A. Rao, and A. Karnik, “Secure v2v communications: Performance impact of computational overheads,” in *INFOCOM Workshops 2008, IEEE*, april 2008, pp. 1–6.

[51] J. Haas, Y.-C. Hu, and K. Laberteaux, “Real-world VANET security protocol performance,” in *IEEE Global Telecommunications Conference. GLOBECOM 2009.*, 30 2009-dec. 4 2009, pp. 1–7.

[52] J. Petit, “Analysis of ECDSA Authentication Processing in VANETs,” in *Proceedings of the 3rd IFIP International Conference on New Technologies, Mobility and Security (NTMS ’09)*, Cairo, Egypt, 2009, pp. 388–392. [Online]. Available: http://dl.acm.org/citation.cfm?id=1790418

[53] J. Petit and Z. Mammeri, “Analysis of Authentication Overhead in Vehicular Networks,” in *Proceedings of the 3rd IFIP Wireless and Mobile Networking Conference (WMNC ’10)*, Budapest, Hungary, 2010, pp. 1–8. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5678756

[54] J. P. Stotz, N. Bißmeyer, F. Kargl, S. Dietzel, P. Papadimitratos, and C. Schleiffer, “PRESERVE D1.1 Security Requirements of Vehicle Security Architecture,” PRESERVE consortium, Deliverable, July 2011.

[55] M. Nowatkowski, J. Wolfgang, C. McManus, and H. Owen, “The effects of limited lifetime pseudonyms on certificate revocation list size in VANETs,” in *Proceedings of the IEEE SoutheastCon 2010 (SoutheastCon)*, march 2010, pp. 380–383.

[56] B. Chaurasia, S. Verma, and S. Bhasker, “Message broadcast in VANETs using group signatures,” in *Fourth International Conference on Wireless Communication and Sensor Networks (WCSN ’08)*, dec. 2008, pp. 131–136.

[57] Z. Ma, F. Kargl, and M. Weber, "Measuring location privacy in V2X communication systems with accumulated information," in *The Sixth IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS'09)*. Macau: IEEE, 2009.

[58] ——, "Measuring long-term location privacy in vehicular communication systems," *Computer Communications*, vol. 33, no. 12, pp. 1414–1427, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0140366410001076

[59] B. Chaurasia and S. Verma, "Optimizing pseudonym updation for anonymity in VANETs," in *IEEE Asia-Pacific Services Computing Conference (APSCC '08)*, dec. 2008, pp. 1633–1637.

[60] D. Chaum, "The dining cryptographers problem: unconditional sender and recipient untraceability," *J. Cryptol.*, vol. 1, no. 1, pp. 65–75, Mar. 1988. [Online]. Available: http://dl.acm.org/citation.cfm?id=54235.54239

[61] A. Wasef and X. S. Shen, "REP: Location Privacy for VANETs Using Random Encryption Periods," *Mob. Netw. Appl.*, vol. 15, pp. 172–185, 2010. [Online]. Available: http://dx.doi.org/10.1007/s11036-009-0175-4

[62] B. C. M. Fung, K. Wang, R. Chen, and P. S. Yu, "Privacy-preserving data publishing: A survey of recent developments," *ACM Comput. Surv.*, vol. 42, no. 4, pp. 14:1–14:53, Jun. 2010. [Online]. Available: http://doi.acm.org/10.1145/1749603.1749605

[63] C. Dwork, "Differential privacy," in *ICALP (2)*, ser. Lecture Notes in Computer Science, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds., vol. 4052. Springer, 2006, pp. 1–12.

[64] S. Rass, S. Fuchs, M. Schaffer, and K. Kyamakya, "How to protect privacy in floating car data systems," in *Proceedings of the fifth ACM international workshop on VehiculAr Inter-NETworking (VANET '08)*. San Francisco, California, USA: ACM, 2008, pp. 17–22. [Online]. Available: http://doi.acm.org/10.1145/1410043.1410047

[65] C. Troncoso, E. Costa-Montenegro, C. Diaz, and S. Schiffner, "On the difficulty of achieving anonymity for vehicle-2-x communication," *Comput. Netw.*, vol. 55, pp. 3199–3210, October 2011. [Online]. Available: http://dx.doi.org/10.1016/j.comnet.2011.05.004

[66] N. Bissmeyer, J. Njeukam, J. Petit, and K. M. Bayarou, "Central misbehavior evaluation for vanets based on mobility data plausibility," in *Proceedings of the ninth ACM international workshop on Vehicular inter-networking, systems, and applications*, ser. VANET '12. Low Wood Bay, Lake District, UK: ACM, 2012, pp. 73–82. [Online]. Available: http://doi.acm.org/10.1145/2307888.2307902

[67] F. Kargl, A. Friedman, and R. Borelli, "Differential privacy in intelligent transportation systems," in *ACM WiSec*, 2013 (to appear).