



Misbehaviour Detection

5GAA Automotive Association

White Paper



CONTACT INFORMATION:
Executive Manager – Thomas Linget
Email: liaison@5gaa.org

MAILING ADDRESS:
5GAA c/o MCI Munich
Neumarkter Str. 21
81673 München, Germany
www.5gaa.org

Copyright © 2022 5GAA. All Rights Reserved.

No part may be reproduced except as authorised by written permission. The copyright and the foregoing restriction extend to reproduction in all media.

VERSION:	1.0
DATE OF PUBLICATION:	13 July 2022
DOCUMENT TYPE:	White Paper
EXTERNAL PUBLICATION:	Yes
DATE OF APPROVAL BY 5GAA BOARD:	31 May 2022

Contents

Scope	04
References	05
1. Preliminaries	07
1.1. Terms and Definitions	07
1.2. Abbreviations	07
2. Background	09
2.1. General	09
2.2. Causes of Misbehaviour	10
3. Architecture of Misbehaviour Management System	13
3.1. Overview	13
3.2. Detailed View of Misbehaviour Management	13
3.3. Local Misbehaviour Management	17
3.3.1. Overview	17
3.3.2. Local Misbehaviour Detection	17
3.3.3. Local Misbehaviour Reporting	18
4. Prior and Related Work	20
4.1. Regulations and Standards	20
4.1.1. Regulations	20
4.1.2. Standards	20
4.2. Misbehaviour Management	21
4.2.1. ITS-S Misbehaviour Management	21
4.2.2. Backend Misbehaviour Management	22
4.2.2.1. Misbehaviour Pre-processing	22
4.2.2.2. Misbehaviour Authority	23
4.2.2.3. Misbehaviour remediation	24
4.2.3. Misbehaviour Organisational Management	24
4.2.3.1. Audit	24
4.2.3.2. Policies	25
5. Policy Questions	26
5.1. Local Misbehaviour Detection	26
5.2. Reporting	27
5.3. Remediation	27
6. Conclusions and Next Steps	28



Scope

In the context of this white paper, 'misbehaviour' refers to the wilful or inadvertent transmission of incorrect data within the V2X network, both to a vehicle and the Misbehaviour Authority.

Note that the definition of misbehaviour is quite specific and narrow here, i.e. strictly related to incorrect content in V2X packets and misbehaviour reports, in contrast to the general literal meaning of misbehaviour as something outside of the norm, or specification.

Therefore, this white paper addresses the following aspects:

- Misbehaviour and attacker models
- Architecture of misbehaviour management system
- Prior and related work
- Policy questions

Aspects that are **not** in scope for this white paper are:

- Physical, network or protocol layer misbehaviour (e.g. incorrect radio frequency, incorrect transmission rate)
- Certain cybersecurity attacks unrelated to message content (e.g. denial of service)
- Research, requirements, design or implementation of a misbehaviour detection system

References

The following documents contain provisions which, through reference in this text, constitute provisions of the present document.

- References are either specific (identified by date of publication, edition number, version number, etc.) or non specific.
- For a specific reference, subsequent revisions do not apply.
- For a non-specific reference, the latest version applies.

[Brecht+2018] Brecht, B., Therriault, D., Weimerskirch, A., Whyte, W., Kumar, V., Hehn, T., Goudy, R. A Security Credential Management System for V2X Communications. IEEE Trans. Intell. Transp. Syst. 19(12): 3850-3871 (2018).

[CP2019] ANNEX 3 to the Commission Delegated Regulation Supplementing Directive 2010/40/EU of the European Parliament and of the Council with regard to the deployment and operational use of cooperative intelligent transport systems. https://cpoc.jrc.ec.europa.eu/data/documents/c-its_certificate_policy_release_preparatory_phase_of_Delegated_Regulation_2019_1789.pdf

[Ghaleb+2019] Ghaleb, F.A., Maarof, M.A., Zainal, A., Al-Rimy, B.A.S., Saeed, F. and Al-Hadhrami, T., 2019. Hybrid and multifaceted context-aware misbehavior detection model for vehicular ad hoc network. IEEE Access.

[GPSVideo] GPS Spoofing Video. Analysing Tesla Under GPS Spoofing Attack - ION GNSS 2020 - Regulus Cyber. https://www.youtube.com/watch?v=_7N-XE0DA-l&t=450s.

[IEEE1609.2] IEEE 1609.2-2016. IEEE Standard for Wireless Access in Vehicular Environments - Security Services for Applications and Management Messages. <https://standards.ieee.org/ieee/1609.2/6038/>

[Heijden+2018] van der Heijden, R.W., Dietzel, S., Leinmüller, T. and Kargl, F., 2018. Survey on misbehaviour detection in cooperative intelligent transportation systems. IEEE Communications Surveys & Tutorials.

[ISO21434] ISO/SAE 21434:2021. Road vehicles - Cybersecurity engineering. <https://www.iso.org/standard/70918.html>

[ISO24089] ISO/DIS 24089. Road vehicles - Software update engineering. <https://www.iso.org/standard/77796.html>

[Kamel+2019] Kamel, J., Haidar, F., Jemaa, I.B., Kaiser, A., Lonc, B. and Urien, P. A misbehaviour authority system for sybil attack detection in c-its. In 2019 IEEE 10th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON).

[Kamel+2020] Kamel, J., Wolf, M., van der Hei, R.W., Kaiser, A., Urien, P. and Kargl, F. VeReMi extension: A dataset for comparable evaluation of misbehaviour detection in VANETs. In 2020 IEEE International Conference on Communications (ICC).

[Michelson+2019] Michelson, D., Noori, H., and Ramsay, Q. Interference detection and reporting in IEEE 802.11 p connected vehicle networks. In: 2019 IEEE 90th Vehicular

Technology Conference (VTC2019-Fall).

[Moalla+2012] Moalla, R., Labiod, H., Lonc, B. and Simoni, N. Risk analysis study of its communication architecture. In 2012 Third International Conference on The Network of the Future.

[Monteuuis+2018] Monteuuis, J.P., Petit, J., Zhang, J., Labiod, H., Mafrica, S. and Servel, A., 2018, September. Attacker model for connected and automated vehicles. In ACM Computer Science in Vehicle Symposium

[Noori+2020] Noori, H., Michelson, D. and Henry, K. Reporting Spectrum Misbehaviour using the IEEE 1609 Security Credential Management System. In: 2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring).

[RFC3647] Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework. <https://datatracker.ietf.org/doc/html/rfc3647>

[Shen+2020] Shen, J., Won, Jun Y., Chen, Z. Drift with Devil: Security of {Multi-Sensor} Fusion based Localisation in High-Level Autonomous Driving under GPS Spoofing. In 29th USENIX Security Symposium (USENIX Security 20), 2020.

[Spectrum2021] Spectrum Misbehaviour Report. Prepared by: ETAS Embedded Systems Canada Inc. Prepared for: Transport Canada Contract No: T8080-180316, March 2021. <https://tcdocs.ingeniumcanada.org/sites/default/files/2021-07/Security%20Credential%20Management%20System%20%28SCMS%29%20%E2%80%93%20Spectrum%20Misbehaviour%20Report.PDF>

[Trauernicht+2019] Trauernicht, J. and Bißmeyer, N. 'Deterministic sybil attack exclusion in cooperative-intelligent transportation systems'. In 2019 17th ESCAR Europe.

[TR103460] ETSI TR 103 460 V2.1.1. Intelligent Transport Systems (ITS); Security; Pre-standardisation study on Misbehaviour Detection; Release 2. https://www.etsi.org/deliver/etsi_tr/103400_103499/103460/02.01.01_60/tr_103460v020101p.pdf

[TS103097] ETSI TS 103 097 V2.1.1. Intelligent Transport Systems (ITS); Security; Security header and certificate formats; Release 2. https://www.etsi.org/deliver/etsi_ts/103000_103099/103097/02.01.01_60/ts_103097v020101p.pdf

[TS103759] DTS/ITS-00561. Intelligent Transport Systems (ITS); Security; Misbehaviour Reporting service; Release 2. https://portal.etsi.org/eWPM/index.html#/home?WKI_ID=59560

[UNR155] UN Regulation No. 155 - Cyber security and cyber security management system. <https://unece.org/transport/documents/2021/03/standards/un-regulation-no-155-cyber-security-and-cyber-security>

[UNR156] UN Regulation No. 156 - Software update and software update management system. <https://unece.org/transport/documents/2021/03/standards/un-regulation-no-156-software-update-and-software-update>

[Zeng+2018] Zeng, K. C., Liu, S., Shu, Y. All Your GPS belong to us: Towards stealthy manipulation of road navigation systems. In 27th USENIX security symposium 2018.

1. Preliminaires

1.1. Terms and Definitions

Many terms used in this document are explained in the WG2 document 5GAA_A-170188_V2XDEF_TR, '5GAA V2X Terms and Definitions'. The following definitions also apply:

GMBD: The act of detecting misbehaviour by the Misbehaviour Authority, by investigating and corroborating misbehaviour reports.

LMBD: The act of detecting misbehaviour locally by the receiver of suspicious V2X packets.

Misbehaviour: Wilful or inadvertent transmission of incorrect data within the V2X network.

Misbehaviour Authority: A component of the Public Key Infrastructure that receives reports of malicious or potentially malicious application activities, analyses them, and determines whether to take mitigating actions.

1.2. Abbreviations

ASN.1	Abstract Syntax Notation One
BSM	Basic Safety Message
CAM	Cooperative Awareness Message
DENM	Decentralised Environmental Notification Message
DoS	Denial of Service
ETSI	European Telecommunications Standards Institute
FW	Firmware
GMBD	Global Misbehaviour Detection
GPS	Global Positioning System
HW	Hardware
IEEE	Institute of Electrical and Electronics Engineers
ISO	International Organisation for Standardisation
ITS	Intelligent Transport Services
ITS-S	ITS Station
LMBD	Local Misbehaviour Detection
MA	Misbehaviour Authority
MBD	Misbehaviour Detection
ML	Machine Learning
OBU	Onboard Unit
OEM	Original Equipment Manufacturer
PHY	Physical Layer
PKI	Public Key Infrastructure
RF	Radio Frequency
RA	Registration Authority
RSU	Roadside Unit
SCMS	Security Credential Management System
SPAT	Signal Phase and Timing
SW	Software
V2X	Vehicle-to-Everything
VSOC	Vehicle Security Operation Centre

2. Background

2.1. General

Information in V2X packets is used by receiving vehicles to make safety critical decisions. It is therefore imperative that such information is correct and trustworthy, representing the true physical reality. It should, however, be noted that information received using radio communication is one source of input, a vehicle has many onboard sensors whose data is fused with received external sensor input to make an informed decision. It should also be noted that using external inputs, such as information received on radio communication, poses a challenge from a functional safety aspect that needs to be considered.

Cryptographic techniques such as digital certificates and signatures are commonly used so a receiver can verify that a sender has valid security credentials (authentication), and the received packets are indeed sent by the claimed sender without being tampered with (integrity). However, that is not enough to guarantee that the information in V2X packets is indeed correct and trustworthy. For example:

- Digital certificates (and the corresponding private signing keys) may be stolen and used by malicious attackers to impersonate the original certificate owner and launch attacks, or
- The software (SW) or firmware (FW) of a legitimate V2X device may be compromised by attackers such that the payload of the V2X packets may be maliciously modified to suit the attackers' needs before being signed and transmitted, or
- A legitimate V2X device may have malfunctioning sensors and therefore may be sending V2X packets with bad data.

For these reasons, the correctness of V2X messages cannot be guaranteed simply by verifying the digital certificate and signature. In the context of this white paper, misbehaviour refers to the wilful or inadvertent transmission of incorrect data within the V2X network. Note that this definition is quite specific and narrow, i.e. strictly related to incorrect content in V2X packets, in contrast to the general literal meaning of misbehaviour as something outside of the norm.

Incorrect information in V2X packets could be caused by a number of things, including but not limited to the malicious attacks discussed earlier. Faulty onboard components, temporary malfunction of a sensor, such as GPS not working inside a tunnel, or cameras failing to detect objects in poorly lit condition, could result in misbehaving packets as well. Whatever the cause or motivation, the net effect is the same – the information sent out in a V2X packet is substantially different from the physical reality, and is beyond the nominal margin of error considered acceptable by the industry.

The general approach is to first detect such misbehaviour locally by receivers of such suspicious V2X packets. This process is called Local Misbehaviour Detection (LMBD). Some of these misbehaviours may also be reported to a central authority in the Public Key Infrastructure (PKI) called a Misbehaviour Authority (MA). The MA may investigate and corroborate such reports and the process is called Global Misbehaviour Detection (GMBD). If deemed necessary, the certificates used by the misbehaving device (and any

future certificates that would be available to the device) may be identified and revoked, permanently negating the influence of such a device on the network.

Misbehaviour detection and reporting may be done by any V2X device receiving a suspicious packet, such as an onboard unit (OBU) in a vehicle or a roadside unit (RSU).

V2X messages contain the full V2X networking stack, from layer 1 (Wireless Physical Layer, PHY) to layer 7 (Specific V2X Applications). Many of these messages are periodic in nature while others are event-triggered. Misbehaviour could mean incorrect information in any layer in the stack, but the primary concern would be the semantic content in the application payload of the packets. Therefore, misbehaviour detection could be done in each of these layers, or at least could leverage all the information available at or within these layers. For example, radio frequency (RF) properties in the PHY may provide important clues that are harder to be tampered with than application payloads and so could be used to assist misbehaviour detection at the higher layer. On the other hand, misbehaviour at the application layer may be very specific to each application. For example, a traffic light management system using V2X may exhibit misbehaviour that is entirely different than a platooning application. In regard to system resources, it is beneficial to detect misbehaviour as early as possible.

Many V2X applications may be built on top of common V2X messages. For example, periodic broadcast of Basic Safety Messages (BSM) or Cooperative Awareness Messages (CAM) enable a broad set of safety applications such as Left-Turn Assist, Lane-Change Warning, Emergency Brake Warning, etc. Therefore, misbehaviour detection for BSM/CAM may be considered a common service, such as at the Facility Layer in ETSI architecture, offered to benefit a broad set of applications.

2.2. Causes of Misbehaviour

Misbehaviour can be caused both by malfunctioning devices and malicious actors. This section focuses on the latter. Overall, there are two kinds of attackers that are of concern in the context of MBD for V2X networks. The first kind are internal attackers, i.e. malicious V2X devices with necessary hardware, software and valid V2X credentials allowing them to inject bogus information into the V2X network and other unsuspecting devices. If other devices accept such bogus information without question, it may negatively affect the safety and efficiency of the overall traffic in the attackers' single- or multi-hop communication range. The second kind of attacker is an external who can manipulate the environment to cause sensor malfunction or for it to misread the data. Both of these attacker types are capable of propagating incorrect information in the V2X network.

Below is a (non-exhaustive) list of attacks to illustrate more concretely the resulting undesirable misbehaviour, i.e. incorrect information in the V2X network. A few commonly discussed attacks in networking security literature, but not within the scope of this discussion, are also highlighted to further distinguish specific threats in the V2X system.

- **Masquerade attacks:** An attacker obtains valid security credentials such as digital certificates by some means (e.g. illegal marketplace that sells stolen security credentials) and then uses them with a HW/SW/FW platform that is under their control and capable of V2X communications with other legitimate V2X devices. Because the platform is fully

Because the platform is fully under the attacker's control, it can generate whatever V2X packets are needed, so whole V2X packets can be infected with bogus information. The platform itself would usually have wireless communication capability to transmit and receive standard V2X packets. Even if the platform is impersonating a vehicle (using its digital certificates) and thus 'mobile' in principle, it could actually be physically stationary.

- o **Sybil attacks:** Sybil attacks are launched by an attacker to emulate multiple or even a large number of devices by concurrently using many certificates valid for a given time period. This attack may be launched to create a fleet of 'ghost vehicles', or to subvert a network service's reputation system by exerting a disproportionately large influence.

- **Data manipulation attacks:** Instead of setting up a full platform with HW/SW/FW, as in the masquerade case, an attacker can also attempt to compromise parts of a victim's platform by exploiting some vulnerabilities or manipulating parts of the V2X packet's content. This kind of attack can be launched via numerous entry points because the entire attack surface of a vehicle or RSU can be exploited for that purpose, and so can be further broken into several notable sub-categories, as discussed below.

- o **Internal sensor spoofing attacks:** A message is manipulated by changing the actual internal sensor readings. For example, this can be done via malware injected into a certain part of the SW/FW stack so that some fields in the V2X packets are manipulated.

- o **External sensor spoofing attacks:** This sub-category of attacks interferes with input from sensors that report external circumstances, e.g. GPS spoofing attack resulting in bogus GPS locations in the V2X packets sent by the victim's vehicle. Spoofing the GPS signal has been seen in shipping, drones, and connected and autonomous vehicles [GPSVideo, Shen+2020, Zeng+2018].

- o **Replay attacks:** An attacker may record packets from a legitimate device and repeat it sometime later. If it simply repeats the packet without any modification, it may be too easy to identify because V2X packets would have timestamps that give it away. But the attacker may be able to manipulate parts such as the timestamp and make it harder to detect. Such attacks could be considered as a special case of Data Manipulation Attacks.

- o **Wormhole/tunnel attacks:** A message from one geographic location is transmitted at a different geographic location by an attacker. This is easy to detect for messages such as BSM/CAM beacons with absolute location data, but may be less obvious for application messages lacking such location fields.

The common nature of these data manipulation attacks is that the attacker may be able to manipulate some parts of the payload but not others (for example, GPS spoofing affects all fields related to or derived from GPS, but not others).

- **Denial of Service (DoS) attacks:** DoS can be launched by a device to exhaust a network resource so that the system stops responding or functioning properly. A naïve but easy to detect example would be for a device to generate and transmit fully formed messages at a rapid rate to potentially congest the channel and overload the network. This is done at the application layer and is relatively easy to detect. DoS can also be launched at the physical layer, i.e. a jamming attack. Because DoS in general does not manifest itself as incorrect data in V2X packets, it is not in the scope of this white paper.

Misbehaving messages can result in a range of responses from vehicles that are fooled by them. For example, misbehaving messages can lead to unnecessary hard-braking

events. Although an attacker has no guarantee that an individual message will lead to that specific outcome, he/she/it will know that it can cause increased disruption to driving (congestion) and potentially the risk of collision. While the ultimate solution would be for the MA to identify and stop (eliminate/revoke) the misbehaving entities, each V2X receiver also needs to be able to mitigate misbehaviour – or the consequences of misbehaving V2X packets – on the ground (locally).

On the other hand, it is important to recognise that most MBD algorithms are not 100% bullet proof, so non-zero false positives and false negatives are to be expected. False positives cause a non-misbehaving device to be falsely accused of misbehaviour, and hence negatively impacted in the system, from its messages being dropped to its certificates being revoked wrongfully. False negatives are real misbehaviours that cause sensors to malfunction or worse, but the malicious attack is not detected. Often MBD algorithms can be tuned to balance the overall impact of false positives and false negatives when the consequences are carefully considered in each case.

In addition, MBD itself also introduces a new line of attack into the V2X system through vectors such as:

- **Slander attacks:** False misbehaviour reports are purposefully created about non-misbehaving devices and reported to the MA to discredit the victims which may lead to the victims' certificates being wrongfully revoked. This could also be done by combining with Sybil attacks to appear as multiple coordinated reporters.
- **Denial of service attacks against the MA:** An attacker can overload MA by just sending bogus MB reports that would force the MA to analyse each report and waste communication and computing resources. This is especially impactful if combined with Sybil attacks to scale up the load against the MA.
- **Evasion attacks:** Once a specific MBD algorithm is known to be used, an attacker may also exploit the specifics of the algorithm to evade detection by keeping it under a certain threshold or purposefully injecting 'noise' to cause the algorithm to fail in a machine learning-based MBD system.

While the above attacks are not the primary concern of MBD itself, they may introduce new kinds of vulnerability into the entire V2X network. Therefore, it is important and necessary to design MBD and reporting with robust protection against these new attacks. For example, reporter accountability may help mitigate against slander attacks and DoS against the MA. However, this must be balanced with concerns that reporters' privacy is not significantly compromised compared with non-reporters, and that reporters are not being penalised inadvertently.

3. Architecture of Misbehaviour Management System

3.1. Overview

Misbehaviour detection is carried out in the context of a Misbehaviour Management System, which has local (on an ITS-S) and backend components [TS103759]. In the Misbehaviour Management System, ITS-Ss create misbehaviour reports and send them to the Backend Misbehaviour Management System. Based on this and other information, the backend system determines what has actually occurred and what, if any, remediation actions to take. A high-level functional architecture of these systems is shown in Figure 1:

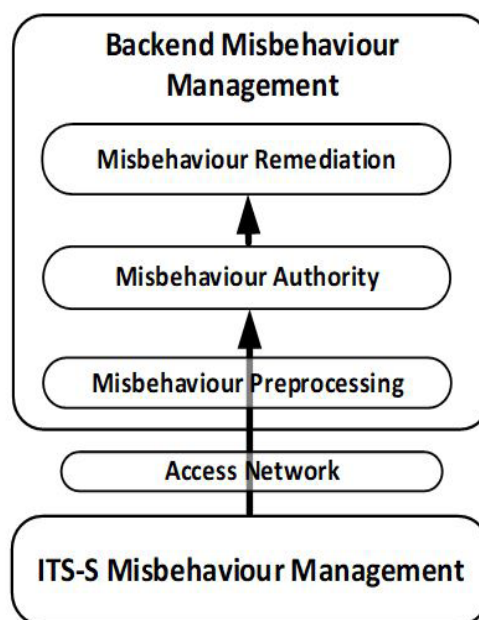


Figure1: Misbehaviour Management

- The ITS-S Misbehaviour Management System is responsible for detecting misbehaviour, generating misbehaviour reports, and sending them to the Backend Misbehaviour Management System. It may also react locally to the detected misbehaviour, for example, by informing other applications on the ITS-S or on other ITS-Ss about the misbehaviour.
- The Access Network provides access from the ITS-S Misbehaviour Management to the Backend Misbehaviour Management System for report uploading. It is not an active 'participant' at the application payload level in misbehaviour management.
- The Misbehaviour Preprocessing component acts on the misbehaviour reports before they are passed to the Misbehaviour Authority component, for example to improve the quality of the information received by the MA or to improve the privacy of the reporters. Examples of components within Misbehaviour Preprocessing are given later in this section.

- The MA component takes in misbehaviour reports and other information, possibly after preprocessing by the Misbehaviour Preprocessing component. It uses these to establish what happened in a reported misbehaviour event and who was responsible for it, and determines what remediation actions (e.g. revocation) should be taken.
- The Misbehaviour Remediation component is responsible for implementing the remediation activity determined to be necessary by the Misbehaviour Authority component.

Figure 1 also shows information flow from ITS-S to the Misbehaviour Authority element, and from the MA element to the Misbehaviour Remediation element. The information flow is shown as going directly from the ITS-S to the MA because the purpose of the flow is to provide report information between the two, even though some preprocessing is carried out en route. The process steps are:

- **Local detection:** An ITS-S runs misbehaviour scans to detect incoming suspicious messages from ITS-S in their neighbourhood.
Note – It is encouraged that ITS-S check their outgoing messages to avoid sending messages that will be classified as misbehaviour.
- **Reporting:** An ITS-S that has detected misbehaviour sends a misbehaviour reporting (MR) message to the Backend Misbehaviour Management System.
Note – It is not assumed that all observed instances of misbehaviour lead to the generation or upload of a report. Whether an ITS-S decides to generate a report on observed misbehaviour is implementation specific.
- **Misbehaviour Authority:** The MA collects and analyses the information from the reports, which may have had preprocessing applied to them. The MA identifies the type/severity of the reported misbehaviour and, if appropriate, the identity of the misbehaving ITS-S, and determines the suitable reaction required to protect the system.
- **Remediation:** A remediation is triggered accordingly (e.g. ITS-S deny listing at the PKI).

3.2. Detailed View of Misbehaviour

Figure 2 provides a more detailed view of the overall Misbehaviour Management System. Each component of the system is briefly described below.

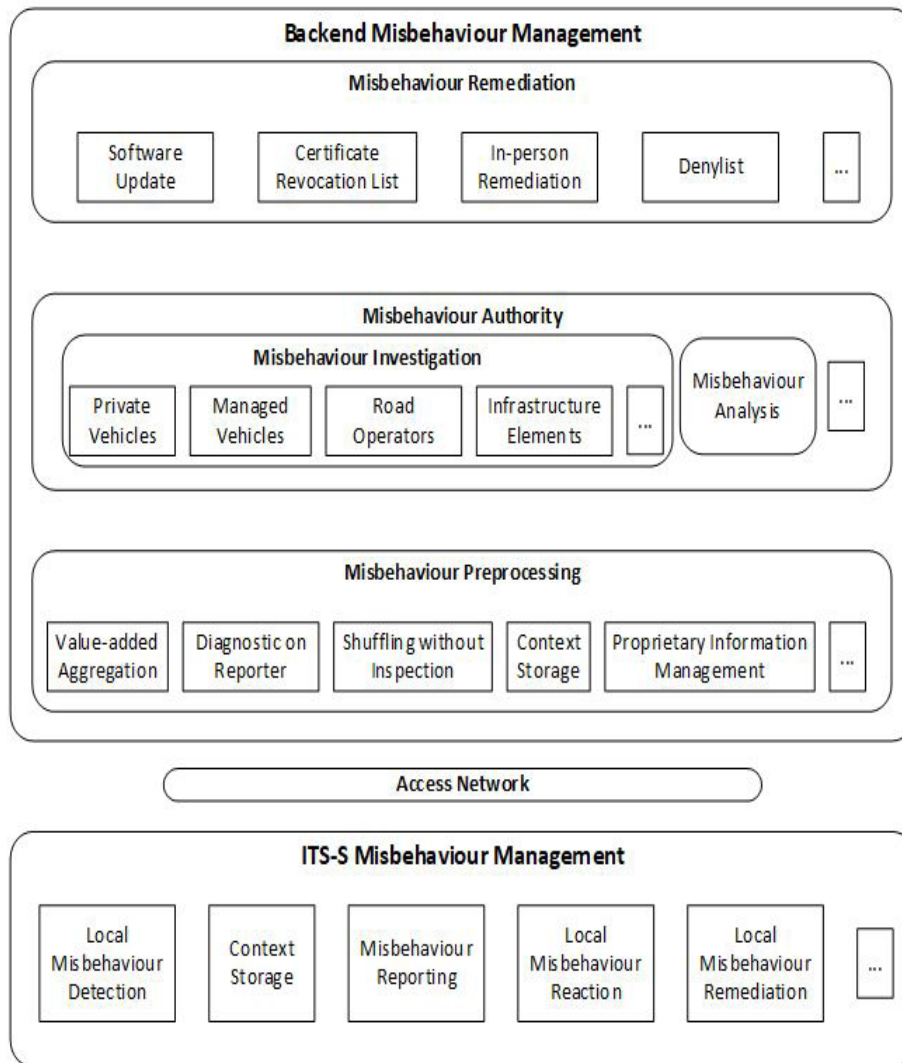


Figure 2:: Misbehaviour Management (detailed)

The ITS-S detects misbehaviour happening in its vicinity. At the ITS-S, misbehaviour management includes the following functional components:

- **Local Misbehaviour Detection:** Responsible for analysing incoming data and detecting any potential misbehaviour. The subject of misbehaviour may include the ego ITS-S. The LMD classifies and categorises the detected misbehaviour for further processing.

Note: It is a good practice for the ITS-S to ensure that the data transmission is reliable, plausible and consistent.

- **Context Storage:** Responsible for storing context information, i.e. information that is relatively long-lived and may be relevant to more than one misbehaviour report.

- **Misbehaviour Reporting:** Responsible for generating, storing, and transmitting reports of detected misbehaviour.

- **Local Misbehaviour Reaction:** Responsible for any reaction to the observed misbehaviour. The reaction does not involve communicating a report to the Backend Misbehaviour Management System.

- **Local Misbehaviour Remediation:** Responsible for any remediation action concerning the misbehaviour, e.g. triggering a self-diagnostic or a software update request procedure.
- **Other:** Components of the ITS-S that may be identified and/or defined in future.

On the backend security side, a full misbehaviour management system architecture includes the following functional components, and may include others:

- **Misbehaviour Preprocessing:** Can sometimes include sub-components, which may or may not be required for the baseline operation of the system (i.e. a full system deployment may contain any or all of them):
 - o Value-added Aggregation: Aggregates reports based on certain parameters/features such as misbehaviour location.
 - o Diagnostic on Reporter: Performs diagnostics on the reporter to establish the reliability of information in its reports.
 - o Shuffling without Inspection: Shuffles reports from multiple reporters to improve privacy.
 - o Context Storage: Stores context information so that reporters can refer to it rather than having to directly include it in their reports
 - o Proprietary Information Management (PIM): Enables routing and processing of proprietary information, i.e. information that is relevant to the diagnosis of misbehaviour but should not be revealed directly to the MA.
 - o Other: Preprocessing components in addition to the ones above.
- **Misbehaviour Authority:** May be considered as containing the following components which (except for 'Other') are necessary for baseline operation of the system:
 - o Misbehaviour Investigation: Determines the ITS-S at fault in reported misbehaviour incidents. This may involve making queries to other parts of the system, e.g. for pseudonym resolution. Operationally, this may be a single system or separated into multiple sub-components, for example for different applications or different ITS-S types.
 - o Misbehaviour Analysis: Determines the facts on the ground for reported misbehaviour incidents, and the severity of the misbehaviour. To analyse the reports along with the outcome of the investigation from the above sub-component. The misbehaviour analysis may be carried out before and/or after the investigation.
 - o Other: Misbehaviour Authority components other than the ones above.
- **Misbehaviour Remediation:** May include sub-components which or may or may not be necessary for the initial operation of the system, such as:
 - o Software Update: Enforces software update.
 - o Certificate Revocation List (CRL): Generates, stores, and distributes CRLs.
 - o In-person Remediation: Implements remediation by taking physical action at the misbehaving ITS-S's location.
 - o Deny List: Generates, stores, and distributes deny lists. These are distinguished from CRLs in that deny lists are distributed to CAs and used to determine which ITS-Ss should not receive certificates, while CRLs are also distributed to ITS-Ss and used to make 'trust decisions' on incoming application messages.
 - o Other: Remediation components other than the ones above.

3.3. Local Misbehaviour Management

3.3.1. Overview

The local misbehaviour reporting module receives notification of an observed misbehaviour. It determines whether to generate a report, assembles the report, signs and encrypts it, and then either sends it or stores it for later sending. Over time, it also manages stored reports to ensure proper prioritisation of storage space.

3.3.2. Local Misbehaviour Detection

A possible reference architecture for the misbehaviour reporting subsystem in the vehicle ITS-S is shown in Figure 4. Each component of the system is briefly described below.

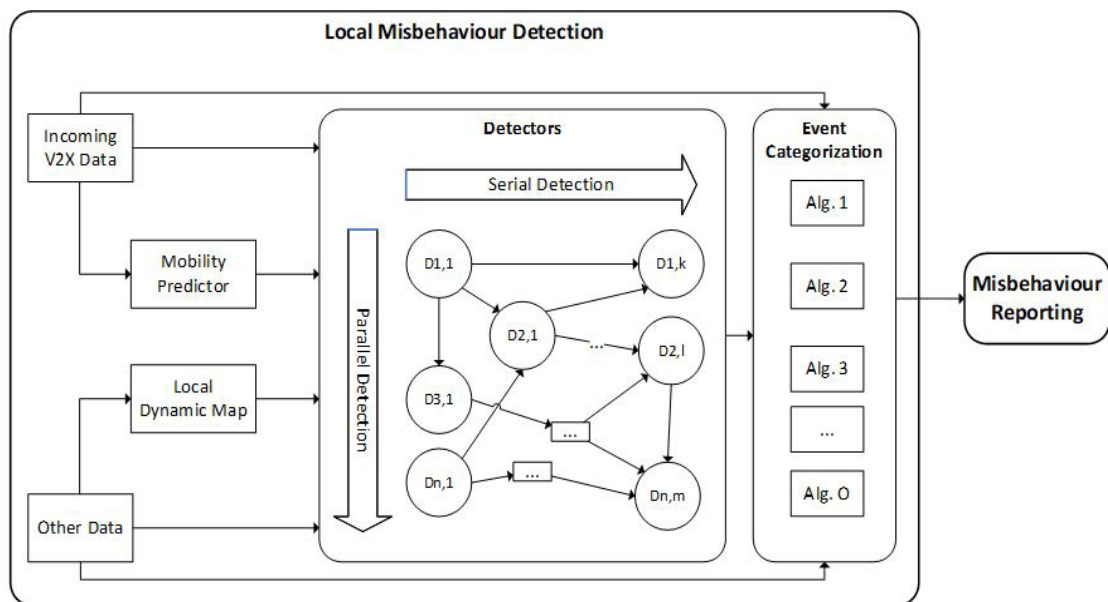


Figure 3: Local Misbehaviour Detection Process

- Incoming V2X Data is any information coming via the V2X communication medium, e.g. Cooperative Awareness Message (CAM) and Decentralised Environmental Notification Message (DENM).
- Other Data is any information other than the V2X Data, e.g. sensor and map data.
- Mobility Predictor estimates the dynamics of vehicles in the ITS communication range prior to the reception of next incoming data.
- Local Dynamic Map (LDM) is a database managed by the ITS-S containing V2X and other data, such as the ego vehicle's position and speed.
- Detectors identify V2X data that are inconsistent with the ego vehicle's 'perception of ground truth' or otherwise impairing the correct operation of the V2X system. Detectors

can be connected among themselves in serial, parallel, or some combination of both, e.g. detectors D1,1 and Dn,1 could be run in parallel and then their outputs fed into D2,1.

- Event Categorisation aggregates the outputs of all individual detectors, along with data from other sources, to determine whether a V2X message is suspicious or not. A suspicious message is classified according to the categorisation results.

3.3.3. Local Misbehaviour Reporting

A possible reference architecture for the local misbehaviour detection subsystem in the vehicle ITS-S is shown in Figure 3. Each component of the system is briefly described below.

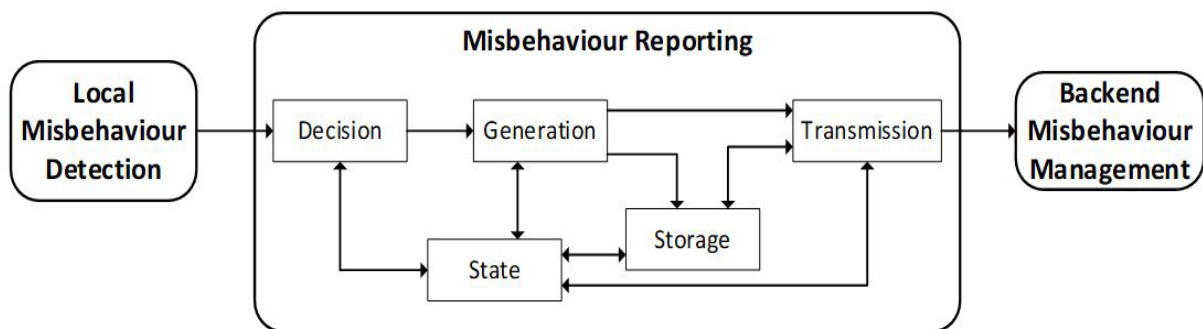


Figure 4: Local Misbehaviour Reporting Process

The rationale for this architecture is that a misbehaviour reporting system may have to manage three distinct 'budgets': one for report creation (as there may be competing demands for access to a signing process, or for processor time in general); one for report storage (as the total volume of reports generated may exceed the storage available or allocated for them); and one for report transmission (as it may not be possible to upload all generated reports due to intermittent connectivity or a data service plan putting a ceiling on the total amount that may be transmitted in a particular time period). The architecture allocates one functional entity to be responsible for managing each of these budgets. In this architecture:

- (Misbehaviour Reporting) Decision decides whether or not to generate a misbehaviour report for an event observed by the local misbehaviour detection system.
- Generation creates the misbehaviour report making use of the information or input provided from the misbehaviour detection subsystem and (optionally) other information obtained from State. After report creation, it also decides – on the basis of classification and categorisation – whether to send the report to Storage and/or Transmission.
- Transmission decides whether to send the reports to the MA. If multiple reports are available for sending, Transmission decides which ones to send and in what order.
- State stores information that may be used by other components inside the misbehaviour reporting subsystem to carry out their functions.
- Storage stores misbehaviour reports provided by Generation, provides them to Transmission ready for the MA, and deletes old reports or reports with low prioritisation

4. Prior and Related Work

4.1. Regulations and Standards

Regulatory bodies and standards development organisations (SDOs) are working on developing security regulations and standards to secure vehicles against misbehaviour attacks. This section provides a brief overview of published (or to be published) documents.

4.1.1. Regulations

Two new UNECE regulations, UN Regulation No. 155 [UNR155] and UN Regulation No. 156 [UNR156], require vehicle manufacturers to implement security measures in the following four categories:

- Cybersecurity risk management for the vehicle
- Detection and the remediation of automotive security incidents
- Enabling software updates without compromising the vehicle's safety and security
- Securing vehicles by design to mitigate risk along the value chain

Accordingly, vehicle manufacturers must implement these security measures to mitigate cybersecurity risks.

4.1.2. Standards

Among all SDOs, the European Telecommunications Standards Institute (ETSI) is leading the standardisation effort for specifying a cybersecurity system against V2X misbehaviour attacks. The outcome of this effort is two documents identified as ETSI Technical Report (TR) 103 460 [TR103460] and Technical Specification (TS) 103 759 [TS103759].

The TR is a document providing prior technical work and open challenges remaining for writing TS 103 759. The latter is a standard under development that defines a V2X misbehaviour detection and reporting system for a subset of V2X message types: Cooperative Awareness Message (CAM) and Decentralised Environmental Notification Message (DENM). The standard content includes the description of the Misbehaviour Management System. Currently, the document focuses on two components of the MMS: the detection and the reporting of attacks by the ITS-S. At the time of writing, the standard specifies a basic set of 15 detectors for the CAM. For each new and existing V2X message type, future releases of this standard are expected to include an updated list of detectors. Note that the document also specifies the structure of the misbehaviour report sent by the vehicle to the MA.

In addition to the documents mentioned above, the following security standards are noteworthy: readers may refer to International Organisation for Standardisation (ISO)

21434 [ISO21434] if they need a framework to assess the security risk of misbehaviour attacks on their V2X system. Additionally, they can read ISO 24089 [ISO24089] for technical information (e.g. requirements and recommendations) related to software updates. Finally, readers can look at [11] for an overview and a description of ITS security standards, such as Institute of Electrical and Electronics Engineers (IEEE) 1609.2 [IEEE1609.2] or ETSI TS 103 097 [TS103097].

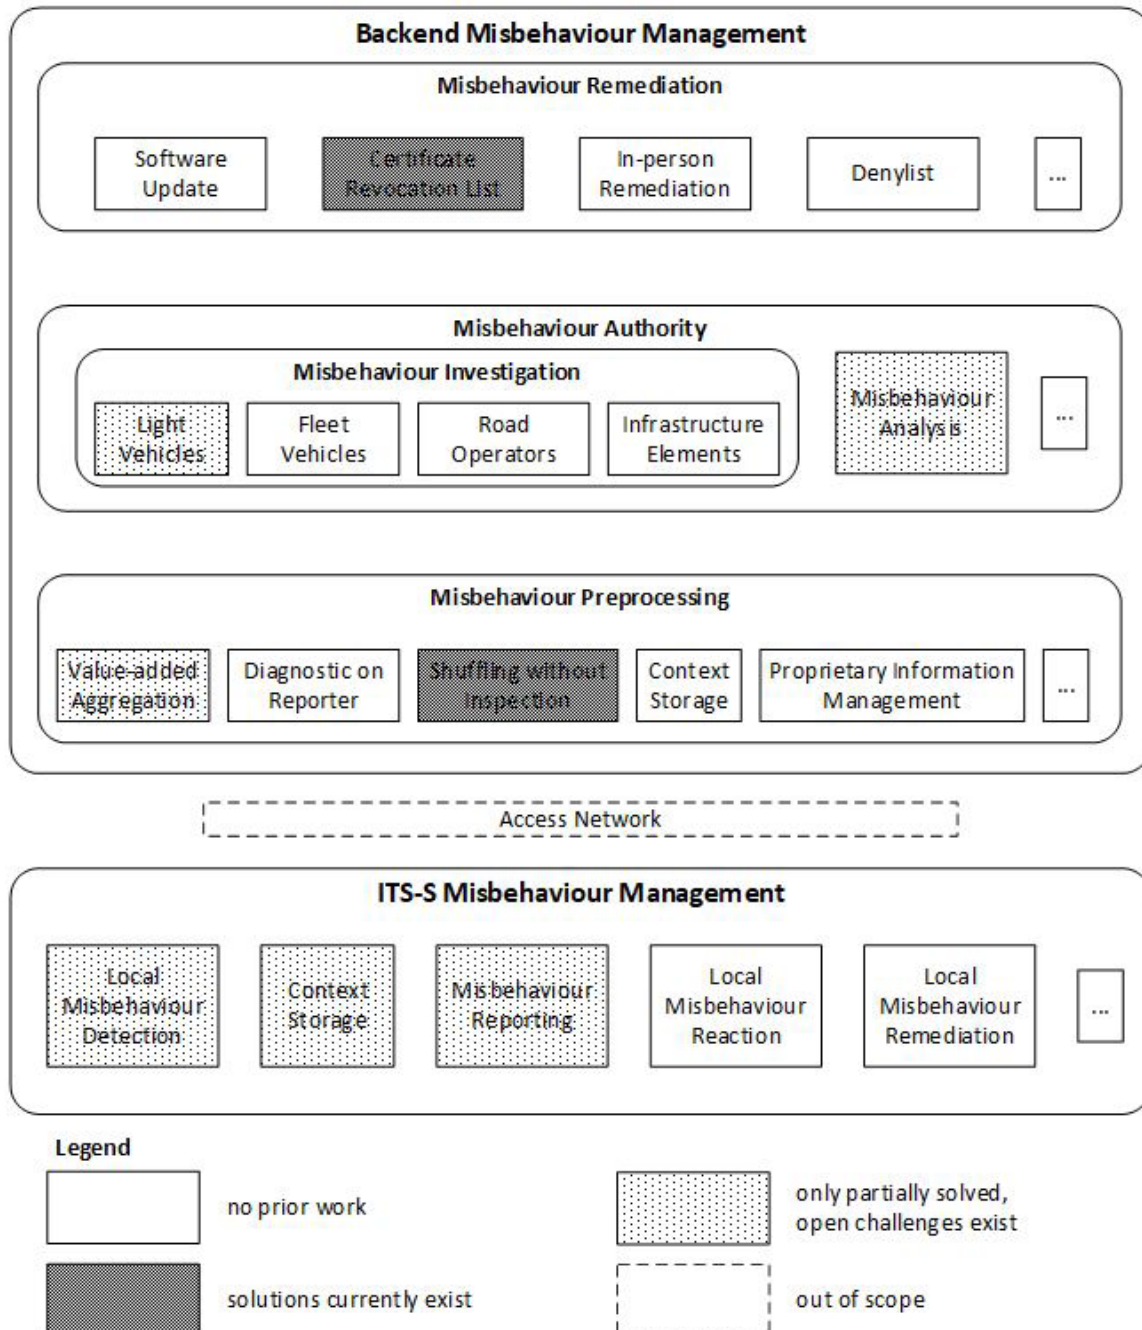


Figure 5: Overview of Prior Work for a Misbehaviour Management System

4.2. Misbehaviour Management

As depicted in Figure 5, several components within the Misbehaviour Management System have not been addressed in the literature. so this section provides a brief overview of the prior work and open challenges.

4.2.1. ITS-S Misbehaviour Management

Misbehaviour management at the level of the ITS-S is a topic that contains a lot of prior work. However, as depicted in Table 1, some components remain open challenges.

Components	References to the prior work
Local Misbehaviour Detection	[Kamel+2020, Heijden+2018]
Context Storage	[Ghaleb+2019]
Misbehaviour Reporting	[Noori+2020, Spectrum2021]
Local Misbehaviour Reaction	
Local Misbehaviour Remediation	

Table 1: Activities Overview for ITS-S Misbehaviour Management

The bulk of the prior work focuses on Local Misbehaviour Detection (LMBD). In addition to the ETSI TR 103 460, an academic survey provides a good overview of the prior work related to LMBD [Heijden+2018]. It is important to note that only 20 attack types have been tested and all relate to a single message type (BSM or CAM) [Kamel+2020]. Therefore, despite the numerous detectors found in the prior work, several open challenges remain in LMBD.

For the remaining addressed components, context storage and misbehaviour reporting, the existing literature is scant. For the first component, one contribution aims to use the mobility data from surrounding vehicles (referred to as context) to determine the plausibility and consistency of a BSM/CAM [Ghaleb+2019]. The ETSI TS 103 097 standard defines the misbehaviour report and describes the reporting process. Additionally, there is one contribution that provides content not defined in the TS [Noori+2020]. This work proposes to extend the current format of a misbehaviour report to include evidence required for detecting spectrum interference attacks. An Abstract Syntax Notation One (ASN.1) definition of this spectrum misbehaviour report can be found in [Spectrum2021].

Potential open challenges should include the extension of the MBR format to all V2X message types. Additionally, some efforts should look at potential local reactions to an ITS-S after detecting a misbehaviour attack. For instance, should the targeted ITS-S discard all future messages from the attacker? Finally, it is likely to be desired that an ITS-S

can update its LMBD system. Despite the short deadline for enforcing the regulation, there is no prior work that addresses this technical requirement.

4.2.2. Backend Misbehaviour Management

4.2.2.1 Misbehaviour Preprocessing

As depicted in Table 2, there is little prior work related to misbehaviour preprocessing. The existing work focuses on two components named ‘value-added aggregation’ and ‘shuffling without inspection.’ The other components remain open challenges to be solved.

Components	References to the prior work
Value-added aggregation	[Kamel+2019] [Michelson+2019]
Diagnostic on reporter	
Shuffling without inspection	[Brecht+2018]
Context storage	
Proprietary information management	

Table 2: Activities Overview for Misbehaviour Preprocessing

Two contributions exist for ‘value-added aggregation’ [Kamel+2019, Michelson+2019]. The first is used for detecting Sybil attacks. The contribution aggregates misbehaviour reports referring to attack(s) involving multiple ghost vehicles. The aggregation consists of a neural network looking for similarities between misbehaviour reports. The neural network will try to determine a set of features such as the Euclidean distance between the reporter and the reported. A second considered feature is the Euclidean distances between the reporter and the reported ITS-S in other misbehaviour reports. Through this aggregation, the neural network can determine if the reports prove the existence of a Sybil attack (or not). The second contribution is used for localising spectrum interference. The MA aggregates misbehaviour reports referring to a spectrum interference. Then, the MA uses an ML technique to correlate the reports and identify the number of interference or congestion events in an area. For each area, the MA can triangulate the precise location of the interference thanks to the reporter’s location.

Looking at ‘shuffling without inspection’, there is no prior work that shuffles reports from multiple reporters to improve their privacy. However, the Security Credential Management System (also known as V2X PKI) [Brecht+2018] has a shuffling process performed by the Registration Authority (RA). Future work could re-use the design and specification of this process for shuffling misbehaviour reports.

4.2.2.2 Misbehaviour Authority

The MA is a topic that has drawn significant attention (Table 3). Both components, hereon referred to as investigation and analysis, have some well-defined use cases in the literature. The next paragraphs will describe a use case per component and how it has been addressed in the literature.

Components	References to the prior work
Misbehaviour Investigation	[Kamel+2019, Trauernicht+2019]
Misbehaviour Analysis	[Michelson+2019]

Table 3: Activity Overview for Misbehaviour Authority Core Functionality

Regarding Misbehaviour Investigation, a common use case is the detection of a Sybil attack. A Sybil attack consists of a malicious ITS-S using its pool of pseudonym certificates to send multiple BSMs with distinct identifiers. From a receiver perspective, the victim will think that there are multiple (ghost) vehicles within its communication range. This attack aims, for instance, to use a fleet of vehicles to block the manoeuvre of a targeted victim. The ability to identify the attacker among all the ghost vehicles is a challenging task for an ITS-S. Thus, it is the role of MA to detect Sybil attacks. Currently, several approaches to misbehaviour investigation have been proposed. A first approach consists of counting the number of misbehaviour reports referring to a ghost vehicle event [Kamel+2019]. The paper assumes that the existence of several ghost vehicles within a narrow timeframe and a precise geographic location could be a Sybil attack. The proposed solution is to report BSMs belonging to a ghost vehicle so that the Linkage Authority (LA) can determine if those BSMs have been transmitted by the same attacker. If it is the case, then the LAs provide the 'linkage seed' of the attacker to be added to the certificate revocation list. Although this solution is valid for the Security Credential Management System (SCMS), this approach is impossible in Europe due to the absence of LAs within the European V2X PKI. Thus, a second paper proposed an approach to revoke all the certificates of an attacker without LA [Trauernicht+2019]. Both approaches allow the permanent revocation of the attacker if necessary.

A common use case, in the prior work, for misbehaviour analysis is the detection of malicious interferences among V2X communication. After detecting the geographical location of the interference, the MA can request an investigation to determine if the cause of this interference was an unintentional (e.g., a faulty hotspot) or intentional interferer. In the second case, the outcome of the investigation could reveal that the interferer is an ITS-S. Therefore, the MA can use the identifier of the interferer for further investigation. For instance, the MA can determine the interferer's involvement in other cases of spectrum interferences by checking the proximity of its geographic location at the time of the interference [Michelson+2019].

Future work could be to expand the list of use cases for both investigation and analysis. For instance, it has yet to be defined if and how the MA can request additional information from an OEM for solving an investigation. Additional topics could try to answer the

analysis. For instance, it has yet to be defined if and how the MA can request additional information from an OEM for solving an investigation. Additional topics could try to answer the following questions. What is the minimal set of evidence needed by the MA to detect a misbehaviour attack? Another aspect would be to determine what is the logic behind an MA's decision. For instance, should the MA's judgement be the same for interference caused by an unintentional interferer compared to an intentional interferer?

4.2.2.3 Misbehaviour Remediation

As depicted in Table 4, misbehaviour remediation is a topic that has been barely explored, with prior work focusing on certificate revocation lists (CRL). Concretely, several regional standards have specified the structure of a CRL and its content. Additionally, several academic papers, ETSI TR 103 460, IEEE Std 1609.2, and certificate policies explain the different processes involving a CRL such as its generation and its distribution [Brecht+2018, Heijden+2018].

Components	References to the prior work
Software Update	
Certificate Revocation List	[Brecht+2018, Heijden+2018]
In-person Remediation	
Deny List	

Table4:ActivityOverviewforMisbehaviourRemediation

However, several open challenges remain when it comes to specifying and implementing the following components: In-person Remediation, Deny List, and Software Update. One approach to meet these challenges could be to define use cases for In-person Remediation. One use case could be the removal from the Deny List of an unintentional misbehaving vehicle. For instance, a vehicle could be identified as misbehaving due to its faulty GPS (e.g. implausible position displacement over time). If the faulty GPS gets repaired or replaced by a certified mechanic, then the vehicle should be removed from the Deny List by the Misbehaviour Management System. Thus, this use case requires a set of processes and protocols for each actor involved in this use case to be defined. Establishing use cases for each component could be a starting point for technical specifications.

4.2.3. Misbehaviour Organisational

4.2.3.1. Audit

As mentioned in UN R155, V2X stakeholders must audit all their V2X systems to prevent or mitigate potential misbehaviour attacks. Current threat assessments found in prior

work focus on attacks targeting BSMs and CAMs [Monteuuis+2018, Moalla+2012]. Thus, the absence of work assessing the security risk related to other V2X message types (e.g. Signal Phase and Timing (SPAT)) may hide some important vulnerabilities affecting vehicle safety.

One way to fill this gap is the extension of prior threat assessment on V2X systems. This extension should include the remaining technical (e.g. missing message type) and organisational gaps (e.g. human maintenance). One outcome of this full threat assessment is a comprehensive list of attacks for each V2X message type.

4.2.3.2. Policies

It is part of good security practices to document all the technical and organisational processes occurring in public key infrastructure (PKI). The literature already contains several instances of this document such as the C-ITS certificate policy in Europe [CP2019]. Current V2X certificate policies do not include any details on processes related to the Misbehaviour Authority. This gap in the documentation may lead to organisational abuses. For instance, a reporter could repeatedly report a misbehaving vehicle with a faulty GPS to speed up its revocation. In terms of In-person Remediation, it is unclear what the interactions between V2X stakeholders and the MA could or should be. For example, a misbehaving vehicle may cause an accident requiring law enforcement authorities to reach out to the MA to understand if the cause is a V2X misbehaviour attack; and if so, they would need the evidence collected by the MA to prosecute the perpetrator.

A solution here is to document all the processes related to misbehaviour in a dedicated misbehaviour policy envelope. Like a certificate policy, the misbehaviour policy would document all the technical and organisational processes related to a Misbehaviour Management System. For instance, the document could explain how an OEM supports the MA in its investigation by providing proprietary information, such as the findings from an OEM Vehicle Security Operation Centre (VSOC). Before writing the actual document, a first step would be to define a framework for defining the scope of the policy, such as RFC 3647 [RFC3647].

5. Policy Question

As is the case with cybersecurity threats, misbehaviour in the V2X environment is likely to evolve over time, so the response to misbehaviour will have to evolve accordingly. Thus, rather than ‘What’s the best misbehaviour detection algorithm?’, we should be asking ‘What do we do with the fact that the best available misbehaviour detection algorithm will differ?’. The best-known algorithm will vary from time to time, and the best algorithm implemented locally may depend on the sensors installed on or in the vehicle. Which of these algorithms will need to be made public, standardised or become uniform? Moreover, there are administrative considerations including privacy implications, incentives, and oversight. Then comes the question: What is the best organisation/group to think about V2X misbehaviour detection and remediation going forward? While some of these questions can be answered satisfactorily using only technical means (algorithm design, computing resources, etc.), most of the central questions around V2X misbehaviour require policy intervention. In the following subsections, we present a non-exhaustive list of such questions with the aim of starting a conversation with policy experts and lawmakers in different jurisdictions.

5.1. Local Misbehaviour Detection

- Should there be some minimum performance requirements for LMBD, just as there are minimum performance requirements for sending BSMs?
- Should we wait and see the level of actual misbehaviour in the system and then design our defences accordingly, or should we assume the worst-case scenario?
- Do local MBD algorithms need to be public?
 - o Yes: allows creation of sensible reporting algorithms
 - o No: allows OEMs to compete
- Do local MBD algorithms need to be standardised?
 - o Yes: allows creation of sensible reporting algorithms
 - o No:
 - allows OEMs to compete
 - may make it hard to change algorithms if attackers work out ways around the current one
- Do local MBD algorithms need to be uniform?
 - o Yes: allows creation of sensible reporting algorithms
 - o No:
 - allows OEMs to compete
 - lowers the bar on developments
- Should there be some self-diagnosis requirements and how will they affect the reporting criteria?
- If there is a lot of misbehaviour in the vicinity, should the vehicle stop participating in the V2X system? It might make sense to shut down particular applications instead of the whole V2X.

5.2. Reporting

- Should the algorithms prioritise misbehaviour reporting based on how critical or severe the misbehaviour is?
- Should there be different algorithms...
 - o depending on reporting vehicle type?
 - o depending on reported vehicle type?
- Can reporting algorithms be proprietary?
- Should highly equipped vehicles be entitled to send more detailed misbehaviour reports than vehicles with more basic equipment (sensors, etc)?
 - o Should they be required to send more detailed misbehaviour report?
- How can the free-rider problem be avoided? For any given OEM the incentive is for their vehicles not to report (i.e. to preserve privacy), thus free-riding on the reporting of others.
- Can a vehicle OEM manipulate reporting by only reporting other OEM vehicles, or by applying different thresholds to different OEM vehicles?

5.3. Remediation

- Should the remediation decision depend on the severity of the misbehaviour?
- Should the remediation decision depend on the MA's confidence that the misbehaviour took place?
- Should the remediation decision depend on whether it is a malfunction or deliberate misbehaviour?
- Can there be an effective oversight mechanism to keep check on the MA's own activities?

6. Conclusions and Next Steps

This document covers different aspects of misbehaviour in the V2X ecosystem: definition and causes of misbehaviour, architecture of a Misbehaviour Management System, prior and related works. The document ends with a list of policy questions. The Security and Privacy working group at 5GAA (WG7) will explore answers and clarify the 5GAA position on these questions in follow-up work. As opportunities arise, it will also engage with relevant people/organisations outside 5GAA.

