



**HAL**  
open science

# Next Generation Vehicles, Safety, and Cybersecurity-The CMX Framework

Jonathan Petit, Gérard Le Lann

► **To cite this version:**

Jonathan Petit, Gérard Le Lann. Next Generation Vehicles, Safety, and Cybersecurity-The CMX Framework. IEEE Transactions on Intelligent Transportation Systems, 2024, 25 (2), pp.1333-1345. 10.1109/TITS.2023.3318376 . hal-04453750

**HAL Id: hal-04453750**

**<https://inria.hal.science/hal-04453750>**

Submitted on 12 Feb 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Next Generation Vehicles, Safety, and Cybersecurity—The CMX Framework

Jonathan Petit\*, Gérard Le Lann†

\*Qualcomm Technologies Inc., USA, petit@qti.qualcomm.com

†INRIA, France, gerard.le\_lann@inria.fr

**Abstract**—Safety, privacy, efficiency and cybersecurity (SPEC) properties are mandatory in vehicular networks. Owing to intrinsic limitations, the V2X framework and related communicating autonomous vehicles are inadequate. We explore the CMX framework (Coordinated Mobility for X = SPEC), Next Generation Vehicles (NGVs), and protocols for safety-critical inter-vehicle communications and agreements that solve problems left open in the V2X framework. Then, we focus on cyberattacks and physical attacks against isolated NGVs and cohorts of NGVs. The cyberphysical security analysis investigates more than 20 attacks and demonstrates that the goal set for CMX is verified. In the presence of attacks, safety is never compromised, possibly at the expense of diminished efficiency.

**Index Terms**—coordinated mobility, V2X, automated vehicle, safety, efficiency, security, cohort, cyberphysical constructs

## I. INTRODUCTION

In this paper, we introduce Next-Generation Vehicles (NGVs) and the CMX framework [1]. CM in CMX stands for Coordinated Mobility, X stands for S, P, E, C, four properties to guarantee:

- **Safety:** collisions entail property damages only. Zero severe injuries and fatalities due to collisions, even without human supervision or interventions.
- **Privacy (passive adversaries):** no personal data can be inferred/extracted from eavesdropped cyber-centric information, and from physical-centric information (tracking of paths and routes followed by vehicles).
- **Efficiency:** road capacity (vehicular densities) is higher than achieved with non-fully automated vehicles for identical velocities.
- **Cybersecurity (active adversaries):** safety is not compromised by internal or external cyber-attacks (e.g., spoofing, Sybil attack, message falsification/suppression, injection of bogus data, Denial-of-Service).

Careful analyses of published work and Vehicle-to-X (V2X) standards (ETSI, SAE, IEEE) reveal that they fail to achieve SPEC properties altogether for two main reasons: reliance on non-deterministic CSMA/CA MAC protocols (hence on an asynchronous time model), and periodic beaconing.

Fundamental results in distributed fault-tolerant computing, time-bounded consensus and safety engineering were ignored. To prove safety, one must consider failures (e.g., slow or unresponsive onboard systems, message losses), and prove that messages are received in due time, deterministically. Unfortunately, this is unfeasible with non-deterministic CSMA/CA MAC protocols. Moreover, it has been proved that consensus

is impossible in an asynchronous time model in the presence of failure(s) [2]. Similar results hold with message losses in a synchronous time model, hence, a fortiori, in an asynchronous model such as current V2X [3]. Consequently, claims of safety in current V2X designs are unfounded.

Therefore, we propose the CMX framework and Next-Generation Vehicles to solve these open issues and achieve SPEC properties. Our main contributions consist of:

- Cyberphysical constructs such as cohorts and related communication protocols. Such constructs serve to “bridge the gap” between plain reality (vehicular networks are asynchronous systems) and time models such as partial synchrony, timed asynchrony [4], or synchrony—thus circumventing well-known impossibility results—at the core of design frameworks needed for meeting the SPEC properties.
- Classification of vehicles based on cyberphysical levels in addition to SAE automated driving levels.
- Specific protocols for inter-vehicular coordination based on time-bounded explicit agreements.
- Solutions for authentication, self-checking capabilities aimed at stopping a vehicle exhibiting abnormal behavior.
- Detailed analysis of cybersecurity in the presence of most severe types of cyber and physical attacks.

CMX addresses open issues that are of concern to many Intelligent Transportation System (ITS) stakeholders, authorities and standards development organizations. One additional objective of this paper is to make stakeholders of the ITS ecosystem aware of an evolutionary path from Connected Automated Vehicles (CAVs) to safe and efficient spontaneous networks of fully automated vehicles.

CAVs/V2X and NGVs/CMX are not antagonistic. In a nutshell, a NGV is a CAV “augmented” with what is missing in V2X for achieving the SPEC properties. Contrary to CAVs, NGVs exchange messages that do not carry GNSS data. As a result, safety properties hold, and this, without enabling eavesdropping and tracking. Protocols and algorithms proper to the CMX framework solve use cases that are currently considered in Maneuver Coordination Service standards (SAE J3186, ETSI TS 103 561).

**The aim of this paper is to verify a major objective of the CMX framework: in the presence of active attackers, S (safety) and C (cybersecurity) are always preserved,** possibly at the expense of reduced E (efficiency)—larger inter-vehicular gaps or cohort splits for example. For the sake of conciseness, the Privacy property will be addressed in

a forthcoming publication. Our focus is on attacks against NGVs and cyberphysical coordination protocols (Longitudinal Join, Lateral Join, Leave, presented in Section V). Attacks on operating systems, system-level software elements, or agents based on artificial intelligence are out-of-scope.

The paper is organized as follows. Section II presents the NGV’s onboard system, its architecture, and its authentication. Section III details the cohort construct, episodic beaconing, and MAC protocols. Section IV presents intra-cohort very short-range directional neighbor-to-neighbor (N2N) communications, and cohort splits. In Section V, we describe cohort-to-cohort (C2C) communications and their usage in safety-critical maneuvers. Section VI details the attacker model and CMX cybersecurity features. Section VII presents a security analysis of NGVs and cohort management protocols along with a summary of the effects of the attacks on SEC properties. A conclusion is provided in Section VIII, along with some future work.

## II. SYSTEM MODEL

CMX conforms to basic Fault-Tolerant Computing principles, e.g., the onboard system is split into two separated and isolated subsystems, triple modular redundancy and fault masking [5], [6]). In fault-tolerant systems, there is a fundamental assumption: any operation attempted by a group of three elements is executed correctly and terminates in time if at most one failure occurs. If more than one failure occur, then the operation attempted is aborted. In the CMX framework, in any set of three (lateral, longitudinal) neighboring NGVs, if at most one NGV is an attacker or fails, desired safety-critical maneuvers are completed correctly and in time. Indeed, under this assumption, a misbehaving NGV is necessarily detected by at least two (longitudinal, lateral) honest NGVs. Majority voting suffices to preserve safety.

### A. NGV’s Onboard System

As depicted in Figure 1, a NGV’s onboard system is comprised of a (safety) critical (C) and a non-critical (NC) subsystems that are physically and logically separated by a redundant gateway (◆), serving to filter data flowing from a NC to a C subsystem and to ensure containment. A NC subsystem cannot write into the memory of a C subsystem. Data (no executable code) can be posted at the interface of a C subsystem, to be read (and checked) by a C subsystem. It follows that native infotainment or applications downloaded by passengers have no access to C subsystems, thus no control over physical maneuvers.

NC and C subsystems host sensing devices (cameras, RADARs, LiDAR, etc.), enabling measurements of distance, direction, velocity, and reading of license plates, as well as communication devices (wireless radio and optical communications such as visible light communication (VLC)). In accordance with the principle of diversified redundancy, devices found in NC and C subsystems are based on different technologies. A NGV knows the type of roadway on which it circulates (e.g., city street, country road, highway) as well as the current lane number (CMX rests on a lane-based system).

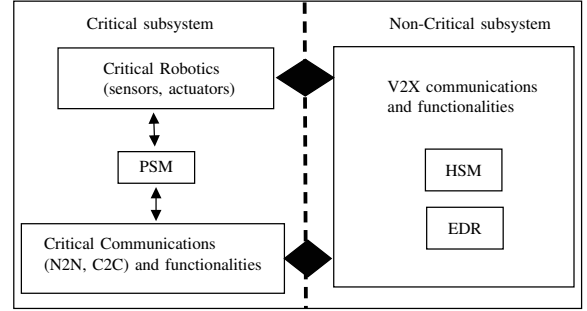


Fig. 1. General architecture of a NGV’s onboard system

A NGV knows the constraints proper to its geolocation (e.g., highest authorized velocity, highest cohort length). SPEC properties are ensured by C subsystem.

1) *NC subsystem*: Processes are run under the control of a regular operating system (e.g., iOS, Android). NC communications are conforming to V2X standards, minus the periodic beaconing of BSM/CAM messages (ruled out in CMX). A hardware security module (HSM) serves to host pseudonym certificates used when NC V2X messages are broadcast. An Event Data Recorder serves to record significant events and states (e.g., driving behaviors, equipment status for predictive maintenance). In CMX, there is a (V2X) cyber-stealth mode detailed in Section VI-B.

2) *C subsystem*: Processes in charge of vital C functions are found in the critical robotics (CR) and in the critical communications (CC) modules. Such processes must meet ultra-high reliability and availability requirements. They are executed on highly reliable processors or on replicated processors. This is commonly referred as *lockstep* CPU, and a requirement for ASIL-D systems. Processes in a C subsystem are run under the control of a real-time kernel similar to the formally verified microkernel named seL4 that was used to protect the critical system of autonomous helicopters, and ensured isolation of its subsystems [7]. We assume a similar system to protect against attacker with physical access aiming at compromising C subsystems. For protections against hardware faults violating modeling assumptions that would “invalidate” seL4, we refer the reader to solutions described in [8]–[10].

C communications are performed with directional very short-range antennas able to adjust transmit and receive power. A NGV has the ability to use radio signal properties to verify the location of a sender.

A C subsystem is equipped with a tamper-responsive proactive security module (PSM)—an extension of Hardware Security Module—that has its own battery (needed for avoiding common-mode failures). A PSM stores a long-term certificate and multiple “short-term” unlinkable certificates delivered at registration for pseudonymous authentication [11]. NGV X’s vehicle profile  $vp(X)$  is recorded in every certificate (see Section II-B).

A PSM contains a set of predicates that are derived from specifications of critical communications protocols and coordination algorithms, expressed in any convenient formalism. They specify forbidden state transitions (i.e., predicate violations), thus ensuring logical safety [12], i.e., “correct

behaviors” in cyber and physical spaces. A PSM monitors executions of C functions and predicate violations. Examples are:

- Relative to MAC protocols: transmission not in assigned channel slot (TDMA), transmission with a non-assigned code (CDMA), and number of consecutive collision exceeds stipulated upper bound (deterministic CSMA).
- Relative to N2N communications (see Section IV): failure to relay a received message that is supposed to be disseminated throughout a cohort.
- Relative to C2C communications (see Section V): failure to respond to a request for joining a cohort.
- Relative to authentication (see Section II-B): delay  $d$  between two signed C2C messages (with different certificates) does not meet the specified lower bound. Prior to engaging safety-critical maneuvers, protocols (LgJoin, LtJoin, Leave) are executed, and are constrained by physics, and therefore,  $d$  is in the order of a few seconds.

Owing to its PSM, a NGV is endowed with self-checking capabilities. When violations of predicates, UTC timestamped, exceed a certain threshold, the NGV is halted, i.e., removed physically from the transportation system. The CR subsystem, informed by the PSM, takes control: flashing warning lights, velocity is reduced, until a safe halt is possible (parking spot or emergency lane). A pseudonymized and encrypted (with the public key from the appropriate organization) V2X message is broadcast by the NC subsystem, aimed at public/private organizations (e.g., police, highway authorities, insurance companies), repeated a small number of times, until the onboard system of the halted NGV is shutdown. This eliminates the possibility of attacks launched by a halted (misbehaving) NGV. That message carries the GNSS location of the halted vehicle and other evidences recorded in an Event Data Recorder, such that accountability is guaranteed. This mechanism achieves “revocation” in the physical space, which is currently missing with revocation of credentials in cyber space (such as in the V2X framework). A discussion of the security of the PSM is provided in Section VII-G.

### B. Authentication of a NGV

NGVs use digital certificates compliant with IEEE 1609.2 [13]. However, in the CMX framework, certificates include information critical for interoperability, access control, and cybersecurity. Each pseudonym certificate attests the vehicle profile  $vp(\cdot)$ , which is composed of SAE level of automation ( $aul(\cdot)$ ) [14], vehicle type ( $vt(\cdot)$ ), vehicle length ( $vl(\cdot)$ ), cyberphysical level ( $cpl(\cdot)$ ), and interoperability set  $Interop(\cdot)$ . A  $cpl$  is a pair  $\{cl(\cdot), pl(\cdot)\}$  of cyber and physical levels respectively, each ranging from 0 to 5, resulting in a matrix of 36 elements (denoted CPL).  $Interop(X)$  is the subset of that matrix which contains  $cp$  levels compatible with  $cpl(X)$ . These security credentials are used to sign messages exchanged in the protocols described in Section V. Therefore,  $vp(X)$  cannot be falsified without invalidating certificates used by  $X$ . Note that, to also provide a physical evidence, license plate displays the vehicle profile. Hence, a NGV profile (subject to standardization) can be read by humans and by digital




Fig. 2. Example of three cohorts and notations

equipment (e.g., cameras, radio receivers). These features are essential for admission control in heterogeneous cohorts, and calculations of smallest safe inter-vehicular gaps by NGVs and by partially automated vehicles, even those without radio capabilities [1].

Short-term certificates are sent along with messages transmitted at unpredictable times, only when necessary (see Join maneuvers, Section V). They can be reused after all the other certificates in the pool have been used once. This limits replay attacks. Compared to periodic beaconing, this reduced storage, computation and communication overhead due to authentication will be appreciated when migrating to post-quantum cryptographic algorithms.

## III. COHORTS

### A. Internal organization

A cohort is a self-organizing vehicular string forming spontaneously on major roads or highways. A cohort has a specification stating its current number of members, denoted  $n$ , bounded by  $n^*$ , its current velocity (denoted  $v$ ) and profile. An isolated vehicle or a cohort head (CH) assigns itself rank 1, while a cohort tail (CT) is ranked  $n^{th}$ . Within a cohort, vehicles are free to accelerate or brake without prior agreement. A cohort is different than a platoon, which is a pre-planned string, where vehicles are usually arranged in specific order for efficiency reasons (smallest safe gaps), ordered by increasing braking capabilities. Vehicles of any SAE level from 0 to 5 ( $aul$ ), and any cyberphysical level (composed of cyber level  $cl$  and physical level  $pl$ )  $cpl = \{cl, pl\}$  from  $\{0, 0\}$  to  $\{5, 5\}$ , may *potentially* become members of the same cohort, in any order. The introduction of cyberphysical levels and interoperability sets, in addition to SAE autonomy levels, is crucial to allow heterogeneous cohorts to self-manage themselves while keeping smallest safe gaps between neighbors of different cyberphysical levels.

In Figure 2, the gap between contiguous members  $X$  and  $Y$  is denoted  $s(v, y|x)$ , where  $v$  stands for current velocity, and is comprised between  $s_{min}(v, y|x)$  and  $s_{max}(v, y|x)$ . As detailed in Section IV, members exchange neighbor-to-neighbor (N2N) messages via a bidirectional N2N link. A N2N message contains name  $\{r, g\}$ , where  $r$  is the rank and  $g$  is the current lane number of the sender. Gaps between consecutive cohorts (in the same lane) are denoted  $S(v, \cdot)$ , lower bound  $S_{min}(v, \cdot)$ ,

such that in case of highest hard braking conditions, a cohort head never hits a cohort tail. Before a cohort head closes the gap with a cohort tail, cohort-to-cohort (C2C) messages are exchanged (see Section V). N2N and C2C communications are based on RF and optical (e.g., VLC) technologies (the latter is not detailed here). RF communications for safety are directional and very short-range, compatible with the C-V2X standards, based on beamforming and small MIMO antennas. N2N and C2C messages carry critical codes such as “set velocity to  $v$ ”, “LtJoin request”, or “immediate lane clearing”.

Every cohort  $\Gamma$  is associated a profile, denoted  $\Pi(\Gamma)$ , that varies with time and comprised of the following parameters:

- Current  $\Gamma$ 's velocity  $v$  (within some bounded tolerance)
- $n$ , current number of  $\Gamma$ 's members ( $n \leq n^*$ )
- Current  $\Gamma$ 's topology  $TP(\Gamma)$ , a global state comprised of  $n$  integers, each giving the size of a “vehicular slot”, denoted  $vslot(y, \Gamma)$  for member  $Y$  with rank  $y$ . Current  $vslot(y, \Gamma) = vl(Y) + s(v, y|x)$ , where  $X$  is  $Y$ 's predecessor. For a cohort head,  $vslot(1, \cdot) = vl(CH)$ . Vehicular slots are ordered by increasing ranks in  $TP(\cdot)$ . The physical span of a cohort (asphalt occupancy) is trivially derived from its topology.
- $\Gamma$ 's interoperability set  $Interop(\Gamma)$  (see Section V-A). Note that contrary to  $Interop(NGV)$ ,  $Interop(cohort)$  is not recorded in a certificate.

### B. Episodic beaconing

This scheme, introduced in Section 3.4.1 of [1], is implemented in the NC subsystem of every NGV, to be activated/deactivated as desired or when necessary. One purpose of this scheme is to contribute to traffic data knowledge. When activated, members of a cohort ( $\Gamma$ ) take turns in broadcasting a pseudonymous V2X message carrying the identifier of the road where  $\Gamma$  circulates,  $\Gamma$ 's velocity and length (computed out of  $\Gamma$ 's topology), and the current GNSS coordinates of  $\Gamma$ 's head and tail. Its merit, compared to the V2X periodic beaconing scheme, lies with savings regarding radio spectrum occupancy and privacy. A cohort of  $n$  members ( $n$  greater than some lower bound) must signal its existence. In that case, a individual member broadcasts at rates  $n$  times smaller than with strict periodic beaconing. Another purpose of episodic beaconing is to ensure safety when visibility conditions are poor on bidirectional roads without a central separation. In this case, episodic beaconing is activated automatically, in order to make every vehicle aware of the existence of incoming traffic. Moreover, NGVs would move at reduced velocities, warning lights on. In such circumstances, overtaking would probably be prohibited by law. That suffices for avoiding heads-on collisions, especially in the presence of cyberattacks (see Section VII).

### C. MAC protocols

MAC protocols are utilized within cyberphysical constructs, where members must access shared radio channels. Contention issues arise with N2N and C2C communications. Only deterministic MAC protocols are eligible, such as TDMA, synchronous CDMA, and CSMA-DCR (deterministic collision

resolution) [15], [16]. Such protocols (that are out-of-scope) must meet a number of requirements, notably timeliness (worst-case access delays have upper bounds, in the order of 20-30 ms [17]) and immunity to cyberattacks (attacks on MAC protocols should be detected instantly).

SWIFT is a deterministic TDMA protocol [18], which is collision-free, contrary to TDMA protocols based on reservations. In SWIFT, there is a bijection from channel slot to cohort rank. Thus, every member trivially infers the start time of its slot from its unique rank. It is necessary to prove that any cyberattack aimed at disrupting such MAC protocols in cohorts necessarily violates at least one of the PSM predicates.

## IV. INTRA-COHORT COMMUNICATIONS

### A. N2N communications

Intra-cohort communications are performed upstream (decreasing ranks) and downstream (increasing ranks). N2N messages are not authenticated (but admission to a cohort (via C2C) is subject to authentication). Members exchange N2N messages which carry C data and contain a sequence number  $sq$  proper to each member. N2N messages are acknowledged, via integers  $asq$  serving to tell which N2N message is being acknowledged ( $asq = sq$ ). N2N communications are defined for two ranges:

- Range-1 send/receive primitive for N2N messages and heartbeats. Heartbeats serve to meet dependability requirements. When there are no messages to be sent on a N2N link, neighbors exchange heartbeats to check liveness of N2N link. A heartbeat carries “I am alive”. A member ranked  $r$  processes heartbeats received only from neighbors with ranks  $r - 1$  and  $r + 1$ . Heartbeats are not acknowledged. A N2N link is diagnosed as broken when more than  $u^*$  consecutive losses are observed. After  $u^*$  misses, the diagnosing vehicle triggers a cohort split (see below).
- Range-2 send/receive primitive to disseminate N2N messages cohort-wide or between consecutive neighbors. In non-lossy conditions, a range-2 N2N message is necessarily heard by range-1 and range-2 neighbors, owing to range-controlled antennas (see Section 8 of [1]). This is essential for coping with message losses and attacks. This primitive is utilized for cohort management, cohort-wide dissemination (CWD) and cohort-wide agreement (CWA) [1]. In CWD and CWA, cohort members relay a N2N message (acceptance conditions checked, see below) received from range-2 and range-1 neighbors to opposite range-1 and range-2 neighbors. With CWA, all members agree on the same knowledge (current cohort status and profile).

### B. Acceptance Conditions (AC)

A range-2 N2N message  $m$  is accepted if the following conditions are fulfilled:

- Body fields of  $m$  received from range-2 neighbor, relayed by range-1 neighbor are identical, delivered within a stipulated bounded time interval.

- Acknowledgment numbers  $asq$  received from range-1 and range-2 neighbors (same directionality) must be returned within specified time intervals.

When AC is not fulfilled,  $m$  is discarded without being acknowledged, which may lead to a retransmission of  $m$  (up to  $u^*$  times). Owing to AC, a range-2 N2N radio communication is an atomic action, even under cyberattack. This means that for any pair of intended recipients, either both receive and process the N2N message, or none of them does.

### C. Cohort split

Given that a cohort must contain correct members only, a scheme called *cohort split* serves to meet safety requirements. Let us consider N2N radio link failures. From the bottom part of Figure 4, consider cohort  $\Gamma$ , vehicles  $P$  (rank  $\gamma - 1$ ) and follower  $Q$ , and assume that a  $P/Q$  link failure occurs. When detected by  $P$ ,  $P$  sends to  $Q$  a cohort split message via the optical N2N link (e.g., VLC), and triggers CWD upstream to send a N2N message that carries “cohort split at rank  $\gamma - 1$ , update cohort topology,  $n = \gamma - 1$ ”. The entire cohort  $\Gamma$  is thus aware of the split maneuver.  $P$  is now tail of the truncated cohort.  $Q$  (assumed to be honest) decelerates until inter-cohort spacing  $S_{min}(v, \cdot)$  is observed with  $P$ , and then triggers a downstream CWD.  $Q$  becomes head of a new cohort (comprising its followers (if any)).  $Q$  may be first to detect the link failure.  $Q$  sends  $P$  a cohort split message via the optical N2N link, and decelerates until becoming head of a new cohort. CWD algorithms are activated, as shown above. Splits can also be caused by malicious behaviors, but are detected by the PSM. Then, the companion CR subsystem is instructed to halt the misbehaving vehicle (see Section II-A), which involves a cohort split. A cohort is thus a linear set of honest NGVs without permanent internal failures. This is fundamental for three reasons. Firstly, the well-known impossibility result of achieving common knowledge in asynchronous and synchronous systems in the presence of partitioning is circumvented thanks to the cohort split scheme, which translates partitioning in cyber space into partitioning in physical space. Consequently, cohort-wide agreements are feasible despite partitioning. Secondly, the cohort split scheme solves the open problem of how to establish a worst-case bound for message losses experienced while executing CWD or CWA. Thirdly, owing to that scheme, CWD and CWA algorithms have known worst-case termination time bounds despite failures.

## V. INTER-COHORT PROTOCOLS AND MANEUVERS

Cohort-to-Cohort (C2C) communications are of two types: short-range between cohorts mutually in line-of-sight (LOS), and short/medium range when not in line-of-sight (when converging towards an unsignalized intersection for example). In this paper, we focus on short-range LOS communication (from a cohort standpoint), especially longitudinal and lateral join protocols (*LgJoin*, *LtJoin*), and *Leave*. Note that C2C acceptance conditions are handled similarly to N2N communications. Signed C2C messages exchanged in the first rounds of *LgJoin* or *LtJoin* protocols carry requestor’s

and responder’s certificate (which contains respective vehicle profile) and respective *Interops*.

### A. Cohort interoperability sets, Join Acceptance Conditions

NGVs of various generations are bound to share physical space and cyber resources. Issues raised by safe interoperability in heterogeneous vehicular networks can be solved by resorting to cyberphysical levels and interoperability sets (see Sections 6.2 and 6.3 of [1]). Cyber levels represent the communication capabilities of the vehicle (e.g., “vehicle is equipped with a C-V2X Release 14 modem”), while physical levels represent their “mechanical” capabilities (e.g., braking power).

For safety reasons, NGVs of low *cpl* shall be kept apart from NGVs with high *cpl*. Moreover, for the sake of efficiency, cohorts that contain NGVs of low *cpl* and those that contain NGVs of high *cpl* should circulate in different lanes (dynamic cohort/lane assignment). Therefore, the purpose of cohort interoperability set, denoted *Interop*, is to enable the formation of cohorts of desired heterogeneity or homogeneity levels. We anticipate the existence of standardized *Interops*, each a dense subset of the 36-element CPL matrix (see Figure 3). A notable example is that of “radioless” vehicles (i.e.,  $cl = 0$ ). Their cohorts would have a specific *Interop* comprised of the 6 elements of the first row of CPL. Figure 3 depicts an example of the CPL matrix. Sets  $\varphi(k)$  are defined as dense subsets of the CPL matrix, where elements are adjacent to each other. In Figure 3, roots  $k$  of sets  $\varphi(k)$  are shown in italic within colored circles:

- $\varphi(14)$  comprises 9 elements
- $\varphi(11)$  comprises 6 elements
- $\varphi(35)$  comprises 4 elements

*Interops* defined by authorities may have empty intersections in the CPL matrix. Accordingly, various acceptance conditions for joining a cohort (JAC) can be proposed. Let us present two simple examples. Let *Interop*(*req*) and

$i \backslash j$	0	1	2	3	4	5
0					4	5
1		7	8	9	10	11
2		13	14	15	16	17
3		19	20	21		
4					28	29
5					34	35

$$\varphi(14) = \{7, 8, 9, 13, 14, 15, 19, 20, 21\}$$

$$\varphi(11) = \{4, 5, 10, 11, 16, 17\}$$

$$\varphi(35) = \{28, 29, 34, 35\}$$

Fig. 3. The CPL matrix and interoperability sets  $\varphi(\cdot)$  for the restricted heterogeneity mode

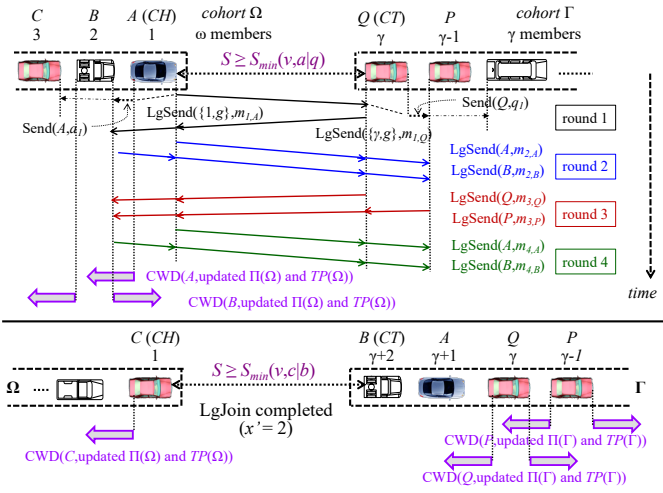


Fig. 4. LgJoin protocol (top) and Join maneuver completed with Cohort-Wide Dissemination (bottom)

$Interop(res)$  stand for the  $Interop$  of a requester and a responder respectively.

- For non-intersecting  $Interops$ , the  $Interop$  of a cohort is the  $Interop$  of its head. JAC:  $Interop(req) = Interop(res)$ .
- For intersecting  $Interops$ , a cohort's  $Interop$  is updated according to  $Interops$  proper to newly accepted members. JAC:  $Interop(req) \subseteq Interop(res)$ . For cohorts that only accept LgJoin, one would obtain dynamic formations that mimic pre-planned platoons (braking powers replaced by  $Interops$ ).

In this paper, we consider non-intersecting cohort  $Interops$ . The outcome of a LgJoin/LtJoin operation (see Section V) depends on JAC. When JAC is not met, a Join protocol and related maneuver are canceled.

Let us now introduce three important network protocols for cohort management. The following protocols are designed to cope with conflicting concurrency (i.e., vehicles that attempt risk-prone maneuvers at about the same time), which is an essential issue, almost totally ignored in published work.

### B. Longitudinal Join (LgJoin)

Figure 4 (top part) depicts the rounds of messages exchanged via directional LgSend that precede a cohort merge (which is slightly more complex than a single vehicle joining a cohort). Let us describe the LgJoin process where cohort  $\Omega$  wants to join the cohort ahead  $\Gamma$ . A detects vehicle  $Q$  with its onboard sensor(s), and measures the distance  $S$  with  $Q$  ( $S$  is kept constant during LgJoin). Recall that  $A$  and  $B$  share the same  $Interop(A)$ .

- Round 1:  $A$  adjusts its transmit power to reach  $Q$  and sends an unsigned “Hello” C2C message  $m_{1,A}$  containing distance  $S$  and  $vl(A) + s(v, A|B)$ .  $Q$  adjusts its transmit power to reach  $A$  and  $B$ , and responds by returning to  $A$  and  $B$  an unsigned “Hello” C2C message containing its size and gap with predecessor  $P$ .

- Round 2: After receiving  $Q$ 's response,  $A$  and  $B$  send a request to join.  $A$ 's pseudonym certificate (which contains the vehicle profile  $vp(A)$ ) is used to sign the join request, which also contains  $Interop(\Omega)$ . Ditto for  $B$ .  $Q$  and  $P$  verify  $A$ 's and  $B$ 's signatures and associated certificates, and check JAC (i.e., is  $Interop(\Omega)$  equal to  $Interop(\Gamma)$ ?).
- Round 3:  $Q$  and  $P$  return a negative response to  $A$  and  $B$  when JAC is not met (LgJoin is canceled). If JAC is met,  $Q$  and  $P$  respond positively to  $A$  and  $B$ .  $Q$ 's response is signed with  $Q$ 's certificate, and contains cohort  $\Gamma$  topology information and  $Interop(\Gamma)$ . Ditto for  $P$ . At the end of round 3, pairs  $(A, B)$  and  $(P, Q)$  are mutually authenticated.
- Round 4:  $A$  and  $B$  send confirmation messages to  $Q$  and  $P$  (merge is initiated), which do not need to be signed.

Figure 4 (bottom part) shows cohorts after the merge and cohort-wide dissemination needed for updating their respective internal states. At the end of LgJoin, symmetrical range-2 knowledge (of vehicle profiles) is established among new members in  $\Gamma$  ( $A$  and  $B$ ), and previous tail and its predecessor ( $Q$  and  $P$ ). In case that LgJoin would need be repeated shortly after a split, and no new vehicles have “appeared” (i.e., detected) between  $A$  and  $Q$ ,  $A$  will only send  $Q$  a hash of its certificate (used in previous LgJoin) along with the latest N2N sequence number (exchanged over link  $Q/A$  for recent proof-of-membership).

### C. Lateral Join (LtJoin)

An LtJoin operation is called to insert a vehicle ( $E$  in Figure 5) between two contiguous members ( $B$  and  $C$ ) of cohort  $\Gamma$  headed by  $A$  in an adjacent lane. Like in LgJoin, LtJoin cyber rounds serve to strike necessary agreements prior to the physical maneuver (insertion requests may be granted or rejected). Directional C2C LtSend messages have a maximum communication range in the order of 10 meters, and latencies in the order of 20-30 ms.  $E$  provides its individual  $Interop(E)$ , even if a member of some cohort. Privacy-preserving unambiguous naming/designation of a specific vehicle within line-of-sight is an open problem. A simple solution exists for cohort structured networks (out-of-scope of this paper). Owing to this solution, vehicles  $E$ ,  $B$  and  $C$  can “talk to each other”, while ignoring messages not designated to them (and vice-versa).

- Round 1:  $E$ 's pseudonym certificate (which contains the vehicle profile  $vp(E)$ ) is used to sign the join request message—which contains  $Interop(E)$ .  $E$ 's length  $vl(E)$  is especially important here.  $B$  and  $C$  verify  $E$ 's certificate, and check JAC (i.e., is  $Interop(E)$  equal to  $Interop(\Gamma)$ ).
- Round 2:  $B$  and  $C$  return a negative response to  $E$  when JAC is not met (LtJoin is canceled). If JAC is met,  $B$  and  $C$  individually compute their vote (yes/no). A negative vote may be due to a number of reasons. One is physical impossibility (e.g.,  $E$  is too large, appropriate deceleration by  $C$  is not feasible at current velocity). Another one is conflicting concurrency. For example,  $B$  is told to decelerate (a cohort-wide decision disseminated

via CWD), or  $C$  has decided to trigger a cohort split, or  $B$  and  $C$  have just granted an LtJoin request from  $F$ .  $B$  and  $C$  send each other a range-2 N2N message carrying their vote. A positive decision is reached only if both votes are positive.  $B$  and  $C$  send  $E$  their common decision via signed directional C2C messages, which contains the vehicle profile of the responder, and  $Interop(\Gamma)$ .

- Round 3: In case  $E$  learns that the decision is negative, LtJoin is aborted.  $\Gamma$ 's membership is unchanged. Prior to receiving decision messages, due to conflicting concurrency,  $E$  may want to abort its request.  $B$  and  $C$  need to know in bounded time whether  $E$  has received their decisions. Therefore, a C2C confirmation message is needed. When confirmation by  $E$  is positive,  $C$  starts decelerating to create an insertion slot with  $B$  large enough for accommodating  $E$ .  $E$  initiates the lane change maneuver, which is performed when physically feasible (under the control of its CR module).
- Round 4 (the insertion case): Once  $E$  is inserted, it inherits  $C$ 's rank (3 in Figure 5), and respective ranks of downstream cohort members are incremented via a CWD.  $B$  and  $A$  send new member  $E$  a N2N message carrying  $vp(A)$ .  $C$  and  $D$  send  $E$  a N2N message carrying  $vp(D)$ .  $E$  can trust  $vp(A)$  and  $vp(D)$  when they are received identically from two members. Ditto for  $Interop(\Gamma)$ , which  $E$  has received from  $B$  and  $C$ . Recall that  $A$  and  $D$  are aware of  $vp(E)$ —see round 2. Symmetrical range-2 knowledge is now available to  $E$ , and  $E$ 's range-1 and range-2 neighbors. As with LgJoin, symmetrical range-2 knowledge (of vehicle profiles) is needed to compute safe gaps between newly inserted neighbors.

#### D. Leave

A Leave maneuver may involve running first the LtJoin protocol in case a cohort occupies the lane targeted by the leaving vehicle. Consider a cohort of consecutive neighbors  $A, B, E, C$  and  $D$  (see Figure 5, after  $E$ 's insertion via LtJoin). Recall that every member knows the vehicle profiles of its range-1 and range-2 neighbors. Member  $C$  decides to leave, moving to lane 1 (now free).  $D$  must learn  $vp(A)$  and  $A$  must learn  $vp(D)$  promptly in order to maintain range-2 knowledge. A Leave operation by  $C$  involves the following rounds:

- Round 1:  $C$  triggers range-2 N2N messages carrying “code = I am leaving”, one aimed at  $D$ , another aimed at  $E$  and  $B$ .
- Round 2:  $C$  leaves and  $D$  closes the gap with  $E$ .  $D$  knows  $vp(E)$  and can compute a safe gap with  $E$ .
- Round 3:  $E$  and  $B$  send each a N2N message to  $A$  carrying “ $vp(D)$ ”, and a N2N message to  $D$  carrying “ $vp(A)$ ”. They both trigger CWD upstream to update cohort-wide status knowledge, thus confirming the Leave maneuver.

No certificates are used in *Leave* operations.  $A$  and  $D$  have received range-2 vehicle profiles from two sources, enabling cross-checking. When they match, they can be trusted (the 1-out-of-3 assumption). A *Leave* operation maintains symmetrical range-2 neighbors' profiles knowledge.




Fig. 5. Example of scenario. NGV  $X\{r,g\}$  has rank  $r$  in a cohort in lane  $g$ .

#### E. Complexity of Join protocols

The complexity of our protocols is measured by their amount of communication between the executing processes, i.e., the number of rounds or messages<sup>1</sup>. Join protocols are variations of well-known  $n$ -phase commit protocols ( $n > 2$ ) in distributed computing/databases [19], complexity  $\mathcal{O}(n)$ . Let  $\beta(n, m)$  stands for the complexity of our Join protocols, where  $m$  stands for the number of messages (post maneuver messages are not included). Exact complexity figures are shown below.

- LgJoin
  - Round 1: 1 C2C and 2 N2N messages
  - Round 2, 3, and 4: 2 C2C messages each
  - $\beta(n, m) = 4, 9$
- LtJoin
  - Round 1: 1 C2C message
  - Round 2: 2 N2N and 2 C2C messages
  - Round 3: 2 C2C messages
  - $\beta(n, m) = 3, 7$

### VI. ATTACKER MODEL AND CMX FEATURES

Prior to embarking on the security analysis (Section VII), we provide an overview of the attacker model, and CMX core features that thwart cyberphysical threats. In addition to threaten safety, an attacker's objective can be to disrupt traffic, or to affect one or multiple targets by forcing them to change route or stop. The consequence could be to steal the asset, to affect the fleet operator's operation (and thus has financial impact), or to endanger the asset or passengers being transported. In Section VII, we will explore a number of attacks and demonstrate how they are unfeasible or handled satisfactorily in the CMX framework.

#### A. Cyber Attacker Model

We consider the following cyber attacker model:

- External/Internal (to the network): An external attacker does not have valid security credentials to send signed messages. An internal attacker has valid security credentials and can send inaccurate information.
- Cohort member/Non-cohort member: A non-cohort member attacker is not member of the same cohort as her target(s). A cohort member attacker has passed the security checks of LgJoin/LtJoin and aims at disrupting the cohort or its member(s).
- Passive/Active: A passive attacker simply eavesdrops messages to threaten sender's privacy (the P property

<sup>1</sup>Indeed, common measures of complexity (worst-case, average-case, asymptotic complexity) are inadequate in this context.



is out-of-scope of this paper). An active attacker sends inaccurate information, manipulates content of messages, or does not follow the network protocols. For example, an active attacker illegitimately rejects a valid LtJoin request.

- Irrational/Rational: A rational attacker aims at gaining benefits from the attack (e.g., shorter time to destination, pay less). An irrational attacker does not gain any benefits but aims at disrupting the system.
- Colluding/Non-colluding: Colluding attackers work together to perform the attack. Non-colluding attackers work individually and do not know the existence of another attacker. In this paper, due to the 1-out-of-3 assumption, we focus on non-colluding attackers.
- Distant/Proximate: A proximate attacker is in line-of-sight of the vehicle under attack. Conversely, a distant attacker is in non-line-of-sight of the vehicle under attack.
- Partial/Full compromise of NGV: A partial compromise implies that the attacker’s NGV’s PSM can prohibit her from launching many of the attacks. However, in this paper, we examine the case of an attacker’s vehicle that is fully compromised, PSM included, to investigate worst case scenarios. However, as further discussed in Section VII-G, compromising the PSM requires physical access and a high level of sophistication.

For conciseness, physical attacks on an immobile NGV are not examined in detail, to the exception of the manipulation of a PSM by a dishonest human (see Section VII-G). Owing to range controlled safety-critical radio communications, an attacker is proximate to its target(s). Consequently, attacks may backfire on its originator, e.g., being involved in a crash due to its own attack. Moreover, honest NGVs can take a snapshot of the event (in the physical and cyber spaces), and report misbehaviors to competent authorities. Attackers would be exposed to judicial retaliations, which is an additional discouraging perspective.

## B. CMX Features

The power of the cyberphysical approach that underlies the CMX framework gives the opportunity of combining cyber and physical solutions to thwart a wide range of attacks.

1) *No distant cyberattacks*: Safety-critical messages are sent and received by very short-range power-controlled directional antennas (longitudinal ranges smaller than 40m approx.). A fraudulent message issued by a distant attacker has no effects on a targeted victim, since that message will not be received. Moreover, the sender must be within LOS (the LgJoin, LtJoin, and Leave cases) or within LOS shortly after having announced its arrival (e.g., intersection crossing). This is a severe limitation that constraints an attacker to be physically in proximity of the target vehicle(s). This is in contrast with the V2X approach, where the “periodic beaconing” scheme is utilized, with GNSS data in medium-range omnidirectional messages, which greatly facilitates tracking and distant cyberattacks.

2) *Proof-of-Presence (PoP) functionality*: The physical presence and motions of nearby vehicles can be monitored/checked via various devices (RADARs, LiDARs, cameras, passive/active optics (VLC), etc.). Optical technologies

are used as a secondary communication channel for N2N communications, and for enabling proof-of-presence approaches in C2C communications. For example, Dolev et al. [20] proposed Optical PUF assisted unforgeable fingerprints to provide a robust vehicle identification using optical and radio channel. Another out-of-band channel for proof-of-presence is to share Inertial Measurement Unit data [21].

3) *Agreement in cyberspace precedes safety-critical maneuvers*: Safety-critical maneuvers can only be undertaken after a specific communication protocol (LgJoin, LtJoin, Leave) has been executed fully and successfully by NGVs involved (at least three in the general case). Owing to the 1-out-of-3 assumption, the presence of an attacker is systematically detected by at least two honest NGVs, which suffices for aborting a risky maneuver in presence of an attacker. Risky situations can also be due to multiple NGVs triggering conflicting concurrent maneuvers. Most published results regarding safety are established for scenarios involving two vehicles only, which is not general enough. Thanks to the 1-out-of-3 assumption, we can consider sets of any number of NGVs, and still demonstrate safety and cybersecurity.

4) *Self-checking capabilities*: In every safety-critical system, self-checks are mandatory. This is especially the case with “unattended” systems. Fully automated vehicles are “unattended” by passengers. The CMX framework, which aims at showing that fully automated driving is realizable, includes the PSM concept, possibly a forerunner of self-checking/self-aware capabilities needed in future NGVs. An attack by a cohort member or a non-cohort member may violate some of predicates stored in a PSM, and hence, the misbehaving NGV will be halted (physically removed) by its CR subsystem, instructed to do so by its PSM.

5) *The cyber-stealth mode for V2X communications*: Contrary to V2X, periodic beaconing is not available in the CMX framework. A NGV cannot be detected from afar if its NC subsystem is inactive (i.e., no V2X messages sent out, no detectable electromagnetic energy). On very few occasions, a NGV must output a V2X message. One shall thus differentiate *core V2X communications* from other short-medium range V2X communications. Core V2X communications include:

- Extremely rare outgoing omnidirectional messages (e-Call, halted vehicle, etc.),
- Incoming I2V messages that carry traffic data and state changes relative to roads not reflected in emaps (e.g. lane closing).

A NGV with only core V2X communications enabled is said to operate in *cyber-stealth mode*. To the exception of outgoing core V2X communication, a NC subsystem is totally mute when in cyber-stealth mode, but it would not be “deaf”. Crowdsourcing services (e.g., Waze, OpenStreetMap) are available via applications uploaded in NC subsystems. Such services do not imply periodic beaconing. Other outgoing V2X communications germane to e-working and e-shopping are enabled only when the cyber-stealth mode is explicitly deactivated by passengers—who should be aware of potential risks in cyber and physical spaces. Passengers have access to the cyber-stealth mode via an on/off option. Recall that C

subsystems are isolated from NC subsystems. When the cyber-stealth mode is on, the P property ensured by a C subsystem cannot be compromised. This mode is detailed in Section 5.4 in [1].

## VII. CYBERSECURITY ANALYSIS

In this section, we follow a system engineering approach and investigate attacks on communications, network protocols, risk-prone maneuvers, and NGV system. Note that any attack based on forging or modifying a certificate or a signature is automatically detected. Therefore, we omit these attacks from the following analysis, because all protocol rounds that use signed messages are intrinsically protected. The case of certificate theft or collusion are also ignored due to their level of sophistication (see Section VII-G). For each attack, we conclude if it affects safety or efficiency.

### A. Attacks on Intra-cohort Communications

#### 1) Spoofing by an external/internal non-cohort member:

An attacker could send messages (N2N, heartbeat) on behalf of a cohort member. However, intra-cohort communications are performed via directional antennas, beams aligned on lane axis. Hence, any heartbeat or N2N message received by a cohort member from an unaligned antenna can only be generated by some vehicle foreign to that cohort. Moreover, recall that a receiver could use RF physical layer features to detect inconsistencies between the message content and the physical location of the sender (see Section II-A). Hence, an attacker that attempts a spoofing attack while not being physically in the cohort would be detected. This attack has no effects on SPEC properties.

2) *Spoofing by a cohort member:* To be successful, a cohort member attacker must be one or two ranks away from the targeted member  $M$  (rank  $r$ ). The rank read in the header of its spoofing message must be  $r \pm 2$ , or  $r \pm 1$ . Moreover, the attacker must provide the last sequence number  $sq$  used by  $M$ , which implies prior eavesdropping. If the attacker and  $M$  transmit at the same time, a (channel) collision may occur, which is a manifestation of a cyberattack. Deterministic MAC protocols mandatory for cohorts shall be collision-free, which implies no channel slot stealing (TDMA) or no code violation (CDMA). If  $M$  is silent, no (channel) collision occurs, and the attacker can transmit successfully. Positive acknowledgment(s) will be returned to  $M$ , which has not sent anything (i.e., the attack is detected). This kind of scenario may be attempted up to  $u^*$  times, after which the targeted victims ( $M$  and maybe  $M$ 's neighbors) would leave the cohort or trigger a cohort split. Moreover, spoofing wrong data is automatically eliminated thanks to the Acceptance Conditions (AC). Spoofing attacks can only generate unnecessary cohort splits, which does not affect safety but only efficiency.

3) *Attacks on N2N protocols:* In this type of attack, a cohort member does not follow the intra-cohort protocols as prescribed. For example, member  $B$  in Figure 5 omits to send a heartbeat, or to acknowledge a N2N message, or to repeat an unacknowledged N2N message, or acknowledges with a wrong  $asq$  number, up to  $u^*$  times. This is observed by its PSM

(violations of predicates) which triggers a physical halt, or by its neighbors  $A$  and  $C$ , who trigger a cohort split (violation of AC). Property S is not compromised. Unnecessary cohort splits only affect efficiency. Note that such attacks affect the attacker too, which makes her irrational.

4) *N2N message tampering/suppression/injection:* Owing to AC, a receiver ( $C$  in Figure 5) is able to detect falsification or suppression of a longitudinal N2N message  $m$  sent by a range-2 neighbor  $A$  and relayed by a range-1 dishonest neighbor  $B$ .  $C$  can accept  $m$  only if seen twice with correct sequence numbers. Hence, a tampering or suppression attack by external/internal, cohort/non-cohort member attacker is detected. If cohort head  $A$  is an attacker that injects a misleading N2N message,  $B$ , who has only one vehicle ahead (i.e., no range-2 upstream neighbor) can observe range-1 neighbor  $A$  physically. N2N messages that matter are those carrying maneuver codes which modify a cohort behavior, e.g., "immediate lane clearing" or "immediate lane change". If  $A$  does not undertake the maneuver itself, then  $B$  trivially detects the inconsistency between the message from  $A$  and  $A$ 's behavior. Consequently, attacker  $A$  cannot gain any benefit from such attack.

Even though attacking a directional antenna from the side-lobe is challenging, an attacker ( $E$  in Figure 5) can suppress downstream message from  $A$  to affect  $B$  and  $C$ 's knowledge. However,  $B$  would again notice a mismatch between  $A$ 's physical behavior and the information received (or missing) from  $A$ , thus triggering a cohort split. The attack affects only efficiency.

Similarly a dishonest cohort tail may pretend having received an (nonexistent) warning or emergency message  $w$  from a distant vehicle. Such a message can only be a medium-range omnidirectional message received by NC subsystems, then posted to C subsystems of NGVs within radio range. Consequently, more than one member must have received  $w$ . This condition is trivially checked via CWA (and majority voting). Every member is invited to declare whether or not it has seen  $w$ . When the majority condition is not satisfied, the event is ignored. The attack fails, even in the presence of multiple attackers (under the 1-out-of-3 assumption).

### B. Attacks on Longitudinal Join Maneuvers (LgJoin)

As stated in Section V-A, JAC must be met to complete LgJoin/LtJoin operations. This condition is supposed to be met in the analysis of attacks on LgJoin/LtJoin.

1) *Fake join:* The objective of such attack is to create unnecessary message exchanges or physical maneuvers. An internal non-cohort member attacker can send unauthenticated Hello C2C message (LgJoin round 1) to the cohort (tail) and ignore the other rounds without actually joining (or as a Denial of Service—see Section VII-F). A cohort tail which would have responded to the round 1 join request would abort its own LgJoin, thus avoiding sending a message and its certificate to the attacker in round 2. In case the attacker would be willing to spend one of its certificates and run the LgJoin protocol fully, without closing the gap with the cohort tail, then nothing ensues. Since the attacker keeps itself  $S_{min}$  away from the tail, safety cannot be jeopardized. This is an irrational attack.

2) *Sybil attack*: A single vehicle approaching a cohort requests a LgJoin and pretends to be a cohort composed of more than one vehicle. One should note that this attack is somewhat irrational as the potential effect can only be bloating the size of the target cohort, and the primary victim is the attacker herself as her request has a higher likelihood of being rejected (new cohort size higher than  $n^*$ ). This attack requires the attacker to send C2C messages on behalf of a ghost follower. If successful, this attack affects efficiency. To detect this attack, a receiver could use physical layer properties such as Received Signal Strength Indicator (RSSI) to identify that all messages are coming from the same exact location [22]. Moreover, the attacker would have to waste certificates to send messages on behalf of the ghost follower, diminishing its benefit further. But more importantly, the attacker's PSM would prevent sending signed C2C messages more frequently than authorized (see delay  $d$  in Section II-A), making this attack impossible.

In the case of "inverse Sybil", a dishonest cohort head claims to be a single vehicle while her cohort is actually composed of at least two members. This is equivalent to an illegitimate accept attack (see below), for the more general case when the targeted cohort is not full. This attack does not affect safety.

3) *Spoofing*: Following Figure 5, assume honest NGV  $G$  sends a *Hello* C2C message (LgJoin round 1) to  $D$ , which is dishonest.  $D$  wants to masquerade as its predecessor  $C$  (see Section VII-A), so as to force a merge with cohort  $\Omega$  headed by  $G$ , without knowing whether the JAC is met ( $Interop(\Gamma) = Interop(\Omega)$ ).  $D$ 's attack can only fail, owing to the LgJoin protocol. In its round 1 response,  $D$  must quote the existence of  $C$ . Therefore, in round 2,  $G$  and  $H$  send their signed C2C messages to  $D$  and  $C$  (antennas have been tuned for reaching  $C$ ). Owing to the 1-out-3 assumption,  $C$  is honest.  $C$  discovers that  $D$  has been silent in round 1 (no N2N message sent by  $D$  to  $C$ ).  $C$  triggers a cohort split immediately.

4) *Illegitimate accept*: Consider a full cohort  $\Gamma$  and a LgJoin by  $G$ . Dishonest  $D$  (of rank  $n^*$ ) wants to illegitimately accept  $G$ 's join request. To avoid detection by  $C$ ,  $D$  must pretend to be alone in its round 1 response (otherwise  $C$  would receive round 2  $G$ 's and  $H$ 's C2C messages). The LgJoin protocol would be successfully executed. When  $G$  closes the gap with  $D$ , CWA is triggered by  $G$  (and  $D$  if honest) to update the new cohort-wide knowledge. Since CWA is based on bidirectional range-2 N2N messages,  $C$ ,  $H$  and  $I$  will hear from  $G$ , learning that  $\Gamma$  now includes more than  $n^*$  members.  $D$ 's attack being uncovered by  $C$  and  $G$ , they immediately trigger a cohort split. Since  $n^*$  is a "pessimistic" integer (smaller than highest acceptable value), adding members in excess of  $n^*$  very briefly does not threaten safety, even in the unlikely case of a simultaneous hard braking event in  $\Gamma$  (in which case,  $D$  is a victim of its own attack). This attack fails, with an unfavorable consequence for the irrational attacker:  $D$  is first isolated, then halted by its PSM (violations of predicates).

5) *Illegitimate reject (no cohort merge)*: Another attack is for  $D$  to ignore LgJoin requests in order to diminish  $G$ 's benefit in joining the cohort (e.g., fuel efficiency, smoother ride, knowledge of upstream events via CWD or CWA). As

the first round of LgJoin only happens between  $G$  and  $D$ ,  $C$  is not aware of the request and cannot detect  $D$ 's misbehavior. This attack only affects efficiency.

6) *Forceful tailgating*: A risk-prone physical attack is when a NGV ( $G$  in Figure 5) tries to get too close to a (honest) cohort tail  $D$  (inter-vehicle distance shorter than  $S_{min}$ ) without having run the LgJoin protocol. This attack is blocked by  $G$ 's sensors (no safety threat). Shall  $G$ 's sensors be faulty,  $D$  and  $G$  could be in danger of a rear-end collision shall an emergency brake happen in  $\Gamma$ . The attack is detected by  $D$ 's sensors. To protect  $\Gamma$ 's safety,  $D$  initiates upstream CWD to inform  $\Gamma$ 's members of the split due to dishonest  $D$ 's follower, and sacrifices its efficiency by splitting with  $\Gamma$  (creating  $S_{min}$  gap between  $D$  and  $C$ ).  $D$  decelerates, forcing  $G$  to decelerate, and then changes lane to stay away from the attacker. This attack affects  $D$ 's efficiency.

### C. Attacks on Lateral Join Maneuvers (LtJoin)

1) *Fake join*: Similarly to *fake join* attack in LgJoin, an attacker aims to create unnecessary message exchanges by executing LtJoin while not actually performing the maneuver. This irrational attack is useless and detected by perception sensors. This attack does not affect safety nor efficiency.

2) *Ghost vehicle attack*: The attacker executes LtJoin on behalf of her ghost vehicle. As above, the ghost vehicle is not physically seen by the honest cohort members. This irrational attack fails.

3) *Spoofing*: To be able to spoof messages on behalf of genuine vehicles ( $B$ ,  $C$ ,  $E$  in Figure 5), the attacker must be physically at the location of either  $B$ ,  $C$  or  $E$ , which is impossible.

4) *Illegitimate reject*: Consider again  $E$  that wants to join cohort  $\Gamma$  between  $B$  and  $C$ . The LtJoin request addressed by honest  $E$  to two neighbors  $B$  and  $C$  is turned down by  $B$  or  $C$ . This could be a Denial-of-Service attack by  $B$  or  $C$ . This could also be a correct reject due to concurrency— $F$  in lane 3 has already been accepted for joining cohort  $\Gamma$  between  $A$  and  $B$ , or between  $B$  and  $C$ , or between  $C$  and  $D$ , and the deceleration by  $B$ ,  $C$  or  $D$  is being started.  $E$  is able to discriminate between these two possibilities because  $E$  observes a physical gap created by  $B$ , or  $C$ , or  $D$ . In case of contradictory responses from  $B$  and  $C$  (a violation of the LtJoin protocol in round 2),  $E$  aborts its LtJoin request. Knowing the existence of an attacker in  $\Gamma$ ,  $E$  will target another cohort to change lane. No safety threats, and marginal efficiency loss for  $E$ .

5) *Forceful join*: Consider a lateral attacker  $E$  (in Figure 5) that is denied a LtJoin because failing JAC or too long for being safely inserted between contiguous neighbors  $B$  and  $C$ . Nevertheless,  $E$  forces an illegitimate insertion. But  $B$  and  $C$  make no room for  $E$ , hence the insertion is physically impossible. Observe that this attack fails even in the presence of a member colluding with the attacker (a violation of the 1-out-of-3 assumption). If  $B$  is the colluder, nothing ensues, since  $C$  will not create an insertion slot for  $E$ . If  $C$  is the colluder,  $C$  could decelerate to enable insertion of the attacker  $E$ . That deceleration and the attacker's lateral motion would be

immediately detected by honest member  $B$  (radar, cameras). Upon detection,  $B$  immediately triggers a cohort split with  $C$  and activates an upstream CWD in order to inform members of cohort  $\Gamma$  about the split, which updates  $\Gamma$ 's profile accordingly. Attacker  $E$  is now a “fake” head of a cohort with  $C$  ranked 2<sup>nd</sup>.  $D$  would necessarily detect the deceleration performed by  $C$ , and expect a CWD involving  $B$ , its range-2 upstream neighbor in (supposedly unchanged) cohort  $\Gamma$ . But  $B$  has moved away. Noticing the lack of such a CWD (and silence from  $B$ ),  $D$  would trigger a cohort split. This attack, which isolates the attacker and the dishonest member, has no effects on safety and efficiency.

6) *Inter-cohort swap*: In this attack, a malicious cohort-member wants to swap her place in a cohort with another lateral vehicle. This attack could happen when two colluding attackers are members of two different (but geographically lateral to one another) cohorts. If this internal/cohort member/active/rational attack is successful it would lead to an insertion of a NGV with poor cyber and braking capabilities, potentially jeopardizing more seriously the cohort's safety in case of a safety-critical event. However, such maneuver must be preceded by the execution of LtJoin and Leave protocols. The genuine follower and predecessor of the attacker will observe a physical change without prior execution of the mandatory protocols. They trigger a cohort split, thus preserving safety by sacrificing efficiency.

#### D. Attacks on Cohort Split Maneuvers

A cohort split between  $X$  and its follower  $Y$  is triggered when the  $X$ - $Y$  N2N link is diagnosed as broken (by  $X$  or  $Y$ ). This may be due to ordinary failures (onboard systems, RF communications) or attacks. Whereas causes are unknown to them,  $X$  or  $Y$  observes that the number of consecutive N2N message or heartbeat losses exceeds  $u^*$ . Splits are mandatory, otherwise gaps between  $X$  and  $Y$  would be smaller than  $S_{min}$ , thus unsafe. An attacker  $X$  could purposefully trigger a series of split-join-split-join in order to affect efficiency and riding experience. However, when  $Y$  wants to “rejoin”  $X$  after a split (failures or attacks may be transient),  $Y$  is able to check whether or not a new vehicle has appeared (within LOS) between itself and  $X$  once the split maneuver was over. If not the case,  $Y$  could rejoin its former cohort (LgJoin messages carry hashes of certificates used when they became neighbors). Conversely,  $Y$  may prefer to stay away from its former cohort since  $X$  may be permanently faulty or an attacker. The effect of this type of attack is a diminished efficiency (in terms of road capacity and riding experience).

#### E. Attacks on Leave

1) *Fake Leave*: A malicious cohort member can announce a Leave maneuver without actually leaving. The objective is to initiate the CWA protocol for creating a wrongful update of the cohort profile. This attack is not effective as a CWA is only executed after the leaving vehicle has physically performed the Leave maneuver. Indeed, the predecessor and follower of the attacker would not detect physical changes with their onboard sensors. Safe gaps and range-2 knowledge are preserved.

2) *Spoofing*: As none of the messages are signed in the Leave protocol, an attacker can attempt to send a message on behalf of a genuine vehicle to fake its departure from the cohort. In round 2, the neighbors of the genuine member wouldn't observe the departure and hence abort the protocol. This attack is unsuccessful.

3) *Trigger unsafe overtaking*: This is a special case of Leave worth investigating in the specific case of bidirectional roads without a central separation (see Section 3.4.2 in [1]). For example, in Figure 5, imagine that  $B$  (honest) wants to overtake  $A$  (attacker).  $B$  informs neighbors  $A$ ,  $C$  and  $D$  (Leave round 1, N2N message carrying “code = I intend to overtake  $A$ ”). Instead of code “do not pass”,  $A$  replies with a N2N message carrying code “you are clear”, despite the fact that there are incoming vehicles.  $A$ 's goal is to create a head-on collision involving  $B$  and other vehicles. Such an attack is feasible if one assumes that the overtaking is authorized. However, at locations identified as dangerous (e.g., uphill roads, where visibility is limited), traffic laws would prohibit the maneuver. In case of temporary poor environmental conditions (e.g., heavy rain, fog, no visibility), the episodic beaconing scheme is activated by all vehicles. Every vehicle is aware of the existence of incoming traffic. That is sufficient for avoiding bad consequences of the attack (safety is not jeopardized). Moreover, the attack would be detected because contradicting the episodic beaconing scheme exerted by incoming vehicles (similar to consistency checks in V2X misbehavior detection system [22]).

#### F. Brute-force Attacks

1) *Denial of Service*: An external attacker can send N2N or C2C messages without valid credentials in order to perform a Denial of Service attack (i.e., waste communication and increase computation overhead at receivers). However, because messages are not signed, the attack would not affect safety. Note that this attack would not violate attacker's PSM predicate because she would use another radio device or a software-defined radio to send invalid messages [23]. This attack would affect efficiency because the receiver would be busy processing bad packets and could miss some deadlines (e.g., executing LgJoin or LtJoin rounds, sending acknowledgements). A receiver under attack refrains from attempting safety-critical maneuvers, shutting off the attacked directional antenna temporarily or until observing a physical change in that direction.

2) *Jamming*: Any wireless communication technologies are sensible to jamming attacks. Such attacks can be conducted over radio channels with readily available devices. Attackers may be static, or mobile. Within jamming range, no messages may be sent/received. In the V2X framework, there cannot be any cooperation or coordination among vehicles, since they are mute and deaf. Safety is threatened, notably by distant attackers. In the CMX framework, safety is ensured by resorting to very short-range directional communications based on diversified redundant technologies (radio and optics (e.g., VLC)). Therefore, NGVs can still coordinate their motions safely, exchanging range-1 N2N messages via VLC. These

messages serve to trigger a cohort split to preserve safety. NGVs may rejoin when exiting the jamming zone. This attack affects efficiency.

### G. Attacks on NGV System

To provide a more complete picture of CMX security, this section looks at the security of the NGV system per se instead of focusing on the cohort construct. Note that the following attacks require a high level of sophistication.

1) *Attacks on sensors*: Such attacks consist of blinding onboard sensors used by a CR subsystem for safe navigation. To be successful, blinding must affect diversified redundant sensors (RADARs, LiDARs, cameras, etc.), and hence, requires significant resources. Moreover, attackers must be within LOS of a targeted vehicle, and properly aligned with it. When the CR subsystem or the PSM of the attacked NGV observes a redundancy level below a minimum level, then that NGV is halted. Only efficiency is affected.

2) *Attacks on sensor diversity*: NGVs navigate in a cyberphysical space and use sensor diversity to detect cyber and physical presence of other NGVs. Therefore, an attacker can inject inconsistencies between the two modalities. For example, an attacker can display a different  $vp(\cdot)$  on her license plate (read by the camera) than the one reported in the C2C message. The attacker will simply not be accepted in the cohort. So there are no benefits for the attacker. Therefore, as discussed previously, this attack would only affect efficiency.

3) *Attacks on certificate*: Attackers might be tempted to modify their vehicle's attributes (contained in their certificate) in order to affect LgJoin/LtJoin. For example, an attacker might want a larger physical space during an LtJoin, thus claiming a larger vehicle's length. Another example is if an attacker wants to join a more "efficient" cohort (i.e., in which cohort members with higher  $cpl$  demonstrate better braking power, faster reliable communication than the attacker) without fitting in the interoperability set of this cohort. Then she has to modify her  $cpl$  and her *Interop*. Thankfully, it is impossible to modify certificates as digitally signed by a trusted authority. One could also think about colluding attackers that share their  $cpl$  (i.e., certificate) or sign messages on behalf of each other (e.g., via tunneling). The viability of these two use cases depends on the security of the PSM (see below). However, sharing or tunneling only works in the cyberspace and does not provide the physical attribute required (e.g., license plate sticker, visual cues). Hence, this forces the attacker to also create fake physical evidence, increasing the level of effort. Thanks to the protections provided by physical evidences and the PSM, these attacks are unsuccessful.

4) *Attacks on PSM*: As one might have noticed from the previous sections, the PSM plays an important role in ensuring the resilience of the system.

In addition to storing security credentials, PSM stores predicates that enforce correct behaviors<sup>2</sup>. However, due to the impossibility of proving completeness of predicates in changing environments, predicates cannot encompass every possible

abnormal behavior in cyberspace (violations of "correctness" while executing protocols and algorithms), or in physical space (refusal to perform mandatory maneuvers, e.g., cohort split). Obviously, the NGV system does not preclude the update of predicates to cover new vital protocols. Over-the-air updates are secured following best practices such as ones described in UN Regulation No.156 [24].

The physical attacker model for the PSM is a human who has access to an (immobile) NGV, and aims at exfiltrating, infiltrating, inserting, or deleting content in the PSM. One example of such attacker is an evil mechanic [25], which could swap devices or perform side-channel attacks. To protect against this type of attacker, we have to assume the PSM to be a FIPS 140-2 Level 4 device (and under Common Criteria Protection Profile to be precise). This means that the PSM is tamper-responsive. An example of solution to detect complete physical extraction of the PSM is to use Physical Unclonable Function (PUF), such as Optical PUF, coating PUF, or PUF-inside-a-PUF [26], to identify modification of its environment. Other solutions exist in the field of self-aware computing systems [27]. After detection of attacks on a PSM, the NGV is immobilized and alert messages are sent to authorities. Later, the NGV must be re-enrolled and provisioned with security credentials.

Defeating the PSM would be a significant win for an attacker. Nevertheless, similarly to attacking certificates, this attack only gives escalated privileges in cyberspace. The attacker still has to modify the physical evidence (e.g., sticker on license plate) or appearance of the vehicle to match the cyber-attributes of the target vehicle. Hence, compromising solutions that exist to secure PSM wouldn't give full privileges. That being said, attacks on PSM will be detected, rendering the NGV inoperable before it can hit the road.

## VIII. CONCLUSION

In order to guarantee safety, privacy, efficiency and cybersecurity (SPEC) properties in networks of partially and fully automated vehicles, a new framework called Coordinated Mobility for X (CMX) has been proposed. The CMX framework encompasses cyberphysical constructs (cells, cohorts, flocks) endowed with proven SPEC properties. From a network perspective, CMX relies on protocols and distributed algorithms for timed-bounded inter-vehicular communications, reliable message dissemination, and trusted explicit agreements/coordination. From a security perspective, CMX relies on onboard proactive security module and proof-of-presence.

We performed a safety and cybersecurity analysis of CMX. We considered rational and irrational attackers, and analyzed the most essential attacks on intra and inter cohort communications and agreements, safety-critical maneuvers, and NGVs. Our security analysis highlighted that, as intended, safety is not compromised by attacks, which only affect efficiency at best (see Table I).

To conclude, this paper demonstrated the strong security benefits of CMX and its robustness against many attacks. Next-Generation Vehicles that are derived from CMX may prefigure future fully automated vehicles. In light of the benefits provided by the CMX framework and NGVs, which appear

<sup>2</sup>This is a major difference compared to commonly used Hardware Security Module.

TABLE I  
EFFECT OF ATTACKS ON SEC PROPERTIES (X= NO EFFECT; ✓= AFFECTED)

Attack target	Attack	Effect		
		S	E	C
Intra-cohort	Spoofing by non-cohort member	X	X	X
	Spoofing by cohort member	X	✓	X
	Attacks on N2N protocols	X	✓	X
	N2N tampering	X	✓	X
LgJoin	Fake join	X	X	X
	Sybil attack	X	✓	X
	Spoofing	X	✓	X
	Illegitimate accept	X	✓	X
	Illegitimate reject	X	✓	X
	Forceful tailgating	X	✓	X
LtJoin	Fake join	X	X	X
	Ghost vehicle attack	X	X	X
	Spoofing	X	X	X
	Illegitimate reject	X	✓	X
	Forceful join	X	X	X
	Inter-cohort swap	X	✓	X
Cohort split	Unnecessary split	X	✓	X
Leave	Fake leave	X	X	X
	Spoofing	X	X	X
	Trigger unsafe overtaking	X	✓	X
Brute force	Denial of Service	X	✓	X
	Jamming	X	✓	X
NGV	Sensor	X	✓	X
	Sensor diversity	X	✓	X
	Certificate	X	X	X
	PSM	X	X	X

to match what is expected from fully automated driving, we encourage the ITS community to revisit current V2X solutions.

Future work will investigate the privacy property in the CMX framework, the cybersecurity of other CMX protocols for safety-critical maneuvers (e.g., for lane merge, zipper merge, and efficient crossings of unsignalized intersections of arbitrary topologies), and safe interoperability in heterogeneous vehicular networks. We will also perform simulations in order to quantify the performance of CMX protocols and algorithms that have been previously proven correct.

## REFERENCES

- [1] G. Le Lann, "Cyberphysical constructs and concepts for fully automated networked vehicles," *INRIA Research Report 9297*, Oct. 2019.
- [2] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of distributed consensus with one faulty process," *Journal of the ACM (JACM)*, vol. 32, no. 2, 1985.
- [3] U. Schmid, B. Weiss, and I. Keidar, "Impossibility results and lower bounds for consensus under link failures," *SIAM Journal on Computing*, vol. 38, no. 5, 2009.
- [4] F. Cristian and C. Fetzer, "The timed asynchronous distributed system model," *IEEE Transactions on Parallel and Distributed systems*, vol. 10, no. 6, 1999.
- [5] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE transactions on dependable and secure computing*, vol. 1, no. 1, 2004.
- [6] D. Powell, "Failure mode assumptions and assumption coverage," in *IEEE Symposium on Fault-Tolerant Computing (FTCS)*, 1992.
- [7] K. Fisher, J. Launchbury, and R. Richards, "The HACMS program: using formal methods to eliminate exploitable bugs," *Philosophical Transactions of the Royal Society of London: Mathematical, Physical and Engineering Sciences*, vol. 375, no. 2104, 2017.
- [8] T. Murray, D. Matichuk, M. Brassil, P. Gammie, T. Bourke, S. Seefried, C. Lewis, X. Gao, and G. Klein, "seL4: From General Purpose to a Proof of Information Flow Enforcement," in *IEEE Symposium on Security and Privacy (S&P)*, 2013.
- [9] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish *et al.*, "seL4: Formal verification of an OS kernel," in *ACM Symposium on Operating systems principles (SIGOPS)*, 2009.
- [10] J. M. McCune, Y. Li, N. Qu, Z. Zhou, A. Datta, V. Gligor, and A. Perrig, "Trustvisor: Efficient tcb reduction and attestation," in *IEEE Symposium on Security and Privacy (S&P)*, 2010.
- [11] J. Petit, F. Schaub, M. Feiri, and F. Kargl, "Pseudonym schemes in vehicular networks: A survey," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, 2014.
- [12] B. Alpern and F. B. Schneider, "Recognizing safety and liveness," *Distributed computing*, vol. 2, no. 3, 1987.
- [13] IEEE Standards Association, "IEEE Std 1609.2-2016, Standard for Wireless Access in Vehicular Environments—Security Services for Applications and Management Messages," 2016.
- [14] SAE, "J3016: Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles," June 2018.
- [15] M. Haddad, P. Muhlethaler, A. Laouiti, R. Zagrouba, and L. A. Saidane, "TDMA-based MAC protocols for vehicular ad hoc networks: a survey, qualitative analysis, and open research issues," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 4, 2015.
- [16] U. Roedig, A. Barroso, and C. J. Sreenan, "f-MAC: A deterministic media access control protocol without time synchronization," in *European Workshop on Wireless Sensor Networks (WSN)*, 2006.
- [17] K. Antonakoglou, N. Brahmhi, T. Abbas, A. E. Fernandez Barciela, M. Boban, K. Cordes, M. Fallgren, L. Gallo, A. Kousaridas, Z. Li *et al.*, "On the needs and requirements arising from connected and automated driving," *Journal of Sensor and Actuator Networks*, vol. 9, no. 2, 2020.
- [18] G. Le Lann, "A collision-free MAC protocol for fast message dissemination in vehicular strings," in *IEEE Conference on Standards for Communications and Networking (CSCN)*, 2016.
- [19] D. Skeen, "A quorum-based commit protocol," Cornell University, Tech. Rep. 82-483, 1982.
- [20] S. Dolev, L. Krzywiecki, N. Panwar, and M. Segal, "Optical puf for non-forwardable vehicle authentication," *Elsevier Computer Communications*, vol. 93, 2016.
- [21] J. Han, M. Harishankar, X. Wang, A. J. Chung, and P. Tague, "Convoy: Physical context verification for vehicle platoon admission," in *ACM Workshop on Mobile Computing Systems and Applications (HotMobile)*, 2017.
- [22] S. So, J. Petit, and D. Starobinski, "Physical layer plausibility checks for misbehavior detection in v2x networks," in *ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2019.
- [23] J. Petit and R. Ansari, "V2X Validation Tool," in *BlackHat USA*, 2018.
- [24] UNECE, "UN Regulation No. 156 - Software update and software update management system," E/CE/TRANS/505/Rev.3/Add.156, Mar. 2021.
- [25] J. Petit, M. Feiri, and F. Kargl, "Revisiting attacker model for smart vehicles," in *IEEE International Symposium on Wireless Vehicular Communications (WiVeC)*, 2014.
- [26] U. Rührmair, J. Martinez-Hurtado, X. Xu, C. Kraeh, C. Hilgers, D. Kononchuk, J. J. Finley, and W. P. Bursleson, "Virtual proofs of reality and their physical implementation," in *IEEE Symposium on Security and Privacy (S&P)*, 2015.
- [27] A. Jantsch, N. Dutt, and A. M. Rahmani, "Self-awareness in systems on chip—a survey," *IEEE Design & Test*, vol. 34, no. 6, 2017.