

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/337023484>

Cyberphysical Constructs and Concepts for Fully Automated Networked Vehicles

Technical Report · October 2019

CITATIONS

0

READS

56

1 author:



Gerard J Le Lann

National Institute for Research in Computer Science and Control

83 PUBLICATIONS 1,112 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Safety-Critical Communications in Intelligent Vehicular Networks [View project](#)



Safety, privacy, trust, and immunity to cyberthreats in future vehicular networks [View project](#)



Cyberphysical Constructs and Concepts for Fully Automated Networked Vehicles

Gérard Le Lann

**RESEARCH
REPORT**

N° 9297

16/10/2019

Project-Team RITS

Cyberphysical Constructs and Concepts for Fully Automated Networked Vehicles

G rard Le Lann¹

Project-Team RITS

Research Report N  9297 — 16/10/2019

Abstract: Human lives are at stake in networked systems of automated vehicles. Drawing from mature domains where life/safety critical cyberphysical systems have already been deployed as well as from various scientific disciplines, we introduce the SPEC (Safety, Privacy, Efficiency, Cybersecurity) problem which arises in self-organizing and self-healing networks of fully automated terrestrial vehicles, and CMX functionalities intended for vehicular onboard systems. CM stands for Coordinated Mobility, X stands for S, P, E and C. The CMX framework encompasses cyberphysical constructs (cells, cohorts) endowed with proven properties, onboard proactive security modules, unfalsifiable cyberphysical levels, protocols and distributed algorithms for timed-bounded inter-vehicular communications, reliable message dissemination, trusted explicit agreements/coordination, and privacy preserving options that insulate passengers from illegitimate internal cyber-surveillance and external eavesdropping and tracking. We establish inter alia that safety and privacy can be obtained jointly, by design. The focus of this report is on SE properties. Notably, we show how to achieve theoretical absolute safety (0 fatalities and 0 severe injuries in rear-end collisions and pileups) and highest efficiency (smallest safe inter-vehicular gaps) jointly, by design, in spontaneous cohorts of vehicles. Results conveyed in this report shall open new opportunities for innovative research and development of high societal impact.

Key-words: Intelligent transportation systems, fully automated driving, vehicular networks, onboard robotics, inter-vehicular wireless communications, communicating autonomous vehicles, cyberphysical constructs, cells, cohorts, cyberphysical levels, safety, privacy, efficiency, cybersecurity, multi-access protocols, lossy channels, message dissemination, distributed agreements, worst-case time bounds, gap control.

¹ gerard.le_lann@inria.fr

Cyberphysical Constructs and Concepts for Fully Automated Networked Vehicles

Résumé : Les vies humaines sont en jeu dans les réseaux de véhicules automatisés, à l'instar de domaines matures où des systèmes critiques en matière de sécurité-innocuité ont déjà été déployés. Les connaissances acquises dans ces domaines ainsi que dans diverses disciplines scientifiques permettent de définir le problème SPEC (Safety, Privacy, Efficiency, Cybersecurity) qui se pose dans les réseaux auto-organisés et auto-réparateurs de véhicules terrestres à conduite entièrement automatisée. On introduit CMX, un ensemble de fonctionnalités destinées aux systèmes bord. CM est l'abréviation de Coordinated Mobility, et X signifie S, P, E et C. L'ensemble CMX repose sur des constructions cyberphysiques (cellules, cohortes) dotées de propriétés prouvées, les concepts de module de sécurité proactif et de niveaux cyberphysiques infalsifiables, des protocoles et des algorithmes distribués pour communications inter-véhiculaires en temps borné, dissémination fiable de messages, coordination et accords explicites dignes de confiance, ainsi que sur des options de protection de la vie privée qui permettent aux passagers d'interdire la cyber-surveillance illégitime interne et externe (écoutes radio et pistage des trajets). On établit qu'il est possible de garantir conjointement sécurité-innocuité (safety) et respect de la vie privée (privacy), par conception. Ce rapport est consacré aux propriétés SE. En particulier, on montre comment obtenir la sécurité-innocuité absolue théorique (taux nul de mortalité et de graves blessures en cas de collisions longitudinales) et maximiser l'efficacité (espaces inter-véhiculaires minimaux) conjointement, par conception, dans les cohortes spontanées de véhicules. Les résultats contenus dans ce rapport devraient ouvrir de nouvelles perspectives de recherche et développement à fort impact sociétal.

Mots clés: Systèmes de transport intelligents, conduite entièrement automatisée, réseaux véhiculaires, robotique embarquée, communications inter-véhiculaires sans fil, véhicules autonomes communicants, constructions cyberphysiques, cellules, cohortes, niveaux cyberphysiques, sécurité-innocuité, respect de la vie privée, efficacité, cyber-sécurité, protocoles d'accès multiple, canaux non fiables, dissémination de messages, accords distribués, bornes de temps pires cas, contrôle d'espacement.

0. Executive Summary	6
1. Introduction.....	8
2. Overview of Our Approach	11
3. System Model	17
4. Cells	22
5. Cohorts.....	24
6. Cyberphysical Levels.....	36
7. Cohort Admission Control, the LgJoin Operation	45
8. Gap Control, Safety and Efficiency Properties	48
9. Human Factors, Cyber-surveillance and Liabilities	55
10. Conclusions and Perspectives	56
Glossary.....	58
Bibliography.....	60

0. Executive Summary

A glimpse of history

The Intelligent Transportation Systems (ITS) community explores issues related to partially and fully automated driving. Early works on autonomous vehicles (AVs) in platoon formations appeared in the late 1970s, aimed at improving safety and efficiency (SE) properties on the roads, namely, ratios of fatalities and severe injuries as well as inter-vehicle gaps smaller than those experienced with human driving. Works conducted by the robotics community are so far the unique solid foundations regarding SE properties. Results (sensing capabilities and behavioral algorithms) are rooted in scientific and engineering disciplines—along with proofs and quantified SE figures for various risk-prone scenarios. Existing AVs (in 2019) are equipped with onboard systems that build on these results. Fatalities and severe injuries reported with AVs, along with related lawsuits, bring to light a major ambiguity. AV users are told that they do not have to drive “most of the time” whereas, contradictorily, they must commit themselves to intervene appropriately whenever needed, knowing that they will probably be held liable for accidents, no matter what.

For a number of technical reasons—sensors are inoperative in non-line-of-sight conditions, and non-technical reasons—e.g., charging toll fees automatically, the concept of connected AVs (CAVs) has emerged later on. CAVs have onboard systems augmented with message passing capabilities based on wireless radio communications. SE issues in strings (spontaneous linear formations) and networks of CAVs have been addressed by the Cooperative ITS (C-ITS) community since the turn of the century. In broad terms, the C-ITS community is comprised of the robotics community and the communication/networking community, herein referred to as CC-ITS, whose main goal is to define vehicle-to-everything (V2X) functionalities intended to enhance SE properties achieved by onboard robotics capabilities.

Where do we stand?

From the beginning, the CC-ITS community has opted for Wi-Fi technologies referred to as DSRC (Dedicated Short Range Communications), and companion protocols now translated into IEEE and ETSI standards. DSRC-V2X functionalities are backed by the Car 2 Car Communication Consortium (C2C-CC). The 5G Automotive Association (5GAA) which is promoting cellular wireless technologies (such as 4G LTE and 5G) advocates for Cellular V2X (C-V2X) functionalities. The automotive industry is divided on those choices. Some members of C2C-CC are also members of 5GAA.

Careful analyses reveal that V2X functionalities fall short of the expectations. None of the most elementary principles in safety engineering for life/safety critical cyberphysical systems are met with V2X functionalities. Notably: No bounded communication delays, no demonstrated high reliability/availability in the presence of realistic rates of message losses or other failures, lack of solutions for explicit inter-vehicular agreements and coordination, no quantified figures regarding savings in deaths and severe injuries, no demonstrated “augmented” efficiency. In the CC-ITS “cooperative driving” world, CAVs do not “cooperate”. Rather, they exchange messages unreliably and move according to computations performed unilaterally. Neither periodic beaconing (PB) nor local dynamic maps (LDMs)—see below—help in guaranteeing that such moves do not conflict with decisions made concurrently by other CAVs, a central issue with regard to safety. Proofs of properties for worst-case conditions are mandatory in life/safety critical systems. Contrary to solutions coined by the robotics community, designs that underlie V2X solutions come with no proofs. Simulations, testing and experiments cannot replace proofs for worst-case scenarios. Moreover, some V2X functionalities violate well-known impossibility results established since the 1980s in the area of Distributed Computing.

Among unfortunate design choices, PB and LDMs deserve a special mention. Local dynamic maps are built out of periodic broadcasts of messages (beacons) that carry unencrypted vehicles’ geolocations and other characteristics. Beacons shall be issued at frequencies ranging between 1 Hz and 10 Hz. Besides being useless regarding SE properties (this can be proven), periodic beaconing is harmful, exposing CAVs to cyberattacks. Unknown remote static or mobile attackers located up to approx. 300 m away from broadcasters can provide themselves with local dynamic maps and launch cyberattacks aimed at specific zones or/and CAVs. V2X functionalities can undermine SE properties achieved by onboard robotics capabilities.

Finally, V2X functionalities facilitate illegitimate cyber-surveillance. Despite reliance on Public Key Infrastructure services and certificate-based pseudonymous broadcasts, it has been shown that passengers’ privacy can be compromised via cyber eavesdropping and physical tracking of vehicles. V2X functionalities are counter-productive as regards privacy and cybersecurity (PC) properties. Public awareness is growing and unlimited cyber-surveillance is increasingly considered to be ethically and morally unacceptable. This has not gone unnoticed. In 2017, the UNECE World Forum for Harmonization of Vehicle Regulations (WP 29) published a Resolution on Data Protection in Automated and Connected Vehicles, calling “all relevant parties involved to fully respect the users’ rights to the protection of their personal data and privacy”. The C2C-CC has

objected to these recommendations, asserting that it is impossible to deliver road safety without breaching passengers' privacy. To be valid, that statement should be backed with an impossibility proof. Such a proof has not appeared yet and will never appear for the simple reason that safety and privacy properties can be achieved jointly, by design, proofs given.

CAVs conformant to V2X functionalities are equivalent to smartphones-on-wheels. Smartphones-on-wheels are susceptible to cyberattacks. They may kill. The automotive industry has entered into competition (cooperative competition) with the digital giants and the telecommunications industry. Monetization of personal data inferred from partially obfuscated information collected wirelessly (a common goal) is greatly facilitated by V2X functionalities.

A broader and positive perspective

An open question is whether *safety and privacy and efficiency and cybersecurity* (SPEC) properties can hold in spontaneous networks of *fully automated (unmanned) vehicles* (SAE level 5). If it is possible to answer in the affirmative, then some of the solutions may be applied to partially automated driving (SAE level 4 and below), and not the opposite. In this report, we consider self-organizing and self-healing networks of next generation vehicles, denoted NGVs in order to avoid confusion with CAVs equipped with V2X functionalities. We introduce CMX functionalities intended for NGVs onboard systems. CM stands for Coordinated Mobility, and X stands for S, P, E, and C.

CMX functionalities encompass cyberphysical networking constructs, such as cells and cohorts (strings with a specification), endowed with proven properties, onboard proactive security modules, unfalsifiable cyberphysical levels, protocols for timed-bounded inter-vehicular communications, distributed algorithms for reliable message dissemination, for trusted explicit agreements/coordination, and privacy preserving options that insulate passengers from illegitimate internal cyber-surveillance and external eavesdropping or tracking. PB/LDM schemes are discarded. The focus of this report is on SE properties. Notably, we show how to achieve theoretical absolute safety (0 fatalities and 0 severe injuries) in rear-end collisions and pileups as well as highest efficiency (smallest safe inter-vehicular gaps) jointly, by design, in spontaneous and heterogeneous cohorts of NGVs. CMX-based solutions for avoiding head-on and lateral collisions in various settings (e.g., unsignalized intersections, zipper merging) and for achieving PC properties in networks of NGVs will appear in forthcoming publications.

5GAA members may have a decisive advantage over the C2C-CC community, owing to the potential of cellular radios. However, superiority in communication capabilities does not suffice. For achieving the SPEC properties, higher-level solutions drawn from or inspired by the CMX framework shall be implemented atop cellular radio channels. Future users, who are increasingly aware of the risk of eavesdropping and tracking with V2X functionalities, in addition to the risk of being injured and/or held liable for accidents experienced with CAVs, may prefer to buy, to rent, to share next-gen vehicles, and to ride next-gen taxis.

To summarize:

Contrary to widespread belief, safety and privacy can be achieved jointly, by design.

In this report, we demonstrate that it is possible to achieve theoretical absolute safety (0 fatalities, 0 severe injuries) in rear-end collisions and pileups, as well as highest efficiency (smallest safe inter-vehicular gaps) jointly, by design, in cohorts of fully automated vehicles.

Safety and efficiency properties achieved with onboard robotics can be degraded by V2X functionalities, owing to cyberattacks. Furthermore, privacy threats inexistent with onboard robotics become a concern with V2X functionalities.

CMX functionalities augment safety and efficiency properties achieved with onboard robotics significantly (theoretical absolute safety is feasible). Safety is immune to remote and local cyberattacks. There are no privacy risks other than those incurred via direct vision (by humans, cameras, etc.) with contemporary vehicles.

We have to decide now: In which type of motorized society do we want to live?

1. Introduction

Approximately 40 years ago, the concept of distributed computing has emerged in the wake of the computer networking revolution, prompted by technological advances. The spotlight would be definitely turned on communication protocols, cyber constructs, and distributed algorithms rather than on computers and centralized control. “The network is the computer” was a famous motto. Replace computer with vehicle. The same evolution is taking place before our eyes. There is however an essential prerequisite for the deployment of vehicular networks: safety proofs. In ordinary distributed cyber systems, human lives are not endangered, whereas vehicular networks are life/safety critical distributed cyberphysical systems that may cause fatalities and inflict severe injuries.

If history is any guide, solutions for making automated driving a reality can only rest on cyberphysical constructs and concepts, enriched with protocols and algorithms based on a number of specific disciplines, among which one finds robotics, distributed computing, and safety engineering.

1.1. From platoons to spontaneous autonomic vehicular networks

In the late 1970s, truck platooning has been studied by the robotics community [1-2]. Primary goals were higher safety and efficiency (small gaps between trucks at medium/high velocities). Work on Automated Highway Systems led to distinguishing 4 categories: autonomous vehicles, low cooperative vehicles, high cooperative vehicles, and platooned vehicles [3]. To the exception of autonomous vehicles (AVs), all vehicles are assumed to be equipped with radio communication capabilities. These technologies supplement onboard systems based on robotics (sensors, actuators, and kinematics/motion control laws). Approximately 20 years later, partially automated driving was studied in the context of ad hoc networks of autonomous cooperative vehicles (VANETs), leading to the concept of connected automated vehicles (CAVs). More recently, conditions and settings for partially/fully automated driving have been refined through the SAE automated driving levels ranging from 0 to 5 [4].

Wireless radio communications and network protocols at the core of Vehicle-to-Everything (V2X) functionalities for CAVs are defined in documents produced by the Cooperative Intelligent Transportation Systems (C-ITS) community [the automotive industry (manufacturers, OEMs, tier-1 suppliers), the Car 2 Car Communication Consortium (C2C-CC), the 5G Automotive Association (5GAA), and standardization bodies such as IEEE and ETSI committees in charge of wireless vehicular communications]. We assume familiarity with this documentation. Some C2C-CC’s members are also members of 5GAA. The automotive industry is divided on the choice of wireless radio technologies (Wi-Fi or cellular).

“Connected” is an ambiguous attribute, almost inevitably understood as “connected to Internet, clouds and Web-based services”. That is a satisfactory interpretation if one is mostly interested in providing vehicles and passengers with, e.g., traffic data, infotainment, automated toll collect, remote assistance, mobile e-working and e-shopping environments. When safety and efficiency are primary goals, “communicating” rather than “connected” would be a more appropriate qualifier, since safety implies that vehicles shall be able to interact with one another directly, deterministically and quasi-instantaneously.

An open question is whether *safe and efficient fully automated driving* is feasible. Safety can be trivially demonstrated assuming very large inter-vehicular distances and very cautious behaviors (e.g., low velocities), leading to congested roadways and travel times higher than with human driving. Efficiency (e.g., smallest inter-vehicular distances at highest velocities) is a fundamental requirement.

In this report, we consider next generation vehicles—denoted NGVs in order to avoid confusion with CAVs equipped with V2X functionalities—forming self-organizing self-healing short-lived or long-lived networks referred to as autonomic vehicular networks (AVNs). NGVs could prefigure future SAE level 5 vehicles. Vehicular networks examined by the C-ITS community have no structuring constructs. We introduce two cyberphysical constructs, communication protocols, coordination algorithms, and the concept of cyberphysical levels which supplement SAE automated driving levels. Issues that arise in heterogeneous AVNs comprising vehicles of diverse cyberphysical levels are also addressed.

A more general question is whether *safety and privacy and efficiency and cybersecurity* (SPEC) properties can be achieved altogether in AVNs. The answer is no with V2X functionalities as currently envisioned in the C-ITS community. The answer is yes with CMX functionalities introduced in this report, CM standing for Coordinated Mobility, X standing for S, P, E, and C.

Fully automated driving is a reachable goal. Exploitation of knowledge spanning a number of specific disciplines suffices for demonstrating the SPEC properties. This report is focused on safety and efficiency properties. Privacy and cybersecurity properties are established in forthcoming publications and reports.

1.2. The SPEC problem

The number of fatalities in road accidents amounts to approx. 9 per day to date in France. US figures are approx. 10 times higher. The challenge is thus to know whether and how NGVs may coordinate safely (and efficiently). Since radio communications are prone to eavesdropping and favor cyberattacks, we also need to show how to thwart privacy threats and cyberattacks that would compromise safety. Therefore, the following four properties shall hold:

- **Safety**: Ratios of crashes are significantly smaller than those achieved with human driving and partially automated driving (human supervision and interventions). Ideally, fatalities and severe injuries are eliminated. Crashes should entail property damage only.

- **Privacy** (passive adversaries): No personal data can be inferred or extracted from cyber-centric information (wireless communications), from physical-centric information (paths and routes followed by vehicles).

- **Efficiency**: Vehicular densities (resp., inter-vehicular gaps) are significantly higher (resp., smaller) than with non-fully automated vehicles for identical velocities. (This is antagonistic with the safety property.)

- **Cybersecurity** (active adversaries): Safety is not compromised by remote or/and local cyberattacks (masquerading, Sybil attacks, man-in-the-middle attacks, message falsification/suppression, injection of bogus data, intrusions, replay attacks, DoS such as, e.g., jamming).

Those four properties define the SPEC problem.

Our overall objective is to solve the SPEC problem as it arises in every setting, from superhighways to urban settings, where NGVs, scooters, bicyclists, and pedestrians happen to share streets and roads. For example, we have devised solutions for safe and efficient crossings of unsignalized intersections in the presence of pedestrians, as well as for explicit NGVs-pedestrians coordination.

In this report, we examine AVNs comprised of NGVs only. All kinds of NGVs are considered (privately owned cars, shared/rented cars, taxis, commercial vehicles, trucks, emergency vehicles, etc.), forming spontaneous and heterogeneous AVNs that share highways and roads with a central separation, without any restrictions (no pre-assigned lanes, no forbidden lanes). NGVs and AVNs are life/safety critical cyberphysical systems of systems. Such systems have already been designed and deployed in various domains (e.g., nuclear or chemical plants, air transportation, manned space missions). Therefore, in addition to automation control, control theory, robotics, kinematics, wireless communications, network protocols, and distributed fault-tolerant algorithms (for example), basic principles in safety engineering have been our design drivers.

1.3. V2X functionalities versus CMX functionalities

Broadly speaking, the C-ITS community is comprised of the robotics community, referred to as RC-ITS in this report, and the communication/networking community, herein referred to as CC-ITS, whose main goal is to define V2X functionalities, X standing for V (vehicle), I (infrastructure), P (pedestrian), and N (network). The CC-ITS community has followed a bottom-up approach, focusing first on Wi-Fi technologies and standards such as IEEE 802.11p and ETSI ITS-G5 (MAC protocol), IEEE 1609-4 (channel management), referred to as DSRC (Dedicated Short Range Communications), backed by C2C-CC's members. Cellular technologies (4G LTE and 5G) are promoted by 5GAA's members. The automotive industry is currently divided on which technology should be adopted: DSRC-V2X or C-V2X. Various enhancements are under current examination, among which one finds IEEE 802.11bd and 5G NR V2X.

A number of application-level V2X functionalities have been defined and turned into standards, such as SAE J2945/1-2.2, ETSI TS 102 637-2 (Cooperative Awareness Messaging) and SAE J2735 (Basic Safety Messaging), ETSI TS 102 637-3 (Decentralized Environmental Notification service), IEEE 1609-2 and ETSI TS 102 940 (security services), to name a few. The CC-ITS approach rests essentially on short/medium range broadcast communications, road-side units and remote services reachable via telecommunication networks. V2X functionalities are based on hybrid mobile edge/cloud computing.

Despite intrinsic limitations, works conducted by the RC-ITS community are so far the unique solid foundations regarding safety and efficiency properties. In agreement with a number of experts, we came to the conclusion that these properties are not improved in worst-case hazardous scenarios with the adjunction of V2X functionalities. Furthermore, periodic beaconing of CAMs or BSMs that carry GNSS space coordinates (in particular) is harmful. Since CAMs/BSMs are broadcast over medium ranges, they can be heard by unknown remote entities (vehicles and static nodes). This favors privacy threats—which are increasingly considered ethically and morally unacceptable, and safety achieved by onboard robotics can be ruined by remote cyberattackers. Shortcomings of V2X functionalities are examined in a separate report.

To solve the SPEC problem, one must address the four properties altogether at every design stage, which implies following a top-down approach. CMX functionalities encompass cyber-physical constructs (vehicular cells and cohorts in this report), aperiodic event-triggered messaging, communication protocols, and distributed coordination algorithms based on very short range directional LOS and NLOS communications. Antennas are power/range controlled. Beamforming and beamsteering capabilities derive from sensor data (radars, ultrasonic sensors and cameras, for example). Optical communications (ignored so far by the CC-ITS community) are an integral part of the CMX framework. The crucial merit of very short range directional radio/optical communications lies with the possibility of instantiating reciprocal and explicit inter-vehicular coordination quasi-instantaneously. In contrast with V2X, CMX functionalities are based on mobile edge computing, making safety immune to hazards and threats incurred with accessing and utilizing cloud-based services.

Designers of functionalities “added” to onboard robotics meant to solve the SPEC problem shall meet the following requirements:

- R1: Demonstrate quasi-instantaneous distributed inter-vehicular coordination in worst-case conditions, giving analytical expressions of time bounds for coordination in the presence of highest radio channel contention and highest failure patterns (to be characterized).

- R2: Demonstrate (e.g., logical proofs, analytical calculus) that these functionalities improve safety and efficiency properties achieved with onboard robotics. Establish reduction ratios for fatalities, severe injuries and safe inter-vehicular gaps (reduced asphalt occupancy), under clearly stated assumptions, of high coverage.

- R3: Demonstrate that with these functionalities, (1) remote eavesdropping, tracking and cyberattacks are impossible, (2) nearby eavesdropping, tracking and cyberattacks are inoffensive (neither privacy nor safety can be compromised).

The above requirements are not met with V2X functionalities. Consequently, they cannot provide life/safety critical services. Requirements R1 and R2 are addressed in this report. Solutions for meeting requirement R3 are given in forthcoming publications. However, the potential of CMX functionalities with regard to privacy and cybersecurity concerns can be briefly outlined as follows (see also Subsection 5.4). Firstly, with very short range aperiodic directional communications, eavesdropping and cyberattacks by remote or local adversaries are more complicated, thus less likely, than with medium range periodic broadcast communications. Secondly, as will be seen, messages in the CMX framework are of no interest to would-be eavesdroppers. Metadata are non-reversible, and message bodies do not carry sensitive data, such as GNSS space coordinates. CMX communications are privacy preserving. Thirdly, owing to diversified redundancy and specific designs, notably message acceptance conditions based on range-2 communication primitives, safety is immune to cyberattacks, including brute-force cyberattacks such as physical jamming of radio channels.

As regards local adversaries, the usual adversary challenge changes radically. Whatever its computational power, a local adversary must travel alongside targeted NGVs while they exchange aperiodic messages, long enough for its spying to be of any interest or for launching cyberattacks against these vehicles. It follows that such adversaries are observable by vehicles under potential spying or cyberattacks. Spying is not “anonymous” any longer, and possibly reciprocal. Moreover, eavesdropping on very short range radio communications does not aggravate spying as it can be practiced today. (Humans and cameras can read vehicles plates, see passengers in nearby vehicles and collect pictures.) Local cyberattacks aimed at creating catastrophes (crashes) may impact the attacker(s) themselves. Likely, they would belong to the category of irrational attacks.

Some members of the CC-ITS community assert that safety and privacy cannot be achieved altogether.

That would be a universally true statement if backed with an impossibility proof. Such a proof has not appeared yet and will never appear for the simple reason that the SPEC problem as it arises in AVNs can be solved (proofs provided).

In fact, asserting that safety and privacy cannot be achieved altogether is tantamount to acknowledging that V2X functionalities do not solve the SP (safety, privacy) problem, let alone the SPEC problem.

Research work conducted by various teams has generated important results. Innovative solutions have paved the way to a global solution to the SPEC problem. Combining knowledge accumulated for almost half a century with novel radio/optical communication technologies opens up very promising perspectives.

In this report, we demonstrate that it is possible to achieve theoretical absolute safety (0 fatalities, 0 severe injuries) in rear-end collisions and pileups, and highest efficiency (smallest safe inter-vehicular gaps) jointly, by design, in linear formations of next-gen vehicles.

1.4. Contributions

- A presentation of some fundamental design principles in life/safety critical system domains and in safety engineering,

- A perspective on how robotics, computer science and cyberphysics supplement each other in the C-ITS area.

Regarding CMX functionalities:

- An introduction to cyberphysical constructs—the vehicular cell and the vehicular cohort,
- A brief presentation of a partitioned architecture for onboard systems,
- An introduction to the concept of proactive security modules,
- A detailed presentation of the cohort construct and cohort splits,
- A presentation of very short-range intra-cell, intra-cohort and inter-cohort communication primitives needed for safe and efficient inter-NGV coordination,
- Examples of time bounds achieved with deterministic MAC protocols and distributed coordination algorithms in the presence of highest channel contention and highest message loss rates,
- An introduction to cyberphysical levels and to the concept of cyberphysical level interoperability for NGVs and cohorts,
- How to enable the spontaneous formation of cohorts endowed with desired heterogeneity degrees,
- How to perform cohort-lane assignment dynamically,
- Analyses of safe gap control schemes in strings and in cohorts,
- How to achieve theoretical absolute safety and high efficiency in heterogeneous cohorts.

This report is organized as follows. Our approach and relationships between robotics, distributed computing and cyberphysics are presented in Section 2. A system model is introduced in Section 3, encompassing a schematic onboard system architecture, the concept of proactive security modules, and critical communications contrasted with non-critical communications. We state our assumptions regarding failures, cyberattacks, loads, and we recall essential dependability requirements to be met in life/safety critical systems. In Section 4, we present the cell construct. The cohort construct is detailed in Section 5, notably intra-cohort and inter-cohort communication primitives, message formats, naming, MAC protocols, distributed coordination algorithms which can sustain high message loss frequencies, and cohort splits. In Section 6, we introduce the concepts of cyberphysical levels, vehicle and cohort profiles, and sets of interoperable cyberphysical levels, as well as cohort operative modes, from unrestricted heterogeneity to strict homogeneity. Controlled cohort admission is described in Section 7 with the longitudinal join cyberphysical operation. In Section 8, we present gap control schemes for theoretical absolute safety and highest efficiency, and we give lower bounds of directional antennas ranges. Section 9 briefly addresses human factors, liability concerns, and privacy issues arising with illegitimate cyber-surveillance. Conclusions and perspectives are exposed in Section 10, followed by a glossary of terms and a bibliography.

2. Overview of Our Approach

Let us first survey some essential principles and well known results that underlie the CMX framework. Then, we elaborate on how robotics, distributed computing and cyberphysics cross-fertilize each other.

2.1. Basic principles in safety engineering and life/safety critical system designs

Safety engineering is a mature field [5-7]. A fundamental design principle states that life/safety critical functions and non-critical functions must be defined and specified separately. Their respective implementations (hardware, firmware, software) are physically separated, in isolation from each other.

As regards critical functions, fundamental design principles of interest to this report are as follows:

- Diversified functional and physical redundancy is mandatory; otherwise response times are infinite in the presence of failures.

- No reliance on external help, since an external entity may malfunction, may fail, may be compromised, or too slow. GNSS data (time and space coordinates) are the only exception. In systems operated via control centers (e.g., ships and maritime transportation, planes and air transportation), constituent entities are bound to

follow pre-planned trajectories and must broadcast their space coordinates at regular intervals, so that control centers can check whether trajectories are correctly followed by constituents.

- Designs shall be conducted top-down, starting with the overall problem under consideration (bottom-up designs get quickly outdated due to continuous technological advances).

- Assumptions relative to types of failures and cyberattacks as well as their worst-case occurrence densities must be stated. Figures of reliability and availability achieved under worst-case conditions shall be given.

- Assumptions relative to worst-case computational and communication loads (event arrival densities) must be stated. Analytical expressions of relevant worst case response times/reaction delays shall be given (timeliness proofs).

- Safety properties sought must be defined unambiguously. Proofs must be given for worst-cases.

- Validation of an overall system design precedes validation of software designs/implementations, which precede experimentation, simulation, testing and certification.

Applied to NGVs and AVNs, the above principles can be restated as follows:

a) Life/safety critical functions (C functions) and non-critical functions (NC functions) aimed at onboard systems are designed and specified separately. They are implemented on distinct onboard subsystems (processors, storage, operating systems, etc.) that are physically separated, in isolation from each other.

For C functions:

b) Diversified redundancy as regards onboard sensors, processors, storage, communication devices (for example) is mandatory.

c) No reliance on entities external to a NGV such as, e.g., road-side units, cloud-based services or control centers. Access to GNSS data is the only exception: a NGV must be able to locate itself in 2D space (within some accuracy) as well as be aware of current UTC (known out of GNSS time). Contrary to constituents of systems operated via control centers, NGVs in AVNs do not follow pre-planned routes. Therefore, they are not bound to broadcast their respective space coordinates at regular intervals. NGVs shall rely solely on their ability to sense each other and to communicate directly among each other in order to make coordinated safety-preserving decisions, very much like TCAS (Traffic Collision Avoidance System) in air transportation.

d) Designs shall be conducted top-down, starting with the overall SPEC problem. Designs for high-level functionalities must not depend on specific communication technologies, contrary to bottom-up approaches. Those approaches restrict the choice of appropriate solutions artificially, they lead to costly revisions of designs and they necessitate redoing certification, due to emerging communication technologies.

e) Assumptions relative to types of failures and cyberattacks as well as their worst-case occurrence densities (onboard systems, inter-NGV communications) must be stated. Quantified dependability properties (reliability, availability, etc.) shall hold under worst-case conditions.

f) Assumptions relative to worst-case computational and communication loads (C event arrival densities) relative to onboard systems and inter-NGV communications (notably MAC-level contention) must be stated. Timeliness properties (upper bounds on response times/reaction delays) are essential. Their analytical expressions shall be given for worst-case conditions.

g) Safety must be proven for worst cases. Quantified ratios of harmful crashes (fatalities and severe injuries) experienced or avoided shall be given. Since safety trivially holds with large inter-vehicular distances (poor efficiency), savings in inter-vehicular gaps (asphalt occupancy) shall be expressed along with corresponding ratios of harmful crashes. Ideally, theoretical absolute safety shall be proven, under realistic assumptions, for inter-NGV gaps smaller than inter-CAV gaps in vehicular strings.

Safety has to do with *proofs* for *worst-case conditions*. Figures relative to harmful crashes derived from discrete event simulations or experimental testing are insufficient. Authors have raised the issue of the credibility of discrete event simulation studies [8]. Worst cases cannot “emerge from” techniques and tools based on statistical approaches. Such approaches are useful for providing *average figures* obtained with a model (protocol, algorithm, setting, vehicular configuration, collision patterns, etc.) for *non-worst case conditions*, which model must have been validated and verified beforehand. Figures of average delays, average message losses, average crashes obtained by simulation cannot be trusted unless *simulated models are proven to match real objects* (demonstrations or sufficient evidence). Most often, this essential requirement is ignored. Ditto with testing and experiments (see Subsection 2.4).

Validation, verification and certification issues are not within the scope of this report. That they play a pivotal role in critical systems has been known for decades (see standards such as ISO 26262 and ASILs). With modern systems, verification and certification shall partially be performed at runtime, since certification has evolved from compliance with standards to the construction of assurance cases, which requires appropriate frameworks [9].

Finally, too often, no clear distinction is made between system engineering and software engineering. When a system fails, it is commonplace to blame the software. In fact, in many cases, roots of failures can be found in early phases of a system lifecycle. Initial overall requirements are necessarily expressed in some natural language, which is prone to errors and misinterpretations. Therefore, overall specifications of “what is desired”, of designs that meet “what is desired” (system design specifications), may differ (possibly significantly) from specifications matching “what is *really* desired”. See for example the Ariane-5 launcher failure in 1996 [10] and the Boeing 737-Max crashes in 2018 and 2019. An incorrect specification of application-level requirements leads to a system design specification which can only result in a failure or a catastrophe, even if correctly implemented in software, formally verified or not (that is not the issue).

2.2. Cross-fertilization of robotics and distributed computing

Distributed cyber systems could not have come into existence without the concepts of communicating processes, protocols and synchronization/coordination algorithms for fair, reliable, deadlock-free mutual exclusion, serializable concurrent executions, resource sharing, etc. Unsurprisingly, similar concepts and solutions transposed into physical space are needed for safe and efficient fully automated driving.

2.2.1 Resource sharing—Reactive and proactive safety

Asphalt (2D systems), asphalt and air space (3D systems) are the shared resources of interest with terrestrial and aerial vehicular networks. There are three classes of algorithmic solutions for resource sharing: detection-and-recovery, prevention, avoidance. The former class is inapplicable (one cannot “roll back an accident”). Prevention is aimed at prohibiting the emergence of hazardous (unsafe) or deadlock-prone (unsafe and inefficient) conditions. Solutions are the province of distributed algorithms (computer science). Avoidance is relied on for maintaining non-hazardous conditions while making progress (also, in case some of the assumptions that underlie prevention schemes would be violated). Solutions are the province of automation control and robotics.

Proactive safety—Prevention schemes

Whenever needed, vehicles run (cyber) algorithms or protocols in order to preclude the emergence of hazardous conditions, prior to executing physical motions (collision-free trajectories) safely. Akin to processes in networked cyber systems, NGVs influence their respective behaviors, directly and reciprocally via e.g., x-way handshakes, 2-phase locking, atomic commit, and they strike explicit agreements prior to engaging risk-prone maneuvers or whenever some hazardous condition develops.

Reactive safety—Avoidance schemes

With avoidance schemes based on sensing technologies and robotics capabilities, vehicles within LOS monitor each other, keeping safe inter-vehicular distances while behaving according to explicit agreements settled beforehand via prevention schemes.

This is how computer science and robotics can be “married” consistently. Both types of safety are needed. That is the rationale for our work, which builds upon results established by the RC-ITS community.

2.2.2 Guesses and hypothetical predictability

Safety problems in open spontaneous vehicular networks are more complex than those solved in physical space postulating advance knowledge of future physical events or state transitions. This observation applies in particular to vehicles’ trajectories computed in advance. Unsignalized intersection crossing is a typical example. Claims of safety can be found in publications where vehicles located in the same multilane entrance cross an intersection following pre-computed trajectories that do not intersect. It is thus postulated that vehicles do not have to move as necessary so as to cope with untold intentions and/or to unexpected hazards. Moreover, prior to entrance, vehicles are assumed to be ideally positioned in appropriate lanes for avoiding intersecting trajectories. Clearly, such assumptions have a very poor coverage. Reality is more complex.

These observations apply to all designs based on assuming predictability, a classical hypothesis. Vehicles are supposed to sense, plan and instantiate “safe” trajectories, silently. Consequently, every vehicle has to make guesses relative to future behaviors of other vehicles, has to make provisions for events that would force a vehicle to deviate from a predicted behavior, and must “dare” entering risk-prone maneuvers. For proving that

designs based on assuming predictability are safe, it is necessary to anticipate and to analyze all cases where a vehicle might be forced to undertake an evasive maneuver. That is a formidable challenge with realistic scenarios, where uncoordinated evasive maneuvers may be undertaken at about the same time by multiple vehicles, i.e. in the presence of generalized conflicting concurrency—an issue largely ignored so far.

With approaches based on “guessing”, safety is conditioned on assuming that if a vehicle deviates from some predicted trajectory, such a deviation must be instantiated with a long enough distance for other vehicles to adjust their own trajectories. This hypothesis is problematic. Firstly, how may a vehicle “know” which assumptions have been made by other vehicles relative to its future trajectory, in every possible scenario? Secondly, how to prove that vehicles will never deviate from postulated trajectories too late, that is when the “safe distance” requirement is not met? Thirdly, in order to specify the “safe distance” requirement in realistic settings, i.e. where more than 2 or 3 vehicles come into play, designers must identify and specify every possible chain of safety-critical remote (NLOS) and nearby (LOS) events and related causal dependencies that lead a vehicle to deviate from some given trajectory, furthermore accounting for transient and permanent failures.

That seems unrealistic. Think also of a distracted young child who suddenly runs across a street, forcing a vehicle V to swerve into nearby vehicle W for avoiding a fatality, or a fast moving misbehaving vehicle X forcing V to swerve into W for avoiding a fatal collision with X . In both cases, roots of causal chains may remain unknown (the culprits may disappear promptly), whereas V would be declared guilty. And that is not necessary, notably under approaches where vehicles tell each other, with feed-back if so desired, which are their intended behaviors, rather than being “mute”, thus doing away with “guessing” and “daring”. The mathematical model for safety assurance presented in [11] is a recent example of approaches that achieve reactive safety. (In fact, rather than proving safety, the formal model of blame behind the Responsibility-Sensitive Safety approach is aimed at knowing who might be held liable for an accident.) It is not too difficult to demonstrate that safety holds in worst case scenarios when inter-vehicular distances are “large enough”, in the absence of failures and cyberattacks. The real challenge is to prove safety without sacrificing efficiency, under realistic assumptions. Thanks to radio and optical communications, and explicit inter-vehicular agreements that achieve proactive safety, we do not have to keep relying on guesses.

2.3. Distributed computing and cyberphysics

Seminal work in Distributed Computing appeared in the late 1970s/early 1980s. Various timing models have been defined, ranging from full synchrony to full asynchrony, along with optimality and impossibility results regarding problems in the presence of concurrency, variable delays and failures (e.g., fair mutual exclusion, leader election, atomic commit, distributed consensus) [12].

Note that “consensus” has three different meanings of relevance here. In distributed systems and discrete time models, consensus is the problem of deciding on some unique value out of multiple proposals. In vehicular robotics and continuous time models, consensus is a problem at the core of string stability (no shock waves due to repeated accelerations and decelerations). In multi-agent networked systems (swarms and flocks), consensus is the problem of achieving collective collision-free motion of agents out of local knowledge relative to neighbors’ behaviors.

It might also be useful to clarify a possible ambiguity regarding “safety”. In critical physical systems, safety means “no catastrophes”, while in computer science safety means “no violations of predicates”. In the latter case, safety is trivially ensured if no state transitions take place (thus the additional requirement of liveness [13]), which is exactly the converse of what must happen in the physical space, where “doing nothing” (e.g., not braking) is unacceptable. Proving “logical” safety and liveness is necessary, albeit not sufficient when “physical” safety is sought. Failure models essential to demonstrations of safety and principles for Dependable Computing were established at the turn of the century [14].

2.3.1 Spontaneous open networked systems and privacy-preserving naming

Spontaneity implies openness, i.e. lack of advance knowledge regarding AVN membership, names or identities of vehicles, geolocations, velocities, intentions, and occurrences of hazardous conditions, to name a few. Knowledge needed to achieve safety can only be gained on-line and quasi instantaneously by onboard processes. We are thus facing problems significantly more complex than problems solved in Distributed Computing where one traditionally considers a set of n processes numbered from 1 to n , or carrying priorities numbered from 1 to n .

In open systems, no process has advance knowledge of n , of other processes names or/and priorities. One may be tempted to follow a “cut-and-paste” approach, by re-using algorithms designed for distributed cyber systems to solve “similar” problems in cyber-physics. Most often, that leads to invalid solutions. For example, in numerous publications, one finds “solutions” to various problems of collision-free motions (e.g., crossings of unsignalized intersections) based on distributed consensus algorithms designed for cyber space. That such

algorithms are proven correct for closed systems, and not for open systems, is a recurrent overlook. Solutions that rest on postulating (i) the existence of n competing vehicles, (ii) that each vehicle owns a unique name (from 1 to n), are of no (theoretical, practical) interest, since the naming problem in open systems is harder than consensus. From a theoretical perspective: One cannot *assume* advance knowledge of n and vehicles' names, and then assert that one has solved consensus in vehicular networks, for this amounts to *postulating* that a problem harder than consensus is solved in order to solve consensus. From a practical perspective: Postulates (i) and (ii) do not match reality.

This observation applies to every distributed computing problem \mathcal{P} arising in spontaneous cyberphysical networked systems. The naming problem must be solved prior to addressing \mathcal{P} . The more complex *privacy-preserving naming (PPN)* problem arises in AVNs where, by definition, vehicles have transient neighbors. Registration plates can be read (cameras for example). For privacy reasons, such identifiers shall not be used within messages. Similarly, vehicles shall not reveal IP/MAC addresses of onboard devices. Hence the questions: Which identifiers/names shall be used by communicating vehicles? In AVNs, is it possible to provide neighboring vehicles that have never exchanged any messages beforehand with knowledge of their respective privacy-preserving identifiers/names *at first encounter*? Informally, the *PPN* problem for univocal inter-vehicular communications is as follows.

Without resorting to GNSS space coordinates or non-reversible names/identifiers:

- Show how a vehicle V can “talk to” neighbor W and let other neighbors know that V does not want to “talk to” them, when V is unaware of which name shall be used for designating W , right here, right now,
- Reciprocally, show how W can “talk to” V (and let other neighbors know that W does not want to “talk to” them) when W is unaware of which name shall be used for designating V , right here, right now,
- Show how V can be certain that its responder is W (no masquerading).

An instance of the *PPN* problem arises with lane changes (lateral join maneuvers). Assume that V wants to move to an adjacent lane, inserting itself ahead of Q which follows P . This is the optimal choice for V 's insertion (minimizing time and energy needed to instantiate the lane change). V , P and Q may reveal their non-reversible names/identifiers only after they have established an unambiguous triangular “connection”.

In a weaker version of the *PPN* problem, V provides W with a non-reversible name/identifier (the second requirement in the above is trivially met). The weak *PPN* problem arises with longitudinal join maneuvers, when a NGV or a cohort is caching up with a NGV or a cohort. Note that geocasting cannot be considered due to reliance on GNSS space coordinates. Moreover, geocasting is sub-optimal: since recipients of m cannot tell right away which one of them is targeted by V , additional message exchanges are necessary. Naming issues are addressed in Section 5 and Section 7.

2.3.2 MAC protocols and impossibility results in fully asynchronous networked systems

Designers of MAC protocols face algorithmic issues germane to the naming problem. The aim of MAC protocols is to dynamically assign names (integers referred to as “collision-free access times” or “back-off delays”) to vehicles contending for channel access. MAC protocols that generate *unique names in strictly bounded finite time* are the exception.

That is not an issue of minor importance. MAC protocols play an essential role with regards to possibility and impossibility results. With collision-prone probabilistic MAC protocols (CSMA-CA for example in the V2X framework), there are no finite upper bounds on channel access delays—see Subsection 3.4. At best, there are finite upper bounds but their values remain unknown. Systems resting on such protocols belong to the category of *fully asynchronous systems*.

Designers of V2X functionalities are thus faced with numerous impossibility results regarding the construction of global sates (local dynamic maps built out of periodic beaconing cannot be consistent) [15-17], reliable broadcast [18], and distributed consensus, in the presence of failures, to name a few examples. In [19], it is shown that distributed consensus is impossible in the presence of a single failure (an onboard software process or a processor in our case) even though communications are fully reliable. The impossibility of consensus in the presence of message losses is demonstrated in [20] for loss rates beyond modest thresholds. To the best of our knowledge, no solutions aimed at circumventing these impossibility results can be found in the CC-ITS literature. These roadblocks are eliminated in AVNs where NGV onboard systems are equipped with deterministic MAC protocols whereby channel access delays have proven upper bounds under worst-case contention conditions.

2.3.3 Pseudonymous communications

From the above presentation of the *PPN* problem, it follows that *V* and *W* must own unambiguous names prior to exchanging proofs of authentication. Pseudonym schemes [21] serve to obfuscate identities of message senders—as well as to check whether messages are issued by honest authenticated vehicles. However, they do not allow for spontaneous univocal designations of specific vehicles, as required in the *PPN* problem. Virtues of pseudonym schemes are undermined when coupled with periodic beaconing, i.e. broadcasts of CAMs/BSMs that contain GNSS space coordinates, which favors tracking, thus privacy threats [22]. The rationale for pseudonymous communications is to eliminate such threats. They are not with periodic beaconing.

Random pseudonym changes do not suffice to avoid tracking [23]. Silent periods make tracking attacks more complex [24]. However, silence is antagonistic with safety. A number of authors have studied tradeoffs between durations of silent periods and crash occurrence. In [25], an analysis of intersection crash avoidance establishes that silent periods shall be shorter than 2 seconds, highlighting the need for designs that achieve privacy and safety properties jointly.

This matches our overall approach to the SPEC problem: no tradeoffs, both properties shall hold, fully, by design. Pseudonym schemes are an integral part of CMX functionalities (periodic beaconing is not).

2.3.4 Dependability and coverage

A “cut-and-paste” approach to dependability would also be inappropriate. Highest densities of failures or cyberattacks are location-dependent and time-varying. A major difficulty lies with utilizing as such results established for non-open cyber systems/networks where correctness and termination of protocols or algorithms are proven to hold for up to f failures (faulty senders/receivers and message losses) or/and cyberattacks. Beyond f , all bets are off. Consequently, in AVNs, it is impossible to claim—even less to prove—safety unless a finite worst-case upper bound f^* is established for f . In the CC-ITS literature, where can we find designs that establish the existence of bounds f^* ? What is a “good” value for f or f^* in a spontaneous open system? That value cannot be “frozen” in advance. That value cannot be assumed “agreed upon” dynamically by vehicles, for this is the very problem under consideration. Something else is needed.

A fundamental principle in safety engineering states that application-level functionalities in charge of safety must meet ultra-high reliability and availability requirements. More precisely, failures of such functionalities shall have probabilities of occurrence well below $\varepsilon = 10^{-x}$ per hour, $x > 3$ [26], page 35 in [27] and [28]. Moreover, assumptions under which a safety critical functionality matches these probabilities should be closely scrutinized, in order to check whether they reflect reality. Designs/solutions that rest on assuming some idealized reality are of no interest. Events and/or conditions that are not considered at design time shall be clearly stated. They are called “residual assumptions”, deemed to have probabilities of occurrence much smaller than ε . The concept of assumption coverage introduced in [28] can be generalized and applied to properties other than dependability as follows: The coverage of an assumption, of a demonstrated property, is defined as the probability that the said assumption or property matches operational reality. In our case, this applies to functionalities based on wireless communications that are “added” to onboard robotics. In the V2X framework, local dynamic maps and periodic beaconing are application-level functionalities in charge of safety. Where in the CC-ITS literature can we find proofs or evidence that packet delivery ratios meet the ε requirement, under assumptions of very high coverage?

2.4. On training, learning, AI, and testing

Just like human-driven vehicles, NGVs will form spontaneous networks. To some extent, spontaneity and safety are not antagonistic in the case of contemporary vehicles since humans must learn and follow driving rules aimed at safety. The following question arises with NGVs: Is it necessary to rely on advance training and learning for mastering unexpected hazardous scenarios safely? Algorithmic learning is quite useful for pattern matching, accurate positioning, object-centric recognition and disambiguation—no false negatives, no false positives. Despite known weaknesses (deep learning techniques can be fooled), algorithmic learning is efficient when applied to static or semi-static data: Accumulated knowledge is not nullified due to changes in the universe considered. That is not the case with data collected by observing behaviors of existing vehicles on today’s roadways or during experimental testing.

Firstly, such data are quickly outdated, since technologies (e.g., onboard software, sensors, and driving assistants) are in constant evolution. If learning is not a convergent process, learning cannot play a pivotal role regarding safety, notably safety proofs for worst-case scenarios.

Secondly, assuming convergence, what do we get? Driving rules found in existing manuals is one possible outcome. If the case, why don’t we implement these deterministic algorithms in onboard systems? Another imaginable possibility is to infer how human-driven vehicles mixed with AVs and CAVs manage to avoid

collisions “almost always” in every possible setting and under every possible condition, as well as why and how collisions do occur. That is a very unlikely outcome. We still do not understand how birds or fish avoid hitting each other in dense flocks or swarms (see Section 3). Something fundamental appears to be overlooked. More than the level of “intelligence” imparted to an onboard system, *collective* “intelligence” matters in AVNs, determined by explicit interactions and reciprocal influences between NGVs—the networking fabrics. It is hardly imaginable that appropriate distributed coordination algorithms needed for safety and efficiency, along with their correctness proofs, can “emerge” out of data mining or deep neural networks.

Thirdly, how many crashes and fatalities shall be experienced before asserting that learning is “good enough”?

Fourthly, given that no automated vehicles of SAE level 5 can be seen yet, what can be learned or inferred *today* about *future* safe driving rules that will apply when these vehicles are on the roads? Most existing vehicles are human-driven or supervised by humans (CAVs). Humans do not always obey driving rules. On the contrary, NGVs will follow behavioral rules (protocols and coordination algorithms) that, very likely, will differ from current ones.

Fifthly, conclusions reached by a number of authors relative to the testing of AVs or CAVs must be emphasized. In [29], authors show that the number of miles involved with testing is way too high for demonstrating that AVs/CAVs reliability in terms of fatalities and severe injuries is significantly higher than achieved with human driving. Their conclusion: “Under even aggressive testing assumptions, existing fleets would take tens and sometimes hundreds of years to drive these miles—an impossible proposition if the aim is to demonstrate their performance prior to releasing them on the roads for consumer use.”

It follows that learning out of currently observable scenarios cannot help much for safety concerns. Trustworthy advances that build on AI are bound to emerge in the future. For example, tools based on AI techniques “smarter” than current design assistants, theorem provers or model checkers shall help in designing protocols and algorithms at the core of onboard systems as well as in proving their properties. Regarding on-line decision making, algorithmic learning has a role to play in specific sensing problems where data sets are not subject to unexpected modifications. How to dynamically adjust values of physical parameters postulated at design time to real values experienced on the roads is an important topic. See Subsection 8.3 for an example.

One crucial question is how to transition from current situation to fully automated driving, while addressing concerns documented in numerous surveys, where potential users have voiced apprehension about traveling in driverless vehicles piloted by AI-based algorithms. In the interim, it might be wise to adopt design frameworks and solutions that minimize the gap (regarding safety related decisions) between human cognition/reasoning and algorithmic/artificial intelligence, by devising and proving protocols and behavioral algorithms that can be understood, checked, and certified by humans assisted with appropriate tools. That is the premise at the core of the CMX approach. CMX functionalities are presented in the remainder of this report.

3. System Model

3.1. Next-gen vehicles

3.1.1 Onboard systems—An overview

Vehicular networks are heterogeneous, comprised of NGVs endowed with diverse capabilities and rights (SUVs, private/shared/rental cars, ambulances, police cars, fire brigade, trucks, trailers, and so on). A NGV is equipped with an onboard system—a multiprocessor, which receives inputs from diverse sensors, notably radars, lidars, ultrasonic sensors and cameras, a GNSS device, a speedometer that delivers current velocity denoted v , radio antennas and optical devices. Sensors serve to measure inter-vehicular gaps. Actuators serve to change velocities and to enforce safe gaps, computed by C processes fed with cyberphysical data which are essential to safe driving.

An onboard system runs various firmware/software processes in charge of C and NC functions. A subset of C functions, referred to as vital functions, must meet ultra-high reliability and availability requirements (see Subsection 2.3.4). Unique names assigned to arterials in the physical world appear in human-readable (paper printed) maps. The same goes for e-maps. They show unique names assigned to arterials, to lanes in arterials, starting from 1 for rightmost lane (leftmost lane in some countries). Regarding geolocation, we assume lane-level accuracy for lateral coordinates [30] and inaccurate longitudinal coordinates. Thanks to geo-positioning capabilities and emaps, a NGV knows the type of roadway where it circulates (e.g., city street, country road, highway) and the corresponding highest authorized velocity, denoted v^* . Thanks to onboard sensors, a NGV knows weather and surface conditions (e.g., dry, wet, snowy). In some countries, weather dependent authorized velocities are displayed on road-side panels which can be read by cameras.

Under the CMX approach, “short-term” certificates can be reused. It is thus not necessary for a NGV to import new “short-term” certificates via remote Public Key Infrastructures when running out of certificates, since the notion of “running out” does not make sense. For this reason, such certificates are referred to as *reusable* certificates. It might nevertheless be required to proceed with a renewal of certificates from time to time for *administrative* reasons, unrelated with (inexistent) certificate exhaustion.

Besides data linked with pseudonymous authentication, long term and reusable certificates contain a vehicle profile, denoted $vp(X)$ for vehicle X (see Subsection 6.1.5). It follows that $vp(.)$ data read in a valid certificate can be trusted.

Vehicle X 's profile $vp(X)$ found in a certificate cannot be falsified without invalidating that certificate.

External NGV interfaces such as human readable and digital plates (subject to standardization) provide nearby neighbors with knowledge relative to NGV profiles. These features are essential as regards admission control in heterogeneous cohorts (see Section 7) and calculations of smallest safe inter-vehicular gaps by NGVs and by vehicles without radio capabilities (see Section 8).

The number of reusable certificates stored in a PSM at registration time is small (a few hundreds). They are mainly used in conjunction with event-based C operations such as, for example, lane changes, join maneuvers (admissions in cohorts), and crossings of unsignalized intersections (the UX problem) or unsignalized roundabouts (the UR problem). Why should a vehicle “prove” that it has been authenticated each time it sends a message? That is unnecessary if, as commonly assumed, most vehicles are not malicious. With vehicular networks that lack structuring constructs and that rely on V2X communications, there might be no alternative. That is not the case with cohorts which are trustworthy linear AVNs (see Section 7).

A NHTSA study establishes an average of 1 lane change per 2.8 miles traveled [33]. Extrapolating these findings (which hold for human-driven vehicles) and considering other C maneuvers performed at comparable rates (a NGV utilizes 1 certificate per join maneuver, be it a requestor or a responder), one arrives at a rough estimate of at most 100-120 certificates consumed per day assuming daily journeys of 80 miles. With approximately 900 certificates stored in a PSM, at least one week would elapse between two consecutive utilizations of the same certificate. It follows that certificates can be reused without exposing NGVs to tracking. Reusable certificates are sent along with aperiodic messages transmitted over very short ranges at unpredictable times. These messages carry unencrypted data that have absolutely no value for an eavesdropper (see further).

It follows that privacy cannot be compromised with N2N or C2C communications. Visual spying (plate reading, photos of nearby car passengers) is vastly more efficient than eavesdropping on radio channels. A preview of privacy and cybersecurity properties can be found in Subsection 5.4. Thanks to pieces of information disseminated throughout this report, an attentive reader may infer some solutions aimed at privacy and cybersecurity based on the CMX framework.

3.2. Introduction to very short range N2N and C2C critical communications

Onboard sensors may fail or be under cyberattacks. Diversified redundancy (sensor fusion) serves to cope with such undesired events. Obviously, radio communications cannot improve safety properties achieved by onboard robotics unless they also rest on diversified redundancy, in conformity with **principle b**).

In the CMX framework, C communications are based on radio and optics (Visible Light Communications, passive optics). Optical communications are essential for coping with jammed radio channels. In this report, we focus on radio communications exclusively. N2N and C2C communications are performed via diverse antennas and channels, accessed through channel specific MAC-level protocols, thus providing for the necessary redundancy above the physical level.

Cohorts are strings of NGVs that follow each other in the same lane. An isolated NGV is a particular case—a cohort of size 1. The number of NGVs members of a cohort is denoted n , with a limit denoted n^* . Only those vehicles in close longitudinal and lateral proximity may hit each other.

Therefore, in conformity with **principle c**), *direct* very short-range inter-NGV communications suffice for achieving safety. Directional N2N communications are used by members of the same cohort and by NGVs that want to perform lane changes. Directional C2C communications are used by NGVs members of consecutive cohorts in the same lane. In the case of cohorts that converge towards an unsignalized intersection or roundabout, collision prevention rests on explicit coordination shortly before entrance, based on very short-range semi-omnidirectional C2C communications.

We have devised the following communication primitives:

- Range-1 longitudinal unicast, aimed at reaching 1 closest member, via the *send primitive*, for exchanging heartbeats and N2N management messages.
- Range-2 longitudinal multicast, aimed at reaching 2 closest contiguous members, via the *Send primitive* (intra-cohort N2N messages).
- The range-2 longitudinal *Receive primitive* which mirrors the *Send primitive*.
- Range-2 and range-3 longitudinal multicast, aimed at reaching up to 3 closest NGVs that belong to two cohorts that follow each other, via the *LgSend primitive* (inter-cohort C2C messages).
- Range-1 lateral multicast (resp., unicast) aimed at sending a N2N message to closest lateral neighbors (resp., to a specific lateral neighbor) in an adjacent lane, via the *LtSend primitive* (not detailed here).
- Very short-range C2C semi-omnidirectional broadcast triggered via the *SBcast primitive* (not detailed in this report). This primitive serves to achieve inter-cohort coordination at unsignalized intersections and roundabouts.

In conformity with **principle d**), designs of communication primitives, message formats, contents, and protocols have been driven by the SPEC problem. Thus the decision to have no GNSS space coordinates quoted in N2N messages or C2C messages, contrary to CAMS, BSMs, and DENMs.

The decoupling of safety and privacy is an essential benefit of the top-down approach that underlies the CMX framework.

3.3. Failure assumptions, cyberattacks and dependability

In conformity with design **principle e**), we state our failure assumptions, aligned with the terminology defined in [14] and by the Distributed Computing community. As for radio antennas, transceivers, and message handlers, we consider the following failures: stop, send omission (a message to be sent out is not transmitted) and receive omission (an incoming message is not delivered). Regarding radio channels, we consider stop and omission failures (message losses). For onboard processors, we consider stop failures. These transient failures are recoverable since C functions rest on diversified redundancy.

Incorrect executions of vital functions—introduced in Subsection 3.1.1—due to transient hardware faults shall not compromise safety. Consequently, a vital function must be executed on highly reliable processors or instantiated as a set of replicated software processes executed on 2 or 3 processing units (e.g., a multicore processor with no common-mode failure), assuming no Byzantine behavior [34], depending on whether fault detection or fault masking is required. In this report, we consider that Byzantine failures are vanishingly rare, i.e. they belong to the set of residual assumptions.

Loss of a vital function is a fatal failure. Vital functions must meet ultra-high reliability and availability requirements: probabilities of occurrence of fatal failures shall be well below 10^{-x} per hour, $x > 3$ (see Subsection 2.3.4). A fatal failure is promptly detected at system level (thanks to heartbeats), leading to a cohort split and to the halting of a failing NGV (see Subsection 5.2.4). It follows that a cohort is free from fatal failures. It should thus be possible to prove that densities of fatal failures occurrence meet the above stated reliability/availability requirements.

3.3.1 The 1-out-of-3 assumption

Communications shall be immune to cyberattacks, message suppression or/and falsification, in particular. All solutions aimed at thwarting cyberattacks rest on assuming adversary coalitions of limited cardinality, such as “only a minority of vehicles may behave maliciously”. What is a minority of an undefined ensemble? In the absence of a precise characterization of “vehicular network” (boundaries, highest number of vehicles, etc.), such an assumption is meaningless. Moreover, assuming “only a minority...” on the average does not eliminate the possibility of having to face short-lived local coalitions where a majority of vehicles are malicious. Unambiguous hypotheses shall be stated if one seeks to prove that cyberattacks compounded with failures cannot jeopardize safety. That is feasible with cohorts.

Inspired by the classical “at most n -out-of- m ” faulty processes assumption in Dependable Computing, we have worked out solutions based on assuming “at most 1 malicious or faulty member among any 3 consecutive members”. Under this assumption, N2N message relaying, cohort-wide message dissemination and cohort-wide agreement are immune to N2N message falsification and/or suppression (cyberattacks can be detected

instantly). We have not found published work where safety is proven to hold in vehicular stings where up to one third of the members may fail and/or may launch cyberattacks. This assumption can be generalized so as to cope with larger adversary coalitions (e.g., 2-out-of-5 assumption), leading to more costly solutions. This assumption applies as well to C2C messages exchanged in the course of LtJoin and LgJoin operations.

3.3.2 Loads and timeliness

In Subsections 5.1.4, 5.2.1 and 5.2.2, in conformity with design **principle f**, we give analytical expressions of worst-case time bounds δ , Δ and Ξ achieved with protocols and coordination algorithms that are based on C communications. Loads that must be sustained by onboard processors and radio channels shall be kept as low as possible. The triggering of C functions is event-driven, hence aperiodic. We establish arrival loads (messages, failures) under which bounds δ , Δ and Ξ hold true (see Subsection 5.2.3). These load conditions are trivially fulfilled in real AVNs. Merits of event-based solutions become evident when compared to periodic CAMs/BSMs broadcasts performed at 1-to-10 Hz frequencies by dozens or hundreds of vehicles. The resulting cyber pollution is unjustified, leading to wasting huge amounts of computational and communication resources.

3.4. Short/medium range V2X communications

In charge of NC functions, NC sub-systems will be equipped with V2X communication capabilities, radio ranges in the order of 300 m. (Such ranges are unnecessarily large for fulfilling safety goals.) Highest numbers of radio channel contenders are unknown since they depend on local time-varying conditions (topology, traffic density). Moreover, CSMA-CA is a collision-prone probabilistic MAC protocol. It follows that worst-case channel access delays are unbounded. Demonstrations have been given long ago [35-36], corroborated in numerous well respected publications. They hold a fortiori in vehicular networks that rely on CSMA-CA. This is confirmed in [37], page 40, where one reads “the latency generally increases with the number of stations, and a maximum latency cannot be guaranteed”. Furthermore, average-case channel access delays may be unacceptably large, due to contention [38], and dishonest executions [39-42].

Since safety and efficiency properties are obtained with very short-range N2N and C2C communications, NC subsystems do not need to perform periodic CAM/BSM beaconing. A NC subsystem cannot write in a C memory space (see Fig. 1). A NC subsystem that receives an event-driven V2X message of interest to a C subsystem writes that message in a specific part of the memory space shared with a CC subsystem, and posts an event. Upon notification, the CC subsystem imports the V2X message and processes its contents, which cannot include executable code.

Besides requests for assistance (e-Call), examples of V2X services commonly cited are:

- (a) Contributions to and awareness of traffic data/conditions
- (b) Notification of emergency conditions
- (c) Infotainment and mobile access to office/work environments.

3.4.1 Service (a)

Periodic beaconing is not needed. Estimates (statistical data) suffice for route planning and mobility management purposes. The following variation of crowdsourcing is proposed. The NC subsystem of every member of a cohort (Γ) runs an algorithm Alg serving to tell whether it must broadcast a pseudonymous V2X message carrying the identifier of the arterial where Γ circulates, Γ 's length, and the current GNSS space coordinates of Γ 's head and Γ 's tail, all of them computed out of cohort topology $TP(\Gamma)$ —see Subsection 6.3.5. When Alg outputs “your turn”, a certificate is retrieved from the (NC) HSM and sent along with the signed V2X message. Receivers learn about the existence of a cohort of n NGVs, its span (asphalt occupancy) and boundaries of geolocations occupied by non-identifiable NGVs.

It is not necessary to refresh statistical traffic data at high frequencies. Pseudo periods in the order of 20 seconds should suffice. Besides preserving privacy, this scheme would induce light channel loads: exactly 1 broadcast (deterministic Alg) or close to 1 broadcast (probabilistic Alg) per cohort per algorithmic round. Rates for a cohort member would thus be in the order of $3/n$ broadcasts per minute, rather than F broadcasts per second with beaconing frequency F . Assuming daily journeys in the order of 80 miles, i.e. approx. 2 hours of driving, and average $n \approx 18$, approximately 18 certificates would be extracted from a HSM every day. Like certificates in a PSM, certificates consumed for NC messaging can be safely reutilized on a weekly basis.

3.4.2 Service (b)

Emergency conditions can be handled safely without resorting to V2X messaging. Let us examine two classical examples used for “justifying” reliance on V2X messaging with regard to safety.

When an accident occurs, V2X messages shall be broadcast for announcing “accident at geolocation G ”. It is customary to consider that such messages are C messages, which amounts to making no distinction between “collision prevention” and “a collision has occurred”. Post-crash messages are NC messages comparable to NC messages that carry “congestion ahead” (for example). Aimed at mobility management and optimized journey planning, these NC functions shall not be equated with C functions. If geolocation G is sufficiently distant, there is enough time for deciding what to do. A NGV moving at 108 km/h has 10 seconds for making a decision when G is 300 m away (no risks of collisions). Otherwise, we are in the presence of a C event, which we address explicitly as a hard-braking scenario (see Section 8). Our terminology is reasonably precise: a V2X message that carries “accident” data is an *a posteriori* declaration of failure (the safety property has been violated), contrary to N2N or C2C messages aimed at preventing accidents *a priori*.

Consider now bidirectional roads without a central separation, and avoidance of head-on collisions due to overtaking maneuvers. Overtaking is forbidden in the absence of sufficient visibility (hilly/mountainous roads), a rule enforced by NGVs onboard systems. Poor visibility may also be due to weather conditions (heavy fog for example), in which case NGVs would move at reduced velocities, flashing warning lights turned on if necessary, overtaking prohibited. In good visibility conditions, overtaking shall be allowed only if there is no approaching vehicle (say V) in an oncoming lane. Overtaking by follower F implies the presence of at least one leader L ahead, and free space between L and V . When L detects oncoming V , L sends F a N2N message with C code “do not pass”. This is similar to the “stay in lane” scenario depicted in Subsection 5.2.1. Anyhow, F is not blind. Its onboard sensors can detect oncoming V and abort an attempted overtaking maneuver.

3.4.3 Service (c)

Passengers can have access to those services via their smartphones. One may question the usefulness of duplicating such *non-critical* services. This question is especially appropriate with shared/rented vehicles: Why would a passenger prefer to use a smartphone-on-wheels unfamiliar to her/him rather than a personal smartphone customized to her/his needs (profiles, passwords, browsers and preset navigation options, etc.)? Drawbacks are obvious (for example, traces of activities can be recorded). Connectivity features suffice, such as, e.g., powering and smartphone screen displayed on a dashboard. This question arises as well with privately owned vehicles. Moreover, the automotive industry cannot keep pace with technological advances that are more or less continuously harnessed by the smartphone industry, processing hardware/firmware in particular. Despite updated applications downloads, smartphone-on-wheels may look outdated more or less rapidly. In any case, given that a C subsystem is isolated from a NC subsystem, *non-critical* services commonly available on smartphones can be offered to passengers of NGVs which can operate in cyber stealth mode.

3.5. The cyber stealth mode

Contrary to CAVs, periodic beaconing is neither utilized nor available in NGVs. NGVs shall not behave as smartphones-on-wheels (CAVs with V2X functionalities). As amply demonstrated, smartphones-on-wheels are susceptible to remote and local cyberattacks. They may kill. A NGV cannot be detected from afar if its NC subsystem is inactive (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:

- Infrequent *outgoing* messages (service (a)),
- Extremely rare *outgoing* messages (e-Call, halted vehicle, etc.),
- *Incoming* messages that carry infotainment data (radio, movies, traffic data, news, etc.).

To the exception of outgoing core V2X communications, a NC subsystem is totally mute (thus no remote tracking or/and eavesdropping), but it would not be “deaf”. Thanks to partitioned NGVs onboard system architectures, infotainment has no access to CC subsystems (see Fig. 1), thus no control over physical motions. A NGV with only core V2X communications enabled is said to operate in cyber stealth mode. Other outgoing V2X communications such as e.g., e-working and e-shopping are enabled only when the cyber stealth mode is explicitly deactivated by passengers. Passengers aware of potential cyber threats shall have access to the cyber stealth mode via an on/off option (see Subsection 9.2).

4. Cells

Consider dense traffic conditions and a set of NGVs spanning 3 adjacent lanes in stationary conditions, as shown in Fig. 2. Thanks to sensors (cameras in particular), a NGV is able to see surrounding neighbors within 360° line-of-sight (LOS). Imagine that some NGV intends to undertake a risk-prone maneuver (within its lane or across adjacent lanes). This NGV can (1) exchange (send, receive) messages with its two nearest

longitudinal predecessors and successors in the same lane and with its nearest lateral neighbors, (2) strike *explicit agreements* with some or all of them as regards appropriate modifications of trajectories and velocities.

Let us rephrase the above, considering human-driven vehicles. A driver is able to see surrounding vehicles within 360° LOS. Imagine that a human driver intends to undertake a risk-prone maneuver. Thanks to turn signals, hand gesture, eye contact and human cognition, a driver can communicate with nearest neighbors and strike *implicit agreements* with some or all of them as regards appropriate modifications of trajectories and velocities. (Think about zipper merging, which humans handle reasonably well.)

By definition, a NGV may hit or may be hit only by nearest NGVs. Therefore, if a NGV has means of checking and influencing as desired the behaviors of its nearest neighbors, with reciprocal feedback, risk-prone maneuvers cannot result in collisions. The vehicular cell concept originates from this first observation.

To some readers, this may sound familiar. Indeed, in dense swarms and flocks of fish or birds, collisions never occur, except in the case of predator attacks. Predator attacks in birds/fish flocks are brick wall events in AVNs. Self-organizing networked flocking systems studied in automation control are comprised of locally interacting agents equipped with dedicated sensing, communication and computing/cognitive capabilities (e.g., linear/non-linear dynamics, AI-based algorithms) which serve to organize the collective motion of agents that align their behaviors with their neighbors, despite noisy information relative to dynamics [43-46].

With next-gen vehicles, can we expect “doing better” than with human-driven vehicles if we are not able to do “as well” in the first place? Then, if we have found smart solutions for doing “as well”, we should be able to show that we can “do better”.

Flocks and AVNs share important features. Fish or birds communicate via sounds and optics. NGVs communicate via radio and optics. Structures, sense of direction, single or multiple destinations are other examples of common features. They differ on degrees of freedom, durations (an AVN may be short-lived) and heterogeneity (not all members of an AVN look alike). According to most recent results in bird flocking networks, the number of neighbors monitored by a bird in a flock is in the order of 7. Collective collision-free motion is achieved by transitivity and sense of direction.

Let us pursue the analogy. Birds or fish spontaneously form flocking networks without having to change their externally visible interfaces. That is not the case for humans. In fact, NGVs with appropriate exposed interfaces would be “clothes” that humans have to put on in order to move safely throughout 2D physical space. In forthcoming publications, we explore these “appropriate externally visible interfaces”. The vehicular cell concept originates from this second observation.

Coordination schemes at the core of fish/bird flocks are not very well understood yet. Simulations and AI techniques are used in naturalistic sciences to gain better insight into this area. Conversely, NGVs and AVNs are human-made constructs. Therefore, coordination schemes that govern vehicular flocks are known (rather than inferred out of learning), since these schemes must be designed and proven correct prior to deployment.




Fig. 2. The cell construct illustrated

A cell with vehicle Y as nucleus is shown in Fig. 2. Depending on sizes of NGVs, vehicular density, relative alignments, and velocities, a cell may comprise between 1 (a cell nucleus is an isolated vehicle) and 13 members (nucleus, 2x2 longitudinal neighbors, 2x4 lateral neighbors). Cell(Y) illustrates a quasi-maximal cell: inter-NGV gaps equal to smallest gaps, NGVs of identical size, 3 lateral neighbors on the right, and 4 lateral neighbors on the left. Coincidentally, 7 (the average number of observed neighbors in a bird flock) is the midpoint of interval $\{1, 13\}$.

Very short range intra-cell N2N communications are *directional*. They rest on optics and radio technologies. Longitudinal and lateral N2N communications are allocated different radio channels. Longitudinal (resp., lateral) N2N communications have ranges of 80 m (resp., 10 m) approximately—see Subsection 8.6 for lower bounds of radio ranges. One could also consider very short-range *omnidirectional* N2N communications, knowing the drawbacks: poor spatial reusability and higher channel access delays.

N2N messages do not carry GNSS space coordinates (such data are useless for vehicles in mutual LOS). Rather, they carry codes of C maneuvers, such as “lane change requested”, “lane change request rejected”, “lane change request granted”, “lane merging ahead”, “set velocity to ...”, “clear lane” (emergency vehicle), “emergency braking”. Standard-making bodies will define such codes.

Cell membership varies with time. Cell members must be able to designate each other unambiguously. Consider Fig. 2. How may Y “tell” that a specific message is aimed at “this” particular left-side neighbor, when a name for “this” is unknown? Besides being unambiguous, “this” must not be reversible, i.e. no personal data should be inferred from “this” (the *PPN* problem in Subsection 2.3.1).

General planar constructs are obtained by combining cells as desired. In Section 5, we focus on the cohort construct, which is a linear formation of cell nuclei that follow each other in the same lane, as shown in Fig. 2. In upcoming publications, we examine lateral C scenarios and the merits of vehicular flocks (formations of adjacent cohorts) for collision-free entrance and crossing of unsignalized intersections and roundabouts.

5. Cohorts

CMX functionalities proper to cohorts are now detailed. The cohort construct was introduced in [47] and refined in [48]. A cohort is a formalization of platoons and strings, where membership is not planned in advance (contrary to platoons) and where members assign themselves consecutive ranks. The original motivation which led to the cohort concept was to demonstrate the possibility of eliminating rear-end crashes occurring at high relative velocities, while keeping very small inter-vehicular gaps. This implies proving theoretical absolute safety (no fatalities, no severe injuries) under assumptions of high coverage.

We can nowadays witness long vehicular strings on the roads. It is necessary to set upper bounds n^* on the number of NGVs that may follow each other in a spontaneous cohort, n^* valued for every setting by authorities in charge of road safety. This is needed for optimizing traffic fluidity. In urban settings, delays experienced for entering and crossing an unsignalized intersection or roundabout must be small, meeting a fairness constraint (no arterial is favored a priori), which implies short cohorts, thus small bounds n^* . This is also needed on highways and major roads, where bounds n^* shall meet antagonistic requirements. They should be high for optimizing efficiency, given that inter-cohort gaps are larger than intra-cohort gaps at medium and high velocities. Conversely, bounds n^* must be small in order to minimize disturbances due to lane changes. Indeed, lateral insertions *between cohorts* shall be preferred to lateral insertions *between cohort members* (cut-in maneuvers) whenever possible.

Consider the case where multiple lane changes shall be performed consecutively by a NGV for reaching an exit ramp. With cohorts of small size, that would be feasible rapidly most of the time. One could have $n^* \approx 12$ in urban settings, and $n^* \approx 25$ for highways. Choosing some n^* for cohorts moving at low velocities v (cities, congested highways, etc.) has a minor impact on efficiency, since inter-cohort gaps are not much higher than intra-cohort gaps when v is small. One may also anticipate the possibility of adjusting n^* according to local parameters such as, for example, the cardinality of an intersection (n^* defined as an inverse function of the number of arterials).

5.1. Intra-cohort communications

N2N messages that carry C data are generated by CC subsystems upon the occurrence of aperiodic C events which may originate from within a cohort or from the outside. The range-1 *send* primitive serves to trigger heartbeats and N2N management messages. The range-2 *Send* and the range-2 *Receive* primitives introduced in Subsection 3.2, detailed in Subsection 5.1.3, serve to perform local N2N messaging. The CWD (cohort-wide dissemination) and the CWA (cohort-wide agreement) primitives, detailed in Subsection 5.2.1 and 5.2.2 respectively, serve to perform upstream and downstream cohort-wide coordination.

5.1.1 N2N messages and acknowledgments

Any 2 contiguous members exchange N2N messages (one way) and acknowledgments (piggybacked on N2N messages sent the other way) on a bi-directional N2N dual link, based on radio and VLC technologies. N2N messaging is handled via the sliding window scheme found in the Internet TCP protocol.

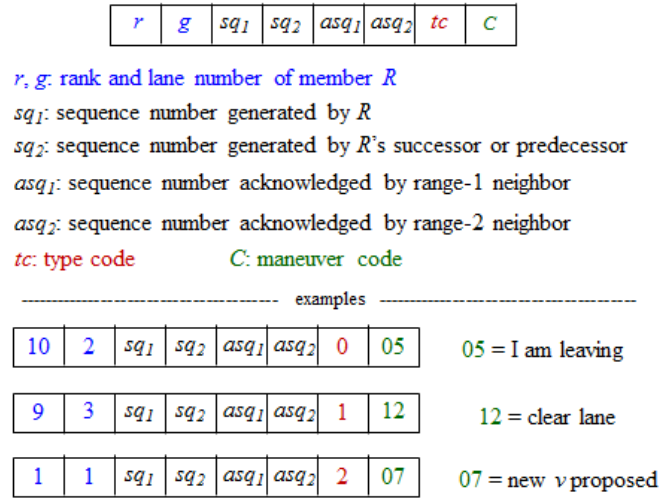


Fig. 3. N2N message header format

Thanks to consecutive sequence numbers assigned to messages sent in a given direction, N2N messages can be acknowledged unambiguously. Waiting for an acknowledgement is non-blocking (while waiting for an acknowledgement, subsequent Send or send operations are possible). When there are no messages to be sent on a N2N link, neighbors exchange heartbeats. No GNSS space coordinates appear in heartbeats.

N2N messages carry C data (codes of risk-prone maneuvers or notifications). In addition to examples given in Section 3, one would have: “no lane change” (fast approaching vehicle on adjacent lane), “current cohort topology $TP(.)$ ”. In addition to its rank r , $r \in [1, n]$, a cohort member knows the number of its lane (g). The header of a N2N message comprises 8 fields (the CRC/FEC field is ignored here), 1 byte long each, as illustrated in Fig. 3. The source field is a pair $\{r, g\}$. Owing to the range-2 Send primitive and N2N message types, destination names are implicit. An integer denoted sq_1 (resp., sq_2) appears in the 3rd (resp., 4th) field. These sequence numbers are computed by a Sender and by a relaying neighbor, respectively. An integer denoted asq_1 (resp., asq_2) appears in the 5th (resp., 6th) field. They serve to acknowledge messages received from range-1 and range-2 neighbors. Type codes denoted tc appear in the 7th field, with $tc = 0$ meaning range-2 Send, $tc = 1$ meaning “cwd” (dissemination), $tc = 2$ meaning “cwa” (agreement), $tc = 3$ meaning range-1 send (“heartbeat” or “management”), $tc = 4$ meaning “empty” message. C codes are given in the 8th field.

Consider 3 neighbors P , Q and R (rank r), P predecessor of Q and R successor of Q . Let us examine an upstream Send by R , assuming no failures. R Sends Q and P a message $m(r)$ generated spontaneously. In $m(r)$, field $sq_1 = \{153, r\}$ and field $sq_2 = \{\text{nil}\}$. Upon reception of $m(r)$, Q relays $m(r)$ with fields $sq_1 = \{147, r-1\}$ and $sq_2 = \{153, r\}$. P receives $m(r)$ and sequence number $\{153, r\}$ twice.

At some later time, P Sends a downstream message $m(p)$ where $asq_1 = \{147, r-1\}$ and $asq_2 = \{153, r\}$. Upon reception, Q relays $m(p)$, with $asq_1 = \{153, r\}$ (range-1 downstream neighbor's sq_1) and $asq_2 = \{\text{nil}\}$. R learns that its message $m(r)$ has been correctly delivered to Q and P , sequence number $\{153, r\}$ acknowledged twice. If message $m(r)$ is lost, matching fields asq are set to $\{\text{nil}\}$ by P and Q . It may also happen that no downstream N2N message is received by R in at most δ time units after $m(r)$ has been issued. (Upper bound δ is examined in Subsection 5.1.4.) R would then repeat $m(r)$, up to u^* times. Beyond threshold u^* , R triggers a cohort split (see Subsection 5.2.4).

Whenever needed, an “empty” N2N message (field $tc = 4$) is generated for the purpose of acknowledging a received N2N message. If a message m is put to wait due to some other N2N message waiting for transmission prior to the arrival of m , the matching field asq is set to $\{\text{wait}\}$. Highest rates of N2N message arrivals are such that only 1 wait may be experienced by a message in the course of cohort-wide message dissemination (see Subsection 5.2.3). It follows that at most 1 wait may be experienced by a message in the course of a round of

highest duration δ . Therefore, a N2N message put to wait when sent in a given direction within round x is acknowledged during round $x+1$, via a message circulating in the opposite direction.

Most N2N message bodies are very short. Longest N2N messages carry a list of proposals—see Subsection 5.2.2, or a cohort profile $\Pi(\cdot)$ —see Subsection 6.3.5. With n^* in the order of 25, their size is in the order of 100 bytes, header and CRC/FEC field included. Let θ stand for the largest transmission duration of a N2N message, which is smaller than 0.14 ms with 6 Mbits/s channels (DSRC-V2X or C-V2X technology). Adding 0.86 ms for framing and message processing time (relaying for example), we can set θ to 1 ms. (θ shall not be confused with δ , the worst-case channel access delay.)

The cohort construct is illustrated in Fig. 4. The gap kept by Y which follows X is denoted $s(v,y|x)$. $S(v,ch|ct)$ stands for the inter-cohort gap kept by a cohort head CH with CT , tail of a preceding cohort.




Fig. 4. The cohort construct illustrated

5.1.2 Naming, pairs $\{r,g\}$

Fundamental properties of the cohort construct result from a spontaneous ranking of members. An early example of ranking can be found in [49]—concepts of convoy, authenticated members and vehicle sequence numbers. An isolated vehicle assigns itself rank 1. A cohort member learns its rank r upon admission in a cohort. Cohort membership is updated every time a member leaves or joins in, longitudinally or laterally.

When a member leaves a cohort, its successors (if any, not the case if that member is a cohort tail) decrement their respective ranks. When a vehicle joins a cohort at rank r , members previously assigned ranks equal to or higher than r (if any, not the case if the new member is cohort tail) increment their respective ranks. A pair $\{r,g\}$ is a non-reversible identifier, built upon admission in a cohort, which admission is conditioned on verified authentication. According to Subsections 2.3.1 and 2.3.3, the *PPN* problem must have been solved beforehand. That is the case with *LtJoin* and *LgJoin* operations (see Section 7 for the latter).

Eavesdropping on N2N communications is of no interest since no personal data can be inferred from pairs $\{r,g\}$, sequence numbers, types or codes. That is why no additional tracking/privacy risks are incurred by a NGV due to N2N communications, other than those risks existing with human-driven vehicles and cameras (physical tracking, reading of plates, and so on).

Despite the fact that names $\{r,g\}$ are not unique in lane g , designation of a specific vehicle with rank r is unambiguous. Consider the case where X and follower Y are assigned the same name $\{1,g\}$, i.e. when X and Y are isolated vehicles (cohort heads) separated by inter-cohort gap $S(v,\cdot)$. Any other vehicle V which can observe X and Y must be in lane g or in a lane adjacent to g . If V is in lane g , the only case of interest arises when V is positioned between X and Y . Therefore, V cannot mistake X for Y , or vice-versa. (If V is ahead of X or behind Y , then V cannot see X and Y .) If V is in a lane adjacent to g , V cannot mistake X for Y due to the size of inter-cohort gap $S(v,\cdot)$. When $r \neq 1$, any two NGVs assigned the same name $\{r,g\}$ belong to different cohorts. No confusion is possible.

Designations of vehicles shall be distinct from identifiers that appear in N2N messages. For the sake of simplification, we keep using R as the name of a NGV that activates a communication primitive, keeping in mind that $\{r, g\}$, rather than R appears in the source field of N2N messages or heartbeats issued by R . A NGV must provide a valid certificate to be admitted in a cohort. A misbehaving member is promptly removed, thanks to predicate checking and PSMs (see Subsection 3.1.2). It follows that cohort members can trust each other. They are thus relieved from the need to be re-authenticated every time they send a N2N message.

5.1.3 The *Send* and the *Receive* primitives

Range-2 communication primitives available to cohort members are as follows:

- *Send* (R, m): member R ranked r^{th} sends message m to neighbors with ranks $r+1$ and $r+2$ (downstream *Send*) or to neighbors with ranks $r-1$ and $r-2$ (upstream *Send*), if any.

Send (R, m) is a longitudinal multicast endowed with the atomic range-2 reachability property. Atomicity is provided by radio technologies such as DSRC-V2X or C-V2V PC5 mode 4 [50]. A neighbor receives or does not receive m correctly. Byzantine behavior [34] of a Sender does not need to be considered. In the opposite case (e.g., VLC, mmWave 5G [51]), a range-2 span would be instantiated as two consecutive range-1 *send* operations, with relaying. Man-in-the-middle cyberattacks by a relaying member (e.g., forgery or suppression of m) shall then be addressed (not within the scope of this report).

- *Receive* (R, m): triggered by the arrival of message m , serves to decide whether m should be processed (for acceptance checking); member R ranked r^{th} processes messages received from downstream members with ranks $r+1$, $r+2$, and from upstream members with ranks $r-1$, $r-2$ exclusively. Other messages are discarded.

The range-2 *Receive* primitive enforces the basic distinction between physical reachability of an antenna and the actual processing of the contents of an incoming message. This distinction is essential for thwarting cyberattacks (see Subsection 5.4). Since N2N message headers and bodies carry life/safety critical data, contents of a received message must be trustworthy.

Consider again the 3 consecutive neighbors P , Q and R examined in Subsection 5.1.1, where Q relays message $m(r)$, sequence number $\{153, r\}$, issued by R . In the absence of permanent losses, P should receive $m(r)$ tagged with sequence number $\{153, r\}$ twice, which suffices to detect forgery or suppression of $m(r)$ by Q . In such a case, neither R nor P is malicious (owing to the 1-out-of-3 assumption). These cyberattacks are defeated with the following *AC* condition.

Acceptance Condition (AC)

Upon reception, a N2N message m is accepted for further processing if and only if:

- Body fields of m Received from range-1 and range-2 neighbors are identical, delivered within a stipulated bounded time interval (derived from the MAC protocol in use).
- Field asq_1 in m and field asq_2 in m Received from range-1 and range-2 neighbors, respectively, are equal.

When *AC* is not fulfilled, m is discarded without being acknowledged, which may lead to a retransmission of m (up to u^* times).

When a N2N message is generated spontaneously by some (unique) member, say U , U 's neighbors cannot apply *AC*. This is a source of vulnerability in the presence of malicious members (U could generate bogus N2N messages). Injection of bogus data is examined in forthcoming publications.

5.1.4 MAC protocols

Few MAC protocols ensure the existence of finite upper bounds δ for channel access delays in worst-case contention conditions (a property not within the grasp of existing IEEE/ETSI standards for V2X communications). Such timeliness proofs, which are mandatory for safety, imply knowing the worst-case number of contenders at any time. These numbers remain unknown in vehicular networks without "structures". Given our goals, only deterministic MAC protocols are eligible.

TDMA protocols designed for DSRC channels can be found in [52]. SWIFT is a deterministic collision-free TDMA protocol [53], which can be implemented on DSRC channels and on cellular radio channels. SWIFT is more efficient (smaller δ) in the case of directional antennas/communications. Channel slot duration θ has been introduced in Subsection 5.1.1. Let integer h stand for the number of contiguous cohort members within radio range, sender included. Integer h is a measurement of spatial reuse. Members h ranks away from each other can send a message in the same direction successfully at the same time (no interferences, no collisions).

Members are serviced downstream and upstream, alternatively, every $h\theta$, according to increasing ranks, then according to decreasing ranks. With most message-passing protocols, acknowledgements are handled as ordinary messages, inducing additional contention, thus augmenting worst-case time bounds. With MAC protocols designed for directional communications, acknowledgements can be piggybacked on N2N messages, entailing a very modest overhead.

Let t stand for the time at which a N2N message is submitted for transmission by member R , rank r , and t_{swift} stand for the start time of the earliest channel slot allocated to R . Assuming no queuing, smallest channel access delay $\varepsilon \approx 0$ is incurred when $t = t_{\text{swift}} - \varepsilon$. Highest channel access delay δ is incurred when $t = t_{\text{swift}} + \varepsilon$. A SWIFT round consists of h consecutive channel slots in each direction. Therefore:

$$\delta = 2h\theta \quad (\delta \text{ does not depend on } n).$$

With the channel bandwidth offered by Wi-Fi or cellular radio, θ is of marginal importance. Bound δ depends essentially on the multiplication factor ($2h$). This holds for any MAC protocol. One does not achieve a small δ simply by assuming high bandwidth channels.

Ideally, $h = 3$ suffices for instantiating a range-2 Send primitive, whatever the locations of send and receive antennas in vehicles. ($h = 2$ would suffice if directional send and receive antennas in a vehicle do not interfere.) Since range-controlled antennas may span distances larger than desired, it is necessary to assume a worst-case $h > 3$ when computing performance figures. For the example in Subsection 5.2.1, we consider $h = 5$. A postulated $h > 3$ may nevertheless be violated at run-time. Channel collisions would then occur. May safety be compromised? This question is addressed in [53]. The answer is no, owing to the cohort split scheme. Collisions materialize as message or acknowledgement losses. When more than u^* consecutive losses are experienced on a N2N link, a cohort split is undertaken. Members at both ends of a lossy N2N link end up being separated by safe inter-cohort gap $S(v, ch/ct)$. Therefore, safety holds true, whichever values are assigned to h and u^* . In other words, the assumption coverage issue does not arise here.

SWIFT is intrinsically immune to various kinds of cyberattacks. For example, a malicious member may not obey the TDMA slot allocation scheme and “steal” a channel slot owned by another member, de facto impersonating that member. With SWIFT, this masquerading attack is detected instantly. When the number of such attacks exceeds some threshold, a malicious vehicle is halted.

Channel slot timing is provided by GNSS devices. An onboard back-up clock keeps itself aligned with GNSS timing. SWIFT utilizes this clock in case of GNSS outages (losses of GNSS signals caused by poor reception conditions or by deliberate jamming). Affordable clocks have intrinsic drifts small enough for maintaining accurate slot timing in the presence of signal losses lasting several minutes. It follows that the coverage of assuming channel slot timing aligned with GNSS timing is very close to 1. Recall that VLC is also utilized for performing message passing on a N2N link. Directional VLC does not need GNSS timing.

With very short range directional communications, MAC protocols other than TDMA can be considered for intra-cohort messaging, such as synchronous CDMA and deterministic CSMA. For example, time-bounded C-V2V communications are feasible over cellular radio [54] with MAC protocols based on collision detection and deterministic collision resolution similar to “deterministic Ethernet” protocols [55]. In order to cope with radio interferences due to simultaneous transmissions by NGVs in adjacent lanes, different radio channels or different access times can be allocated on a lane number basis. Note however that such interferences are significantly minimized in the case of very short range directional communications of very low electromagnetic energy level. Some deterministic MAC protocols need a single radio channel (to appear). Lane-dependent channel allocation is preferable with cellular radio. Resilience against privacy threats and/or cyberattacks (a SWIFT property) is one of the most important challenges to be addressed with the design of novel MAC protocols.

Non-contiguous members can communicate as desired via N2N message relaying (successive activations of the Send primitive). Two important cases for relaying stem from the need to achieve cohort-wide coordination. That is the purpose of the CWD and the CWA algorithms.

5.2. Cohort-wide coordination

Owing to its complexity, the general problem of reliable, time-bounded and trustable inter-vehicular coordination in vehicular networks has not been solved yet. One major obstacle, whose identification has led to the CMX framework, is the lack of constructs endowed with intrinsic properties that lay the ground for solutions to the SPEC problem, in the same way as lemmas lay the ground for theorems. Thanks to the cohort construct, it is possible to devise algorithms—presented below—that achieve reliable, fast, time-bounded and trustable inter-vehicular coordination. One solves the SE problem by compounding eligible MAC protocols, CWD and CWA algorithms with cyberphysical levels (see Section 6).

5.2.1 Cohort-wide dissemination and common knowledge—the CWD primitive

When the originator of a message m typed “cwd” ($tc = 1$) triggers the CWD primitive, this activates two Send (\cdot, m), one downstream and another one upstream. Destination names are implicit (cohort head and tail). A neighbor that is delivered a message m typed “cwd” relays m in the direction opposite to its arrival. Cohort head and cohort tail do not relay. The case $n = 2$ can be ignored, since direct message exchanges suffice (there is no need for a CWD algorithm). Time-bounded cohort-wide message dissemination matters since members must share common knowledge relative to various time-dependent variables and C data. Updates of cohort membership and propagation of a C event generated from within or from outside a cohort are typical examples.

Numerous protocols have been proposed for message dissemination in VANETs and strings/platoons, most of them based on plain broadcasts of V2X messages (BSMs/CAMs), the case with Cooperative Adaptive Cruise Control (CACC), where V2X messages may be relayed according to various policies (e.g., leader-predecessor following, two-vehicle look-ahead, bidirectional control, and one-vehicle look-ahead topology). Prima facie, it may seem preferable to rest on dissemination algorithms based on short/medium range V2X broadcasts and relaying rather than on very short range N2N inter-member relaying. That is not necessarily the case, as shown in [56-59].

For example, average packet delivery ratios (PDRs) measured during tests or obtained by simulation in different settings (urban, rural, intersection, etc.) are in the order of 50% at distances approximately equal to 100 m [57]. Slightly better figures for distances roughly equivalent to V2X radio ranges are reported elsewhere. In any case, the problem is that packet loss ratios are equal to 1-PDR, i.e. orders of magnitude higher than the 10^{-x} per hour figure ($x > 3$) to be met by life/safety critical functionalities (see Subsection 2.3.4). Local dynamic maps (LDMs) are the essential scheme for safety in the V2X framework, which LDMs are built out of periodic broadcasts of BSMs/CAMs that are undelivered at exorbitant rates. The conclusion is obvious.

Observe further that a PDR measures a *throughput*. For safety, *delays* matter, notably *packet inter-reception (PIR) time*, i.e. intervals of time elapsed between two successful beacon deliveries [60]. The likelihood of inconsistent and divergent LDMs, thus collision risks, increases with large and unbounded PIRs.

Recall requirement R2 in Subsection 1.3. Imagine that radars, lidars or cameras would fail at rates comparable to 1-PDR. Would that be deemed acceptable? Clearly, periodic broadcasting of V2X messages impaired by failure rates orders of magnitude higher than failure rates of onboard robotics cannot improve safety obtained by onboard robotics. Furthermore, dissemination delays in the presence of malicious members that suppress V2X messages to be relayed may not have finite or acceptable bounds with algorithms based on short/medium range broadcasts.

To be eligible, a dissemination algorithm shall meet the following four requirements:

- Reliability: every relaying of a message must be acknowledged,
- Time bounded termination: dissemination terminates in finite time bounds under worst-case conditions regarding contention, failures, message suppressions and falsifications, under clearly stated assumptions,
- Symmetry: worst-case dissemination time bounds are identical for downstream and upstream executions,
- Promptness: worst-case dissemination time bounds in strings or cohorts are sublinear in $n\delta$.

In [53], one can find a presentation of a SWIFT-based CWD algorithm. Cohort-wide dissemination time bound $\Delta(n, f)$ is given for worst-case contention and message loss conditions, where integer f stands for the number of losses (messages or acknowledgments) experienced in the course of CWD. Owing to the range-2 Send primitive and the 1-out-of-3 assumption, the first three above-mentioned requirements are met. See Subsection 6.1.4 for the promptness requirement. Bound $\Delta(n, f)$ is reached when CWD is initiated by a cohort head or a tail. In the absence of queuing, we have:

$$\Delta(n, f) \leq \delta \left\{ 1 + f + \lceil (n-1)/h \rceil \right\} \quad (\text{Eq. 1})$$

Let us consider $h = 5$, $n = 20$ and $f = 5$ (a pessimistic value, to account for very poor radio conditions). We have seen that $\theta = 1$ ms is a fair estimate. Thus $\delta = 10$ ms and $\Delta(20, 5) \leq 100$ ms.

Fig. 5 shows an example of the usefulness of SWIFT-based CWD. V moving at 160 km/h is about to overtake a slow cohort Γ (40 km/h, $n = 20$). No cohort member shall undertake a lane change before V passes.

In case V would broadcast a V2X “warning” message, not all members may receive that message or that message may incur too high a channel access delay to be of any usefulness. V moves too fast for being detected in time by backward-looking sensors of members other than cohort tail CT and, possibly, CT 's predecessor. CT detects V when V is a few meters away and triggers upstream CWD (“stay in lane” N2N message).

In 100 ms, V moves by 4.44 m while CT and Γ move by 1.11 m. Therefore, every cohort member (CH included) has received the N2N warning message by the time V has gained 3.33 m over CT . Collision prevention is achieved. A detailed presentation of the case where multiple vehicles such as V are about to overtake Γ as well as conditions under which Γ 's members may undertake lane changes anew is not within the scope of this report. Our objective here is to show that strictly time-bounded CWD is feasible in the presence of message loss densities way higher than those commonly assumed in the published literature (a 50 Hz message loss rate in our case).

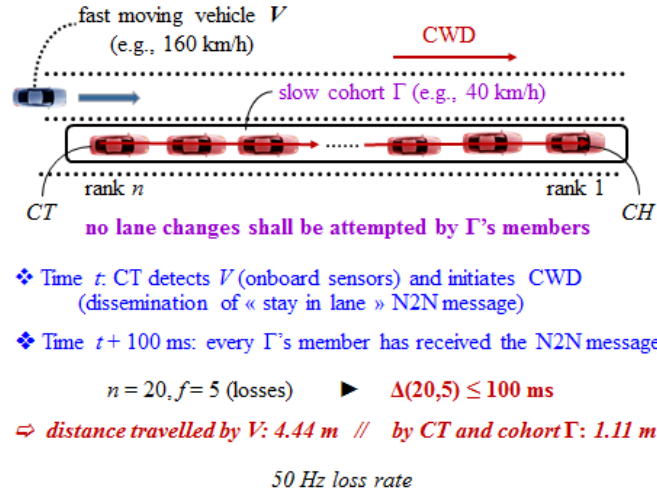


Fig. 5. An example of how to prevent lateral collisions with CWD

CWD matters in numerous scenarios where C events are generated by emergency vehicles endowed with specific types $vt(\cdot)$ and related special rights. Immediate lane clearing is a typical example. Consider a police car or an ambulance or a fire truck, denoted Z , that catches up with cohort Γ in the same lane, Γ 's CT denoted Q . Along with its C2C message which carries C code "immediate lane clearing", Z provides Q with a certificate that contains profile $vp(Z)$. Z 's type $vt(Z)$ is found in $vp(Z)$. When learning that Z is an emergency vehicle, Q triggers CWD upstream. Every Γ 's member onboard CR subsystem is posted the "immediate lane clearing" command by its companion CC subsystem. Another interesting case arises when Z is a police car that wants to stop an isolated vehicle Q or a complete cohort Γ . Z 's C2C message would then carry the C code for "stop within the next ... meters". There are two possibilities. Z may want to join Γ first (see Subsection 7.1), and then trigger CWD as a Γ 's member, or proceed as above (tell Q to trigger CWD).

Cyberphysical operations prompted by emergency vehicles are monitored by PSMs and their associated predicates (see Subsection 3.1.2). A refusal to obey an emergency command is a predicate violation, leading to halting an unruly vehicle. In both cases, Z 's intention is fulfilled (lane is cleared; a vehicle that does not stop is halted). As regards non-emergency vehicles, refusals to obey a "clear lane" decision reached by consensus on the occasion of a LgJoin (see Section 7) are recorded in EDRs.

Rather than some random ordering, possibly with unfair distributions of successful channel accesses, SWIFT induces orderings of "winners" (members that succeed in accessing the radio channel) that are congruent to member rankings (upstream or downstream). Furthermore, within a SWIFT round, contiguous members can send a N2N message every θ . SWIFT is a MAC protocol *and* an efficient symmetrical CWD algorithm. This might be an indication that with top-down design approaches, one finds protocols or algorithms that solve multiple problems at once.

With SWIFT-based CWD, every member knows in bounded time whether or not a N2N message has made it throughout a cohort, thus instantiating the equivalent of a global end-to-end acknowledgement. However, N2N acknowledgments and retransmissions do not suffice for achieving cohort-wide consensus, which has to do with the *semantics* of N2N message contents. To put it simply, CWD does not solve consensus, CWA does.

5.2.2 Cohort-wide agreement and conflicting concurrency—the CWA primitive

The originator of a message m typed "cwa" ($tc = 2$) triggers the CWA primitive, which activates two CWD primitives, one downstream and another upstream. Exact agreement, a.k.a. consensus [12], is needed for coping with conflicting proposals disseminated in N2N messages generated at about the same time by distinct cohort

members in response to C event arrivals. As stressed in Subsection 2.2.2, safety issues related to general conflicting concurrent scenarios that involve multiple vehicles have not received due attention so far. Let us give a few examples.

A message typed “cwa” carries “reduce velocity to (.)” and another one carries “clear lane”. A message typed “cwa” carries “accept lateral Join for left-side neighbor X between ranks r and $r+1$ ”, and another one carries “accept lateral Join for right-side neighbor Y between ranks r and $r+1$ ”. Multiple “cwa” messages carry lateral Join insertions requested in various ranks in a long cohort, only a subset of them being physically feasible due to the multiple quasi-simultaneous deceleration maneuvers that would ensue. A message typed “cwa” carries “accept lateral Join at rank r ” while at the same time members with ranks smaller than r hard brake. Safety requires that all members shall decide on a unique proposal. Otherwise, chaos would develop (first example), a lateral collision would occur (second example), multiple lateral or rear-end collisions would occur (third and fourth examples).

Apart from consensus, approximate agreement may be necessary or sufficient, the case with real numbers [61]. Safety mandates that all members shall agree on some function of the values proposed (e.g., average, midpoint). For example, when weather/environmental conditions change, new velocity v would be agreed upon via some approximate agreement algorithm AA , keeping discrepancies within bounded tolerance. AA would be run only when needed, given that onboard robotics handles mild velocity variations quite well.

Besides well-known impossibility results regarding consensus [19-20], tight conditions for possibility have been established, notably that unanimity or absolute majority are impossible unless the number of message losses does not exceed small bounds [62]. Thanks to the linear structure of cohorts and the cohort split scheme, various distributed consensus algorithms can be considered for CWA. Let us mention two examples.

An algorithm designed for a fully asynchronous system model augmented with unreliable failure detectors, where processes may crash, can be found in [63]. As shown in Subsection 5.2.4, message losses may lead to a cohort split, which amounts to “crashing” a correct NGV (a split separates a correct NGV from its cohort). Failure detectors serve to circumvent impossibility results.

Another algorithm (Eligo), designed for a synchronous system model, is presented in [64]. Eligo is a 2-phase SWIFT-based CWD algorithm. Every “cwa” message carrying a proposal is disseminated upstream and downstream, accumulating other proposals on its way, recorded in a list (1st phase). A cohort head CH or a cohort tail CT that receives a list of proposals collected one way disseminates the received list the other way via CWD (2nd phase, part 1). A member that receives both lists from its two opposite neighbors computes their union and the decision value, which value is then propagated via CWD in both directions (2nd phase, part 2). Under the 1-out-of-3 assumption, this computation shall be performed by 2 contiguous members. Worst-case upper bound $\Xi_{LgDA}(n,f)$ for reaching agreement with Eligo is given in [64]. One finds:

$$\Xi_{LgDA}(n,f) \leq \delta \{ 1 + f + 2 \lceil (n-1)/h \rceil \}.$$

In case a cohort split event occurs while some consensus algorithm is being run, either consensus is aborted (possibly resumed once the split is over) or consensus terminates prior to the triggering of the split maneuver, ditto for approximate agreement. Problems of conflicting concurrency cannot be solved without adopting a network-centric perspective, hence the relevance of cyberphysical constructs.

5.2.3 Worst-case time bounds, failure rates and loads

With upper bounds δ , $\Delta(n,f)$ and $\Xi_{LgDA}(n,f)$, the requirement embodied in **principle f** is met. However, according to **principle e** and **principle f**, one must express worst-case N2N message loads under which δ holds true, as well as worst-case N2N message loads and failure rates under which $\Delta(n,f)$ and $\Xi_{LgDA}(n,f)$ hold true. They depend on n^* and on f^* (the highest value of f). Parameter n^* is known and enforced (controlled cohort admission and cohort splits). In life/safety critical systems, the need for proving bound f^* shall not be ignored (f^* shall not be “guessed”).

This is another unsolved problem in open spontaneous systems in general, in the V2X framework in particular. It is possible to prove the existence of bound f^* in the case of cohorts. Owing to the 1-out-of-3 assumption, the highest number η of N2N lossy links is $\eta = \lceil (n-1)/3 \rceil$. Let u^* stand for the highest number of consecutive losses that may be experienced on a N2N link without leading to a cohort split. Trivially, $f^* = \eta u^*$. Integer u^* can be valued as desired (e.g., by standard-making bodies). Therefore, worst-case bounds $\Delta(n^*, f^*)$ and $\Xi_{LgDA}(n^*, f^*)$ are known.

Realistic values for f^* are much smaller than worst-case bound ηu^* . Assume $u^* = 3$. With $n = 20$, one finds $f^* = 21$. With the numerical values considered in Subsection 5.2.1, one would have $\Delta(20,21) \leq 260$ ms. A total

of 21 losses over an interval of 260 ms amounts to a message loss rate in the order of 80 Hz. Clearly, smaller rates should be assumed in practice, still meeting the extremely high coverage requirement.

$\Delta(n,f)$ as per Eq. 1 holds in the absence of queuing. Queuing may occur with high N2N message arrival rates, yielding worst-case bound $\Delta(n,f,q)$ when a N2N message is put to wait q times in the course of CWD. A detailed schedulability analysis is out of the scope of this report. However, observing that a latency of $2h\theta$ is experienced in the case of queuing, one can write:

$$\Delta(n,f,q) = \Delta(n,f) + 2hq\theta.$$

In Subsection 5.1.1, we have stated that only 1 wait may be experienced by a message in the course of cohort-wide message dissemination. This amounts to postulating rates of N2N message arrivals at most equal to $2/\{\Delta(n,f) + \delta\}$. With the same numerical values as above, one finds that highest sustainable loads are equal to $2/110$, that is approx.18 messages/s. This matches reality. Indeed, aperiodic event-triggered N2N messages have arrival rates vastly smaller than 10 Hz.

5.2.4 Episodic heartbeats and cohort splits

Heartbeats (N2N messages typed $tc = 3$) serve to meet *dependability requirements* (see Subsection 3.3). Stop failures of onboard systems and N2N link stop or omission failures (message losses) must be detected within bounded latency π . Unless failed, a member's onboard system must send at least 1 message every π to range-1 neighbors, across their common N2N link. In case no N2N message has been submitted for transmission during π , an onboard system generates an empty N2N message called a heartbeat.

Transmission of heartbeats ceases when (real) N2N messages are being exchanged. A heartbeat's body is empty. A heartbeat's header carries code C meaning "I am alive". Heartbeats are sent via the range-1 unicast *send* primitive. A member ranked r^{th} processes heartbeats received only from neighbors with ranks $r-1$ and $r+1$. Other heartbeats, which could be received from farther members, are discarded.

N2N messages or heartbeats generated by an onboard system may be incorrectly transmitted, i.e. lost, over a N2N link. An onboard system may fail. A N2N link is diagnosed failed when more than u^* consecutive losses (message or heartbeat) are experienced during π ($\pi > 2h\theta u^*$). Therefore, timers suffice for detecting permanent failures within a cohort.

Given that a cohort must be free from failed elements, a scheme called *cohort split* serves to meet *safety requirements*. Consider neighbors X (rank x) and Y following X , and assume that a X/Y link failure occurs, detected by X . X triggers CWD upstream for sending a N2N message that carries "cohort split at rank x , update cohort topology, $n = x$ ". The entire cohort is thus aware of the split maneuver. X is now tail of truncated cohort. CWD is also triggered downstream by Y , which decelerates until inter-cohort spacing $S_{\min}(v,.)$ is instantiated with X . Y becomes head of a new cohort (comprising its followers). Would the N2N link between X and its predecessor W and the X/Y link fail simultaneously, the same process takes place, with W replacing X in the above example.

Splits are caused by failures and malicious behaviors. A malicious behavior is detected by an onboard CC subsystem. The companion CR subsystem is instructed to halt the misbehaving vehicle, which involves a cohort split. Splits caused by failures can be made as infrequent as desired via an appropriate tuning of parameter u^* . A cohort is thus a linear AVN without permanent internal failures. This is fundamental, for three reasons. Firstly, the impossibility 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 worst-case bound f^* (see above). Thirdly, owing to that scheme, CWD and CWA algorithms have known worst-case termination time bounds despite failures

Noting that failures of optical links are not correlated with failures of radio links and that N2N message passing remains possible via VLC communications, an attentive reader may wonder why it is necessary to trigger a cohort split when a N2N radio link fails. First reason: redundancy is lost. Second reason: with VLC, there are no atomic range-2 N2N communications. The range-2 Send primitive must rest on range-1 relaying. In the presence of a malicious cohort member that suppresses a N2N message to be relayed, the acceptance condition AC (see Subsection 5.1.3) is inoperative. It is possible to thwart such internal cyberattacks. Anyhow, since cohorts shall not include malicious members, cohort splits are utilized.

5.3. Inter-cohort communications and coordination

Besides intra-cohort coordination, inter-cohort coordination is necessary, as illustrated in Fig. 6. Originators of LtJoin and LgJoin requests as well as responders shall demonstrate that they can trust each other.

This implies verified pseudonymous authentication. Pseudonym certificates [21] are part of the CMX framework. In the case of LtJoin or LgJoin operations, past the first contact round, a C2C request or response M is sent accompanied with a pair $\{S, C\}$. S is M 's signature computed out of the private key found in certificate C retrieved from a PSM. In case of incorrect authentication (S and C do not match), a receiver discards M . In the remainder of this report, pair $\{S, C\}$ is sometimes referred to as “pseudonym” and incorrect authentication referred to as “invalid pseudonym”. See Subsection 6.3.6 for Leave operations.

5.3.1 Lateral N2N and longitudinal C2C communication primitives

The coordination of cohorts that are within mutual LOS rests on the following primitives:

- LtSend: range-1 lateral N2N multicast or unicast of maximal ranges in the order of 10 m. A LtSend is triggered by a NGV that undertakes a LtJoin operation in order to insert itself between two contiguous neighbors located in an adjacent lane. LtJoin cyber rounds serve to strike necessary agreements (insertion requests may be granted or turned down).

- LgSend: range-2 and range-3 C2C longitudinal multicast of maximal ranges in the order of 120 m. A longitudinal cyberphysical LgJoin operation is triggered by a NGV or a cohort head (A in Fig. 6) that catches up with an isolated NGV or a cohort tail (cohort Γ in Fig. 6) in the same lane. These NGVs must agree on what to do: create a unique cohort (full merge), or merge partially, or do not merge.

A LgJoin operation triggers a number of LgSend primitives which serve to exchange C2C messages longitudinally in order to reach agreement. See Section 7 for a detailed presentation of LgJoin.

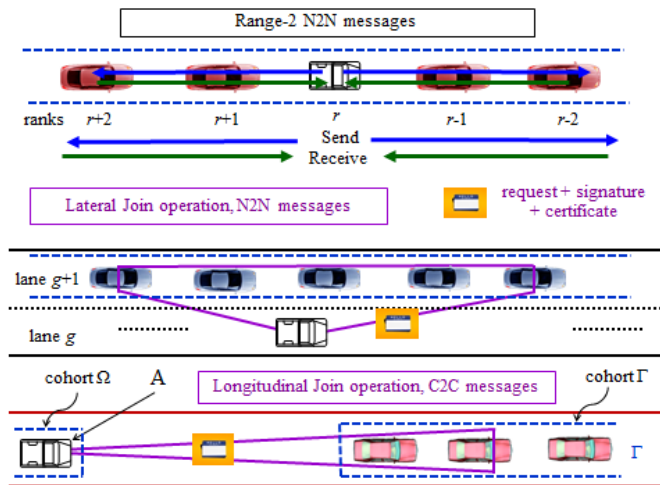


Fig. 6. Intra-cohort N2N communications and inter-cohort LtJoin and LgJoin operations

In the case of unsignalized intersections (UX) and roundabouts (UR), cohorts located in converging arterials may not see each other prior to being on the verge of entering (occlusions, foggy weather, etc.). They cannot “hear” each other either with range-2 directional communications. Cohorts must nevertheless reach consensus (which cohorts can enter safely, which ones have to wait). Time-bounded consensus algorithms rest on the SBcast primitive and semi-omnidirectional C2C communications of ranges in the order of 120 m (braking distances + UX or UR diameter) mentioned in Subsection 3.2. In forthcoming publications, we show how to solve the general UX problem (unsignalized intersections with any number of arterials, any number of lanes per arterial) with MAC protocols that instantiate the equivalent of virtual traffic lights.

5.3.2 Inter-cohort MAC protocols

Distinct radio channels are allocated to lateral N2N communications and to longitudinal C2C communications. Worst-case MAC-level access delays for these communications are necessarily smaller than bounds δ , since contention degrees (concurrent LtSend, concurrent LgSend) are significantly smaller than those experienced with intra-cohort N2N communications.

We have devised novel deterministic MAC protocols based on principles radically different from those underlying MAC protocols for intra-cohort communications (to appear). One of them is CSMA augmented with a deterministic time-bounded collision resolution scheme, which bears some resemblance with the deterministic CSMA-CD protocol described in [65], which became a French Navy standard implemented in real-time local area networks deployed in submarines, surface vessels, and at the European Ariane Launchpad in French Guiana (to name a few examples) in the late 1980s and in the 1990's.

5.3.3 Naming, triples $\{r,g,<R\psi>\}$

Data (contents) and metadata (headers) of inter-cohort messages shall preserve privacy. We have seen that no GNSS space coordinates appear in C2C messages. As regards metadata, one must distinguish designations of vehicles from identifiers. Let ψ stand for a pseudonym linked with a reusable certificate. For conciseness, we use triple $\{r,g,<R\psi>\}$ as the name of member R ranked r in a cohort on lane g that issues a C2C message accompanied with pseudonym $R\psi$. Pair $\{r,g\}$ appears in the message source field. Hereafter, we keep using R as a parameter of a LgSend primitive, keeping in mind that R stands for $\{r,g,<R\psi>\}$. Pseudonym ψ° stands for a condensed version (e.g., a hash) of a pseudonym ψ recently used in a LgJoin operation. Identifier $R\psi^\circ$ serves for a quick re-authentication of R , supplemented with secret sharing. Pseudonyms ψ° are of particular interest in LgJoin operations (round 4, see Subsection 7.1) and in case of temporary cohort splits.

5.4. On privacy, tracking and cyberattacks

Cohorts and fast reliable time-bounded longitudinal and lateral inter-vehicular coordination based on very short range communications are key elements for safety and efficiency in the CMX framework. These properties are established and quantified in the remainder of this report.

Prior to this, a few observations are in order regarding privacy and cybersecurity. Recall that in this report we do not address issues related to privacy threats or cyberattacks aimed at onboard sensors and actuators managed by a CR subsystem. Our focus is on privacy/tracking threats and cyberattacks against V2X functionalities (managed by CC subsystems), which could jeopardize safety, by creating hazards and collisions. These issues have been briefly addressed in [66], where the cyber stealth mode (see Subsection 3.5) was first envisioned. Let us expand on our early analyses.

In the CMX framework, privacy and cybersecurity issues of interest are those arising with NGVs that operate in cyber stealth mode. Otherwise, NGVs are equivalent to CAVs (smartphones-on-wheels), i.e. exposed to well-known threats—to the exception of threats proper to periodic beaconing, since periodic beaconing is not considered in the CMX framework. These threats which have been and are being studied by a number of experts are not within the scope of this report.

The electromagnetic level of very short range directional N2N and C2C communications is very low, compared to that of short/medium range omnidirectional V2X communications. Consequently, a NGV that operates in cyber stealth mode in the midst of multiple sources of high electromagnetic energy levels is hardly detectable by a remote adversary. Core V2X communications are the only exploitable attack angles.

5.4.1 Remote adversaries

In cyber stealth mode, given that outgoing NC V2X messages are aperiodic and extremely rare, tracking is practically unfeasible. Moreover, these messages carry data that are unlikable with personal information, thus of no interest to an eavesdropper. Privacy threats, hardly feasible, are inoffensive.

Remote cyberattacks are performed via short/medium range communications. These communications are handled by specific antennas managed by NC subsystems, in total isolation from antennas and CC subsystems in charge of safety (see Subsections 3.1.1 and 3.2). NGV antennas utilized for C communications listen to specific radio channels, over very short ranges in specific directions, I/O handlers implemented in CC subsystems. No malevolent message issued by a remote adversary can “contaminate” a CC subsystem—which would otherwise order a CR subsystem to engage in a harmful maneuver. Safety cannot be compromised by remote cyberattacks. Nevertheless, being consistent with the obligation of considering worst cases, let us imagine that a malevolent message M issued by remote adversary Z is correctly decoded by an antenna devoted to C communications. Issues to be addressed are thus related to local adversaries.

5.4.2 Local adversaries

Local adversaries are physically within LOS or almost within LOS of targeted NGVs. Life/safety critical N2N and C2C messages are event-triggered, thus aperiodic, making eavesdropping and tracking more complicated than with periodic beaconing. Moreover, we have seen that no personal data can be inferred from the reading of data or metadata proper to N2N messages, heartbeats, or C2C messages. Consequently, local privacy threats, hardly feasible, are inoffensive. More precisely, privacy threats and risks of tracking by local

observers other than those existing anyway with human-driven vehicles are inexistent. Direct visual spying (cameras, humans) is easier than eavesdropping on radio channels dedicated to N2N and C2C messaging.

Regarding cyberattacks, one must distinguish physical reachability (message deliveries) and actual processing (instantiation of behaviors determined by message contents).

Let us begin with N2N messages. Owing to the range-2 Receive primitive and the Acceptance Condition (see Subsection 5.1.3), malevolent message M —assumed to be received locally by cohort member R —would simply be rejected without being processed, unless it is demonstrated that R may consider remote adversary Z as a range-1 or range-2 neighbor. This is an impossible proposition, owing to controlled cohort admission rules. Consequently, Z can only launch a cyberattack on the occasion of a Join operation (LtJoin, LgJoin), i.e. via C2C messaging (see below). A cohort member R may misbehave, attempting a cyberattack from inside a cohort. We have seen that this would be detected by R 's onboard PSM (R is halted). Nevertheless, considering the worst case where detection would not be effective or would be performed “too late”, local cyberattacks detrimental to safety shall be examined. We must question the motivations that would lead of a cohort member to instigate a possibly harmful collision, since that NGV would be involved into that collision. Very likely, such malicious behaviors belong to the category of irrational cyberattacks.

Let us now examine the case of cyberattacks launched via C2C messages. Z 's cyberattack can only succeed if performed via a suite of C2C messages in conformance with the specification of a Join operation. A Join operation is comprised of a number of rounds of C2C message exchanges (messages signed with valid pseudonyms ψ) that involve the participation of NGVs other than Z . See for example the LgJoin operation (Section 7) that comprises 4 cyber rounds, which precede the physical join maneuver(s) in case of acceptance. Firstly, Z must generate more than one message M . Secondly, Z must correctly guess every parameter that should appear in every (fake) C2C message. A vehicle profile $vp(\cdot)$ is an example (see Subsection 6.1.5). Consequently, Z can only be a vehicle that has been registered. Clearly, such remote cyberattacks can only fail, unless it is demonstrated that (1) remote vehicle Z can be perceived as being located ahead of a cohort head or behind a cohort tail, (2) Z can execute the LgJoin cyber rounds successfully, which implies explicit cooperation from (honest) cohort members involved in the LgJoin. These are impossible propositions.

The same reasoning applies a fortiori to LtJoin operations. In addition to cyber rounds based on radio messaging, optical communications play a crucial role in LtJoin operations. A NGV targeted by remote Z cannot perceive Z as being located in an adjacent lane.

Consequently, even if a malevolent message M issued by Z may reach some NGVs antennas devoted to C communications, M will simply be ignored and rejected by these NGVs. No harmful maneuvers can result from the processing of M 's contents. The only cases where cohort members perform risk-prone motions ordered by nearby vehicles arise with emergency vehicles. In its first C2C message, such a vehicle must provide a certificate which contains its profile $vp(\cdot)$, where parameter $vt(\cdot)$ states “emergency vehicle/type” (type = police car, fire truck, ambulance, etc.). Dishonest attacker Z cannot provide a valid certificate that would carry the above $vt(\cdot)$.

5.4.3 Brute force cyberattacks

We have shown that cyberattacks based on the processing of message contents are unfeasible (to the exception of irrational local cyberattacks). What about other cyberattacks that do not rest on the processing of message contents or that can be launched by non-registered adversaries? The physical jamming of radio channels is a brute force cyberattack which can be launched by any vehicle or static adversary equipped a jamming device. PSMs and predicates are inoperative vis-à-vis such cyberattacks. Safety is definitively lost whenever V2X communication channels are jammed. That is not the case with CMX functionalities.

Range-1 N2N communications remain feasible, via VLC. A member that detects the loss of a N2N radio link can trigger a cohort split. Members under cyberattack would then become isolated from each other, separated by *safe* inter-cohort gap $S(v, \cdot)$. Velocity could also be reduced. Therefore, with cohorts, efficiency may be compromised in the event of jammed N2N radio channels, whereas safety is preserved.

A similar conclusion holds with C2C radio channels. In the presence of jamming, no LtJoin or LgJoin operations are feasible. That may only hamper efficiency. In UX or UR settings, jamming cyberattacks can create deadlocks, whereby NGVs and cohorts do not attempt risk-prone maneuvers, such as crossing an unsignalized intersection or roundabout. This could lead to congestion. Such cyberattacks conducted from a vehicle are highly unlikely (unless launched by irrational adversaries), for the reason that an attacker will suffer the consequences of its own cyberattack (blocked in a congested area).

Anyhow, it is possible to thwart such cyberattacks, including jamming performed by a static adversary. NGVs would move at low velocities, keeping longitudinal and lateral distances larger than allowed in the

absence of jamming, while exploiting features and schemes specific to vehicular cells (see Section 4) and vehicular flocks (to appear). Traffic flow would be reduced, and congestion avoided.

Finally, observe that safety is not compromised in the presence of GNSS signal spoofing or jamming, conditions under which safety is definitively lost in the V2X framework. Firstly, GNSS space coordinates do not appear in N2N or C2C messages, which messages determine fully life/safety critical motions and maneuvers undertaken by NGVs. Secondly, onboard CC subsystems maintain correct knowledge of GNSS time coordinates despite GNSS jamming, due to onboard back-up clocks (see Subsection 5.1.4).

5.4.4 Conclusions

Thanks to CMX functionalities, the privacy of passengers in NGVs that operate in cyber stealth mode cannot be compromised in the presence of eavesdropping or/and tracking. Even when the cyber stealth mode is deactivated, motions and trajectories of NGVs remain unaffected by remote or rational local cyberattacks. Therefore, safety cannot be compromised by such cyberattacks. It remains to be seen whether irrational cyberattacks shall be addressed explicitly by designers of NGVs and AVNs, or whether they can be ignored, on the ground that they are highly unlikely, considered to belong to the set of residual assumptions (similar to brick wall events). Preservation of safety when radio channels are jammed is an important merit of CMX functionalities.

One can thus conclude that cohorts and their CMX functionalities have invaluable advantages regarding privacy and cybersecurity, in addition to safety and efficiency. Privacy and cybersecurity issues and solutions in the CMX framework, briefly surveyed in the above, are thoroughly examined in forthcoming reports and publications.

Let us now introduce cyberphysical concepts that play a pivotal role in a number of problems and solutions, notably gap control and demonstrations of safety and efficiency properties (examined in Section 8).

6. Cyberphysical Levels

In platoons, strings and cohorts, safety and efficiency properties depend on inter-vehicular gaps. Our focus is on rear-end crashes (pairwise collisions and pileups) that may occur within a formation and between formations in the same lane. Stability issues are out of the scope of this report. The severity of a rear-end collision is highly correlated with relative velocity at collision time, denoted v_{rel} , hence with braking powers (deceleration rates) and cyber reactive capabilities (in replacement of human reactivity).

In customary presentations of gap control schemes in platoons or strings [67], one considers two contiguous vehicles denoted L (leader) and F (follower), respective braking powers b_l and b_f , moving at velocity v , as well as a detection/reaction delay, denoted λ . F detects L 's velocity change via its onboard sensors or through inter-vehicular communications (a message received by F , issued by L or by some member ahead of L).

Smallest safe gaps $sp_{min}(v,f|l)$ in platoons or strings are computed out of the standard (simplified) formula, ignoring braking actuation latencies:

$$\begin{aligned} sp_{min}(v,f|l) &= \sigma^\circ(v) + [1/b_f - 1/b_l] v^2/2, & b_l > b_f. \\ sp_{min}(v,f|l) &= \sigma^\circ(v), & b_l \leq b_f. \end{aligned} \quad (Eq. 2)$$

Spacing $\sigma^\circ(v) = \lambda v + c^\circ$, c° standing for the inter-vehicular gap at velocity 0 (vehicles at a standstill).

Work on standard leader/follower LOS models has been generalized to CACC formations where a follower is fed with deceleration or hard-braking data carried in a V2X message issued by a platoon/string leader or by any other (NLOS) member, in addition to its predecessor. Since we investigate the role of cyber capabilities in gap control schemes, we consider the latter possibility. Variable λ depends on the MAC protocol and the CWD algorithm utilized by platoon/string members.

Variables b_f and b_l are subject to uncertainties, despite reliance on classical techniques such as Kalman filters for example. Deviations between postulated values for b_l or b_f and real values (on the roads) may be significant. Can they be bounded? Additional difficulties arise with b_l in Eq. 2. Is b_l a *highest* braking power? How may F "guess" b_l ?

These difficulties must be resolved, otherwise violent rear-end crashes can occur, entailing fatalities and severe injuries. Let us show how to address these difficulties by harnessing the potential of essential CMX functionalities, namely, cyberphysical levels.

6.1. Cyberphysical levels for NGVs

Neither braking powers nor cyber reactivity are referred to in the SAE categorization of automated driving levels. The ordering of braking powers and the ordering of cyber capabilities are unrelated. It seems thus necessary to devise some common referential of cyberphysical capabilities if we want to quantify safety and efficiency properties sufficiently accurately in heterogeneous AVNs. Herein, we introduce a referential based on a 2-dimension categorization of cyberphysical levels denoted $cpl(\cdot)$.

The Richter scale for earthquakes rests on a categorization in levels of magnitude ranging from 1 to 9. Vehicles are human-made constructs. Why wouldn't be possible to do for cyber delays and braking powers what we do regarding nature-made phenomena? Just like pollution grades and mileage figures, highest cyber delays and highest *nominal* braking powers can be categorized. The nominal braking power that characterizes a given model of NGV would be measured under specific standard conditions defined in an official referential \mathcal{R} . Measurements would be performed and certified by professional organizations designated by authorities in charge of ground transportation. Every NGV would then be assigned a cyber level and a physical level. Inspired by the SAE categorization in 6 levels, we suggest a partitioning of cyber levels in 6 intervals denoted $cl\ i, i \in [0, 5]$, and a partitioning of physical levels in 6 intervals denoted $pl\ j, j \in [0, 5]$.

6.1.1 Cyber levels

Two delays are of primary importance, namely worst-case upper bound δ (channel access delays) and worst-case upper bound $\Lambda(n)$ of reliable message dissemination delays in a cohort of n members. Assume for a moment that members have identical braking capabilities, gaps equal to $\sigma^\circ(v)$ as per *Eq. 2*. Consider the case of hard-braking by a cohort head *CH*. *CH* issues a N2N message *M* typed "c wd", disseminated downstream and acknowledged upstream via some CWD algorithm. Ideally (and unrealistically), with $\Lambda(n) = 0$, all members would start hard braking at exactly the same time, gaps $\sigma^\circ(v)$ kept unchanged. What could be an optimal $\Lambda(n)$ enabling highest efficiency?

■ Delays δ_i

MAC level access delays δ are established for worst-case contention conditions. According to figures found in published work, desired values of bounds δ range approximately between 20 ms [3] and 100 ms (as per the CC-ITS documentation). Consequently, δ belongs to $cl\ i$ if $\delta_{i+1} < \delta \leq \delta_i, \quad \delta_i = (6-i) 20$.

Inexistent δ_6 is set to 0. Extreme values are as follows:

- $i = 5 \Leftrightarrow \delta_5 = 20 \Leftrightarrow \delta \leq 20$
- $i = 0 \Leftrightarrow \delta_0 = 120 \Leftrightarrow 100 < \delta \leq 120$

MAC protocols that do not guarantee $\delta \leq \delta_0$ are said ineligible, i.e. not considered for NGVs. Vehicles equipped with such protocols are assigned symbol \perp rather than $cl\ i$.

■ Delays λ_i

Delays λ in *Eq. 2* hold for two contiguous members. It should not be necessary to let contiguous members enter into mutual competition for accessing a radio channel devoted to N2N communications. Indeed, $\lambda = \delta$ is not desirable if we look for an optimal $\Lambda(n)$. In line with the definition of δ as a multiple of 20, we define λ as a function linear in $\sqrt{20}$. Delay λ belongs to $cl\ i$ if $\lambda_{i+1} < \lambda \leq \lambda_i, \quad \lambda_i = (6-i) \sqrt{20} = \delta_i \sqrt{5}/10$.

Inexistent λ_6 is set to 0. Extreme values are as follows:

- $i = 5 \Leftrightarrow \lambda_5 = 2 \sqrt{5} \Leftrightarrow \lambda \leq 2 \sqrt{5}$
- $i = 0 \Leftrightarrow \lambda_0 = 12 \sqrt{5} \Leftrightarrow 10 \sqrt{5} < \lambda \leq 12 \sqrt{5}$

■ Delays $\Lambda_i(\tilde{n})$

Delays $\Lambda(n)$ linear in λ can be categorized provided that some nominal bound for n , denoted \tilde{n} , is defined in referential \mathcal{R} . Delay $\Lambda(\tilde{n}) = (\tilde{n}-1) \lambda_i$ belongs to $cl\ i$ if $\Lambda_{i+1}(\tilde{n}) < \Lambda(\tilde{n}) \leq \Lambda_i(\tilde{n}), \quad \Lambda_i(\tilde{n}) = 2 (\tilde{n}-1) (6-i) \sqrt{5}$.

Inexistent $\Lambda_6(\tilde{n})$ is set to 0. Extreme values are as follows:

- $i = 5 \Leftrightarrow \Lambda_5(\tilde{n}) = 2 (\tilde{n}-1) \sqrt{5} \Leftrightarrow \Lambda(\tilde{n}) \leq 2 (\tilde{n}-1) \sqrt{5}$
- $i = 0 \Leftrightarrow \Lambda_0(\tilde{n}) = 12 (\tilde{n}-1) \sqrt{5} \Leftrightarrow 10 (\tilde{n}-1) \sqrt{5} < \Lambda(\tilde{n}) \leq 12 (\tilde{n}-1) \sqrt{5}$

CWD algorithms that do not guarantee $\Lambda(\tilde{n}) \leq \Lambda_0(\tilde{n})$ are said ineligible. Vehicles equipped with such algorithms are assigned symbol \perp rather than $cl\ i$. A reasonable choice would be to have \tilde{n} equal to the supremum of n^* values for all settings, that is $\tilde{n} = 25$ (see Section 5).

6.1.2 Physical levels

Highest braking/deceleration capabilities of currently commercialized vehicles have values ranging between 4 m/s^2 and g approximately.

■ Braking powers β_j

Increments of 1 m/s^2 lead to a partitioning in 6 intervals of highest braking powers denoted β_j (in m/s^2). Consequently, measured β belongs to $pl\ j$ if $\beta_{j-1} < \beta \leq \beta_j$, $\beta_j = 4 + j$.

Inexistent β_{-1} is set to 0. Extreme values are as follows:

- $j = 5 \Leftrightarrow \beta_5 = 9 \Leftrightarrow 8 < \beta \leq 9$
- $j = 0 \Leftrightarrow \beta_0 = 4 \Leftrightarrow \beta \leq 4$

Contrary to cyber levels, some β_j is assigned to every NGV (no symbol such as \perp may appear in pl field j).

6.1.3 Cyberphysical levels

A cyberphysical level is a pair of integers $\{i, j\}$. Boundaries of cyberphysical levels are shown in Table I which instantiates a partially ordered set of pairs $\{i, j\}$. This set can be converted into a totally ordered set comprised of integers $k = 6i + j$ represented as a 6×6 CPL matrix, as shown in Fig. 7. Any given cp level may be indifferently referred to as pair $\{i, j\}$ or element k .

Table I. A categorization of cyberphysical levels

cyber level i	0	1	2	3	4	5
δ_i	120	100	80	60	40	20
λ_i	$12\sqrt{5}$	$10\sqrt{5}$	$8\sqrt{5}$	$6\sqrt{5}$	$4\sqrt{5}$	$2\sqrt{5}$

Assuming $\tilde{n} = 25$:

cyber level i	0	1	2	3	4	5
$\Lambda_i(25)$	644	536.7	429.3	322	214.7	107.3

physical level j	0	1	2	3	4	5
β_j	4	5	6	7	8	9

Consider two contiguous cohort members L and follower F . Let $j(l)$ stand for the physical level assigned to member L . Let $i(f)$ and $j(f)$ be the cyber and the physical levels assigned to L 's follower F . Smallest safe inter-neighbor gaps in cohorts, computed out of the standard (simplified) formula (ignoring braking actuation latencies), are as follows:

$$s_{min}(v, f/l) = \sigma^\circ(v) + [1/\beta_{j(f)} - 1/\beta_{j(l)}] v^2/2, \quad \beta_{j(l)} > \beta_{j(f)},$$

$$\sigma^\circ(v) = \lambda_{i(f)}v + c^\circ.$$

For clarity, $s_{min}(v, f/l)$ will be expressed as follows:

$$s_{min}(v, f/l) = \sigma^\circ(v) + [1/\beta_f - 1/\beta_l] v^2/2, \quad \beta_l > \beta_f,$$

$$\sigma^\circ(v) = \lambda_f v + c^\circ.$$

Symbol \perp is assigned to vehicles not equipped with communication technology, like vehicles equipped with ineligible MAC protocols or CWD algorithms. Nevertheless, they are elements of AVNs, which leads to interesting safety and efficiency problems with heterogeneous cohorts (see further).

More complete categorizations in cyberphysical levels would include other physical parameters, such as vehicle length or highest acceleration rates, which rates appear in analyses devoted to cohort stability. Should cyberphysical levels be adopted, MAC protocols and CWD algorithms would be translated into standards quoting their cyber levels. Authorities and the automotive industry could influence future trends regarding the types of vehicles to be built, seeking homogeneity in both dimensions (cyber and physical levels). For example, premium cars would match cl 4 or 5, and pl 4 or 5, forming compact cohorts (high efficiency) able to move safely at high velocities. Conversely, low-end vehicles (cl 0, 1 or 2, pl 0, 1 or 2) would form sparse and/or slow cohorts which should not interfere with fast moving cohorts. See Subsection 6.3.4 for more on this topic.

6.1.4 On the relevance of cyber levels and sublinear time CWD

The following example shows the relevance of the cyber dimension in hard braking scenarios. For δ , let us examine extreme boundaries 20 ms and 100 ms. With $n = 25$, a CWD algorithm linear in δ yields a worst-case termination time bound ranging between 480 ($i = 5$) and 2,400 ($i = 1$), whereas $\Lambda(25)$ ranges between 107.3 and 536.7, respectively. Consider $i = 5$, thus favoring CWD algorithms linear in δ , and $v = 30$ m/s. A cohort tail would perform the hard braking command carried in a message issued by a cohort head 11.18 m (372.7 ms) “earlier” with a CWD algorithm linear in $\sqrt{20}$. That is comparable to differences in braking distances d travelled by vehicles that have slightly different braking powers. Applying the simple formula $d(v) = v^2/2b$, with $v = 30$ m/s and considering braking powers equal to 6 m/s^2 and 7 m/s^2 , one finds $d(30) = 10.7$ m. Cyber capabilities and physical capabilities shall deserve equal attention in networks of NGVs.

Another example can be given with a scenario where an emergency vehicle E requires lane clearance. Its C2C message is received by a cohort tail that activates a CWD algorithm for disseminating an upstream message carrying C code 12 (“clear lane” in Fig. 3). Desired physical motions are obtained out of cyber capabilities, and $\Lambda(n)$ is the parameter that determines how quickly vehicles ahead of E learn that they must move to some other lane.

With an eligible CWD algorithm, messages are acknowledged. Consequently, repetitions due to message losses shall be accounted for in $\Lambda(n)$, leading to $\Lambda(n,f)$. That would complicate matters since this is predicated on assuming that authorities and professional organizations may agree on some worst-case bound f^* . It seems more realistic to let designers of CWD algorithms be responsible for providing delays $\Lambda(n,f)$ achieved by their algorithms when $n = \tilde{n}$. Then, it would suffice to compare some $\Lambda(\tilde{n},f)$ with $\Lambda_i(\tilde{n})$ boundaries in Table I to know which cl i is matched by a CWD algorithm under some given f or, reciprocally, which f^* can be sustained for a given bound $\Lambda(\tilde{n},f)$ or cl i .

For example, with SWIFT-based CWD, $\tilde{n} = 25$, $\theta = 1.1$ ms and $h = 5$, one finds $f \leq 13$ for a match with cl 4 (bound $\Delta(25,13) = 209$), and $f \leq 3$ for a match with cl 5 (bound $\Delta(25,3) = 99$).

Let us now demonstrate that cyber level 5 is not an unattainable utopia with MAC protocols and CWD algorithms similar to SWIFT. In Subsection 5.1.4, we have stressed that $h \geq 3$, irrespective of n , and that safety is ensured whichever value is assigned to h . Given that range control is imperfect, antennas tuned for $h = 3$ may span higher ranges. For a good spatial reusability (members can transmit at the same time without incurring a collision), h shall not be too high. Let us assume $h \leq 7$ (high coverage, very poor power/range control). This assumption only serves to compute performance figures.

Numerical examples, with $n = 25$ and $f = 0$ (as with bound $\Lambda(n)$):

- Case $\theta = 1$ ms and $h = 4$: $\delta = 8 \Rightarrow$ cl 5 $\Delta(25,0) = 56 \Rightarrow$ cl 5
- Case $\theta = 1.25$ ms and $h = 7$: $\delta = 17.5 \Rightarrow$ cl 5 $\Delta(25,0) = 87.5 \Rightarrow$ cl 5

Let us establish the general conditions on δ under which SWIFT-based CWD matches cl 5. This is done by showing that the promptness requirement for eligibility stated in Subsection 5.2.1 is met: bound $\Delta(n,.)$ is sublinear in $n\delta$. Since $\Lambda(n)$ is linear in λ , which is linear in $\sqrt{20}$, let us compare $\Delta(n,.)$ and $2\Lambda(n)$. Factor 2 is due to the fact that, contrary to $\Delta(n,.)$, $\Lambda(n)$ does not include acknowledgement latencies. For a fair comparison, integer f should be set to 0 in $\Delta(n,f)$. Recall also that $\Delta(n,0)$ includes 1 delay δ for the first channel access. Such a delay is not accounted for in $\Lambda(n)$. Therefore, linearity in $\sqrt{\delta}$ holds under the following condition:

$$\delta \lceil (n-1)/h \rceil \leq 2(n-1) \sqrt{\delta} \quad (\text{Eq. 3})$$

Consider integer $p > 0$. Integer $\lceil (n-1)/h \rceil = p$ when $(p-1)h+1 \leq n-1 \leq ph$. Let p° stand for $(p-1)h+1$. Eq. 3 holds true if it is valid for the smallest value of the term on the right, that is for $n-1 = p^\circ$.

The case $p = 1$ must be addressed separately from the general case, since h matters in the general analysis. With $p = 1$, $2 \leq n \leq h+1$. The case $n = 2$ shall be ignored, since CWD is not utilized (see Subsection 5.2.1). For $n = 3$, Eq. 3 yields $\delta \leq 16$. Trivially, Eq. 3 holds for $n > 3$, yielding bounds on δ higher than 16.

In the general case ($p > 1$), Eq. 3 writes as follows: $\sqrt{\delta} \leq 2[(p-1)h+1]/p$

With $p = 2$, one finds $\delta \leq (h+1)^2$. Since $h \geq 3$, lower bounds to be matched by δ are at least equal to 16. Clearly, lower δ bounds are higher than 16 for $p > 2$.

We have shown that irrespective of p , linearity in $\sqrt{\delta}$ holds whenever $\delta \leq 16$. This condition is slightly tighter than the criterion for cyber level 5 ($\delta \leq 20$).

NGVs equipped with SWIFT-based CWD are categorized cl 5 when $\Delta(\tilde{n},0) \leq \Lambda_5(\tilde{n})$. For a given \tilde{n} , $\Delta(\tilde{n},0)$ is highest when h is smallest ($h = 3$). Highest $\Delta(25,0) = 9\delta$. Therefore, $\delta \leq 11.92$ is the condition for a match with cl 5. When $\delta \leq 16$ (SWIFT matches cl 5), SWIFT-based CWD yields bounds $\Delta(25,0)$ which may or may not be smaller than $\Lambda_5(25)$, depending on h . For example, with lowest $h = 3$, $\Delta(25,0) = 144$, which matches cl 4, while $\Delta(25,0) = 96$ with $h = 5$, which matches cl 5.

Cyber level 5 is not an unattainable utopia.

6.1.5 Vehicle profiles

Every vehicle is associated a profile denoted $vp(\cdot)$: $vp(\cdot) = \{cpl(\cdot), aul(\cdot), vt(\cdot), vl(\cdot)\}$.

Notation $aul(\cdot)$ stands for SAE driving level, $vt(\cdot)$ for vehicle type (e.g., car, mid-size commercial vehicle, emergency vehicle, truck, trailer), and $vl(\cdot)$ for vehicle length. Vehicle profiles may include additional parameters, not detailed here. Profile $vp(X)$ is recorded in the long-term certificate and in reusable certificates allocated to X at time of registration, stored into its PSM. Vehicle profiles also appear on vehicles' plates.

In this report, we restrict our attention to parameters $cpl(\cdot)$ and their companion interoperability sets. Parameters other than $cpl(\cdot)$ also have associated interoperability sets. For example, regarding $vt(\cdot)$, a car may not be interoperable with trucks, i.e. they shall not be members of the same cohort. Principles and rules presented below for interoperability sets associated to cp levels may lay the ground for other parameters.

6.2. Vehicle cyberphysical level interoperability sets $\varphi(k)$

The interoperability set associated to a vehicle X tagged with element k , denoted $\varphi(k)$, defines which vehicles are deemed "compatible" with X , i.e. endowed with comparable cyber capabilities and braking powers. Sets $\varphi(k)$ are stored in onboard systems of NGVs assigned cp level k .

Such sets matter with regards to safety. Numerous non-fatal collisions involving autonomous cars reported to date appear to be caused by human-driven vehicles. Consequently, for the sake of safety, NGVs tagged with low cp levels and those tagged with high cp levels shall be kept apart. Sets $\varphi(k)$ also matter with regards to efficiency (see Section 8).

The definition of sets $\varphi(k)$ is under the responsibility of the automotive industry and authorities in charge of road safety. Sets $\varphi(k)$ may evolve with vehicular technology and policies derived from legislation updates. Authorities and the automotive industry will define cp level *interoperability* accordingly. That would just be an extension of policies and legislations resorted to for adopting speed limits.

We anticipate the need for three cohort operative modes, namely unrestricted heterogeneity, restricted heterogeneity, and strict homogeneity. The CPL matrix is essential for differentiating between these modes. To show how sets $\varphi(k)$ might be defined and used, we offer an simple instantiation which has the merit of enforcing a clean and dynamic separation between the three operative modes. In other words, a NGV that operates under either of these modes will never find itself "mixed with" vehicles that operate under either of the other two modes.

Every member of a cohort that operates under the unrestricted heterogeneity mode has an interoperability set that includes all elements of the CPL matrix, from 0 to 35, irrespective of its own cp level. Members of a cohort that operates under the strict homogeneity mode have the same set $\varphi(k)$, comprising element k only. The definition of vehicle interoperability sets for the restricted heterogeneity mode is slightly more involved.

Sets $\varphi(k)$ are defined as dense subsets of the CPL matrix, where elements are adjacent to each other. Every element $k' = \{i', j'\}$ in a set $\varphi(k)$ is adjacent to $k = \{i, j\}$, adjacency defined as follows:

$$\{|i - i'| = 1 \text{ or } |j - j'| = 1\} \quad \text{and} \quad \{|i - i'| = 1 \text{ and } |j - j'| = 1\}.$$

Consequently, under the restricted heterogeneity mode, a set $\varphi(k)$ has an infimum (lowest k) and a supremum (highest k), and may comprise 4, 6 or 9 elements. Examples are given in Fig. 7.

$i \backslash j$	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. 7. The CPL matrix and interoperability sets $\varphi(\cdot)$ for the restricted heterogeneity mode

Roots k of sets $\varphi(k)$ are shown in italic within colored circles:

- Set $\varphi(14)$ comprises 9 elements, 16 such sets are possible,
- Set $\varphi(11)$ comprises 6 elements, 16 such sets are possible,
- Set $\varphi(35)$ comprises 4 elements, 4 such sets are possible.

Cyberphysical levels and the concept of interoperability sets play a pivotal role in a number of problems and solutions, such as controlled cohort heterogeneity and dynamic cohort/lane assignment strategies (explored below and illustrated in Section 7), and gap control aimed at optimizing efficiency (see Section 8).

6.3. Cyberphysical level interoperability sets $\Phi(\cdot)$ for cohorts and controlled cohort heterogeneity

Cohort/lane assignment and controlled cohort heterogeneity are inter-dependent. In urban settings where highest authorized velocities and/or the number of lanes per arterial are small, it might be useless, impossible or counterproductive to restrict the spontaneous formation of heterogeneous cohorts. Some lanes may be reserved for specific vehicles (e.g., buses, taxis). In multilane highways and major roads, lane reservation may be practiced, notably at rush hours (carpool lanes). Access to other lanes is unrestricted. In the general case examined hereafter, heterogeneity control and cohort/lane assignment must be fully dynamic for the sake of optimal efficiency.

Contrary to sets $\varphi(k)$ tied to vehicles, interoperability sets tied to cohorts change according to modifications in membership. Let $\Phi(\Gamma) = \bigcap_{X \in \Gamma} \varphi(X)$ stand for the cpl interoperability set tied to cohort Γ at some given time. By definition, all elements associated to Γ 's members belong to set $\Phi(\Gamma)$. They are referred to as active elements. In Fig. 8, active elements are roots of sets $\varphi(\cdot)$ shown within colored circles (as in Fig. 7) and elements that are adjacent to one of those roots, shown within squares. Other elements are said inactive, meaning that vehicles that could be tagged with such elements are not present.

Controlled cohort heterogeneity rests on sets $\varphi(\cdot)$ and $\Phi(\cdot)$. A new cohort may be created when an isolated vehicle joins with another isolated vehicle, when an isolated vehicle joins with a cohort, or vice-versa, and when a cohort joins with another cohort. Let ios stand for an interoperability set $\varphi(\cdot)$ or $\Phi(\cdot)$. Consider two sets ios_1 and ios_2 representative of NGVs or cohorts that attempt a Join operation. If successful, a new cohort Γ is created, with $\Phi(\Gamma) = ios_1 \cap ios_2$. A Join operation is subject to conditions.

6.3.1 Intersecting interoperability sets and CPL conditions for a Join

Let $minsize$ stand for the smallest of ios_1 size and ios_2 size (in number of elements).

CPL (cyberphysical level) conditions for a successful Join attempted by NGVs or cohorts characterized by ios_1 and ios_2 leading to new $\Phi(\cdot) = ios_1 \cap ios_2$ are defined after the following rules:

- LR (limited reduction): size of new $\Phi(\cdot) > minsize/2$,
- ZE (zero exclusions): all elements in ios_1 and all elements in ios_2 belong to new $\Phi(\cdot)$.

■ Unrestricted heterogeneity and strict homogeneity

Owing to CPL conditions, sizes of sets ios cannot change, kept to 36 (unrestricted heterogeneity) or to 1 (strict homogeneity). Indeed:

- Either $ios_1 = ios_2$, and a Join operation succeeds, with new $\Phi(.) = ios_1 = ios_2$,
- Or $ios_1 \neq ios_2$, and a Join operation is denied. When ios_1 is of size 36, any $ios_2 \neq ios_1$ is at most of size 9. When ios_1 is of size 1, any $ios_2 \neq ios_1$ is of size larger than 1. In both cases, rule LR is violated.

■ Restricted heterogeneity

Rules LR and ZE are illustrated in Fig. 8, where overlapping interoperability sets $\Phi(.)$ are shown in yellow. Sets ios_1 and ios_2 are shown in blue and green, fully or partially hidden by yellow set $\Phi(.)$.

Figure 8a

Two vehicles X and Y, associated sets $\varphi(4) = \{3, 4, 5, 9, 10, 11\}$ and $\varphi(5) = \{4, 5, 10, 11\}$ join together, forming cohort Γ . Interoperability set $\Phi(\Gamma) = \varphi(4) \cap \varphi(5)$ of size 4 fulfills both rules:

- LR: Size of $\Phi(\Gamma) = 2/3$ size of $\varphi(4)$, and size of $\Phi(\Gamma) =$ size of $\varphi(5)$,
- ZE: Elements 4 and 5 belong to $\Phi(\Gamma)$.

Figure 8b

Two cohorts Γ and Ω , set $\Phi(\Gamma) = \{13, 14, 15, 19, 20, 21, 25, 26, 27\}$ in blue, active elements 15 and 20, and set $\Phi(\Omega) = \{14, 15, 16, 20, 21, 22, 26, 27, 28\}$ in green, active elements 15, 21 and 26, join together, forming new cohort Γ . New interoperability set $\Phi(\Gamma) = \{14, 15, 20, 21, 26, 27\}$ of size 6 fulfills both rules:

- LR: Size of new $\Phi(\Gamma) = 2/3$ of initial sizes of $\Phi(\Gamma)$ and $\Phi(\Omega)$,
- ZE: Elements 15, 20, 21 and 26 belong to new $\Phi(\Gamma)$.

Figure 8c

A rejected Join is illustrated with cohorts Γ and Ω , set $\Phi(\Gamma)$ in blue, size 9, active element 21, and set $\Phi(\Omega)$ in green, size 6, active element 17. Interoperability set $\Phi(.) = \Phi(\Gamma) \cap \Phi(\Omega) = \{16, 17\}$ fulfills none of the rules:

- LR: Size of $\Phi(.) = 2$,
- ZE: Elements 17 and 21 do not belong to $\Phi(.)$.

Figure 8d

Same as Fig. 8b (successful Join), except that cohort Ω includes a member tagged with element 22, which leads to a rejected Join. Rule LR is fulfilled, whereas rule ZE is not.




Fig. 8. Examples of interoperability sets $\Phi(.)$ in the restricted heterogeneity mode

6.3.2 Mutating sets $\Phi(\cdot)$ and restricted heterogeneity

Cohort membership may decrease on the occasion of a Leave operation or a split. Sets $\Phi(\cdot)$ remain unchanged, whatever the heterogeneity/homogeneity mode. Let us demonstrate that owing to rule LR, a vehicle that operates under the restricted heterogeneity mode will never find itself “mixed with” vehicles that operate under either of the other two modes.

A cohort created with an associated set $\Phi(\cdot)$ of size 9 keeps that set unchanged as long as Joins are attempted by cohorts or vehicles with associated sets of size 9 in perfect match with $\Phi(\cdot)$. Otherwise, according to rule LR, a set $\Phi(\cdot)$ of size 9 can only mutate into a set $\Phi(\cdot)$ of size 6. (A direct 9-to-4 mutation is ruled out.)

Identical conclusions hold for cohorts created with an associated set $\Phi(\cdot)$ of size 6, or that have an associated set $\Phi(\cdot)$ resulting from a 9-to-6 mutation. If not kept unchanged, a set $\Phi(\cdot)$ of size 6 can only mutate into a set $\Phi(\cdot)$ of size 4.

Regarding cohorts created with an associated set $\Phi(\cdot)$ of size 4, or that have an associated a set $\Phi(\cdot)$ resulting from a 6-to-4 mutation, it is obvious that such sets can only be kept unchanged. Rule LR prohibits a Join in other cases (with sets of size other than 4, with sets of size 4 not in perfect match with $\Phi(\cdot)$).

Since the reduction process ceases when an associated set $\Phi(\cdot)$ is of size 4, restricted heterogeneity can neither mutate into strict homogeneity (sets $\Phi(\cdot)$ of size 1) nor mutate into unrestricted heterogeneity (sets $\Phi(\cdot)$ of size 36). This completes the demonstration.

Dominating interoperability sets $\Phi(\cdot)$

An interoperability set $\Phi(\cdot)$ has a supremum (highest k), in the same way as sets $\varphi(\cdot)$. Consider cohorts of restricted heterogeneity and strictly homogeneous cohorts. When two interoperability sets are not strictly identical, one of them owns the highest supremum. In the case of unrestricted heterogeneity, the supremum of an interoperability set (which comprises 36 elements) is defined as the highest active element k .

A set that owns the highest supremum is referred to as the dominating set. This dominance property is utilized in cohort/lane assignment, lane clearing in particular. A NGV or a cohort must give way to a NGV or a cohort that owns a dominating set, rather than forcing the latter to engage in overtaking maneuvers. For details, see Policies, Subsection 6.3.4.

Lane clearing ordered by an emergency vehicle (Z) is handled differently. In Subsection 5.2.1, we have briefly discussed the scenario where Z orders lane clearing via a C2C message. Z provides a certificate with $vt(Z)$ inside. CJ conditions presented below do not apply. They are superseded by inspection of $vt(\cdot)$. When $vt(\cdot) = \text{“emergency vehicle”}$, lane clearing is mandatory.

6.3.3 CJ conditions for a Join

A Join operation implies physical displacements and is conditioned on verified authentication. Apart from CPL conditions, PF and VA conditions shall be fulfilled.

PF conditions (physical feasibility)

In addition to meeting bound n^* (a Join with a cohort that is full is denied), PF conditions serve to eliminate hazards in conflicting concurrent maneuvers. For example:

- A LgJoin is attempted with cohort Γ while a “clear lane” message is disseminated throughout Γ ,
- A LgJoin is attempted with cohort Γ when Γ 's tail has started a Leave operation, or when Γ 's tail is being halted (due to malicious behavior),
- Two conflicting LtJoin requests are received quasi simultaneously by cohort member R , one issued by a left-side neighbor and the other one issued by a right-side neighbor. One of these requests shall be turned down, for preventing a lateral collision,
- Multiple LtJoin insertions are requested at about the same time at distinct ranks in a long cohort. Only a subset of them is physically and safely feasible due to the multiple quasi-simultaneous deceleration maneuvers that would ensue from acceptance.

VA conditions (verified authentication)

Pseudonyms utilized in the first rounds of a LtJoin or LgJoin operation must be valid (see Subsection 7.1 for LgJoin).

CJ conditions are the union of CPL, PF and VA conditions.

Other conditions for a Join may encompass interoperability sets that are related to parameters such as $aul(\cdot)$, $vt(\cdot)$, and $vl(\cdot)$ found in vehicle profiles $vp(\cdot)$. They are not detailed in this report.

6.3.4 Policies and CJ conditions

CJ conditions are at the core of policies that govern behaviors of NGVs and cohorts in the physical space. LtJoin is not explored in this report. Let us focus on policies that may be followed by requestor A in lane g (an isolated NGV or a cohort head) whenever a LgJoin request is turned down by a responder (an isolated NGV or a cohort tail). For example:

- A overtakes the responder,
- A stays in lane g , $S(v, \cdot)$ away from the responder,
- If A owns the dominating set, the responder must clear lane g , giving way to A (local topological conditions permitting).

The combination of a “clear lane/cohort with dominating set behind” policy maximizes asphalt utilization ratios (efficiency) and passengers’ satisfaction conjointly since NGVs and cohorts with high cp levels will spontaneously circulate in fastest lanes. It turns out that this policy matches driving rules in many European countries, where slower vehicles are supposed to circulate in rightmost lanes.

There is thus no need for restricting the circulation of highly automated vehicles (or any other category) to specific lanes designated a priori. In addition to being antagonistic with the efficiency property, such a constraint might be a source of frustration. In a multilane roadway, vehicles may enjoy fluid traffic conditions while congestion would develop in a lane preassigned to NGVs of SAE level 5 or of $cpl\{5,5\}$, for example. Similarly, thanks to parameters such as $vt(\cdot)$ or $vl(\cdot)$ and their associated interoperability sets, long vehicles (trucks, trailers, etc.) will form homogeneous cohorts spontaneously in lanes not occupied by cohorts of small size NGVs (e.g., sedans, small vans). They may however share the same lane since cohorts of small size NGVs and cohorts of long vehicles are separated by safe inter-cohort gaps $S(v, \cdot)$.

The circulation of vehicles not equipped with radio communication devices or equipped with ineligible protocols or algorithms (see Subsection 6.1.1) will be constrained. Their vehicle profiles $vp(\cdot)$ has argument $cpl(\cdot)$ set to $\{\perp, j\}$. (Profiles are also displayed on vehicles’ plates.) Whenever necessary, they will have to give way to NGVs of any cyber level other than \perp . They will not be able to issue a LgJoin or a LtJoin request, thus kept away from cohorts of NGVs. They will form inefficient cohorts, where *intra-cohort gaps* are *inter-cohort gaps* $S(v, \cdot)$. Unable to issue C2C messages, they will be kept out of unsignalized intersections and roundabouts, only allowed to cross intersections and roundabouts equipped with traffic lights. Since behaviors are recorded in EDRs, violators (refusals to give way, unauthorized crossings) will face penalties.

6.3.5 Cohort profiles and cohort-wide common knowledge

Cohort membership and topology vary with time. Topology $TP(\Gamma)$ of a cohort Γ is 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. Every cohort is associated a profile denoted $\Pi(\cdot)$ that varies with time. Profile $\Pi(\Gamma)$ is comprised of the following parameters:

- Current Γ ’s velocity v (within some bounded tolerance)
- n , current number of Γ ’s members ($n \leq n^*$)
- Current Γ ’s interoperability set $\Phi(\Gamma)$
- Current Γ ’s topology $TP(\Gamma)$.

A cohort profile $\Pi(\cdot)$ is copied in members’ onboard CC subsystems, and updated whenever necessary via CWD or CWA, thus instantiating cohort-wide common knowledge. Complete cohort profiles include interoperability sets related to other parameters found in profiles $vp(\cdot)$. They are not examined in this report.

6.3.6 Scope-limited common knowledge

In addition to cohort-wide knowledge, cohort members share symmetrical scope-limited knowledge.

Symmetrical range-1 neighbors’ profiles knowledge

Symmetrical range-1 neighbors’ profiles knowledge is obtained out of N2N management. In a cohort, every member knows its predecessor’s and successor’s profiles $vp(\cdot)$, notably their cp levels. Recall that profiles are

trustworthy since falsification of a $vp(\cdot)$ implies falsification of a certificate. This knowledge plays an essential role in trusted gap control schemes. Assume that this symmetrical range-1 property holds within a cohort. Let us briefly explain how this property is persistent in the presence of Join or Leave operations.

Consider three consecutive neighbors U , W and X (ordered by increasing ranks). A LtJoin request is issued by NGV V for an insertion between U and W . To be granted by U and W , this request must be accompanied by a valid certificate C_V —with $vp(V)$ inside. Consequently, when V 's insertion is about to start, U knows its new successor's profile and W knows its new predecessor's profile. See Section 7 for details in the case of a LgJoin operation, where a new cohort tail (A in Fig. 9) learns the profile of previous tail (Q in Fig. 9)—now its predecessor, and Q learns the profile of its new successor A . Join operations maintain symmetrical range-1 neighbors' profiles knowledge.

Consider now a Leave operation. In the above set of consecutive neighbors U , V , W and X , three neighbors (U , V , and W) know profile $vp(V)$. Member W decides to leave. X must learn $vp(V)$ promptly, in order to calculate a safe gap with V . V does not know $vp(X)$. A Leave operation by W involves the following cyber rounds:

- W triggers two range-1 send of a N2N message carrying “I am leaving, $vp(V)$, $vp(X)$ ”, one aimed at V and another aimed at X ,
- V triggers a range-2 Send of a N2N message aimed at X , which message carries “This is $vp(V)$, your new predecessor's profile”,
- X triggers a range-2 Send of a N2N message aimed at V , which message carries “This is $vp(X)$, your new successor's profile”.

Vehicle profiles received from two sources can be cross checked, $vp(V)$ by X and $vp(X)$ by V . When they match, they can be trusted (the 1-out-of-3 assumption). No certificates must be consumed. Leave operations maintain symmetrical range-1 neighbors' profiles knowledge.

Symmetrical range-2 gap knowledge

Symmetrical range-2 gap knowledge is another example of knowledge sharing based on N2N management. This knowledge is useful for computing ranges of N2N directional antennas, in support of range-2 Send and Receive primitives. Let us briefly explain how this knowledge is maintained in the face of changing gaps. Assume that this property holds within a cohort: every member is aware of the two upstream and the two downstream gaps. Consider five contiguous neighbors A , B , C , D and E (by increasing ranks). Let $s(v,y/x) \geq s_{min}(v,y/x)$, where Y is X 's successor (notations introduced in Subsection 6.1.3). C is aware of upstream $s(v,b/a)$ and $s(v,c/b)$, of downstream $s(v,d/c)$ and $s(v,e/d)$. Whenever a gap changes by a certain proportion (some bounded tolerance shall be accommodated), that is detected quasi-instantly by both neighbors, thanks to onboard forward looking and backward looking sensors. Range-2 neighbors must be notified.

Assume that gap $s(v,c/b)$ changes, for example due to C (gap increased in case C decelerates, decreased in case C accelerates). That is detected by C and B . C (resp., B) would then send D (resp., A) a range-1 N2N message carrying updated $s(v,c/b)$. Moreover, C (resp., B) would then Send A (resp., D) a range-2 N2N message carrying updated $s(v,c/b)$. A and D can cross check the data received from two sources. When they match, A and D update their $s(v,c/b)$ knowledge. Similarly, C would be notified by D (range-1 N2N message) and by E (range-2 N2N message) in case gap $s(v,e/d)$ would change significantly.

7. Cohort Admission Control, the LgJoin Operation

The maneuver under examination is sometimes referred to as gap closing. Consider 2 cohorts Ω and Γ in the same lane g , Ω moving at velocity v' that catches up with Γ moving at velocity $v < v'$. In the general case, cohorts are of size 2 at least, ω members in Ω and γ members in Γ . A and B , respectively head of cohort Ω and A 's successor, coordinate with Q and P , respectively tail of cohort Γ and Q 's predecessor—see Fig. 9. Both cohorts must agree on whether some or all members of Ω may or may not leave Ω for joining Γ .

The LgJoin operation involves C2C communications and LgSend primitives (the cyber rounds), which serve to convey the data needed to check whether CJ conditions are fulfilled, and to strike explicit agreements. Physical motions follow suite. There are various possible instantiations of LgJoin, which depend on assumptions regarding failures and malicious behaviors. In this report, we introduce LgJoin as a 4-way handshake for the cyber part. Round 1 is the contact round, round 2 the request round, round 3 the reply round, and round 4 the decision round, where a positive or negative decision is made. Our main goal here is to expose basic principles for the general case. Failures are not considered. Falsifications of vehicle profiles are addressed.

Radio ranges

Without loss of generality, let us assume that forward-looking (resp., backward-looking) send and receive antennas are located in the front (resp., back) of vehicles. Let $D(X/Y)$ stand for the radio range separating X 's and Y 's C2C antennas, $d(X/Y)$ stand for the radio range separating X 's and Y 's N2N antennas. These ranges are detailed in Subsection 8.6. Inter-cohort gap $S_{\min}(v, \cdot)$ is given in Subsection 8.4. A $X \leftrightarrow Y$ C2C message is also received by vehicles located between X and Y .

7.1. The four cyber rounds

7.1.1 Round 1 ($\Omega \leftrightarrow \Gamma$): contact round

Purpose: To make Q and P aware of distance $L_{B/A} = s(v, b/a) + vl(A)$, and to make A and B aware of distance $v\text{slot}(\gamma, \Gamma) = vl(Q) + s(v, q/p)$.

Via its sensors (radars, cameras, etc.), A detects Q and keeps itself $D(A/Q) \geq S_{\min}(v, \cdot)$ away from Q , while setting its velocity to v . Message $m_{1,A} = \{\text{hello}, D(A/Q), L_{B/A}\}$ is the C2C message A sends to Q , where *hello* is the C code for “data for LgJoin round 1”. A tunes its forward-looking C2C antenna according to $D(A|Q)$ and activates $\text{LgSend}(A, m_{1,A})$, A standing for $\{1, g\}$. Then, A creates $a_1 = \{\text{hello}, D(A|Q)\}$, tunes its backward-looking N2N antenna to $d(A|B)$ and triggers downstream $\text{send}(A, a_1)$ aimed at its successor B .

Upon receiving $m_{1,A}$, Q computes $D(Q|B)$, tunes its backward-looking C2C antenna according to $D(Q|B)$ and activates $\text{LgSend}(Q, m_{1,Q})$ aimed at A and B . Q stands for $\{\gamma, g\}$ and message $m_{1,Q} = \{\text{hello}, v\text{slot}(\gamma, \Gamma)\}$. Then, Q creates N2N message $q_1 = \{\text{hello}, D(Q|B)\}$, to be received by its predecessor P . Q tunes its forward-looking N2N antenna to $d(Q|P)$ and triggers upstream $\text{send}(Q, q_1)$.

Pseudonyms shall not be utilized in this round. The weak *PPN* problem arises here, to be solved beforehand (see Subsection 2.3.1). Owing to front/rear sensors, as well as to topological alignment of A and Q , Q cannot mistake A with another NGV, and vice-versa. Owing to the cohort construct, Q and P know in advance that names of A and B can only be $\{1, g\}$ and $\{2, g\}$, respectively. Finally, at the end of round 1, A and B have learned Q 's name $\{\gamma, g\}$, thus P 's name $\{\gamma-1, g\}$.

Certificates are exchanged in rounds 2 and 3. Since vehicle profiles are recorded in certificates, the LgJoin operation presented herein can thwart falsifications of vehicle profiles (LgJoin denied).

By the end of round 1, A , B , Q and P have computed radio range lower bounds $D(\cdot)$ for every subsequent round. They know how to tune their respective C2C antennas when a LgSend primitive must be activated. Given the small duration of a LgJoin operation (see Subsection 7.3), these bounds are valid for the entire execution of a LgJoin.

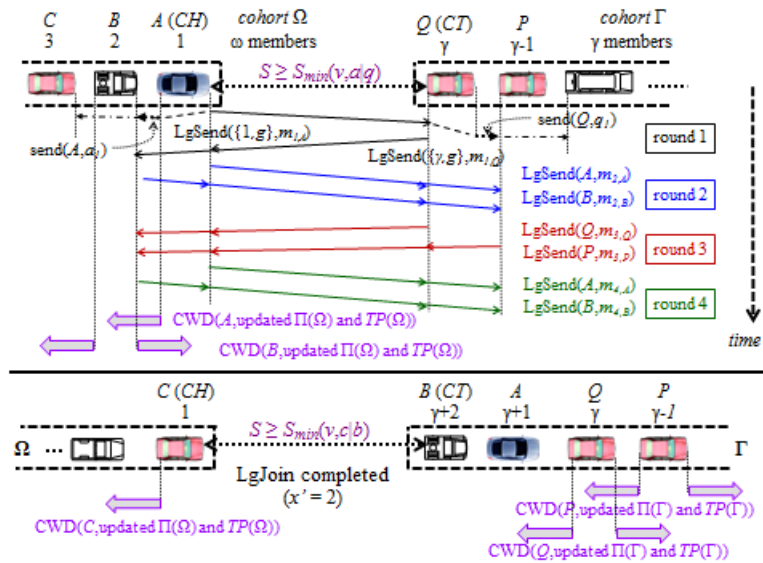


Fig. 9. Illustration of a LgJoin cyberphysical operation

7.1.2 Round 2 ($\Omega \rightarrow \Gamma$): request round

Purpose: A and B ask permission to join cohort Γ .

Upon receiving $m_{1,Q}$, A creates message $m_{2,A} = \{req, \omega, \Phi(\Omega), TP(\Omega)\}$, where req is the C code for “request for a Join”. Then, A triggers $LgSend(A, m_{2,A})$ accompanied with certificate C_A , aimed at Q and P . A stands for $\{1, g, \langle A\psi \rangle\}$. Similarly, upon receiving $m_{1,Q}$, B creates message $m_{2,B} = \{req, \omega, \Phi(\Omega), TP(\Omega)\}$ and triggers $LgSend(B, m_{2,B})$ accompanied with certificate C_B , aimed at Q and P . B stands for $\{2, g, \langle B\psi \rangle\}$.

Under aforementioned assumptions, the C2C message $m_{2,B}$ issued by B is identical to $m_{2,A}$. According to notations introduced in Subsection 5.3.3, pseudonym $\langle A\psi \rangle$ (resp., $\langle B\psi \rangle$) is linked with certificate C_A (resp., with certificate C_B).

At the end of round 2, P and Q have received $m_{2,A}$, C_A , $m_{2,B}$ and C_B . They can compute reply rep .

7.1.3 Round 3 ($\Gamma \rightarrow \Omega$): reply round

Purpose: Q and P reply to A and B .

Let $\Phi_\Gamma(\Gamma) = \Phi(\Gamma) \cap \Phi(\Omega)$ computed by P and by Q .

- Reply rep is set to “no/motive” to prohibit the $LgJoin$ operation. Possible motives are as follows:
 - pseudonym $\langle A\psi \rangle$ or pseudonym $\langle B\psi \rangle$ is found invalid (not both—the 1-out-of-3 assumption)
 - unmet CJ conditions/no lane clearing ($\Phi(\Gamma)$ is the dominating interoperability set)
 - unmet CJ conditions/please clear lane g ($\Phi(\Omega)$ is the dominating interoperability set).
- Reply rep is set to “yes” in all other cases.

Let $m_{3,Q}$ stand for the C2C message to be issued by Q . If $rep \neq$ “yes”, $m_{3,Q} = \{\text{“no/motive”}\}$. Otherwise, $m_{3,Q} = \{\text{“yes”, } x, \Phi(\Gamma), TP(\Gamma), \Phi_\Gamma(\Gamma)\}$, where $x = \min \{n^* - \gamma, \omega\}$ stands for the number of Ω 's members admitted in Γ . Q triggers $LgSend(Q, m_{3,Q})$ along with certificate C_Q . Q stands for $\{\gamma, g, \langle Q\psi \rangle\}$.

Under aforementioned assumptions, the C2C message $m_{3,P}$ issued by P is identical to $m_{3,Q}$. P triggers $LgSend(P, m_{3,P})$ accompanied with certificate C_P . P stands for $\{\gamma - 1, g, \langle P\psi \rangle\}$. Pseudonym $\langle Q\psi \rangle$ (resp., $\langle P\psi \rangle$) is linked with certificate C_Q (resp., with certificate C_P).

At the end of round 3, A and B have received $m_{3,Q}$, C_Q , $m_{3,P}$ and C_P . They can compute decision dec .

7.1.4 Round 4 ($\Omega \rightarrow \Gamma$): decision round

Purpose: A and B return decision dec to Q and P .

Let $\Phi_\Omega(\Gamma) = \Phi(\Gamma) \cap \Phi(\Omega)$ computed by A and by B .

- Decision dec is set to “cancel/motive” in either of the following cases:
 - $rep =$ “yes” but pseudonym $\langle Q\psi \rangle$ or pseudonym $\langle P\psi \rangle$ is found invalid
 - $rep =$ “yes” but $\Phi_\Omega(\Gamma) \neq \Phi_\Gamma(\Gamma)$
 - $rep =$ “no/unmet CJ conditions/no lane clearing
 - $rep =$ “no/unmet CJ conditions/ please clear lane g .
- Otherwise, decision dec is set to “join”.

Let us denote $m_{4,A}$ the C2C message sent by A to Q and P . If $dec \neq$ “join”, then $m_{4,A} = \{\text{“cancel/motive”}\}$. Otherwise, $m_{4,A} = \{\text{“join”, } x, \text{new } \Phi(\Gamma) = \Phi_\Omega(\Gamma)\}$. A triggers $LgSend(A^\circ, m_{4,A})$. Here, A° stands for $\{1, g, \langle A\psi^\circ \rangle\}$, where $\langle A\psi^\circ \rangle$ is a condensed version of $\langle A\psi \rangle$ utilized in round 2.

Under aforementioned assumptions, the C2C message $m_{4,B}$ issued by B is identical to $m_{4,A}$. B triggers $LgSend(B^\circ, m_{4,B})$. B° stands for $\{2, g, \langle B\psi^\circ \rangle\}$, where $\langle B\psi^\circ \rangle$ is a condensed version of $\langle B\psi \rangle$ utilized in round 2.

In addition to verifications of condensed pseudonyms $\langle A\psi^\circ \rangle$ and $\langle B\psi^\circ \rangle$, P and Q may challenge A and B , asking for pair-wise secrets constructed during rounds 2 and 3, thus defeating a potential intruder that would attempt a masquerading cyberattack at the start of round 4. Numerous possibilities exist for building pair-wise secrets. Owing to features proper to N2N messaging, simple schemes can be devised for cohorts.

At the end of round 4, A , B , Q and P share the same decision.

7.2. Termination, the physical part

When $dec = \text{"cancel/motive"}$, Ω and Γ behave according to Policies described in Subsection 6.3.4.

When $dec = \text{"join"}$:

- If $x \neq \omega$, a N2N message carrying {LgJoin granted for x members, updated profile $\Pi(\Omega)$ } is disseminated within Ω . As a result, Ω 's member at rank $x+1$ triggers the split maneuver while members ranked 1, ..., x accelerate until catching up with Γ .

- If $x = \omega$ (full merge), a N2N message carrying {LgJoin granted for Ω } is disseminated. All Ω 's members accelerate until catching up with Γ .

A learns Q 's cp level at the end of round 3. At the end of round 4, A catches up with Q , leading the physical join maneuvers, $S_{min}(v, ch/ct)$ replaced by $s_{min}(v, a/q)$.

Once new members are inserted in Γ , their respective ranks are updated to $\gamma+1, \gamma+2, \dots, \gamma+x$. Q and P activate the CWD primitive within new Γ in order to disseminate a N2N message carrying {LgJoin completed, x new members, updated profile $\Pi(\Gamma)$ }. Fig. 9 shows the particular case where $x = 2$.

7.3. Discussion

Worst-case time bounds for executing the four cyber rounds depend on pseudonym/cryptographic schemes considered (delays incurred with signature verification) and on worst-case MAC-level access delays δ_c for C2C communications (8 channel accesses for LgJoin presented above). Since MAC-level contention experienced with C2C messaging is much lower than contention caused by N2N messaging, bounds δ_c are significantly smaller than bounds δ . Worst-case time bounds for executing the four cyber rounds are in the order of 150 ms. Consequently, distances travelled by NGVs while cyber rounds are being executed are very short, approx. less than one average vehicular slot, for velocities below 180 km/h.

When $dec = \text{"join"}$, once the LgJoin operation is completed, A, B, Q and P are mutually aware of their respective (trustworthy) profiles, thus the symmetrical range-2 neighbors' profiles knowledge property (see Subsection 6.3.6), which is essential for defeating various cyberattacks, masquerading in particular.

Analyses of LgJoin for cases other than those considered in the above (cohorts of size 1, failures and various cyberattacks) are to appear in forthcoming publications.

Mutual trust in cohorts

A LgJoin operation is denied in the presence of malicious requestors or/and responders. That is desired, since no cohort shall admit a misbehaving NGV, and no NGV shall join a cohort that comprises a dishonest member. LtJoin operations are denied for the same reasons. It follows that cohort members can trust each other, provided that members do not behave abnormally *in cyber space or in physical space* after being admitted. An abnormal behavior in cyber space or in physical space is promptly detected by the NGV's PSM (predicate violation—see Subsection 3.1.2). A misbehaving NGVs is halted (emergency lane, safe parking spot) by its CR subsystem, instructed by the CC subsystem. The whole process of physically removing a NGV from the transportation system is described in forthcoming publications.

Since cohort members can trust each other, there is no need to authenticate senders of intra-cohort N2N messages. Consequently, intra-cohort communications do not exhaust certificates stored in a PSM.

8. Gap Control, Safety and Efficiency Properties

Safety proofs consist in demonstrating that the numbers of crashes (experienced over a year, or in specific scenarios) have upper bounds in worst-case conditions—see **principle g**). Our work on CMX functionalities aims at eliminating fatalities and severe injuries caused by all kinds of crashes. The focus of this report is on fatal and severe rear-end crashes (pairwise collisions and pileups). It seems therefore useful to elaborate on the relevance of such crashes.

8.1. On rear-end collisions and pileups

Statistics published by NHTSA for year 2017 in the USA are as follows:

- Property damage: 4,530,000 - Injury: 1,889,000 - Fatality: 34,247

Frontal crashes predominate, followed by lateral crashes, rear-end crashes ranked 3rd. We conjecture that those ratios established for human-driven vehicles, with or without ABS and ADAS capabilities, will change with the emergence of NGVs. Inattention and drowsiness are the major causes of head-on and lateral crashes.

With NGVs, a significant fraction of such crashes shall be eliminated thanks to onboard robotics, N2N/C2C communications, and coordination algorithms. Given that inter-NGV gaps are bound to be small, one may speculate that rear-end crashes due to hard braking may occur more frequently than nowadays if such scenarios are not properly mastered. That is the rationale for the work presented hereafter.

The sliding brick wall paradigm

The “brick wall” paradigm is an extreme model of worst-case hazards in vehicular networks: sudden obstruction by an obstacle (e.g., tree fall, rock avalanches) leading to abrupt stopping, where a vehicle reaches velocity 0 in 0 time unit. Such hazards, considered vanishingly rare, are ignored in published work. In other terms, they belong to the set of residual assumptions (see Subsection 2.3.4). More realistically, it is assumed that any vehicle may be hard braking (irrespective of its rank in a cohort) at any time, leading to multiple brick walls sliding at different velocities.

Theoretical absolute safety

We posit that it is impossible to fully eliminate rear-end crashes. However, crashes that occur at *relative* velocities not higher than some modest threshold v_{rel} shall entail property damage only. Theoretical absolute safety holds when it is proven that *no fatalities, no severe injuries*, can be experienced on the occasion of rear-end crashes, under assumptions of high coverage.

The challenge addressed herein is to express smallest inter-vehicular gaps (highest efficiency) while achieving theoretical absolute safety.

Safety properties in linear formations depend on gap control schemes. Building on gap control devised for pre-planned platoons, we analyze these schemes for strings (spontaneous platoons) and cohorts (strings with ranking and pairwise knowledge of authenticated cp levels). Next, we examine efficiency properties.

8.2. Analyses of safe inter-neighbor gaps in strings and cohorts

The role of cyber levels has been examined in Subsection 6.1. In this subsection, the focus is on physical levels. One must distinguish design time (analyses conducted off-line) from operational time (decisions made on-line, vehicles on the roads). In early publications devoted to truck platooning, gaps denoted $\sigma^\circ(v)$ in Eq. 2 in the order of 2 or 3 meters were considered, since trucks could be arranged in non-decreasing braking powers. That was identified in early NAHSC studies [68]. Such ideal arrangements cannot be assumed with strings or cohorts, due to spontaneity.

8.2.1 Strings

Standard gap control algorithms rest on postulating some braking powers b_l for a leader L and b_f for a follower F , yielding smallest safe gaps $sp_{min}(v,f/l) = \sigma^\circ(v) + [1/b_f - 1/b_l] v^2/2$, $b_l > b_f$.

Rear-end crashes occur when postulated deceleration rate b_l is lower than the real deceleration rate exercised by L or/and when postulated deceleration rate b_f is higher than F 's real deceleration rate. Variable v_{rel} has a coverage that varies according to the following uncertainties:

- Postulated interval $[b_{min}, b_{max}]$ that contains rates b_l and b_f . This interval shall be “small enough” in order to predict “small enough” gaps $sp_{min}(v,f/l)$. Conversely, this interval shall be “large enough” for having a high coverage.

- Postulated approximations ϕ relative to rates b . It is assumed that rate b chosen in interval $[b_{min}, b_{max}]$ may in fact vary from $b-\phi$ to $b+\phi$,

- Approximation error ω made by F relative to b_l . Since one cannot assume that F knows b_l , this rate must be approximated. Error ω stands for the highest difference between b_l assumed to be known by F and highest real rate b_l . As regards theoretical absolute safety, the only case of interest is $\omega > 0$.

Consequently, stricto sensu, smallest safe gaps in extreme conditions are as follows:

$$xsp_{min}(v,f/l) = \sigma^\circ(v) + [1/(b_f - \phi) - 1/(b_l + \phi + \omega)] v^2/2, \quad b_l > b_f.$$

8.2.2 Cohorts

Just like strings, cohorts have their members arranged in arbitrary order of braking powers. Every member is aware of its assigned cyberphysical level $cpl(\cdot)$. We have shown that F knows the cyberphysical level $cpl(L)$ of its predecessor L . Therefore, in addition to λ_f and b_f , F knows b_l . Highest braking power b_l can be trusted, since it has been read in a (valid) certificate provided by L . Notations have been introduced in Subsection 6.1.3.

In the V2X framework, F could be made aware of b_l via a V2X message sent by L . Unfortunately, a falsified b_l can be read in a V2X message generated by a malicious vehicle. Reliance on authentication (signatures of V2X messages are checked to be valid) does not help since bogus data can be found in the body of a V2X message issued by a correctly authenticated vehicle.

In cohorts, unfalsifiable bounds β_j replace hypothetical rates b considered for strings, yielding smallest safe gaps

$$s_{min}(v,f/l) = \sigma^\circ(v) + [1/\beta_f - 1/\beta_l] v^2/2, \quad \beta_l > \beta_f, \quad \sigma^\circ(v) = \lambda_f v + c^\circ.$$

Variable v_{rel} has a coverage that varies according to the following factors:

- Postulated interval $[\beta_0, \beta_5]$ which contains rates β_f and β_l . By definition, there are no uncertainties here.
- Approximations ϕ relative to rates β . By definition, for a NGV assigned pl j , the highest braking power β considered at design time is nominal β_j . Therefore, $\phi = 0$ for β_l . Regarding F , we must consider the lowest braking power β which matches pl j . It follows that approximation ϕ is the error made when equating β with β_f , which is at most 1 m/s^2 , whichever boundary β_f from 4 to 9 is considered.
- F knows β_l . There is no approximation error ω made by F relative to β_l .

Consequently, stricto sensu, smallest safe gaps for extreme conditions are as follows:

$$x_{s_{min}}(v,f/l) = \sigma^\circ(v) + [1/(\beta_f - 1) - 1/\beta_l] v^2/2, \quad \beta_l > \beta_f.$$

8.2.3 Discussion

- Gaps $x_{sp_{min}}(v,f/l)$ are larger than gaps $x_{s_{min}}(v,f/l)$

This is due to variables ϕ and ω . A detailed analysis is out of the scope of this report. Briefly, assuming the same values for β and b , it suffices to consider $\phi = 1$ to gain an insight into a demonstration. Trivially, we have $1/(\beta_f - 1) - 1/\beta_l < 1/(b_f - \phi) - 1/(b_l + \phi + \omega)$, since $1 + \omega > 0$. Moreover, uncertainty ϕ can be greater than 1.

- Types of safety property achieved

With gaps $x_{sp_{min}}(v,f/l)$, variable v_{rel} has an upper bound that depends on postulated rates b_{min} , b_{max} , b_f , b_l , and unbounded approximations ϕ and ω . Therefore:

$$\textit{Probabilistic safety} \text{ is achieved with gaps } sp_{min}(v,f/l) = \sigma^\circ(v) + [1/b_f - 1/b_l] v^2/2, \quad b_l > b_f.$$

With gaps $x_{s_{min}}(v,f/l)$, variable v_{rel} has an upper bound that depends on unquestionable boundaries β_0 , β_5 , trustworthy highest braking powers β_f and β_l , and approximation ϕ bounded by 1. Therefore:

$$\textit{Theoretical absolute safety} \text{ is achieved with gaps } s_{min}(v,f/l) = \sigma^\circ(v) + [1/\beta_f - 1/\beta_l] v^2/2, \quad \beta_x > \beta_y.$$

- Assumptions coverage

Underlying assumptions UA relative to braking powers are summarized in Table II.

Table II. Braking power assumptions

<i>Assumptions</i>	interval considered	approximation of rates	error relative to L 's braking power guessed by F
UAS (strings)	$[b_{min}, b_{max}]$	$\pm \phi$ for b_f , $\pm \phi$ for b_l	ω
UAC (cohorts)	$[\beta_0, \beta_5]$	$< 1 \text{ m/s}^2$ for β_f , 0 for β_l	0

The coverage of assuming interval $[b_{min}, b_{max}]$, $b_{min} > \beta_0$ and $b_{max} < \beta_5$, is necessarily lower than the coverage of assuming the whole $[\beta_0, \beta_5]$ spectrum. The coverage of a nominal bound β_j , measured in well specified conditions, is necessarily higher than the coverage of any postulated braking power b . Neither approximation ϕ nor error ω has guaranteed bounds. Clearly, the coverage of assumptions UAS is necessarily lower than the coverage of assumptions UAC.

The coverage of assumptions that underlie theoretical absolute safety is higher than the coverage of assumptions that underlie probabilistic safety.

This being established, let us revert to standard gaps $sp_{min}(v,f/l)$ and $s_{min}(v,f/l)$ for studying operational time calculations of gaps and efficiency properties.

8.3. On-line calculations of safe inter-neighbor gaps

Simplified formulae for gap control considered in the above can be replaced by more elaborate analytical expressions, such as in [69-72] for example, where stability issues arising in homogeneous and heterogeneous CACC strings are explored. In the CMX framework, analyses aimed at stability issues would be useful for expressing highest inter-neighbor gaps $s_{max}(v,y/x) = (1+\alpha) s_{min}(v,y/x)$, $\alpha > 0$.

Braking powers b and β which appear in analytical gaps $sp_{min}(v,f/l)$ and $s_{min}(v,f/l)$ may not match operational reality. Let \mathcal{B} stand for b or β . Real braking power $\mathcal{B}(t)$ exercised at any time t depends on parameters proper to a vehicle (e.g., tire wear, weight transported) as well as to environmental parameters (e.g., roadway surface, declivity, weather), which parameters can be sensed and calculated on-line by onboard CR subsystems. Let us denote $cs(t)$ the set of parameters sensed and calculated at time t and let us write $\mu(t) = |\mathcal{B}(t) - \mathcal{B}|$.

Algorithmic learning and AI-based techniques could cross fertilize classical techniques aimed at expressing $\mathcal{B}(t)$ out of $\{\mathcal{B} \otimes cs(t)\}$. Assuming that learning about $\mathcal{B}(t)$ is asymptotically sufficiently accurate, uncertainty range $\mu(t)$ would tend towards a sufficiently low value which would yield an accurate bound for $v_{rel}(t)$. This applies equally to CAVs in strings ($\mathcal{B} = b$) and to NGVs in cohorts ($\mathcal{B} = \beta$).

Nonetheless, cohorts have three advantages over strings. One is error ω . In the event of a rear-end crash, relative velocity v_{rel} is necessarily lower when $\omega = 0$. Another one is linked with parameters proper to a vehicle. A time-dependent vehicle status $vs(t)$ can be maintained by an onboard system, which would serve to compute $\mathcal{B}(t)$ out of $\{\mathcal{B} \otimes vs(t)\}$. A vehicle status shall be provided in requests for LtJoin or LgJoin operations, along with a certificate. As a result, every member can be aware of the current highest braking power which can be exercised by a predecessor. A falsified status could be provided by a member behaving as a local adversary. Such malicious behaviors belong to the category of irrational cyberattacks (see Subsection 5.4.2). Finally, when environmental conditions change significantly (detected by onboard sensors), cohort members can run a CWA algorithm in order to agree on values taken by environmental parameters of interest to an entire cohort at current time t . They would therefore compute their respective real $\beta(t)$ out of nominal β consistently. It appears that none of these advantages may exist with V2X functionalities.

For all these reasons discussed in Subsection 8.2 and in this subsection, the probability of having velocity $v_{rel}(t)$ kept below a given value, small enough for entailing property damage only, is higher with cohorts than with strings. The “theoretical” attribute in theoretical absolute safety can be dropped when it is demonstrated that the relationship binding $\mu(t)$ and $v_{rel}(t)$ yields a modest value v_{rel} matching a bound set by public and private bodies (authorities, insurance companies, etc.).

8.4. Safe inter-cohort gaps

It is not sufficient to establish safe intra-cohort gaps. Pileups involving cohorts that follow each other shall be avoided. Stated differently, a cohort must be immune to cyber and physical events that may occur within any other cohort. This is known as the isolation/atomicity property in Distributed Computing. For guaranteeing that CH , head of a cohort, can always stop before hitting CT , tail of a preceding cohort, a smallest *inter-cohort* gap $S(v,ch/ct)$ is enforced by CH 's onboard system fed with data from forward-looking sensors. Since CH is not aware of which cp level is assigned to CT , $S(v,ch/ct)$ is computed assuming *extreme* hard-braking conditions, i.e. CT endowed with some hypothetical braking power hb , higher than bound β_5 given in the CPL matrix. Therefore:

$$S_{min}(v,ch/ct) = \delta_c v + [1/\beta_{ch} - 1/hb] v^2/2,$$

where δ_c stands for the worst-case reaction latency with forward-looking sensors, or for the worst-case MAC-level delay experienced with C2C communications.

8.5. Efficiency properties

Efficiency is derived from the total length “occupied” by $n-1$ gaps in a formation of n members (vehicle lengths not included), denoted SG for strings and CG for cohorts. Our goal is to establish the amount of asphalt space “saved” with cohorts, compared to strings, via the following efficiency ratio reached while achieving theoretical absolute safety:

$$\rho(v) = 1 - CG(v)/SG(v).$$

For fairness, we shall assume the same knowledge at design time in both cases (strings and cohorts). Rates b are replaced by bounds β_j , b and β_j belonging to interval $[\beta_0, \beta_5]$. Moreover, we shall ignore uncertainties proper to strings—thus favoring strings over cohorts—in order to express “pessimistic” ratios $\rho(v)$. (Real ratios can only be higher.)

Efficiency can thus be analyzed by comparing gaps $sp_{min}(v,y/x)$ and $s_{min}(v,y/x)$, Y standing for X 's follower. Since we ignore the role of cyber levels, the same $\sigma^\circ(v) = \lambda v + c^\circ$ appears in gaps $sp_{min}(v,y/x)$ and $s_{min}(v,y/x)$.

In the case of pre-planned platoons, highest efficiency is achieved by having platoon members ranked by non-decreasing β values, leading to $G(v) = (n-1) \sigma^\circ(v)$. In strings and cohorts, a NGV assigned any physical level j comprised between 0 and 5 may “appear” anywhere. A detailed analysis of all possible arrangements—out of the scope of this report—would be of interest conditioned on the availability of sufficiently accurate statistics regarding braking powers of commercialized vehicles, which is not the case. In other words, data is missing for computing an accurate probability of occurrence for every possible pattern. We shall thus focus on best cases and worst cases, for each of the three modes introduced in Subsection 6.2.

8.5.1 Unrestricted heterogeneity

Any β ranging from $\beta_0 = 4$ to $\beta_5 = 9$ may appear in a formation. Let $\sigma^*(v) = \sigma^\circ(v) + [1/\beta_0 - 1/\beta_5] v^2/2$ stand for the gap between a (string, cohort) member labelled β_0 which has a predecessor labelled β_5 .

- Best cases for cohorts

In best-case patterns, members are ranked by non-decreasing β values, thus $CG_{min}(v) = (n-1) \sigma^\circ(v)$. Best-case patterns are unbounded (suites of the same consecutive β of any length may be considered). Best-case patterns range from those where all members are assigned β_0 to those where all members are assigned β_5 . There are approx. n° possible best-case patterns for a cohort of size n . With strings, calculations of gaps for achieving theoretical absolute safety must rest on assuming that every member has a predecessor labelled β_5 . There are two extreme patterns:

- All members are assigned β_0 . Such patterns are worst cases for strings: $SG_{max}(v) = (n-1) \sigma^*(v)$.

In such cases: $\rho(v) = \rho_{max}(v) = 1 - CG_{min}(v)/SG_{max}(v) = 1 - \sigma^\circ(v)/\sigma^*(v)$.

- All members are assigned β_5 . These are the only best case patterns for strings (for any given n , there is only 1 instantiation of this pattern), where $SG(v) = SG_{min}(v) = (n-1) \sigma^\circ(v)$.

In such cases: $\rho(v) = \rho_{min}(v) = 0$.

- Worst cases for cohorts

In a worst-case pattern, members are ranked by decreasing β values. Owing to the impossibility of having 2 neighbors assigned the same β , such patterns are bounded, including at most 6 members (1 member tagged with β_5 , followed by 1 member tagged with β_4 , ..., ending with 1 member tagged with β_0). These longest patterns may be repeated as such or any subset thereof, to encompass n members.

The absolute worst-case pattern is instantiated with consecutive pairs $\{\beta_0, \beta_5\}$ (β_5 first), length wg . The proof is by contradiction. Consider such a pattern and pretend that some other pattern has length $wg' > wg$. In any pair, let us replace β_5 by $\beta' < \beta_5$. The gap between β' and downstream β_0 is reduced, since $1/\beta' > 1/\beta_5$. The gap between β' and upstream β_0 remains unchanged, equal to $\sigma^\circ(v)$. Thus, resulting $wg' < wg$. Let us now replace some β_0 by $\beta' > \beta_0$. The gap between β' and downstream β_5 remains unchanged, equal to $\sigma^\circ(v)$. Given that $1/\beta' < 1/\beta_0$, the gap between β' and upstream β_5 is reduced, leading to $wg' < wg$. Other patterns can only be built out of similar replacements. End of proof.

Therefore, in absolute worst-cases, highest $CG(v) = CG_{max}(v) = \lceil (n-1)/2 \rceil \sigma^*(v) + \lfloor (n-1)/2 \rfloor \sigma^\circ(v)$.

With strings, the same patterns of consecutive pairs $\{\beta_0, \beta_5\}$ lead to the same result. Consequently:

$$\rho(v) = \rho_{min}(v) = 0.$$

- Conclusion and numerical illustration

Numerical results are shown in Table III. Real ratios $\rho(v)$ for every possible pattern range between bounds higher than $\rho_{min}(v)$ and $\rho_{max}(v)$ since we have expressed “pessimistic” ratios $\rho(v)$.

Assume $c^\circ = 1.2$ m and $n = 17$. Since $\rho_{max}(v)$ is minimized with highest $\sigma^\circ(v)$, let us choose highest λ which favors strings, i.e. $\lambda_0 = 12 \sqrt{5}$, which is slightly higher than $\delta = 20$, a value commonly considered for strings. Results in Table III are given for high velocities, $v = 30$ m/s, and for somewhat congested traffic conditions, where $v = 15$ m/s. CG and SG are in meters (centimeters are ignored).

One finds $\sigma^\circ(30) = 2$ m, $\sigma^*(30) = 64.5$ m, $\sigma^\circ(15) = 1.6$ m, $\sigma^*(15) = 17.2$ m.

Table III. Efficiency ratios $\rho(v)$ in the unrestricted heterogeneity mode

$v = 30$	Absolute worst case for cohorts	Best cases for cohorts (β_0 pattern for strings)	$v = 15$	Absolute worst case for cohorts	Best cases for cohorts (β_0 pattern for strings)
cohorts	$CG_{max} = 532$	$CG_{min} = 32$	cohorts	$CG_{max} = 150.6$	$CG_{min} = 25.6$
strings	$SG_{max} = 532$	$SG_{max} = 1,032$	strings	$SG_{max} = 150.6$	$SG_{max} = 275.6$
$\rho(30)$	$\rho_{min} = 0$	$\rho_{max} = 0.97$	$\rho(15)$	$\rho_{min} = 0$	$\rho_{max} = 0.9$

One could suspect that high ratios ρ_{max} result from the need to consider the whole β spectrum from $\beta_0 = 4$ to $\beta_5 = 9$. That is not the case, as shown below.

8.5.2 Restricted heterogeneity and strict homogeneity

With cohorts, restricted heterogeneity and strict homogeneity can be enforced dynamically via controlled admissions, the result being that all β s belong to a specific bounded interval $[\beta_{low}, \beta_{high}]$ which is known to all members. Strings examined in published work operate under the unrestricted heterogeneity mode. Since the other two operative modes explored here are inaccessible to string formations, calculations of efficiency ratios ρ_{max} are of theoretical interest. It may however be useful to show that those ratios are far from being negligible, even when considering a small subset of the whole β spectrum.

According to the definition of restricted heterogeneity, two cases shall be considered (see Fig. 7):

- Case 1: Sets $\Phi(\cdot)$ of size 4 and vertical sets $\Phi(\cdot)$ of size 6, where $\beta_{high} - \beta_{low} = 1$,
- Case 2: Sets $\Phi(\cdot)$ of size 9 and horizontal sets $\Phi(\cdot)$ of size 6, where $\beta_{high} - \beta_{low} = 2$.

With the same numerical values as above for $\sigma^\circ(v)$, we have:

- Smallest β (case 1.1): $\beta_{low} = \beta_0 = 4 // \beta_{high} = \beta_1 = 5$. One finds $\sigma^*(30) = 24.5$ and $\sigma^*(15) = 7.2$.
- Highest β (case 1.2): $\beta_{low} = \beta_4 = 8 // \beta_{high} = \beta_5 = 9$. One finds $\sigma^*(30) = 8.25$ and $\sigma^*(15) = 3.2$.
- Smallest β (case 2.1): $\beta_{low} = \beta_0 = 4 // \beta_{high} = \beta_2 = 6$. One finds $\sigma^*(30) = 39.5$ and $\sigma^*(15) = 11$.
- Highest β (case 2.2): $\beta_{low} = \beta_4 = 7 // \beta_{high} = \beta_5 = 9$. One finds $\sigma^*(30) = 16.3$ and $\sigma^*(15) = 5.2$.

Corresponding ratios $\rho_{max}(v) = 1 - \sigma^\circ(v)/\sigma^*(v)$ are given in Table IV. Ratios $\rho_{max}(v)$ for other pairs $[\beta_{low}, \beta_{high}]$ range between values shown for .1 and .2 sub-cases.

Table IV. Efficiency ratios $\rho_{max}(v)$ in the restricted heterogeneity mode

$v = 30$	Case 1.1	Case 1.2	$v = 15$	Case 1.1	Case 1.2
$\rho_{max}(30)$	0.92	0.76	$\rho_{max}(15)$	0.78	0.50

$v = 30$	Case 2.1	Case 2.2	$v = 15$	Case 2.1	Case 2.2
$\rho_{max}(30)$	0.95	0.88	$\rho_{max}(15)$	0.85	0.69

We have assumed a unique λ to establish the above results. A unique λ and a unique β are features proper to strictly homogeneous cohorts. Since they are characterized by interoperability sets $\Phi(\cdot)$ that contain only one element k , all gaps are equal to $\sigma^\circ(v)$, whatever λ , whatever β . No safe gaps can be smaller when the goal is theoretical absolute safety. Strictly homogeneous cohorts instantiate best case patterns *dynamically*.

8.5.3 Expected ratios $\rho(v)$

As shown above, gap saving ratios achieved with cohorts compared to strings can be quite significant. Inter-neighbor gaps may vary between $s_{min}(v,y/x)$ and $s_{max}(v,y/x) = (1+\alpha) s_{min}(v,y/x)$, $\alpha > 0$. Lowest values of $\rho(v)$ are reached with inter-neighbor gaps $s_{max}(v,y/x)$. A fraction of uncertainties that are germane to on-line calculations of braking powers $\beta(t)$ could also be “encapsulated” in parameter α .

Average ratios $\rho(v)$ could be computed out of the exact proportions of commercialized vehicles endowed with highest braking powers from β_0 to β_5 . It does not seem unreasonable to speculate that those highest braking powers would match a Poisson distribution. Then, besides accurate calculations of ratios $\rho(v)$, one could also compute the coverage of Poisson distributions truncated as per the $[\beta_{\min}, \beta_{\max}]$ intervals. Without that knowledge, one may only speculate that the population of vehicles is biased towards low j values (more low cost vehicles than expensive ones) such as, e.g., $j = 2$ ($\beta = 6$).

Intuitively, this indicates that values of $\rho(v)$ most frequently achieved on the roads are closer to $\rho_{\max}(v)$ than to $\rho_{\min}(v)$. This remark is strengthened by the fact that probabilities of occurrence of 1 out of approx. n^6 best-case patterns for cohorts of n members are significantly higher than the probability of occurrence of the single best-case pattern for a string of n members. Consequently, expected savings in asphalt occupancy in the order of 60% seem to be fair estimates.

8.6. Antennas range control

C2C messages and N2N messages are sent and received via forward-looking and backward-looking antennas located in the front and in the back of vehicles. Ranges $D(\cdot|\cdot)$ and $d(\cdot|\cdot)$ given below are lower bounds of radio ranges, derived from inter-vehicular gaps and vehicle length (found in vehicle profiles).

C2C communications and LgSend primitives

Notations have been introduced in Section 7. Via its sensors (radars, cameras, etc.), A detects Q and sets its velocity to v , standing $D(A|Q)$ away from Q . $D(A|Q)$ is monitored by Q 's backward-looking sensors.

- Round 1, range-1 $A \rightarrow Q$: $D(A|Q) \geq S_{\min}(v, a|q)$.
- Round 1, range-2 $Q \rightarrow B$: $D(Q|B) = D(A|Q) + L_{B/A}$. Q learns $D(A|Q)$ and $L_{B/A}$ via $m_{I,A}$.
- Round 2, range-2 $A \rightarrow P$: $D(A|P) = D(A|Q) + \text{vslot}(\gamma, \Gamma)$. A learns $\text{vslot}(\gamma, \Gamma)$ via $m_{I,Q}$.
- Round 2, range-3 $B \rightarrow P$: $D(B|P) = L_{B/A} + D(A|Q) + \text{vslot}(\gamma, \Gamma)$. B learns $\text{vslot}(\gamma, \Gamma)$ via $m_{I,Q}$.
- Round 3, range-2 $Q \rightarrow B$: $D(Q|B)$ as in round 1.
- Round 3, range-3 $P \rightarrow B$: $D(P|B) = \text{vslot}(\gamma, \Gamma) + D(Q|B)$. P learns $D(Q|B)$ via message q_I .
- Round 4, range-2 $A \rightarrow P$: $D(A|P)$ as in round 2.
- Round 4, range-3 $B \rightarrow P$: $D(B|P)$ as in round 2.

In Subsection 7.2, we have seen that when LgJoin is completed, A , B , Q and P are mutually aware of their respective (unfalsified) profiles, thus the range-2 shared knowledge of neighbors' profiles (Subsection 6.3.6).

N2N communications and Send primitives

Consider X followed by Y . Gap $s(v, y/x) \geq s_{\min}(v, y/x)$, which is enforced by Y , is monitored by X 's backward-looking sensors. For range-1 N2N communications, $d(Y|X) = s(v, y/x)$. For example, in LgJoin round 1, for range-1 send $A \rightarrow B$, $d(A|B) = s(v, b/a)$, and $d(Q|P) = s(v, q/p)$ for range-1 send $Q \rightarrow P$.

Range-2 N2N communications can be computed thanks to the symmetrical range-2 gap knowledge property (Subsection 6.3.6). For example, considering contiguous neighbors W , X and Y as in Fig. 4, Y can reach W 's N2N antenna by tuning its own N2N antenna according to range-2 $d(Y|W) = s(v, y/x) + \text{vslot}(x, \cdot)$.

8.7. Conclusion

The cyberphysical evidence can be illustrated with three examples. *Cyber collisions* on radio channels that may lead to *physical collisions* on the roads are eliminated with deterministic MAC protocols. Owing to the cohort split scheme, *failures in cyber space* (message losses) do not result in *failures in physical space* (accidents). Thanks to PSMs and predicates, *misbehaviors in cyber space* result in *removals from the physical space* (details to appear).

Quantified safety and efficiency properties established in this report rest on the following CMX pillars:

- The cohort construct (a single NGV is a particular case),
- Cyberphysical levels,
- PSMs and vehicle profiles,
- Controlled cohort admission.

These pillars cross fertilize results established by the robotics community. Vehicle cyberphysical levels are recorded in certificates stored in PSMs, thus unfalsifiable. The PSM concept was introduced first in [66] under the name of safety-critical tamper proof device. Cyberphysical levels also appear on vehicle plates. They can be read by vehicles not equipped with radio capabilities.

We have fulfilled our initial goal. Theoretical absolute safety and high efficiency can be achieved, under assumptions of coverage higher than the coverage of assumptions underlying probabilistic safety.

In cohort-structured vehicular networks, theoretical absolute safety and remarkable savings in asphalt occupancy can be achieved jointly, by design.

The cyberphysical perspective is ignored in the V2X framework where, to the best of our knowledge, none of the basic requirements listed in Subsection 2.1 are met. In the V2X “world”, there is no real inter-vehicular “cooperation”. Vehicles send messages. Vehicles receive some or all of them and decide what to do, unilaterally. There are no explicit requests, no explicit feedback, no explicit inter-vehicular agreements. There are no demonstrated worst-case finite bounds crucial for safety and efficiency, such as time bounds for channel access, time bounds for fault-tolerant coordination, bounds on string length, and bounds on the number of failures to be sustained, no demonstrated theoretical absolute safety. Schemes devised for privacy and cybersecurity are fragile. Pitfalls of the current V2X framework regarding privacy have also been pointed out by the UNECE World Forum for Harmonization of Vehicle Regulations (WP 29) in [73]. The CMX framework appears to be in line with WP 29’s recommendations (see below).

Some topics not addressed fully in this report are shortly discussed hereunder.

9. Human Factors, Cyber-surveillance and Liabilities

9.1. Human factors

With existing AVs and CAVs, human vigilance is mandatory. A driver/passenger must be prepared to intervene “whenever necessary”. In critical scenarios, acceptable reaction latencies are in the order of 1 or 2 seconds. It is highly unlikely that within such latencies, a distracted human *driver* might (1) understand “what is going on”, (2) make a correct decision. By definition, *passengers* are out of the driving loop. However, in a patent filed in 2015, one reads: “If the passenger identifies an emergency situation, the passenger may take control of the vehicle immediately. For example, passenger may see an obstacle which computer has not identified...”. Many questions remain unanswered. How can a “passenger” know for sure that an obstacle has not been identified by the onboard “computer” (onboard system)? What if a “passenger” intervention results in an accident, when inspection of the EDR reveals that in fact the computer was in full control of the situation?

Issues arising with human factors and authority sharing (should a human or some automaton be in charge of handling exceptional conditions) are extremely varied and complex, most of them still open, even in mature domains such as air transportation and defense systems. With AVs and CAVs, we eventually have to face the same problems that have surfaced with automated flying: “automation addiction” has eroded pilots flying skills to the point that pilots may not recall how to recover from a loss of control by a flight management system.

When humans are in the driving loop (progressive approaches to automated driving), it remains to be seen whether safety can be proven to hold. Theoretical absolute safety is beyond reach. A lesson drawn from this report is that disruptive approaches (humans out of the driving loop) are worth considering. Fully automated driving is feasible and theoretical absolute safety is a reachable goal. These results may help in shortening the transition period during which CAVs below SAE level 5 are a source of concerns regarding safety.

9.2. Protection against undesired external and internal cyber-surveillance

Privacy threats and illegitimate cyber-surveillance are being increasingly considered morally and ethically unacceptable. Early work on privacy and security issues in vehicular networks [74-76] has laid the ground for solutions aimed at combating privacy threats conducted via short/medium range broadcast communications. The use of pseudonyms and issues of pseudonymous authentication were addressed shortly after (see [77] for example). The cyber stealth mode serves to protect NGV passengers from *external* privacy threats and illegitimate cyber-surveillance performed through V2X communications.

That is not sufficient. It is necessary to show how to defeat *internal* privacy threats and illegitimate cyber-surveillance performed within vehicles via various sensors (e.g., microphones, cameras). The usual argumentation (“this is for drivers/passengers safety”) is more than questionable with AVs or CAVs of SAE

levels below 5, since humans are unable to react *correctly* in sufficiently small delays (see above). This argumentation is obviously flawed with NGVs (SAE level 5). *Fully automated vehicles do not require, do not enable, human interventions*. There is thus no need to keep passengers of NGVs under watch for safety motivations. Internal privacy protection shall be offered to passengers, whereby internal sensors are deactivated after control checks have been passed (protection against thefts, etc.).

Internal cyber-surveillance is justified in the case of public transportation (buses, etc.). As regards rented vehicles, one may speculate that rental companies prepared to operate privacy-preserving vehicles would enjoy a competitive advantage. Ditto with shared vehicles and taxis that would offer pre-paid rides (users are not anonymous). Consequently, for ethical and commercial reasons, NGVs shall be equipped with the following two options similar to the classical stop/start option, which passengers may select as desired:

- *External cyber-surveillance on/off, off meaning V2X communications in cyber stealth mode.*
- *Internal cyber-surveillance on/off, off meaning that internal privacy protection is on.*

Those *on/off* options are in line with the European GDPR ruling [78], whereby no data of direct or indirect relevance to privacy can be collected without explicit consent. Users shall be able to tell in a straightforward manner how they want to set these two options, forever or on a trip basis. The auto industry shall not restrict users' freedom of choice. Insurance companies may charge additional fees for enabling the *off* option(s). That is fair. Detection of events harmful to passengers such as, e.g., heart failure or hijacking, is another common justification for unlimited cyber-surveillance. In such cases, the *on* state would automatically override the *off* state (details are out of the scope of this report). These *on/off* options are also in line with a number of recommendations from lawyers, such as, e.g., "privacy issues must be systematically addressed in advance—before autonomous vehicles become consumer products" [79].

Social acceptability may be another reason for a standardization of these two options. Public awareness of privacy threats is increasing with time [80]. Users of NGVs aware of the risk of eavesdropping and tracking with V2X communications may be reluctant to buy, to rent, or to ride a vehicle without the *on/off* options.

9.3. Liabilities

It might be worth noting that the proposed categorization in cyberphysical levels reveals an essential difference with standard results established by the roboticists' community, in relationship to liabilities. Rear-end crashes may occur due to imprecise knowledge of braking powers. Implicitly, roboticists "embed" these approximations into their (correct) formulae for gap control. In the event of a harmful crash (involving fatalities or/and severe injuries), they could be held liable for having made inappropriate assumptions regarding braking powers. De facto, they bear responsibilities that should be shared with or assumed by others, namely, the automotive industry. Vehicle manufacturers shall be requested to provide the highest β attached to a vehicle, submitted to certification by independent bodies. It would then be possible to assign responsibilities appropriately whenever necessary, since actual conditions that lead to a crash are recorded in EDRs.

Legislation is undergoing profound changes. No one really knows yet who may/shall be held legally responsible in the event of a harmful crash. Given that a non-negligible fraction of the automotive industry knows the limitations of V2X functionalities, stakeholders are unwilling to bear financial losses incurred with accidents caused by AVs or CAVs. Partially automated assisted driving (under human vigilance) is safer from a financial perspective: when an accident occurs, drivers/passengers can be blamed, no matter what. The most agile players are nevertheless getting prepared for the days when "blaming drivers/passengers no matter what" is not considered acceptable any longer. Work on CMX functionalities reported herein may help those agile players to shape up the future of safe fully automated driving.

10. Conclusions and Perspectives

Throughout this report, we have shown how CMX functionalities based on a distributed mobile edge computing model solve the SE (safety, efficiency) problem as it arises in heterogeneous cohorts of NGVs. The need for design frameworks, protocols and algorithms based on solid scientific foundations, as well as worst-case analyses, has already been stressed by various authors. In upcoming publications, we show how to solve some important problems, such as:

- Naming. It has been shown that naming is harder than consensus in asynchronous systems/networks [81]. As pointed out in Subsection 2.3.1, the privacy preserving naming problem that arises in AVNs is even harder.

The power of UPFs (unambiguous physical functions), simple variations of PUFs [82], is particularly interesting.

- Privacy and cybersecurity in cohort-structured AVNs. A preview can be found in Subsection 5.4.
- Safe lateral Join (LtJoin) operation, an instance of Non-Blocking Atomic Commit [83] in cyberphysics.
- Deterministic MAC protocols for lateral N2N radio communications and for semi-omnidirectional C2C radio communications.
- Theoretical absolute safety and high efficiency in contexts and scenarios prone to lateral collisions, such as in fully automated zipper merging.
- The SPEC problem as it arises with entrances and traversals of UX settings of arbitrary topologies.

Solutions to the above-mentioned problems do not rest on “guessing” or “daring”. Numerous issues remain open. For example, in our on-going work, we address the SPEC problem posed by entrances and traversals of UR settings of complex topologies, such as Place de l’Etoile-Charles-de-Gaulle in Paris (12 arterials). Efficient collision-free crossings in dense traffic conditions at rush hours are reasonably well mastered by human drivers. No solutions have been proposed yet for fully automated vehicles. The vehicular flock construct may prove to be relevant here.

The CMX framework lends itself to integration of advances in various fields, such as encryption chips (for PSMs) and wireless radio technologies, which are in constant evolution [84]. Furthermore, fully automated driving involves specific moral and ethical challenges, notably in urban settings where adult pedestrians, young children, impaired persons, bicyclists, animals in some cities and countries, co-exist with NGVs [85]. The SPEC problem in 3D settings (free flocking or organized skies for networks of automated aerial vehicles?) must also be addressed.

Thanks to the pioneering work conducted by the ITS community and important results established by the scientific community over the last 40 years, as well as significant advances accomplished by the computing and communication industries, we now have a better understanding of how to address the overall SPEC problem as it arises with fully automated vehicles in autonomic vehicular networks.

Scientists, engineers, lawyers, experts in social and human sciences, the automotive industry, the hardware and software industry, the telecommunication industry, certification bodies, standard-making organizations, insurance companies, and national/supranational judicial administrations shall join forces for translating into reality solutions drawn from or inspired by the CMX framework.

Safety and efficiency properties achieved with onboard robotics can be degraded by V2X functionalities, due to cyberattacks. Furthermore, privacy threats inexistent with onboard robotics become a concern with V2X functionalities.

CMX functionalities augment safety and efficiency properties achieved with onboard robotics significantly (theoretical absolute safety is feasible). CMX functionalities are immune to remote cyberattacks. Nearby cyberattacks are inoffensive. There are no privacy risks, other than existing risks incurred with direct vision (by humans, cameras, etc.).

We have to decide now: In which type of motorized society do we want to live?

Acknowledgments

I would like to thank Jonathan Petit and Fawzi Nashashibi for their constructive comments and for our discussions on cybersecurity and gap control issues. I am grateful to Steve Shladover for bringing early NAHSC studies to my attention. I also want to give credit to the scientists and engineers whose pioneering works have been valuable keystones for my research. In return, hopefully, along with new generations of innovators, they may find the CMX framework helpful for their future endeavors.

Glossary

AV: autonomous vehicle CAV: connected automated vehicle
 NGV: next-generation vehicle AVN: autonomic vehicular network
 SAE: Society of Automotive Engineers C-ITS: Cooperative Intelligent Transport Systems
 RC-ITS: C-ITS sub-group which defines robotics functionalities
 CC-ITS: C-ITS sub-group which defines V2X functionalities
 C2C-CC: Car 2 Car Communication Consortium 5GAA: 5G Automotive Association
 V2X: vehicle-to-everything DSRC-V2X: Wi-Fi V2X C-V2X: cellular V2X
 SPEC: safety, privacy, efficiency, cybersecurity CMX: coordinated mobility for S, P, E, C
 LOS: line-of-sight NLOS: non line-of-sight
 GNSS: Global Navigation Satellite System UTC: universal time coordinates
 CAM: cooperative awareness message BSM: basic safety message
 DENM: decentralized environmental notification message
 PB: periodic beaconing LDM: local dynamic map
 C: critical NC: non-critical EDR: event data recorder
 N2N: neighbor-to-neighbor C2C: cohort-to-cohort VLC: visible light communications
 HSM: hardware security module PSM: proactive security module
 MAC: multi-access control CWD: cohort-wide dissemination CWA: cohort-wide agreement
 UX: unsignalized intersection UR: unsignalized roundabout
 CH: cohort head CT: cohort tail
 v : velocity v^* : highest authorized velocity
 n : number of members in a cohort n^* : highest number of members in a cohort
 r : rank in a cohort g : lane occupied
 Identifiers of vehicle R :
 - Pair $\{r, g\}$ when R is the sender of a N2N message
 - Triple $\{r, g, \langle R\psi \rangle\}$ when R is the sender of a C2C message
 $R\psi$: pseudonym used by R
 Send(R, m): R sends m to 2 nearest (upstream or downstream) neighbors
 Receive(R, m): R processes m received from 2 nearest (upstream or downstream) neighbors
 send(R, m): R sends m to its nearest (upstream or downstream) range-1 neighbor
 LtJoin: operation for joining in a cohort laterally
 LgJoin: operation for joining in a cohort longitudinally
 LgSend(R, m): primitive for sending a C2C message longitudinally, activated in the course of a LgJoin
 θ : channel slot duration δ : worst-case channel access delay
 h : spatial reuse (in number of contiguous NGVs within N2N antenna radio range)
 f : number of losses experienced while executing CWD or CWA f^* : upper bound of f
 u^* : highest number of losses experienced on a N2N link without incurring a cohort split
 $\Delta(n, f)$: worst-case bound on CWD termination delays
 cl i : cyber level, $i \in [0, 5]$ pl j : physical level, $j \in [0, 5]$ k : cyberphysical element, $k = 6i + j$

k_X : cyberphysical element assigned to NGV X $cpl(X)$: cyberphysical level of NGV X
 δ_i : channel access boundary for a cl i NGV
 λ_i : inter-neighbor communication delay boundary for a cl i NGV
 $\Lambda_i(\tilde{n})$: delay boundary for a cl i cohort-wide dissemination algorithm
 \tilde{n} : supremum of n^* values for all settings
 β_j : braking power boundary for a pl j NGV
 $vp(X)$: X 's vehicle profile $vp(\cdot) = \{cpl(\cdot), aul(\cdot), vt(\cdot), vl(\cdot)\}$
 $aul(\cdot)$: SAE automated driving level $vt(\cdot)$: vehicle type $vl(\cdot)$: vehicle length
 $\phi(k)$: set of cp levels interoperable with cp level k $\Phi(\Gamma)$: Γ 's current cp level interoperability set
 $TP(\Gamma)$: current topology of cohort Γ $\Pi(\Gamma)$: current profile of cohort Γ
 $vslot(y, \Gamma)$: vehicular slot for Y member of cohort Γ v_{rel} : relative velocity at collision time
 $sp_{min}(v, y/x)$: smallest inter-vehicular gap kept by Y which follows X in a string
 $s_{min}(v, y/x)$: smallest safe inter-vehicular gap kept by Y which follows X in a cohort
 $s_{max}(v, y/x)$: highest inter-vehicular gap kept by Y which follows X in a cohort
 $S_{min}(v, ch/ct)$: smallest safe inter-cohort gap kept by CH which follows CT at velocity v
 $\rho(v)$: ratio of asphalt space "saved" with cohorts, compared to strings
 $D(X|Y)$: lower bound of antenna radio range for C2C messaging
 $d(X|Y)$: lower bound of antenna radio range for N2N messaging

Bibliography

- [1] R.J. Caudill and W.L. Garrard, "Vehicle-follower longitudinal control for automated transit vehicles", *Trans. ASME Journal of Dynamic Systems, Measurement, and Control*, vol. 99, 1977, 241-248.
- [2] S.E. Shladover, "Longitudinal control of automated guideway transit vehicles within platoons", *ASME Journal of Dynamic Systems, Measurement and Control*, vol. 100(4), 1978, 291-297.
- [3] S.E. Shladover, et al., "Automatic vehicle control developments in the PATH Program," *IEEE Trans. on Vehicular Technology*, vol. 40(1), Feb. 1991, 114-130.
- [4] <https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety>
- [5] J.H. Lala and R.E. Harper, "Architectural principles for safety-critical real-time applications", *Proceedings of the IEEE*, vol. 82(1), Jan. 1994, 25-40.
- [6] J.C. Knight, "Safety critical systems: Challenges and directions", *ACM Intl. Conference on Software Engineering*, May 2002, Orlando, USA, 547-550.
- [7] P. Koopman and M. Wagner, "Autonomous vehicle safety: An interdisciplinary challenge", *IEEE Intelligent Transportation Systems Magazine*, vol. 9(1), Spring 2017, 90-96.
- [8] K. Pawlikowski et al., "On credibility of simulation studies of telecommunication networks", *IEEE Communications Magazine*, vol. 40(1), 2002, 132-139.
- [9] J. Rushby, "Just-in-time certification", Best Paper award, 12th IEEE Intl. Conference on the Engineering of Complex Computer Systems (ICECCS), Auckland, New Zealand, July 2007, 15-24.
- [10] G. Le Lann, "An analysis of the Ariane 5 Flight 501 failure – A system engineering perspective", *IEEE Intl. Conference on the Engineering of Computer-Based Systems*, Monterey, USA, March 1997, 339-346.
- [11] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "On a formal model of safe and scalable self-driving cars", *Mobileye Document*, 2017, revised Oct. 2018, 37 p. <https://arxiv.org/pdf/1708.06374.pdf>
- [12] N.A. Lynch, *Distributed Algorithms*, Morgan Kaufmann, March 1996, 907 p.
- [13] B. Alpern and F.B. Schneider, "Recognizing safety and liveness", *Distributed Computing*, vol. 2(3), Springer, Sept. 1987, 117-126.
- [14] A. Avizienis et al., "Basic concepts and taxonomy of dependable and secure computing", *IEEE Transactions on Dependable and Secure Computing*, vol. 1(1), 2004, 11-33.
- [15] G. Le Lann, "Distributed systems—Towards a formal approach", *IFIP Congress*, Toronto, Canada, 1977, North-Holland/Elsevier, 155-160.
- [16] M.J. Fischer, N.D. Griffeth, and N.A. Lynch, "Global states of a distributed system", *IEEE Transactions on Software Engineering*, vol. SE-8(3), May 1982, 198-202.
- [17] K. Mani Chandy and L. Lamport, "Distributed snapshots: Determining global states of distributed systems", *ACM Transactions on Computer Systems*, vol. 3(1), Feb. 1985, 63-75.
- [18] V. Hadzilacos and S. Toueg, "Reliable broadcast and related problems", *Distributed Systems*, ACM Press, New-York, 1993, 97-145.
- [19] M.J. Fischer, N.A. Lynch, and M.S. Paterson, "Impossibility of distributed consensus with one faulty process", *Journal of the ACM*, vol. 32(2), April 1985, 374-382.
- [20] U. Schmid, B. Weiss, and I. Keidar, "Impossibility results and lower bounds for consensus under link failures", *SIAM Journal of Computing*, vol. 38(5), Jan. 2009, 1912-1951.
- [21] J. Petit et al., "Pseudonym schemes in vehicular networks: A survey", *IEEE Communications Surveys and Tutorials*, vol. 17(1), 2015, 228-255.
- [22] D. Eckhoff and C. Sommer, "Driving for big data? Privacy concerns in vehicular networking", *IEEE Security & Privacy Conference*, vol. 12(1), 2014, 77-79.
- [23] B. Wiedersheim et al., "Privacy in inter-vehicular networks: Why simple pseudonym change is not enough", 7th IEEE/IFIP Intl. Conference on Wireless On-demand Network Systems, 2010, 176-183.
- [24] L. Buttyàn et al., "SLOW: A practical pseudonym changing scheme for location privacy in VANETs", 1st IEEE Vehicular Networking Conference, Oct. 2009, Tokyo, Japan, 1-8.

- [25] S. Lefèvre et al., “Impact of V2X privacy strategies on intersection collision avoidance systems”, 5th IEEE Vehicular Networking Conference, Dec. 2013, Boston, USA, 71-78.
- [26] J. Gray and D.P. Siewiorek, “High-availability computer systems”, IEEE Computer, vol. 24(9), Sept. 1991, 39-48.
- [27] J. Rushby, “Critical system properties: Survey and taxonomy”, Computer Science Laboratory, SRI International, Technical Report CSL-93-01, May 1993, Revised February 1994, 64 p. Original version published in Reliability Engineering and System Safety, vol. 43(2), 1994, 189-219.
- [28] D. Powell, “Failure mode assumptions and assumption coverage”, 22nd IEEE Intl. Symposium on Fault-Tolerant Computing (FTCS-22), July 1992, 386-395.
- [29] N. Kalra and S.M. Paddock, “Driving to safety—How many miles of driving would it take to demonstrate autonomous vehicle reliability?”, Rand Corporation Report RR1478, 2016, 14 p.
- [30] R. Toledo-Moreo, D. Bétaille, and F. Peyret, “Lane-level integrity provision for navigation and map matching with GNSS, dead reckoning, and enhanced maps”, IEEE Trans. Intelligent Transportation Systems, vol. 11(1), March 2010, 100-112.
- [31] P.R. Lewis et al., “Architectural aspects of self-aware and self-expressive computing systems: From psychology to engineering”, IEEE Computer, Aug. 2015, 62-70.
- [32] N. Dutt, A. Jantsch, and S. Sarma, “Toward smart embedded systems: A self-aware system-on-chip (SoC) perspective”, ACM Transactions on Embedded Computing Systems, vol. 15(2), Article 22, Feb. 2016, 22/1-22/27.
- [33] S.E. Lee et al., “A comprehensive examination of naturalistic lane changes”, NHTSA DOT Technical Report HS 809 702, March 2004, 213 p.
- [34] L. Lamport, R. Shostak, and M. Pease, “The Byzantine Generals problem”, ACM Transactions on Programming Languages and Systems, vol. 4(3), July 1982, 382-401.
- [35] L. Kleinrock and F. Tobagi, “Packet switching in radio channels: Part I - Carrier sense multiple-access modes and their throughput-delay characteristics”, IEEE Transactions on Communications, vol. 23(12), Dec. 1975, 1400-1416.
- [36] D.P. Bertsekas and R.G. Gallager, Data Networks, Prentice-Hall, 2nd Edition, 1992, ISBN 0-13-200916-1, Multiaccess Schemes, Chapter 4, 271-362. <http://web.mit.edu/dimitrib/www/datanets.html>
- [37] 5G PPP White Paper, “5G Automotive Vision”, Oct. 2015, 67 p.
<https://5g-ppp.eu/wp-content/uploads/2014/02/5G-PPP-White-Paper-on-Automotive-Vertical-Sectors.pdf>
- [38] Y. Yao et al., “Delay analysis and study of IEEE 802.11p based DSRC safety communication in a highway environment”, 32nd IEEE Infocom, Turin, Italy, April 2013, 1591-1599.
- [39] P. Kyasanur and N. Vaidya, “Selfish MAC layer misbehavior in wireless networks”, IEEE Intl. Conf. on Dependable Systems and Networks (DSN’03), San Francisco, USA, June 2003, 173-182.
- [40] M. Čagalj et al., “On selfish behavior in CSMA/CA networks”, 24th Annual Joint Conference of the IEEE Computer and Communications Societies, March 2005, Miami, USA, 2513-2524.
- [41] J. Tang, Y. Cheng, and W. Zhuang, “Real-time misbehavior detection in IEEE 802.11-based wireless networks: An analytical approach”, IEEE Trans. Mobile Computing, vol. 13(1), Jan. 2014, 146-158.
- [42] C. Campolo, A. Molinaro, and A. Vinel, “Understanding the performance of short-lived control broadcast packets in 802.11p/WAVE vehicular networks”, 3rd IEEE Vehicular Networking Conference, Amsterdam, Netherlands, Nov. 2011, 102-108.
- [43] A. Jadbabaie, J. Lin, and A.S. Morse, “Coordination of groups of mobile autonomous agents using nearest neighbor rules”, IEEE Trans. Automation Control, vol. 48(6), June 2003, 988-1001.
- [44] R. Olfati-Saber and R.M. Murray, “Consensus problems in networks of agents with switching topology and time-delays”, IEEE Trans. Automation Control, vol. 49(9), Sept. 2004, 1520-1533.
- [45] W. Ren and R.W. Beard, “Consensus seeking in multi-agent systems under dynamically changing interaction topologies”, IEEE Trans. Automation Control, vol. 50(5), May 2005, 655-661.

[46] W. Ren, R.W. Beard, and E.M. Atkins, “Information consensus in multivehicle cooperative control”, *IEEE Control Systems Magazine*, April 2007, 71-82.

[47] G. Le Lann, “Cohorts and groups for safe and efficient autonomous driving on highways”, 3rd IEEE Vehicular Networking Conference, Amsterdam, Netherlands, Nov. 2011, 1-8.

[48] G. Le Lann, “Safety in vehicular networks—On the inevitability of short-range directional communications”, 14th Intl. Conference on Ad Hoc, Mobile, and Wireless Networks (AdHoc-Now 2015), Athens, Greece, June-July 2015, *Lecture Notes in Computing Systems* n° 9143, Springer, 347-360.

[49] A. Studer, M. Luk, and A. Perrig, “Efficient mechanisms to provide convoy member and vehicle sequence authentication in VANETs”, 3rd Intl. Conference on Security and Privacy in Communications Networks (SecureComm), Nice, France, Sept. 2007, 422-432.

[50] A. Bazzi et al., “Study of the impact of PHY and MAC parameters in 3GPP C-V2V Mode 4”, *IEEE Access*, vol. 6, 2018, 71685-71698, arXiv:1807.10699v2.

[51] S. Rangan, T. S. Rappaport, and E. Erkip, “Millimeter-wave cellular wireless networks: Potentials and challenges”, *Proceedings of the IEEE*, vol. 102(3), March 2014, 366-385.

[52] M. Haddad et al., “TDMA-Based MAC protocols for vehicular ad hoc networks: A survey, qualitative analysis, and open research issues”, *IEEE Communications Surveys & Tutorials*, vol. 17(4), 2015, 2461-2492.

[53] G. Le Lann, “A collision-free MAC protocol for fast message dissemination in vehicular strings”, *IEEE Conf. on Standards for Communications and Networking (CSCN)*, Berlin, Germany, Nov. 2016, 7 p.

[54] B. Toghi et al., “Multiple access in cellular V2X: Performance analysis in highly congested vehicular networks”, *IEEE Vehicular Networking Conference*, Taipei, Taiwan, 2018, 8 p.

[55] G. Le Lann and N. Rivierre, “Real-time communications over broadcast networks: the CSMA-DCR and the DOD-CSMA-CD protocols”, *INRIA Research Report* n°1863, March 1993, 36 p.

<https://hal.inria.fr/inria-00074810>

[56] C. Bergenhem et al., “V2V communication quality: Measurements in a cooperative automotive platooning application”, *SAE World Congress 2014*, Detroit, USA, *SAE Intl. Journal of Passenger Cars – Electronic and Electrical Systems*, vol. 7(2), 2014, 9 p.

[57] F. Martelli et al., “A measurement-based study of beaconing performance in IEEE 802.11p vehicular networks”, 31st IEEE Infocom Conference, Orlando, USA, 2012, 1503-1511.

[58] A.J. Ghandour et al., “Dissemination of safety messages in IEEE 802.11p/WAVE vehicular network: Analytical study and protocol enhancements”, *Pervasive and Mobile Computing*, Elsevier, vol. 11, April 2014, 3-18.

[59] X. Yin et al., “Performance and reliability evaluation of BSM broadcasting in DSRC with multi-channel schemes”, *IEEE Transactions on Computers*, vol. 63(12), Dec. 2014, 3101-3113.

[60] E. Renda et al., “IEEE 802.11p VANets: Experimental evaluation of packet inter-reception time”, *Computer Communications*, Elsevier, vol. 75, Feb. 2016, 26-38.

[61] D. Dolev et al., “Reaching approximate agreement in the presence of faults”, *Journal of the ACM*, vol. 33(3), 1986, 499-516.

[62] N. Santoro and P. Widmayer, “Agreement in synchronous networks with ubiquitous faults”, *Theoretical Computer Science* 384, Elsevier Science Direct, 2007, 232-249.

[63] M.K. Aguilera, G. Le Lann, and S. Toueg, “On the impact of fast failure detectors on real-time fault-tolerant systems”, *DISC 2002*, Toulouse, France, *Lecture Notes in Computing Systems* n° 2508, Springer-Verlag, 354-369.

[64] G. Le Lann, “Fast distributed agreements and safety-critical scenarios in VANETs”, *IEEE Intl. Conf. on Computing, Networking and Communications*, Santa Clara, USA, Jan. 2017, 200-206.

[65] G. Le Lann and P. Rolin, French patent n° 2 597 686 (1984), US patent n° 4,847,835 (1989), and European patents.

[66] G. Le Lann. “Autonomic vehicular networks: safety, privacy, cybersecurity and societal issues”. *IEEE Vehicular Technology Conference Spring 2018 -- First International Workshop on research advances in Cooperative ITS cyber security and privacy (C-ITSec)*, Porto, Portugal, June 2018, 5 p.

[67] J. Carbaugh, D.N. Godbole, and R. Sengupta, “Safety and capacity analysis of automated and manual highway systems”, TRB Annual Meeting, Transportation Research Part C: Emerging Technologies, vol. 6(1), Elsevier, 1998, 69-99.

[68] Section 10.4.1 in AHS C3-Volume 1 report, Sections 10.4-13 and 10.4.14 in AHS C3-Volume 2 report, March 1998.

<https://path.berkeley.edu/research/connected-and-automated-vehicles/national-automated-highway-systems-consortium>

[69] P. A. Ioannou and C. C. Chien, “Autonomous intelligent cruise control”, IEEE Transactions on Vehicular Technology, vol 42(4), Nov.1993, 657-672.

[70] S.E. Shladover, D. Su, and X-Y. Lu, “Impacts of cooperative adaptive cruise control on freeway traffic flow”, 91st TRB Annual Meeting Washington, D.C. Jan. 2012, TRB Report 1868, 16 p.

[71] J. Ploeg et al., “Controller synthesis for string stability of vehicle platoons”, IEEE Transactions on Intelligent Transportation Systems, vol. 15(2), 2014, 854-865.

[72] C. Flores, V. Milanés, and F. Nashashibi, “Online feedforward/feedback structure adaptation for heterogeneous CACC strings”, IEEE Annual American Control Conference, Milwaukee, WI, USA, June 2018, 49–55.

[73] 39th Intl. Conference of Data Protection and Privacy Commissioners, Hong Kong, Sept. 2017, 5p.

https://edps.europa.eu/sites/edp/files/publication/resolution-on-data-protection-in-automated-and-connected-vehicles_en_1.pdf

[74] J.P. Hubaux, S. Capkun, and J. Luo, “The security and privacy of smart vehicles”, IEEE Security & Privacy, vol. 2(3), May-June 2004, 49-55.

[75] B. Parno and A.Perrig, “Challenges in securing vehicular networks”, ACM Workshop HotNets-IV, College Park, USA, Nov. 2005, 6 p.

[76] M. Raya, P. Papadimitratos, and J.P. Hubaux, “Securing vehicular communications”, IEEE Wireless Communications, vol. 13(5), Oct. 2006, 8-15.

[77] G. Calandriello et al., “Efficient and robust pseudonymous authentication in Vanet”, ACM VANET’07, Montréal, Canada, Sept. 2007, 19-27.

[78] General Data Protection Regulation, European Commission

https://ec.europa.eu/info/law/law-topic/data-protection/data-protection-eu_en

[79] D. J. Glancy, “Privacy in autonomous vehicles”, Santa Clara Law Review, vol. 52(4), Rev. 1171, Nov. 2012, 70 p. Available at <https://digitalcommons.law.scu.edu/lawreview/vol52/iss4/3>

[80] IEEE Spectrum, Cars That Think/Transportation/Self-Driving, “The self-driving car is a surveillance tool”, M. Anderson, June 24, 2019, 2 p.

[81] H. Buhrman et al., “On the importance of having an identity or, is consensus really universal?”, Distributed Computing, vol. 18(3), Springer, Feb. 2006, 167-176.

[82] R. Maes and I. Verbauwhede, “Physically unclonable functions: A study on the state of the art and future research directions”, in Towards Hardware-Intrinsic Security, Springer Link, Online 12 Oct. 2010, 3-37.

[83] D. Skeen, “Non-blocking commit protocols”, ACM Sigmod International Conference on Management of Data, Ann Arbor, Michigan, May 1981, 133-142.

[84] 5GAA, “Timeline for deployment of C-V2X—Update—5Gaa”, Jan. 2019, 12 p.

http://5gaa.org/wp-content/uploads/2019/01/5GAA_White-Paper-CV2X-Roadmap.pdf

[85] IEEE, “Ethically aligned design—A vision for prioritizing human well-being with autonomous and intelligent systems”, 1st Edition Overview, 2019, 15 p.

<https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e-overview.pdf>



RESEARCH CENTRE OF PARIS

2 rue Simone Iff

CS 42112

75589 Paris Cedex 12

Tél. : +33 (0)1 80 49 40 00

[View publication stats](#)

Publisher

Inria

Domaine de Voluceau - Rocquencourt

BP 105 - 78153 Le Chesnay Cedex

[inria.fr](#)