



HAL
open science

Safe Automated Driving on Highways – Beyond Today’s Connected Autonomous Vehicles

Gérard Le Lann

► **To cite this version:**

Gérard Le Lann. Safe Automated Driving on Highways – Beyond Today’s Connected Autonomous Vehicles. 8th Complex Systems Design & Management Conference “Towards smarter and more autonomous systems”, Dec 2017, Paris, France. hal-01610957

HAL Id: hal-01610957

<https://hal.science/hal-01610957v1>

Submitted on 5 Oct 2017

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

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

Safe Automated Driving on Highways – Beyond Today’s Connected Autonomous Vehicles

G rard Le Lann

Abstract Safe automated driving rests on safety-critical (SC) inter-vehicular (IV) coordination. Safety criticality is defined unambiguously via the Bounded Move requirements. We show that today’s autonomous vehicles and upcoming connected vehicles fail to meet these requirements by huge margins. We present a cyber-physical construct, IV communication protocols and IV agreement algorithms that achieve SC IV coordination in highway autonomic vehicular networks. Worst-case termination time bounds of protocols and algorithms are given, which allows for checking that the Bounded Move requirements are met. These solutions lay the ground for novel standards specifically aimed at safety. Interestingly, they also meet privacy requirements. Some open problems raised with automated driving are put into perspective.

1 Introduction

The purpose of this paper is to shed some light on where we stand and where we might be heading to regarding *safe* human-less driving. Whether such predictions as “autonomous vehicles will be on the roads by 2020” should be trusted depends on (1) what is meant by “autonomous vehicle”, and (2) which environments are considered (dedicated paths or ordinary traffic conditions). Not long ago, we were told that “autonomous vehicles had been driven thousands of miles without any accident”. Such misleading statements are a thing of the past (existing autonomous vehicles have been involved in dozens of accidents, unreported collisions not taken into account). However, given that fierce competition is the rule in the fast growing market of autonomous vehicles, the “fairy tales” era is not over yet.

Early work was focused on platoons [1]. The original concept (pre-planned linear formations with a lead truck piloted by a human driver) has evolved, now referred to as strings (spontaneous formations of autonomous vehicles). Safe inter-vehicular gaps with human driven vehicles are too large, and humans are responsible for about 90% of severe accidents (severe human injuries and fatalities). Seeking for improvements, the Intelligent Transportation Systems (ITS) community has stated two contradictory goals (overall goal denoted Ω):

- Significant reductions of safe inter-vehicular gaps in strings (a few meters, even at medium-high velocities) so as to minimize travel times and pollution (Ω_1),
- Significant reductions in the number of severe accidents (Ω_2). A reduction ratio of 10 is commonly quoted.

G rard Le Lann
INRIA Paris Rocquencourt
RITS Team – BP 105
F-78153 Le Chesnay cedex, France
gerard.le_lann@inria.fr

Safety issues appear to be the least rigorously studied in the ITS field. That must be corrected, since solving the numerous safety problems that remain open is a prerequisite to the advent of safe automated driving. To this end, we examine safety-critical (SC) scenarios, where severe accidents occur if not handled correctly. Given that human life is at stake, SC scenarios to be examined are those where risks of severe accidents are highest, the case with high velocities. Therefore, we consider multilane highways (outside, around, and inside cities), where vehicles form networks of spontaneous short-lived or long-lived strings of various size. A vast majority of SC scenarios result from ordinary or/and intentional maneuvers (e.g., overtaking, lane changes, on-ramp zipper merging, accelerations, decelerations). Other SC scenarios result from unexpected or undesired events, such as brutal stopping (the “brick wall” paradigm) and irrational behaviors, leading to accidents and sudden lane blocking. Goal Ω must be achieved for all these SC scenarios, the latter contributing to the non-zero residual accident ratio (Ω_2).

The ITS community has adopted the SAE standard that identifies 6 levels of driving automation, from 0 to 5 [2]. Level 0 is assigned to human driven vehicles. To the exception of level 5, referred to as full automation, denoted FullAU herein, human supervision is mandatory—laws and insurance companies mandate the presence of human drivers, legally responsible for taking over whenever needed. As of today, vehicles on the roads have SAE levels ranging from 0 to 3. Arguing against or in favor of FullAU driving is pointless unless there are good reasons to believe that *safe* FullAU driving is achievable, currently an open question. In the sequel, we demonstrate that safe FullAU driving on highways is feasible, achieving goal Ω , provided that self-organizing/autonomic vehicular strings are endowed with necessary SC inter-vehicular (IV) coordination schemes.

Focusing on highways is also particularly interesting from other standpoints. A significant fraction of vehicles that travel on highways daily are occupied by people going to and returning from work, wasting about 1 hour of human time per day per vehicle. Thanks to FullAU driving, this huge time budget can be spent in more rewarding activities (e.g., work, rest, education, e-shopping, infotainment).

In Section 2, we show that goal Ω cannot be achieved with connected autonomous vehicles as envisioned today, and we review the limitations of remedies currently under investigation. Vehicular networks [3] are cyber-physical systems-of-systems [4], and SC IV coordination can only rest on some appropriate cyber-physical construct. The relevance of the cyber dimension, vastly overlooked so far, is highlighted in Section 3, where safety criticality is rigorously characterized via the Bounded Move requirements. In Section 4, we examine longitudinal and lateral SC scenarios and we present the cohort construct as well as IV protocols and distributed agreement algorithms that meet the Bounded Move requirements, along with analytical results and illustrations. Perspectives and some open problems are discussed in Section 5.

2 Today's Connected Autonomous Vehicles

Most results that underlie currently deployed autonomous vehicles originate from control theory, kinematics, and robotics. Autonomy is about to be supplemented with “connectedness”. Starting year 2020 in the USA, new vehicles shall be equipped with radio communication devices conformant to WAVE standards [5]. These connected autonomous vehicles, referred to as AU vehicles, are/will be equipped with on-board (OB) systems based on the following technologies:

- LOS (line-of-sight) *sensing* technologies (e.g., radars, lidars, and cameras), that ensure robotics-centric capabilities (e.g., proximity detection, motion and trajectory control, lane keeping),
- *Geo-positioning* technologies, based on GNSS “sensing” (e.g., GPS, Glonass, Galileo) and e-maps, providing vehicle space-time coordinates,
- Non-LOS *radio communication* technologies conformant to IEEE 802.11p or ETSI ITS-G5 standards, providing omnidirectional medium-range (≈ 300 m) communication capabilities [5]. Radio communications are sometimes viewed as just another type of sensing technology—a vehicle “senses” its non-LOS environment thanks to data read from messages received. This (mistaken) vision is unnecessarily restrictive—see Section 3.

Limitations of the aforementioned technologies are widely acknowledged, often referred to as “shyness”: AU vehicles are not “aggressive” enough. In ordinary traffic conditions (no reserved lanes as for busses, taxis, or carpooling), SC maneuvers are based on “daring” and “guessing” whether other vehicles appear to “understand” and are willing to give way. Besides large IV gaps (Ω_1 is not achieved), this translates into overly cautious lane change maneuvers that are unduly aborted or postponed, clearly not the best strategies for achieving Ω_2 . Human or human-like cognitive capabilities appear to be necessary.

• Authority sharing

With AU vehicles, it is mandated that a human keeps his/her hands on the wheel, ready to intervene whenever necessary. In a SC scenario, acceptable reaction latencies are in the order of 1 to 2 seconds. (See Subsection 4.2, where we show that 100 FullAU vehicles forming a string can make SC decisions in less than 208 ms, in less than 1 second when every message must be repeated once, a performance beyond the reach of humans.) Studies devoted to human-machine interaction demonstrate that it takes between 5 and 8 seconds for a distracted driver to take over control, i.e. to understand “what is going on” and to make a correct decision [6], [7]. Humans are too slow, especially in adverse driving conditions (tiredness, heavy rain, fog, darkness).

Note also that trusting humans as a last resort to compensate for incomplete automation is contradictory with the original rationale for automated driving. If about 90% of severe accidents are caused by humans when they are supposed to be continuously at the wheel, how could we trust humans (for achieving Ω_2) when they are told that they do not have to pay continuous attention to driving? How do we know at design time (e.g., in a R&D laboratory) whether a human should be

trusted more than some automaton, or the opposite, for every possible future SC scenario under every possible future conditions? Authority sharing problems have been under extensive examination in the defense domain and in air transportation. With AU vehicles, 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, too often, pilots do not recall how to recover from a loss of control by the flight management system.

- **Artificial intelligence**

AI (machine/deep learning) should help, undoubtedly the case in urban settings, where there is a need for detecting bicyclists, pedestrians, dogs (in some cities), for scene recognition in general. However, vehicular AI faces limitations as regards responsiveness (large response times due to image processing, data retrieving and semantic understanding). In cities, where velocities are reasonably small—enforced in the future, the latter difficulty may not be an impediment. Conversely, as for the handling of SC scenarios on highways, today’s sensing-and-AI engines are too slow—see Subsection 3.2. Fortunately, highways are settings where goal Ω can be achieved with deterministic time-bounded solutions (see Section 4) that rest on assuming lane recognition and lane numbering (cameras and e-maps suffice). Reliance on accurate GNSS-assisted geopositioning is needed with other solutions and in bad weather conditions (e.g., invisible lane marking due to snow). No matter how smart, decisions made by a vehicular AI engine are also based on sensing, and thus suffer from the same limitations proper to sensing/robotics capabilities (see Subsection 3.1). Also, it remains to be seen how one may *prove* safety when AI is resorted to.

- **The need for new WAVE standards**

The inadequacy of today’s WAVE standards and LTE, 4G and 5G technologies regarding safe AU driving has been pointed at by numerous experts. A brief analysis is given in the Appendix. The conclusion is:

For achieving goal Ω , vehicular networks must be provided with specific radio channels and IV communication protocols distinct from those defined in current WAVE standards or used in public telecommunication networks.

We have been here before. After years of experimentation, it became clear that the Arpanet NCP protocol, inappropriate for the Internet, had to be replaced by some novel protocol, now known as TCP/IP. Similarly, years after their initial inception, it is clear that today’s WAVE standards are inappropriate for achieving SC IV communications. Moreover, they do not take advantage of more recent and affordable technologies (short-range high speed radio, optical communications). Lastly, due to mandated periodic beaconing of messages that carry IP or MAC addresses, they are potential threats to privacy (e.g., anonymity breaches, remote intrusions, spoofing, and vehicle tracking are feasible). Solutions based on certified pseudonyms delivered by a cloud-based trustable Authority are far from being satisfactory [8]. There is a need for novel standards (WAVE 2.0). Examples of solutions for WAVE 2.0 are given in the sequel. Today’s WAVE 1.0 standards should be kept for non-SC IV communications.

Since goal Ω cannot be achieved with today's AU vehicles "augmented" with human intelligence or/and AI, one may be tempted to conclude that *safe* FullAU driving is unfeasible, a fortiori, at least within the foreseeable future. Why such is not the case is explained below.

3 Beyond Today's Connected Autonomous Vehicles

AU driving is based on *vehicle-centric sensing* only. *IV coordination* in autonomic vehicular *networks* is needed for achieving goal Ω . The relevance of the "network" dimension (rather than focusing on "the ego vehicle") should not come as a surprise in the Internet age. To paraphrase the famous dictum "The network is the computer" (J. B. Gage, Sun Microsystems), which was popular in the 80's when the focus was still on "the computer", one may write "The (autonomic vehicular) network is the (automated) vehicle". IV coordination encompasses IV communication protocols and IV explicit agreement algorithms needed to prove that, e.g., asphalt resources are shared in mutual exclusion by nearby vehicles and within strict time bounds (otherwise, collisions occur), that concurrent and conflicting maneuvers such as lane changes are handled fairly and in bounded time, and so on. Scientists and experts in distributed dependable computing systems and networks feel on firm ground here, since resource sharing and consensus/agreement in the presence of concurrency and failures are fundamental problems which have been studied for almost half-a-century [9]. And ditto for real-time computing [10]. Solutions devised for cyber systems achieve well-known logical "safety" and liveness properties [11]. In autonomic vehicular networks, the counterparts of these properties are reactive safety and proactive safety, respectively—see below. Since reactive safety is taken care of by sensing/robotics, we are invited to devise the cyber-physical counterparts of the aforementioned cyber solutions that would achieve proactive safety.

3.1 Reactive safety vs. proactive safety

Sensing/robotics solutions serve to achieve *collision avoidance* (as much as possible) in *hazardous* situations created by nearby vehicles. Unfortunately, *hazards* can only be sensed after they are (being) instantiated, likely too late in SC scenarios. Furthermore, avoidance strategies are inherently speculative, since they are based on *guessing (silently)* other vehicles' behaviors. That is *reactive safety*. When "caught" in a SC scenario, an AU vehicle tries to protect itself. On the contrary, IV coordination achieves *hazard prevention* (as much as possible) by *influencing behaviors of other vehicles a priori*, and by *striking explicit IV agreements*. Prior to being undertaken, a risk-prone maneuver is declared, allowing nearby vehicles to return positive or negative feedback via IV messaging.

That is *proactive safety*. Unless worst-case assumptions that underlie IV coordination solutions would be violated, there is no need for collision avoidance.

Sensing/robotics capabilities are necessary for fine-tuning agreed upon maneuvers, keeping in check vehicles behaviors/trajectories (e.g., making sure that sufficient spacing is kept between the 3 vehicles performing an agreed lateral insertion in a string).

Lane changes/merging, un-signalized intersection or roundabout crossings, are examples of SC scenarios where explicit IV agreements are necessary. Accident warning is another example. Common belief is that it suffices to have emergency messages broadcast by damaged vehicles, quoting their geo-positions, so as to warn approaching vehicles. Surprisingly, none of the numerous publications that examine this usage of broadcasting addresses the essential question of what should be the behaviors of vehicles that receive such emergency messages. Brutal braking or random lane changes can only result in more accidents. SC IV coordination and proactive safety are more than needed in such scenarios.

IV communications and IV coordination take time. An essential question is “how much”. More precisely, how to tell whether a solution to a SC IV communications/coordination problem is appropriate with regards to goal Ω ?

3.2 The Bounded Move (BM) requirements

In the current ITS literature, “safety” is referred to without being precisely quantified. (Safety issues related to interactions among vehicles bear limited resemblance with functional safety examined in vehicle-centric frameworks, such as ISO 26262 [12].) Safety-related problems would receive simple solutions if one could assume that vehicles do not move while messages are being disseminated, acknowledged, and agreements are struck. Since this is unrealistic, the best we can aim at is to bound and quantify distances travelled by vehicles involved in SC scenarios. If such distances are small enough, then goal Ω is achievable.

It is customary to use latencies to characterize “safe” response times. For example, WAVE 1.0 standards state that emergency messages shall be delivered in at most 100 ms. That is meaningless. We are not told whether this bound should hold in the presence or in the absence of message losses. Moreover, distances travelled in 100 ms do matter, and they depend on velocities. Binding together space and time variables is mandatory. Safety proofs must be given for worst-case conditions. Time bounds that appear in the following BM requirements are stated for worst cases regarding radio channel contention, vehicular density, and IV message/ack losses. Let λ stand for the worst-case upper bound of channel access delays, Δ_d for the worst-case upper bound of string-wide acknowledged message dissemination delays, and Δ_a for the worst-case upper bound of string-wide or inter-string agreement delays. Consider vehicles involved in a SC scenario and let σ stand for the smallest “car asphalt slot” ($\sigma = \text{smallest car size} + \text{smallest IV gap}$). Typically, σ is in the order of 7 m (5+2). BM requirements are as follows:

- BM₀: a MAC protocol is acceptable only if the distance travelled in λ time units by any vehicle is significantly smaller than σ .
- BM₁: an algorithm for string-wide acknowledged message dissemination is acceptable only if the distance travelled in Δ_d time units by any vehicle is smaller than σ .
- BM₂: a string-wide or inter-string agreement algorithm is acceptable only if the distance travelled in Δ_a time units by any vehicle is smaller than 2σ .

The rationale is simple: if vehicles switch from their current “slots” to at most 1 (BM₁) or 2 (BM₂) “slots” ahead while, respectively, building common knowledge or reaching agreement, then goal Ω can be achieved. (Other sufficiently constraining bounds may be considered.)

Other BM requirements may be defined for open roads or urban settings, which have specific driving regulations (e.g., highest velocities, how to cross intersections and roundabouts with no traffic lights). Physical zones (recorded on e-maps) and digitized signs/panels provide vehicles with on-line knowledge of zone limits (city limits, highway ends here, etc.). Appropriate SC IV coordination schemes would be activated according to which type of zone is entered.

4. SC IV Coordination in Autonomic Vehicular Networks

In the ITS literature, terms such as strings or platoons have received no precise definitions. As a result, it is impossible to demonstrate that the BM requirements are met. These shortcomings can be corrected. Constructs with specifications are necessary for reasoning about (and proving) properties under worst-case conditions, notably worst-case timeliness properties. A cyber-physical construct that is essential to meeting the BM requirements is presented in the sequel. SC scenarios that arise on highways are enumerable. They can thus be examined fully. Safety-preserving on-line decisions derive from driving rules (learned by humans). They happen to be deterministic. They can be translated in protocols and algorithms, and implemented in OB systems.

4.1 Cohorts: A basic construct for autonomic vehicular networks

Cohorts—strings with a specification [13], [14], appear to be an early example of cyber-physical constructs needed to prove the existence of such time bounds as λ , Δ_d , and Δ_a . The cohort concept was originally devised for minimizing the number of vehicles involved in rear-end collisions in “brick wall” conditions. Essential features of cohorts are as follows:

- Ranking. Members communicate via 1-hop neighbor-to-neighbor (N2N) messages and beacons, which are acknowledged. They assign themselves consecutive integers, called ranks. An isolated vehicle assigns itself rank 1. In case

a vehicle X catches up with a cohort Γ , tail N of rank n , X and N execute a cyber-physical remote procedure call based on N2N messaging. If successful, X assigns itself rank $n+1$ and becomes Γ 's new tail. If denied, X stays away (smallest inter-cohort gap) or overtakes Γ . Member ranking is used in SWIFT, a MAC protocol aimed at achieving time-bounded N2N communications in cohorts (see further).

- Fault-tolerance and cohort splits. Neighbors exchange beacons periodically, enabling time-bounded detection of a N2N link failure, due to stop or timing failures of an OB system, or caused by too many message/ack losses. In case of a P -to- Q link failure, a cohort split is triggered: Q decelerates until a safe gap is created with P . Thus, in cohorts, N2N communication network partitioning can be ignored. An inter-cohort gap is safe when the head of a cohort that follows another cohort can stop without hitting the tail of the cohort ahead (gap $S_{ct/ch}$ in Fig. 1).

- Worst-case faulty N2N link conditions. Let f stand for the number of temporarily faulty N2N links in the course of a message dissemination involving n members, and not leading to a cohort split. Trivially, we have $0 \leq f \leq (n-1)z$, where z stands for the highest number of consecutive tolerated faults impacting a N2N link.

- Highest velocity.size product. For minimizing the number of rear-end collisions under “brick wall” conditions, a cohort that moves at velocity v shall not comprise more than $n^*(v)$ members [14]. Very fast moving vehicles are isolated most often. Conversely, a cohort that moves slowly includes many members, for the reason that slow motion is necessarily due to highly dense traffic conditions (otherwise, members would simply leave a slow cohort in order to move faster), typically the

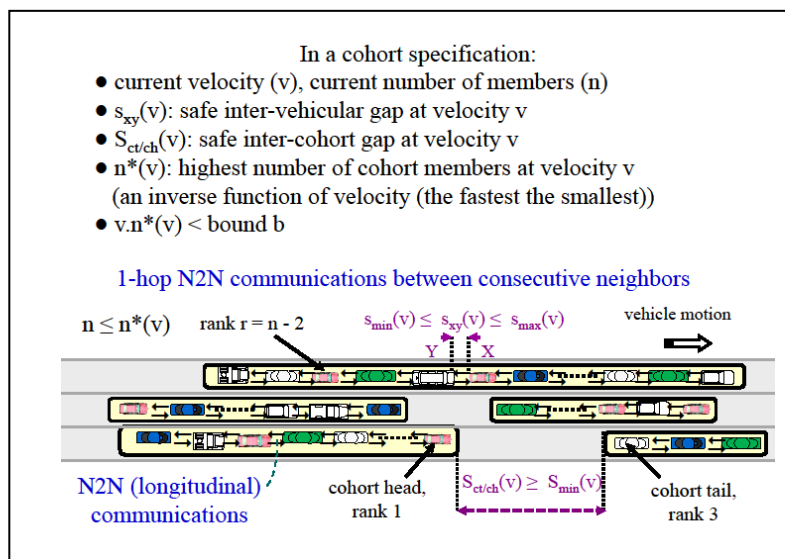


Fig. 1. The cohort construct

case at rush hours or on single lane sections. Based on this observation, the original concept has been refined as follows: $n^*(v)$ must be an inverse function of v (or vice-versa), i.e. the fastest the smallest (or the opposite), which we write $v.n^*(v) \leq b$ [15], [16].

A correct valuation of bound b shall not be arbitrary. Assume that a cohort moving at 250 km/h (highest velocity in WAVE 1.0 standards) shall not include more than 4 members. Thus, with velocities in km/h, we would have $b = 1,000$.

- Naming and privacy. Names appearing in N2N messages are member ranks (rather than IP/MAC addresses in WAVE 1.0 standards). Thanks to consecutive ranking and SWIFT, ranks are authenticated. As a result, masquerading attacks are immediately detected. Neither geo-positioning data nor V2V beaconing is resorted to in the solutions presented below. Therefore, vehicle tracking and anonymity breaches are unfeasible.

4.2 Communication protocols for SC IV coordination

In order to make coordinated behavioral decisions, cohort members shall share some common knowledge, which implies cohort-wide dissemination of SC messages, such as “new n is 11”, “new velocity is 70 km/h”, “lane blocking ahead”, “smallest accepted SAE level is 3” (see further).

- **SWIFT: A deterministic MAC protocol for safety and privacy in cohorts**

In cohorts, one can take advantage of directional antennas. Every vehicle is equipped with a backward looking and a forward looking antenna, small beamwidth (e.g., 25°), short-range (e.g., up to 30 m). A N2N SC message received from a neighbor is acknowledged and forwarded to the opposite neighbor. Acks are piggybacked on messages. With SWIFT [15], $\lambda = 2h\theta$, where h stands for the highest number of contiguous members that may experience interferences due to a transmitter, and θ stands for the largest N2N message transmission duration. Conservative figures are $\theta = 1$ ms, $h = 4$ if malicious or irrational attacks (e.g., suicides) are deemed highly unlikely. Thus, $\lambda = 8$ ms, and BM_0 is met: less than 0.56 m are travelled for $v < 250$ km/h.

- **Worst-case cohort-wide acknowledged message dissemination delays**

With SWIFT, channel accesses match descending and ascending member ranking, alternatively, thus enabling fast symmetrical message dissemination. (This is unfeasible or achieved very inefficiently with conventional CSMA/CA or TDMA protocols.) Assuming no malicious attacks and no concurrency (a single N2N message is disseminated), we have shown in [15] that:

$$\Delta_d(n,f) = 2h\theta \{1+f+\lceil (n-1)/h \rceil\}, \quad n \leq n^*(v).$$

Consider two extreme cases ($b = 1,000$): $v = 180$, $n^* = 5$ and $v = 10$, $n^* = 100$. Let us have $z = 1$ (every N2N message must be sent twice). One finds that BM_1 is met (δ standing for distance travelled in meters):

- $v = 180$: $\Delta_d(5,0) = 16$ ms, $\Delta_d(5,4) = 48$ ms, and $0.8 \leq \delta \leq 2.4$;
- $v = 10$: $\Delta_d(100,0) = 208$ ms, $\Delta_d(100,99) = 1$ s, and $0.58 \leq \delta \leq 2.78$.

Highest bounds Δ_d hold for very pessimistic loss conditions. Tolerated loss frequencies are equal to $z(n-1)/\Delta_d$, i.e. ≈ 100 Hz (99 losses per second) in the above example. Note in passing that periodic beaconing is totally impracticable under such loss frequencies.

- **Acknowledged message broadcast/multicast within and across cohorts**

SWIFT achieves longitudinal acknowledged message broadcast (or multicast) over a cohort or any cohort subset. Compounded with lateral inter-cohort communications (see below), this instantiates 360° time-bounded *acknowledged* broadcast/multicast modes, not available with WAVE 1.0 standards.

- **Autonomic management of heterogeneous vehicular networks**

Thanks to N2N messaging, distinct cohorts of human-driven vehicles and cohorts of AU or FullAU vehicles can form spontaneously, without having to “freeze” specific lanes for either category. Admission in a cohort (in a given lane) may be conditioned on various parameters, SAE level being one of them (see Ranking, Subsection 4.1). This is essential regarding safety and efficiency, since cyber-physical capabilities depend directly on such levels. Let us give an illustration with Δ_d and safe IV gaps $s_{xy}(v)$. Let p stand for $s_{max}(v)/s_{min}(v)$. Reaction delays and p are highest with lowest SAE levels, and p determines parameter h . For the version of SWIFT which copes with malicious or irrational attacks, it can be shown that having $p = 4$ (e.g., SAE level 1) leads to $h = 8$, while having $p = 1.5$ (e.g., SAE level 5) leads to $h = 5$. Bounds Δ_d are smaller with small values of h whenever $f \neq 0$. It is desirable to enable the spontaneous formation of cohorts comprising only AU vehicles of some high SAE level. Indeed, thanks to faster message dissemination/agreements, they would be able to handle SC scenarios more rapidly or/and at higher velocities, while enjoying stipulated safety properties.

4.3 Distributed agreement algorithms for SC IV coordination

Cohort members must also cope with concurrent SC events, such as conflicting lateral lane changes, lane change attempts when a fast-moving vehicle is approaching, or simultaneous steep braking and insertion attempts with a cohort. Agreement algorithms Eligo and LHandshake [16] build on SWIFT.

- Worst-case cohort-wide agreement delays. Let p stand for the number of participants (members that want to disseminate a message or that propose a value—an input to agreement). With Eligo, extending the results presented in [16] and assuming no malicious attacks, we have the following worst-case termination time bound (highest latency after which every member knows that all members know and decide): $\Delta_{swa}(n,f,p) = 2h\theta \{1+p+2[f+\lceil(n-1)/h\rceil]\}$.

Assume $p = \lceil n/10 \rceil$. We have:

- $v = 180$: $\Delta_{swa}(5,0,1) = 32$ ms, $\Delta_{swa}(5,4,1) = 96$ ms; $1.6 \leq \delta \leq 4.8$;

- $v = 10$: $\Delta_{swa}(100,0,10) = 488$ ms, $\Delta_{swa}(100,99,10) = 2.072$ s; $1.36 \leq \delta \leq 5.76$.

BM_2 is met in both cases.

- Worst-case lateral inter-cohort agreement delays. With LHandshake, we have $\Delta_{isa}(g,f,p) \leq 2 \sigma_{max} + \Delta_{swa}(g,f,p)$, where g stands for the number of consecutive

members that must reach agreement, e.g., that receive an insertion request from an adjacent vehicle ($g < 5$ in realistic cases). Variable σ_{\max} stands for the worst-case delay for transmitting a message, MAC access delay included. With WAVE 1.0 standards, σ_{\max} takes unbounded values, thus the need for lateral MAC protocols that would meet BM_0 . With appropriate solutions (to appear in forthcoming publications), LHandshake meets BM_2 .

5 Perspectives

We have shown that it is possible to achieve goal Ω with AU and FullAU vehicles on highways, thus providing some backing to disruptive approaches. The fact that SC IV coordination is feasible while vehicles move by negligible distances opens up new perspectives regarding safe automated driving on highways, notably the prefixing of risk-prone maneuvers in the physical space by explicit IV agreements in cyber space. Cyber-inspired solutions have a significant impact of how to validate OB systems. So far, testing on the roads and intensive simulations are the only solutions considered. Unfortunately, testing on the roads is plagued with serious weaknesses. Key findings in [17] are: “Autonomous vehicles would have to be driven ... sometimes hundreds of billions of miles to demonstrate their reliability in terms of fatalities and injuries... existing fleets would take ... sometimes hundreds of years to drive these miles...”. What about regression? Since technology keeps evolving (e.g., new OB software releases, new hardware, new sensing devices), suites of tests previously performed must be replayed. Likely, testing on the roads is not a convergent process. Cyber-inspired solutions come with proofs, which serve to demonstrate properties for a (usually) very large set of reachable states. Since testing such states is useless, the complexity and the duration of validation phases are significantly reduced. And what about liability issues? As OB hardware and/or software may be frequently modified, which kind of guarantees and liabilities apply to not-so-recent vehicles with OB systems running versions/releases superseded by newer ones [18]?

Research is being pursued by us and others working on the cross-fertilization of robotics and cyber-inspired algorithms. We have designed new versions of SWIFT and IV protocols for cohort join maneuvers and message dissemination that cope with malicious attacks (e.g., impersonation, message falsification/destruction). Other recent results related to optical communications would also contribute to the advent of WAVE 2.0 standards. We are applying the cohort construct to urban and open road settings, such as un-signalized intersections and roundabouts, considering assumptions more realistic than those found in the current literature. For example, we assume that entrance and exit roads have different numbers of lanes, that vehicles have to coordinate extremely quickly since they may have to follow intersecting trajectories while entering or/and crossing an intersection or a roundabout, and they may have to perform lane merging for entering an exit road.

References

1. Shladover, S. Longitudinal control of automated guideway transit vehicles within platoons. *ASME Journal of Dynamic Systems, Measurement and Control*, vol. 100(4), 1978, 291–297.
2. SAE J3016 standard, <https://www.sae.org/news/3544/>
3. Karagiannis, G. et al. Vehicular networking: A survey and tutorial on requirements, architectures, challenges, standards and solutions. *IEEE Comm. Surveys & Tutorials*, vol. 13(4), 2011, 584-616.
4. Rajkumar, R. et al. Cyber-physical systems: the next computing revolution. *ACM Design Automation Conference 2010*, 731-736.
5. IEEE Standard 802.11p. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications amendment 6: wireless access in vehicular environments, July 2010, <http://www.ietf.org/mail-archive/web/its/current/pdfqf992dHy9x.pdf>
6. Gold, C. et al. "Take over!" How long does it take to get the driver back into the loop? *57th Human Factors and Ergonomics Society Intl. Meeting*, Sept. 2013, 1938-1942.
7. Mok, B. et al. Emergency, automation off: unstructured transition timing for distracted drivers of automated vehicles. *18th IEEE Intelligent Transportation Systems Conference*, Sept. 2015, 2458-2464.
8. 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 and Services, (WONS)*, 2010, 176-183.
9. Lynch, N.A. *Distributed algorithms*. Morgan Kaufmann Pub., 1996, 872 p.
10. Moser, H. and Schmid, U. Reconciling fault-tolerant distributed algorithms and real-time computing. *Distributed Computing*, Springer, vol. 27 (3), June 2014, 203–230.
11. Lamport, L. Proving the correctness of multiprocess programs. *IEEE Transactions on Software Engineering*, vol. SE-3 (2), March 1977, 125-143.
12. ISO 26626-1: Road vehicles – Functional safety – Part 10, 2011.
13. Le Lann, G. Cohorts and groups for safe and efficient autonomous driving on highways. *3rd IEEE Vehicular Networking Conference (VNC)*, Nov. 2011, 1-8.
14. Le Lann, G. Safety in vehicular networks—On the inevitability of short-range directional communications. *14th Intl. Conference on Ad Hoc, Mobile, and Wireless Networks (AdHoc-Now)*, June-July 2015, Springer LNCS 9143, 347-360.
15. Le Lann, G. A collision-free MAC protocol for fast message dissemination in vehicular strings. *IEEE Intl. Conference on Standards for Communications and Networking (CSCN'16)*, Oct.-Nov. 2016, 7 p.
16. Le Lann, G. Fast distributed agreements and safety-critical scenarios in VANETs. *IEEE Intl. Conference on Computing, Networking and Communications (ICNC 2017)*, Jan. 2017, 7p.
17. Kalra, N. and Paddock, S.M. *Driving to safety*. Rand Corporation Report n°1478, Oct. 2016, 15 p.
18. Smith, B. W. *Automated driving policy*. Road Vehicle Automation 3, Book Chapter, Springer Lecture Notes in Mobility, 2016, 53-58.
19. Yao, Y. et al. Delay analysis and study of IEEE 802.11p based DSRC safety communication in a highway environment. *IEEE INFOCOM 2013*, 1591-1599.
20. Karlsson, K., et al. Field measurements of IEEE 802.11p communication in NLOS environments for a platooning application. *IEEE VTC Fall-2012*, 1-5.
21. Tonguz, O. K. et al. On the broadcast storm problem in ad hoc wireless networks. *3rd Intl. Conference on Broadband Communications*, Oct. 2006, 11p.

Appendix – Today’s WAVE standards considered inadequate

Today’s WAVE standards for IV communications, such as IEEE 802.11p and ETSI ITS-G5 are based on WiFi technology. They serve to provide so-called “connected vehicles” with access to Internet and cloud services (infotainment, weather data, traffic conditions, etc.), in addition to enabling *best-effort* IV communications. Vehicles can thus be seen as “smartphones on wheels”.

Essential choices are reliance on CSMA/CA as the MAC-level protocol and omnidirectional communications, radio range in the order of 300 m, and interference range in the order of 500 m. There cannot be such bounds as λ with WAVE 1.0 standards. Stochastic channel access delays are exceedingly high in moderate or worst-case contention conditions. This is shown in [19] where average values of MAC delays achieved by the IEEE 802.11p protocol range between 75.3 ms and 211.8 ms, for various channel loads, assuming 1 vehicle every 12 m. At 108 km/h, vehicles travel 6.35 m at least (the 211.8 ms figure is not a strict bound), which is much higher than stated in BM₀.

Mobile radio communications are unreliable [20]. Lack of acknowledgements (acks) in multicast or broadcast modes under WAVE 1.0 standards is another major weakness. Use of acks leads to the broadcast storm problem [21], no time-bounded solution published so far. Since SC message dissemination and IV agreements must be achieved in bounded time despite losses of messages or acks, there cannot be such bounds as Δ_d and Δ_a . None of the BM requirements is met by WAVE 1.0 standards. Ditto for solutions built out of these standards, such as CACC (Cooperative Adaptive Cruise Control), where message broadcast is most often implicitly assumed to be free from contention and fully reliable, in contradiction with numerous impossibility proofs. For example, when less than 2/3 of vehicles targeted by a broadcast do not receive a broadcast message, no agreement/coordination is feasible. Results established ignoring such issues are of little significance for real-world conditions.

A similar conclusion applies to periodic beaconing (broadcasting of CAM messages), frequencies in the 1 to 10 Hz range. Since message losses do occur, “situational maps” are inaccurate, and they may differ significantly for any 2 nearby vehicles, rendering them useless for building a common global view, believed to be necessary for safety (mistakenly). Moreover, in addition to overloading communication channels and OB processors, beaconing amounts to breaching privacy voluntarily, since vehicles reveal their IP/MAC addresses and time dependent geolocations to unknown recipients within ranges in the order of 300 m, making tracking, spying and hacking much easier. Finally, reliance on vehicle-to-infrastructure (V2I) communications via terrestrial nodes, such as road-side units or 3G/LTE/4G/5G relays, can only lead to poorer results in terms of delays. Also, terrestrial nodes may fail and/or be “attacked”, and they can be used for launching all sorts of attacks, man-in-the-middle attacks in particular. V2I communications favor security and privacy threats, hence safety threats.