Integrated Safety and Efficiency in Intelligent Vehicular Networks: Issues and Novel Constructs

Gérard Le Lann

Transport Research Arena – Europe 2012

Integrated Safety and Efficiency in Intelligent Vehicular Networks: Issues and Novel Constructs

Gérard Le Lann*

INRIA, BP105, F-78153 Le Chesnay Cedex, France

Abstract

We present the cohort and the group constructs which are aimed at reconciling safety and efficiency for intelligent vehicular networks on roads and highways, and show how platoons and vehicular ad hoc networks can be structured as cohorts and groups. Fundamental implications of safety requirements are reviewed. A rationale for on-board systems based on diversified functional redundancy is developed, illustrated with a proposal for neighbor-to-neighbor periodic beaconing based on short range unidirectional communications meant to withstand telemetry failures. Worst-case analytical results are given for safe inter-vehicle spacing in cohorts despite inaccurate vehicle space-time coordinates and failing telemetry capabilities. The group construct is based on prefixing usage of sensing-based solutions with omnidirectional communications. Benefits resulting from prefixing vehicle maneuvers with vehicle role assignments are illustrated with the on-ramp-merging safety-critical scenario.

Keywords: Safety; Dependability; Platoons; Vehicular Ad Hoc Networks; Cooperative Automated Driving; V2V Communications.

* Corresponding author. Tel.: +33 (0)1 3963 5364.
E-mail address: Gerard.Le_Lann@inria.fr.

1. Introduction

An intelligent vehicular network (IVN) is an open set of automated vehicles endowed with demonstrated efficiency and safety properties. In this paper, we consider IVNs on roads and highways. Efficiency goals are significant reductions in energy consumption, pollution, and travel times. Such goals rest on reducing human reaction latencies and longitudinal spacing, as well as the durations of risk-prone lateral maneuvers. Safety goals are significant reductions in accident and injury rates, which imply large enough inter-vehicle longitudinal spacing, as well as enforcing reduced velocities in the course of risk prone maneuvers, which are thus time consuming. Clearly, efficiency and safety are antagonistic goals. Platoons, which were introduced in [Caudill & Garrard 1977], and vehicular ad hoc networks (VANETs) are the most studied forms of IVNs. The high compactness requirement, which is met with platoons, is largely ignored in VANETs. In both cases, safety problems arise due to failures impacting on-board (OB) systems or radio channels. Many such problems are yet unsolved. This is a serious impediment to the deployment of IVNs. A global approach building on results that came out of national initiatives in various countries, as well as recent projects such as, e.g., eVALUE, SAFESPOT and HAVEit in Europe, can be envisioned. Limitations proper to platoons and VANETs can be circumvented by "splitting" the platoon construct in two distinct constructs – cohort and group. As a result, IVN complexity can be mastered, and safety proof obligations can be met. Table 1 summarizes the need for "augmenting" platoons and VANETs with novel constructs.

Table 1. A summarized appraisal of platoons and VANETs

Criteria	Platoons	VANETs
Compactness as a design driver	yes	no
Safety in the presence of failures	?	?
Strong features	rigorous definitions of platoons	ad hoc mirrors reality
	platoon is a structuring construct	communication technologies an integral part of VANETs
	proofs (e.g., kinematics)	
Limitations	complexity	no rigorous definitions for "clusters"
	communication technologies are an add-on	lack of structuring constructs
		simulations, few proofs

In Section 2, we review some important implications of safety requirements. The distinction between non safety-critical and safety-critical (SC) scenarios is explored in Section 3. Cohorts and groups are introduced in Section 4. A rationale for neighbor-to-neighbor unidirectional communications in cohorts is exposed in Section 5, along with recent findings on how to keep safe inter-vehicle spacing despite failing telemetry capabilities. In Section 6, we elaborate on the open issue of time-bounded multi-access channel delays in mobile wireless networks.

2. Safety Requirements and Some Fundamental Implications

Safety is a system property which does not reduce to such properties as, e.g., reliability or availability. An instantiation of an IVN may be safe in the presence of vehicles that violate reliability requirements. Conversely, a set of vehicles, each meeting reliability requirements, may instantiate an unsafe IVN. However, a violation of a dependability property may lead to safety hazards. This has been acknowledged

years ago for OB systems, as witnessed by the ISO 26262 standard, an adaptation of the IEC 61508 standard for Automotive Electric/Electronic Systems, which focuses on *functional safety of individual vehicles*. With IVNs, *functional and non-functional safety properties of vehicle networks* are at stake. Currently, there are very few solutions for IVNs that build on state-of-the-art in Dependable Computing [Avizienis et al. 2004]. Let us briefly review three major obligations that derive from safety requirements.

2.1. Diversified functional redundancy

An automated vehicle is equipped with an OB system designed to ensure a number of functions. In addition to processing, I/Os, and storage functions, an OB system provides an automated vehicle with the following three major functions:

- Space-time localization and scene recognition:
  - 360° positioning and global time keeping (e.g., GPS/EGNOS/Galileo devices) referred to as GP; in order to improve the accuracy of space-time coordinates and to withstand GP outages, plain GP is augmented with e.g., dead reckoning and inertial systems for space coordinates, and with clocks for time coordinates; the corresponding function is denoted GP+
  - 360° scene recognition (e.g., cameras, sensor fusion, AI) and lateral detection (e.g., radars/lidars)
- Longitudinal telemetry: directional sensing-based technology (e.g., radars, infrared, free-space optics) serving to enforce safe longitudinal inter-vehicle spacing
- Short-range omnidirectional communications: 360° radio communications (V2V and V2I), such as e.g. those specified in the IEEE 802.11p and IEEE 1609.x standards, based on a stochastic multi-access control (MAC) protocol resting on carrier-sense and collision avoidance (CSMA-CA) as well as seven 6 Mbits/s channels, one of them (referred to as the SC-channel) reserved for safety-critical messages.

For this paper, it suffices to consider that OB function failures fall in two categories, tolerable failures and fatal failures. A failure is tolerable whenever masking or recovery is feasible, in due time. A failure is fatal whenever some function is lost, and no other function can supersede the lost function in due time. Every OB function shall be implemented out of diversified redundant hardware, data, and software capabilities, so as to avoid common cause failures. This is common practice with many commercial airplanes, where the altimeter function is built out of 5 altimeters implemented out of 3 differing technologies. However, this does not suffice, since a function may be lost (fail-stopped or detectably erroneous), transiently or permanently, under adverse conditions. That some other function shall be able to supersede a failed function is mandatory in every safety-critical domain (e.g., air transportation), where system reliability or availability figures have a lower bound in the order of $1\text{-}10^{-9}$ per hour. For example, trajectories of commercial airliners are under the control of such diversified functions as automated navigation, terrestrial radio control, and TACAS [US DoT/FAA 2011]. IVNs are no exception. Many design principles that have been validated in aeronautics over the past 25 years could be usefully applied to IVNs. As a result, we would be able to provide satisfactory answers to such simple questions as: “What if a follower’s longitudinal telemetry function fails when the spacing with its predecessor is 3 m, they both circulate at 120 km/h, and the predecessor decelerates abruptly?” – see Section 5.

2.2. Realistic modeling of the cyber-physical world

Assumptions that underlie a modeling of the IVN cyber-physical world shall match reality. To put it simply, 10^{-9} per hour (see above) could be viewed as the highest acceptable probability of having at least one of such assumptions violated at run time. Among examples of unrealistic modeling that can be found in the published literature (e.g., highways with a non varying number of lanes), let us elaborate on vehicle space-time coordinates. Let γ (resp., τ) stand for the inaccuracy of GP+ longitudinal space (resp., time) coordinates. Numerous published solutions for IVNs rest on postulating that space-time coordinates are as

accurate as desired, e.g. γ smaller than 1 meter and τ in the order of a few nanoseconds. The former figure rests on postulating highly dense populations of DGPS devices, road-side units, or in-road units, which is currently unrealistic, and an optimistic (best-case) assumption regarding the future. The latter figure amounts to equating GP satellite-to-ground signal propagation delay inaccuracies with the discrepancy between the time values read on 2 GP devices in 2 different vehicles. Moreover, both figures rest on postulating the absence of GP outages. Since safety experts would simply discard designs based on such assumptions, more realistic figures need be considered. For example, while GP inaccuracy of space coordinates may exceed 40 m, GP+ inaccuracy γ would be in the order of 15 m. GP inaccuracy of time coordinates is in the order of 1μs. With "good enough" affordable clocks (serving to withstand GP outages), assuming outages last less than 10 s, intrinsic clock drift in the order of e.g. $0.5\ 10^{-5}$, mutual time discrepancy 2τ for any two vehicles would be in the order of 100 μs.

Elements of the terrestrial referential, e.g., road-side units (RSUs), landmarks, as well as their exact space coordinates constitute topological data made available to OB systems via electronic maps. Such data can be combined with vehicle GP+ space coordinates and sensing-based functions for achieving lane-level positioning [Skog & Händel 2009, Toledo-Moreno et al. 2010]. Lane-level positioning has been demonstrated by a number of innovative companies (e.g., Toyota, Google, BMW).

2.3. Proofs for worst-case conditions

Here, "conditions" refer to safety-critical (SC) scenarios considered at design time, meant to reflect adverse conditions experienced in an IVN. They are distinct from *assumptions* relative to cyber-physical world models reviewed above. Meeting safety requirements implies *proving* specific properties for (realistic) *worst-case conditions*. Again, to put it simply, 10^{-9} per hour could be seen as the highest acceptable probability of having at least one of those run-time conditions that underlie a solution and/or its companion proofs violated at run time. The belief according to which V2V beaconing at 10 Hz is always feasible in IVNs is an appropriate illustration of a non worst-case condition. Consider a highway, 8 lanes both ways, and compactness in the order of 1 vehicle every 11 m – a realistic worst-case compact scenario. Due to highway curvatures, a radio interference radius in the order of 400 m would encompass all vehicles circulating on a transmitter's lane at most 600 m away from the transmitter, approximately, both directions, i.e. 110 vehicles per lane. Beacons carry vehicle localizations, velocities and safety related data, yielding a beacon size in the order of 6.8 Kbits, encryption overhead included (privacy mandates encryption). Simple calculations show that 880 vehicles generate a total load which saturates a 6 Mbits/s channel for beaconing frequencies higher than 1 Hz. Moreover, the above calculations rest on assuming a 100% channel utilization ratio, which is unrealistic with CSMA-CA radio channels.

2.4. Shared/final authority and socio-technical issues

IVNs are life-critical socio-technical systems [Rasmussen 1997]. The combined expertise of professionals from diverse areas is needed for addressing *all* issues involved with automated driving, vehicle safety and road safety, ranging from behavioral psychology, to hazard analyses and safety engineering, to resilient distributed real-time systems design and engineering. With IVNs, 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 due to a flight management system failure (pilots have to *guess* what to do). A FAA study found that many accidents and major incidents occur when the pilots and the technology are failing together. If we are to avoid such problems with IVNs, the essential issue of shared/final authority shall be addressed very rigorously, which implies (1) knowing how to decide when it is safer to trust an OB automaton or a

human driver, (2) developing comprehensive human-system interfaces that would provide human drivers with detailed diagnoses and recommendations as regards "what to do" in the event of a fatal OB system failure. Moreover, with hybrid/mixed IVNs, safety and interoperability requirements translate into a highly complex version of the shared/final authority issue. A decision local to a vehicle (is the human or some OB automaton in charge?) may depend on knowing the degrees of automation of surrounding OB systems, which knowledge is not spontaneously available, and highly transient. Although seemingly not closely related, these open questions bear some resemblance with safety and interoperability issues relative to the integration of unmanned airborne vehicles into controlled airspace – see Subsection 2.1.3 in [DeGarmo 2004].

3. Safety-Critical vs. non Safety-Critical Scenarios

In SC scenarios, V2V communication delays shall be significantly smaller than 100 ms. Vehicles involved in a SC scenario can only be reasonably close to each other. Such scenarios may develop far away from RSUs, and RSUs are vulnerable (out of service, e.g. due to intentional attacks). Thus, the handling of SC scenarios shall only be based on 1-hop V2V communications. Contemplating V2I communications would amount to assuming that investments in road-side infrastructures are economically viable in the long term. V2V radio ranges (smaller than interference ranges) are in the order of 250 m. Thus, reliance upon V2I communications and 2-hop V2V relaying would imply having 1 radio-equipped RSU per km, which entails huge construction costs. It might prove uneasy to demonstrate palatable returns-on-investments in the face of fast changing radio technology. (Imagine that some novel V2V communication standard is adopted circa 2020, necessitating radio-equipped RSUs every 20 km only.)

Safety proofs cannot rest on assuming the availability of an operational RSU everywhere anytime. However, RSUs are appropriate communication helpers as regards non SC communications in IVNs, and RSUs are good geo-positioning helpers. Since assuming arbitrarily dense RSUs is unrealistic, the following question arises: what is the optimal tradeoff between inaccuracy bounds γ and τ on the one hand and smallest distances between consecutive RSUs on the other hand?

A SC scenario always has at least one initiator, vehicle denoted Z (Z' added when more than one initiator). A categorization of highway-centric SC scenarios can be built out of four criteria (single-lane or multi-lane, originating event is intentional or unintentional), as illustrated below:

- Intentional single-lane scenario: misbehaving human driver; Z reverted to manual mode, brutal acceleration or deceleration,
- Unintentional single-lane scenario: brutal stopping; Z stops abruptly (the "brick wall" paradigm, e.g. collision with a deer); Z and Z' involved in an accident, 1 lane blocked,
- Intentional multi-lane scenario: on ramp merging; Z is the entrant vehicle,
- Unintentional multi-lane scenario: emergency stopping (individual multi-lane change); Z wants to reach the emergency lane as soon as possible.

Hazard analyses and proofs of properties rest on a set NB which stipulates nominal bounds for such variables as, e.g., acceleration and deceleration rates, velocities, and inter-vehicle spacing, as well as OB function failures to be tolerated (see Subsection 2.1). Consider a finite bounded set V of vehicles. A non SC scenario is defined as follows: V occupies a single lane, every vehicle in V behaves according to NB, experiencing no failures or tolerable failures only. A SC scenario corresponds to either (1) V occupies multiple lanes, one vehicle at least performing lane change maneuvers, or (2) at least one vehicle in V (single lane or multi-lane) violates some bound in NB or experiences a fatal failure. Since most frequent scenarios are non SC, they are also referred to as stationary scenarios. Given that the duration of a SC scenario is very short (few seconds at most), such scenarios are also referred to as transitory scenarios.

It turns out that this distinction between SC and non SC scenarios leads to the concepts of cohort and group, generic constructs that encompass platoons and VANETs.

4. Cohorts and Groups

The cohort construct is concerned with the handling of non SC scenarios, i.e. single-lane stationary scenarios. The group construct, whereby vehicles are assigned specific roles prior to undertaking risk prone maneuvers, is aimed at handling SC scenarios, i.e. single-lane and multi-lane transitory scenarios. In the sequel, our notation for bounds is as follows (b is a mute variable): b° for lower bound, b• for upper bound.

4.1. Cohorts

A cohort is an ad hoc set of $r \leq r^{\bullet}$ contiguous vehicles circulating on a single lane. Contrary to platoons, lane changes are not under the responsibility of cohort management. On a given lane, a vehicle leaves (resp., joins) a cohort simply by decelerating (resp., accelerating). Cohort management is fully distributed (a head plays no particular role), resting on telemetry and some other function(s) – see Section 5. In case some bound in set NB is violated, a SC scenario is triggered without prior approval from the cohort head. Some contiguous members of a cohort may form a pre-planned platoon, while others do not.

Let v_u stand for vehicle U's velocity, and s_{xy} for the spacing between two contiguous vehicles X and Y on the same lane, X preceding Y. Inter-vehicle spacing has been extensively explored [Shladover 1991], under various car-following models [Panwai & Dia 2005] and for mixed vehicle networks [Chakravarthy et al. 2009]. We apply and extend existing work to the cohort construct. Spacing s_{xy} which depends on X's and Y's velocities is such that $s^{\circ} \leq s_{xy} \leq s^{\bullet}$. Bound s° is derived from safety calculations for smallest velocities (e.g., 3 m for velocities smaller than 20 km/h). Bound $s^{\bullet}$, which is derived from efficiency calculations (cohort compactness), is reached when $v_x = v_y = v^{\bullet}$ (e.g., 150 km/h). In case Y would detect that its spacing with X is increasing closer to $s^{\bullet}$, either Y accelerates so as to remain a member of its current cohort, or Y decelerates until s_{xy} reaches value S°, Y becoming head of a cohort – see Fig.1.

In addition to inter-vehicle spacing, it is necessary to define inter-cohort spacing, denoted $S_{ct/ch}$, CH standing for the head of a cohort and CT the tail of the preceding cohort. (Observe that there is no distinction made between variables s and S in the published literature.).

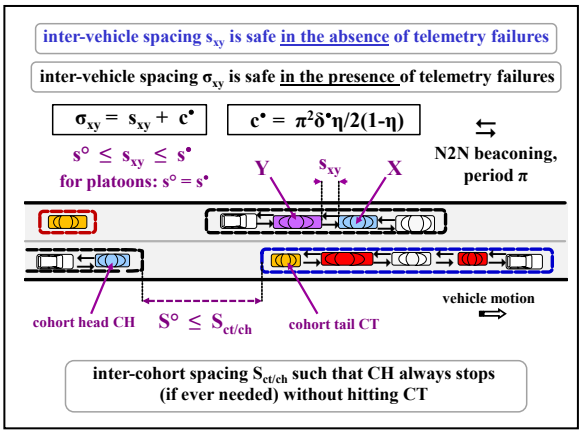

Fig. 1. Cohorts on a 2-lane highway

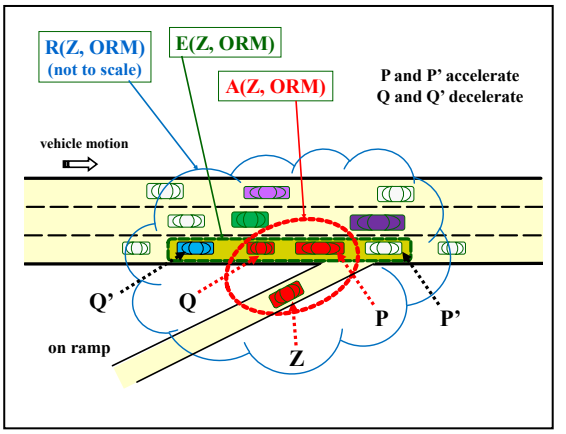

Fig. 2. Groups and the On-Ramp-Merging scenario

$S_{ct/ch}$ has a lower bound S° derived from safety calculations for smallest velocities (e.g., 15 m for velocities smaller than 20 km/h). S• is reached when $v_{ct} = v_{ch} = v•$. Specifying r• and $S_{ct/ch}$ permits to set an upper bound for the number of vehicles that may be involved in a collective rear-end collision would on-board systems of cohort members experience fatal failures simultaneouslyGiven that $S_{ct/ch}$ is enforced, a cohort head always stops before hitting the tail of a preceding cohort (the "brick wall" paradigm). Thus, the occurrence of a tolerable failure impacting a vehicle in a cohort is handled in a non visible manner outside that cohort. This "immunity" property is essential for safety. In the future, safety regulations will stipulate bounds found in set NB, such as s°, s•, S°, r•, v•, which bounds will be monitored, and enforced whenever needed. At times, highest velocity bounds v* may be set lower than nominal v•, due to temporary or local conditions (e.g., weather or highway surface). Cohort management rests on distributed agreement protocols, which protocols serve to maintain common knowledge of current v*, as well as reaching event-driven or repeated agreements [Dolev et al., 1986] on, e.g., new mutual spacing, new beaconing periods. Set NB also comprises variables σ° and σ_{xy}, the counterparts of s° and s_{xy} in the presence of telemetry failures – see Subsection 5.2. Cohorts of 1 vehicle each entail poorest compactness/efficiency. Given that with automated driving, vehicle motions are under the control of OB systems, instantiations of low density patterns can be avoided by enforcing the creation of maximally compact complete cohorts (r• members each) whenever feasible.

4.2. Groups

A type F SC scenario started by an initiator Z is denoted {Z, F}. A velocity upper bound v•(F) is specified for every type F. All messages exchanged in the course of a SC scenario are SC-messages broadcast over the V2V SC-channel. At some unpredictable time, Z broadcasts a SC-message denoted M(Z, F). Three groups and three algorithmic phases are defined with scenario {Z, F}. Created in phase 1, group R(Z, F) comprises vehicles that receive M(Z, F). Created in phase 2, group E(Z, F) comprises vehicles in R(Z, F) that may have to take some active part in scenario {Z, F}; such vehicles are said to be *eligible* (for becoming actors). Group A(Z, F) is created in phase 3. This group comprises *actors*, i.e. vehicles that do have to take an active part in scenario {Z, F}. Every actor plays a specific role. Roles depend on type F. These phases are followed with two kinematic phases. In phase 4, actors undertake coarse grain maneuvers inferred from their roles assigned in phase 3. An actor (as well as neighbors, possibly) may have to decelerate (velocity not higher than v•(F)) or to accelerate, or start changing lane. In phase 5, actors perform fine grain maneuvers, under the control of their sensing-based proximity capabilities (notably, side-looking capabilities). Due to space limitations, we only illustrate multi-lane SC scenarios with On-Ramp Merging – see Fig. 2. (Group forming is detailed in forthcoming publications.)

Z is the entrant vehicle. In phase 2, only those members of R(Z, ORM) which circulate on lane 1 (rightmost or leftmost lane, depending on the country considered) run the eligibility test. E(Z, ORM) includes every vehicle X estimating that it will reach the merging point at a time t(X) comparable to t(Z), estimated time of Z's arrival at merging point quoted in M(Z, ORM). Every eligible vehicle X broadcasts SC-message EM(Z, ORM, X) carrying t(X), its own situational data, as well as neighbors' data (if any). In phase 3, Z runs some optimization function having SC-messages EM(*) as inputs, and chooses 2 contiguous vehicles, denoted P and Q, P (resp., Q) being assigned the role of Z's predecessor (resp., Z's successor). When only 1 SC-message EM(Z, ORM, X) is received, Z's choice is obvious (X = P or X = Q). As soon as phase 3 is over, Z, P and Q start phase 4 by adjusting their respective velocities so as to make them approximately equal when they reach the merging point. Moreover, P and Q adjust their respective velocities so as to create a "slot" between them, permitting Z to get inserted on the highway. When P, Z and Q are in line-of-sight with each other, phase 5 is started, consisting in fine lane "insertion" tuning. Neighbors P' and Q' adjust their respective velocities accordingly. Cohort management is invoked when phase 5 is about to terminate. Z is assigned the rank previously held by Q, and ranks held by Q and its followers are incremented.

Network protocols and coordination algorithms are run in phases 1, 2 and 3, whereas phases 4 and 5 rest on sensing-based control algorithms. Durations of phases 1, 2 and 3 are approximately in the {1-2} seconds range for unintentional scenarios, in the {3-5} seconds range for intentional scenarios. Membership of A(Z, ORM) is "frozen" until scenario {Z, ORM} terminates. Motions of vehicles other than actors are under the control of cohort management. Via role assignments, actors know "what to do" sufficiently ahead of time, prior to invoking sensing-based functions. As a result, risk-prone maneuvers are undertaken under conditions much more favorable, i.e. safer, than in the absence of role assignments.

5. Diversified Functional Redundancy – Rationale for N2N Unidirectional Communications

There is a blatant lack of symmetry as regards functions currently considered with OB systems (Subsection 2.1): 2 functions are related to terrestrial resources, and only 1 function is devoted to radio resources. Stated differently, there is no provision made for "longitudinal communications" that would back longitudinal telemetry.

5.1. N2N unidirectional communications in cohorts

Latencies with telemetry capabilities are in the order of few ms. In platoons or cohorts, if Y follower of X has its telemetry failing at time t, spacing s_{xy} is out of control, an obvious safety hazard in case X would decelerate at t at some "high" rate $\delta_x(t)$. We have seen that another OB function shall supersede a failed telemetry function. V2V communications are not an ideal solution, for at least two reasons. Firstly, V2V communication channels may be jammed (accidentally, intentionally). Secondly, worst-case detection/reaction latencies achievable with V2V communications are much higher than a few ms, due to shared radio channel access delays in the presence of 100s of contenders (MAC access delays). Since the hazardous situation caused by a failing telemetry function ought to be resolved by just two contiguous members, some other communications-based function is needed. Unidirectional communications are feasible with small beamwidth radio antennas [Ramanathan et al. 2005]. With such antennas, restricted to span very short line-of-sight ranges (e.g., 20 m), it is possible to provide any two cohort members with a "private" communication channel, a function referred to as neighbor-to-neighbor (N2N) communications, implemented via at least one couple of front-looking and rear-looking unidirectional antennas, operating on channel(s) other than those allocated to V2V communications. Tunable antennas with transmit power proportional to inter-vehicle spacing help in mitigating radio interferences. By choosing an appropriate technology, N2N communications can be immune to V2V channel jamming.

N2N communications consist of messages and beacons exchanged over N2N links. By the virtue of N2N messaging and "linear" downstream/upstream relaying, a cohort can be structured as a chain or a virtual ring. It is reasonably easy to devise distributed fault-tolerant agreement algorithms out of such features, which algorithms are essential for cohort management – see Subsection 4.1. A N2N beacon is shared by two contiguous cohort members only (no relaying). N2N beaconing is a periodic process, period π – see Fig.1. Vehicles need not have access to the GP+ time referential (good timers/clocks suffice). Let us show how N2N beaconing serves to withstand telemetry failures.

5.2. Safe inter-vehicle longitudinal spacing despite telemetry failures

Consider two neighbors X and Y, Y following X. By definition, when Y's telemetry function fails, spacing s_{xy} is safe. We want to find the expressions of $\sigma°$ and σ_{xy}, the counterparts of $s°$ and s_{xy} in the presence of telemetry failures. Let $\delta•$ stand for the highest value of nominal deceleration rates sustainable by every vehicle (this will be mandated by safety authorities). Since space-time coordinate inaccuracies γ

and τ cannot be held negligible, X's N2N beacon sent at t carry X's deceleration rate $\delta_x(t)$ measured or computed over interval [t-π, t], rather than X's longitudinal space coordinates. Since telemetry failures shall be viewed as tolerable failures, one defines η, $0 < \eta < 1$, such that $\forall t$, $\delta_x(t) \leq \eta\delta\bullet$. Variables η and $\delta\bullet$ belong to set NB. *The distinction between X decelerating at a rate at most equal to $\eta\delta\bullet$, and X forced to violate this bound matches exactly the distinction between a stationary scenario and a SC scenario*. Worst-case delays for transmitting a beacon over a N2N link are in the order of a few ms, and convenient values of π are smaller than 1 s. Such values are feasible with MAC protocols devised for N2N beaconing (a very small number of contenders may be involved), in contrast with V2V beaconing which involves hundreds of contenders under worst-case traffic density and a high number of lanes (both highway directions). In [Le Lann 2011], one can find a detailed presentation of a solution based on periodic N2N beaconing, as well as the derivation of spacing $c\bullet$ such that $\sigma_{xy} = s_{xy} + c\bullet$ is proved to be a safe X/Y spacing in worst-case telemetry failure conditions. The exact formula is: $c\bullet = \pi^2\delta\bullet\eta/2(1-\eta)$. Notice one unexpected outcome of interest:

Worst-case spacing $c\bullet$ is a constant (does not depend on velocities).

We can now answer an open question:

How much should be added to s*, the spacing usually considered in platoons (s* ≈ 2 or 3 m) in order to withstand telemetry failures? Practical answer: No more/much less than ≈ s*.

Indeed, in most realistic cases, numerical values of variables π, $\delta\bullet$ and η are such that $c\bullet$ is smaller than, or comparable to s°. For example, with $\delta\bullet = 7.5$ m/s^2 and $\eta = 0.77$, we have:

$\pi = 100$ ms (10 Hz) ➡ $c\bullet = 0.13$ m ‖ $\pi = 400$ ms ➡ $c\bullet = 2.01$ m ‖ $\pi = 500$ ms ➡ $c\bullet = 3.14$ m.

Since s° is approximately equal to s*, the inter-vehicle spacing in platoons, ensuring a safe inter-vehicle spacing within cohorts despite telemetry failures entails a marginal loss of compactness, (much) smaller than s* for most practical settings. These results can be easily exploited by the transportation industry, regulatory bodies, and safety certifiers. In case the probability of experiencing coincidental failures of telemetry and N2N beaconing could not be considered negligible, triple functional redundancy would be mandatory. In other words, V2V beaconing shall be relied upon as well for safe spacing enforcement. In [Le Lann 2011], we also show how to make use of these results when considering well established spacing algorithms known to enforce safe and stable cooperative adaptive cruise control, such as the algorithm given as Equation 45 in a report from the US Dpt. of Transportation [Fitch et al. 2008].

6. Unresolved MAC Issues with Mobile Wireless Communications

Merits of V2V/V2I communications regarding, e.g., early warnings or collision avoidance are discussed in numerous publications. However, a major problem remains open with mobile wireless networks: how to prove that channel access delays are finite and bounded (non stochastic bounds) in the presence of worst-case contention and hidden nodes? To the best of our knowledge, there is no published "deterministic" MAC protocol, be it based on CSMA, CDMA, or TDMA [Willke et al. 2009], which solves this problem under realistic assumptions. With the RTS/CTS scheme found in the IEEE 802.11p standard, collisions are neither avoided nor resolved deterministically. Various MAC protocols such as location-based or space division multi-access protocols rest on assuming that different vehicles in proximate neighborhood necessarily compute different GP+ positioning data, either at the same time or at times approximately equal. This amounts to assuming that γ and τ are negligible. Since safety mandates making the opposite assumption, such protocols cannot be considered as valid solutions to the time-bounded MAC delays problem whenever safety requirements shall be met. Lack of finite and strictly

bounded MAC delays suffices for disproving safety, regardless of how "smart" the application layer atop the MAC layer. This weakness is a serious impediment to full scale deployments of IVNs.

The belief according to which MAC protocols for mobile wireless networks can only be stochastic appears well founded whenever IVNs are viewed as unstructured "collections" of vehicle "clusters". Conversely, there are no proofs establishing that it is impossible to devise IVN constructs that would make it possible to find "deterministic" MAC protocols. Surprisingly, although not initially devised to that end, cohorts and groups happen to be essential cornerstones for solving this long standing problem with IVNs (with mobile ad hoc networks, more generally).

7. Conclusion and Next Steps

This paper summarizes recent findings related to combined safety and efficiency issues in IVNs, of theoretical and practical value. Our on-going work is focused on harnessing the cohort and the group constructs so as to devise innovative solutions for IVNs, to be presented in forthcoming publications. Among the topics that have been explored, one finds (1) "deterministic" MAC protocols that guarantee strictly time-bounded access delays (despite γ and τ), for V2V and N2N event-driven messages and periodic beacons exchanged in platoons or VANETs, (2) reliable broadcast/multicast protocols which depart from conventional acknowledgment-based solutions, building on "cooperative altruism" made feasible with N2N periodic beaconing – reliability figures achieved with such protocols are much higher than figures published for conventional solutions, (3) protocols able to generate pseudonyms dynamically (privacy requirements) and autonomously, based on V2V communications (no reliance on RSUs).

Adding some structuring (cohorts and groups) to IVNs has profound implications, comparable to those induced by the concepts of atomic transactions and concurrency control algorithms in distributed computing and distributed databases in the 80's, solutions now found in cloud computing. Cohorts (resp., group forming protocols) are the cyber-physical counterparts of atomic transactions (resp., concurrency control algorithms). This fresh vision should open up new horizons.

References

Avizienis A., Laprie J.-C., Randell B., & Landwehr C. (2004). Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable and Secure Computing*, vol.1(1), 11-33.

Caudill R.J., & Garrard W.L. (1977). Vehicle-follower longitudinal control for automated transit vehicles. *Trans. ASME Journal of Dynanic Systems, Measurement, and Control*, vol. 99, 241-248.

Chakravarthy A., Song K., & Freron E. (2009). Preventing automotive pileup crashes in mixed-communication environments. *IEEE Trans. Intelligent Transport. Systems*, vol. 10(2), 211-225.

DeGarmo M.T. (2004). *Issues Concerning Integration of Unmanned Aerial Vehicles in Civil Airspace*. Report MP 04W0000323, MITRE Center for Advanced Aviation System Development, 98 p.

Dolev D., Lynch N.A., Pinter S.S., Stark E.W., & Weihl W.E. (1986). Reaching approximate agreement in the presence of faults. *Journal of the ACM*, vol. 33(3), 499-516.

Fitch G.M. et al. (2008). *Safety benefit evaluation of a forward collision warning system: final report*. NHTSA, US Department of Transportation, HS 810 910, 100 p.

Le Lann G. (2011). *Cohorts and groups for safe and efficient autonomous driving on highways*. 3rd IEEE Vehicular Networking Conference, Amsterdam (NL), Nov. 2011, 1-8.

Lygeros J., Godbole D.N., & Broucke M. (2000). A fault tolerant control architecture for automated highway systems. *IEEE Trans. Control Systems Technology*, vol. 8(2), 205-219.

Panwai S., & Dia H. (2005). Comparative evaluation of microscopic car-following behavior. *IEEE Trans. Intelligent Transport. Systems*, vol. 6(3), 314-325.

Ramanathan R., Redi J., Santivanez C., Wiggins D., & Polit S. (2005). Ad hoc networking with directional antennas: A complete system solution. *IEEE Journal Selected Areas in Communications*, vol. 23(3), 496-506.

Rasmussen J. (1997). Risk management in a dynamic society: a modeling problem. *Safety Science*, vol. 27(2/3) (Elsevier), 183-213.

Shladover S.E. (1991). Longitudinal control of automotive vehicles in close-formation platoons. *ASME Journal on Dynamic Systems, Measurement and Control*, vol.113, 231-241.

Skog I., & Händel P. (2009). In-car positioning and navigation technologies—A survey. *IEEE Trans. Intelligent Transport. Systems*, vol. 10(1), 4-21.

Toledo-Moreo R., Bétaille D., & Peyret F. (2010). Lane-level integrity provision for navigation and map matching with GNSS, dead reckoning, and enhanced maps. *IEEE Trans. Intelligent Transport. Systems*, vol. 11(1), 100-112.

US DoT, Federal Aviation Administration. (2011). *Introduction to Traffic Alert and Collision Avoidance System II Version 7.1*. Doc. HQ-111358, 50 p.

Willke T.L., Tientrakool P., & Maxemchuk N.F. (2009). A survey of inter-vehicle communication protocols and their applications. *IEEE Comm. Surveys and Tutorials*, vol. 11(2), 3-20.