On the Power of Cohorts - Multipoint Protocols for Fast and Reliable Safety-Critical Communications in Intelligent Vehicular Networks

Grard Le Lann

► To cite this version:

Grard Le Lann. On the Power of Cohorts - Multipoint Protocols for Fast and Reliable Safety-Critical Communications in Intelligent Vehicular Networks. ICCVE - International Conference on Connected Vehicles and Expo - 2012, IEEE, TRB, ACM, IFAC, Nov 2012, Beijing, China. pp.35-42. hal-00769133

HAL Id: hal-00769133

<https://hal.inria.fr/hal-00769133>

Submitted on 28 Dec 2012

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 destine au dpt et la diffusion de documents scientifiques de niveau recherche, publis ou non, manant des tablissements d'enseignement et de recherche franais ou trangers, des laboratoires publics ou privs.

On the Power of Cohorts – Multipoint Protocols for Fast and Reliable Safety-Critical Communications in Intelligent Vehicular Networks

Grard Le Lann

IMARA – INRIA Paris-Rocquencourt – France
Gerard.Le_Lann@inria.fr

Abstract—We report on recent findings related to safety-critical V2V multipoint communications in ad hoc networks of fully automated vehicles, in the presence of communication failures. Neither classical failure assumptions nor multipoint protocols at the core of existing communication standards can be considered, since they do not meet the high reliability and strict timeliness requirements set for safety-critical scenarios. We introduce a novel unbounded omission failure model, the concept of proxy sets which builds on the cohort construct, and Zebra, a suite of geocast, convergecast, and multicast protocols specifically designed for safety-critical 1-hop multipoint communications. Analytical expressions of worst-case termination time bounds are given for each Zebra protocol, which is mandatory with safety requirements. These results have a number of practical implications, which are discussed. They should be of interest to safety authorities and to the transportation industry involved in future deployments of intelligent vehicular networks.

Keywords—*Networks of Automated Vehicles, Inter Vehicular Communications, Safety, Reliability, Data Dissemination.*

I. INTRODUCTION

Intelligent vehicular networks (IVNs) are sets of interconnected and automated vehicles resting on combined use of sensing-based solutions (e.g., radars, lasers, cameras, GPS, robotics) and computing/networking-based solutions (e.g., real-time processors, storage capabilities, radio communications). As of now, platoons and VANETs are the most popular examples of IVNs. However, neither of these constructs meets the high reliability and strict timeliness requirements set for safety-critical (SC) scenarios, which has so far prohibited full scale deployment of IVNs. See [1] and [2] for comprehensive surveys. We consider IVNs circulating on roads or highways, and we focus on problems arising from unintentional omission failures (message losses) which impact V2V multipoint communications resorted to in SC scenarios. Omission failures are the dominant instances of failures in mobile wireless communications. Collisions, which are natural phenomena in contention-prone shared channels, are not similar to omissions. Malicious failures (e.g., deliberate jamming, Byzantine failures) and non safety/time critical communications (e.g., Internet access) are out of the scope of this paper.

Achieving reliable multipoint communications in wireless networks has been studied in both static settings (e.g., [3], [4]) and mobile settings (e.g., [5], [6]). However, to the best of our

knowledge, all existing proposals, standardized protocols included, be they based on masking or acknowledgments [7], rest on assuming a bounded failure model, which model fails to capture the very nature of wireless communications in IVNs. Due to hazards potentially induced by such protocols, safety authorities may/shall mandate adoption of novel and trustworthy solutions for the handling of safety/time-critical communications which can withstand high percentages of message losses. This paper introduces such a solution. Distances travelled by vehicles while V2V communication protocols are executed must be bounded. Therefore, demonstrating safety properties implies establishing worst-case upper bounds for protocol termination, which is not doable via simulations. Analytical expressions of worst-case termination bounds are mandatory.

This paper is organized as follows. In Section II, we briefly introduce two novel constructs, cohorts and groups, motivated by the goal of endowing IVNs with combined efficiency and safety properties. Section III is devoted to presenting the multipoint communication phases involved with SC scenarios. Classical omission models and protocols for reliable communications are reviewed in Section IV, and diagnosed inappropriate for withstanding omission patterns proper to mobile wireless short-lived communications. In Section V, we introduce Ω , a novel omission failure model which is more “aggressive”, i.e. more realistic, than traditional bounded models, as well as the concept of proxy sets. The Zebra protocols suite which solves the reliable and timely multipoint communication problem arising with SC scenarios is presented and evaluated in Section VI. The practical significance of our results is discussed in Section VII.

II. NOVEL CONSTRUCTS FOR SAFETY AND EFFICIENCY

SC scenarios may develop far away from a road-side unit (RSU). Consequently, V2I communications cannot be contemplated. Moreover, given the time bounds sought-after, only 1-hop V2V communications can be considered. Safety requirements for future IVNs are bound to resemble those met in commercial air transportation, where safety/time-critical functions must exhibit reliability or availability figures at least as high as $1-10^{-9}$ per hour [8], which implies on-board systems based on diversified functional redundancy. With regard to communication functions, diversification translates into having V2V communications backed by some other technology.

Let us now introduce two fundamental constructs, cohort and group, based on distinguishing *stationary* risk-free scenarios from *transitory* SC scenarios.

A. Introduction to Cohorts

Short range (e.g., 30 m) unidirectional communications are natural candidates for backing failing V2V functions [9]. In [10] and [11], we have developed the concept of neighbor-to-neighbor (N2N) communications and the companion cohort construct. Vehicles are equipped with short-range forward-looking and backward-looking unidirectional antennas. Cohorts are a formalization of the concept of bounded strings of vehicles circulating on a single lane. Safe spacing between cohort member X and follower Y can be small (e.g., 3 m), monitored and enforced by telemetry capabilities (e.g., radars, lidars, lasers). We have established that an additional inter-vehicular spacing of less than 1 m is needed for withstanding failing telemetry functions, thanks to periodic N2N beaconing. Similarly, some safe spacing S (e.g., 40 m) is maintained between any 2 consecutive cohorts, enforced by telemetry capabilities and V2V communications. Let cs stand for the current size of a cohort, and cs^* for the (enforced) highest value of cs . N2N communications offer two major services:

- Event-based messaging, for achieving coordination and membership knowledge. Let $VA(X)$ stand for vehicle X attributes, e.g., X 's length, type (truck, sedan ...), and color. $VA(X)$ is useful in the physical phases of SC scenarios. Whenever some vehicle X joins a cohort at rank r_X (bound cs^* shall not be violated), new member X creates a N2N message carrying $VA(X)$ and r_X , denoted $Nm(X)$, and initiates a global (cohort-wide) hop-by-hop bidirectional dissemination (GD) protocol. Upon receiving $Nm(X)$, every cohort member records $VA(X)$, increments its local copy of variable cs , and forwards $Nm(X)$ to its successor or predecessor, if any. Every member previously assigned a rank equal to or higher than r_X increments its rank,

- Periodic exchange of beacons, notably “do you hear me” beacons, denoted $Fb(*)$, which serve to detect a failing N2N communication link, in which case a cohort split is undertaken by both vehicles (say X and X 's follower Y) involved. X becomes tail of its current cohort and Y becomes head of a new cohort, positioned S away from X , while X and Y broadcast a specific V2V SC “cohort split” message – see Section VI.

B. Introduction to Groups

Groups serve to reason about transitory SC scenarios. A SC scenario is assigned a type, denoted F , and is always triggered by at least one vehicle, denoted Z . In this paper, we focus on SC scenarios involving lane changes. Vehicles participating in a SC scenario shall be assigned specific roles, prior to undertaking risk-prone maneuvers, which is feasible with V2V omnidirectional communications. A SC scenario comprises 3 cyber phases resting on V2V multipoint messaging, followed by 2 physical phases involving coarse grain and fine grain maneuvers, controlled by sensing-based functions. Durations of a SC scenario are in the $\{2-5\}$ seconds range, dominated by physical maneuvers latencies. Thus, total durations of the cyber phases shall be (much) less than 0.5 second.

III. MULTIPOINT V2V COMMUNICATIONS IN A SC SCENARIO

Z and vehicles involved in a SC scenario exchange situational data so as to coordinate their respective behaviors.

A. Situational Data

Let $\Psi_t(X)$ stand for X 's situational data current at time t – see Table I. To the exception of X , all parameters depend on t . Examples of lane type are “reserved for automated vehicles”, “in-road lane”, “merging lane”. Lanes are numbered sequentially. Longitudinal coordinate λ is computed combining e.g., GPS/GNSS data, e-maps, and the distance travelled from the last (downstream) landmark or RSU. This yields small longitudinal location inaccuracy γ . In the case of an isolated vehicle (cohort of size 1), s is assigned symbol $\perp$.

Let $\Phi_\theta(X)$ stand for situational data that shall hold true for X at some future time $\theta > t$, where l^* and j^* are targeted lane type and number, respectively – see Table II. To the exception of X , all parameters depend on θ . The accuracy of a space-time predicate $\Phi_\theta(X)$ is inversely proportional to time horizon $\theta-t$. W.l.o.g., for the sake of clarity, and considering the small durations allowed for executing the cyber phases, we will ignore issues related to positioning inaccuracies in the sequel.

B. Anatomy of V2V Communication Phases in a SC Scenario

The Zebra protocols suite presented in Section VI encompasses the 3 cyber phases of a SC scenario, illustrated here with the on-ramp merging scenario ($F = ORM$), see Fig. 1. Entrant vehicle Z circulating on a on-ramp lane intends to move to highway lane 1. Ignore failures and radio channel access contention for the moment.

1) Cyber phases 1, 2 and 3

Selective Geocast (SGcast) triggered by Z serves to identify vehicles which match predicates stated in $\Psi_t(Z)$ and $\Phi_\theta(Z)$, carried in a message denoted $M(Z,F)$. Thus, in addition to current locations (“where are you”, the classical Geocast [12]), we have predicates related to projected situational data (“where will you be”). Other protocols are Convergecast (Ccast), a union of Unicast (Ucast) protocols [13], and Multicast (Mcast).

Phase 1: At time t , Z initiates SC scenario $\{Z,F\}$ by doing SGcast $\backslash M(Z,F)$. $M(Z,F)$ is received by vehicles forming an unknown group $B(Z,F)$. Let $R(Z,F)$ stand for $B(Z,F) \cap \text{lane } 1$. A $t^\circ \approx t$, every vehicle X member of $R(Z,F)$ runs an e-test, which consists in processing the contents of $M(Z,F)$ in order to evaluate whether, given $F = ORM$ and $\Psi_t(X)$, X might match predicate $\Phi_\theta(Z)$ at time θ , as well as accommodate $VA(Z)$. In particular, velocities at θ have to be approximately identical.

TABLE I. X 'S CURRENT SITUATIONAL DATA

X : vehicle id	t : current time	l^* : lane type	j : lane number
r : rank in cohort	λ : longitudinal coordinate on j	v : velocity	s : spacing with predecessor on j

TABLE II. X 'S PROJECTED SITUATIONAL DATA

X	θ : future time	l^*	j^*	λ^* : longitudinal coord. on j^*	v^*
-----	------------------------	-------	-------	--	-------

Details of the e-test are not shown in this paper. Vehicle X knows it is eligible, denoted e-vehicle, when the e-test is successful. E-vehicles form group $E(Z, F)$. By construction, $E(Z, F) \subseteq R(Z, F)$. Let n_e stand for the number of e-vehicles.

Phase 2: An e-vehicle X responds to Z by doing $Ucast\backslash C(X, Z, F)$. Assume $n_e \geq 2$. Within $E(Z, F)$, 2 contiguous vehicles, denoted P and Q, referred to as actors, will be designated. P and Q will be in charge of creating some sufficient spacing so as to have Z safely “inserted” between them on lane 1. Two options are available regarding the contents of message $C(X, Z, F)$. Under the Z-driven option, $C(X, Z, F)$ is a candidacy message which carries $VA(X)$ and $\Psi_r(X)$. Moreover, message $C(X, Z, F)$ contains the $n_e - 1$ pairs $\{VA(*), \Psi_r(*)\}$ relative to other members of $E(Z, F)$. Z decides on pair $\{P, Q\}$, out of the $C(*, Z, F)$ messages returned by the e-vehicles. Under the E-driven option, the choice of pair $\{P, Q\}$ rests on e-vehicles, via the execution of an internal agreement protocol. Under this option, $C(*, Z, F)$ is a unique choice message which carries $VA(P)$, $\Psi_r(P)$, $VA(Q)$, and $\Psi_r(Q)$. For the sake of conciseness, we consider the Z-driven option only in the sequel. Once its $C(*, Z, F)$ message has been sent, an e-vehicle “freezes” its spacing with its predecessor.

Phase 3: Z does $Mcast\backslash D(Z, F)$ over group $E(Z, F)$. $D(Z, F)$ is a decision message carrying $\{P, Q\}$. P and Q form group $A(Z, F)$, $A(Z, F) \subseteq E(Z, F)$. Upon receiving $D(Z, F)$, P and Q “lock” themselves as actors for scenario $\{Z, F\}$.

Termination of phase 3 coincides with the start of physical phase 4. Every participant has advance knowledge of which role to play, i.e. which maneuvers ought to be undertaken during phases 4 and 5. In the general case, P (resp., Q) adjusts its acceleration (resp., deceleration) rate. Vehicles other than actors behave according to cohort management protocols. The above applies to any SC scenario involving lane change(s).

2) Required properties

Owing to omissions, messages $M(Z, F)$, $C(*, Z, F)$, or/and $D(Z, F)$ may be lost. A multipoint communication can be modeled as a set of n virtual point-to-point links. With SGcast or Mcast, every link originates from a unique sender, and the same message circulates on every link. With Ccast, every link ends at a unique receiver (sink), and a link carries a message generated by a sender, via a Ucast.

Figure 1. Groups and the On-Ramp-Merging SC scenario

Wherever appropriate, we will use Xcast to mean SGcast, or Ccast, or Mcast. Among n links, up to f may experience an omission. An omission may be a send omission (transiently faulty sending antenna), a receive omission (transiently faulty receiving antenna), or a message loss (transiently faulty medium). Omissions are distinguishable from permanent V2V antenna failures, which are not considered in this paper.

Let T_i stand for the duration of cyber phase i , $i \in [1, 3]$. Due to safety requirements, only the expressions of T_i 's worst-case upper bounds matter. Properties required for every cyber phase are as follows (informal presentation):

- High reliability: every message that a vehicle intends to get delivered to a given set of vehicles is received by every such vehicle, under worst-case omission patterns.

- Strict timeliness: successful message delivery occurs in no more than T_1 , T_2 , T_3 , for SGcast, Ccast, and Mcast, respectively, under worst-case asphalt (vehicle density and velocities) and ether (channel contention) traffic conditions.

IV. OMISSION FAILURES AND RELIABLE MULTIPOINT COMMUNICATIONS – CLASSICAL MODELS AND PROTOCOLS

A. Classical Modeling of Omissions and Related Solutions

Inevitably, f is some integer function of n , denoted $\xi(n)$. Typically, $\xi(n) < n/3$ or $\xi(n) < n/2$. See [14] for an exhaustive treatment of reliability issues in distributed systems under link failures. Classical solutions for reliability are round-based protocols, resting on positive acknowledgements (acks) or/and negative acknowledgements, or resting on masking. In the former case, w.l.o.g., let us consider acks only, i.e. positive-acknowledgement-retransmission (PAR) protocols. A SGcast (resp., Mcast) by sender Y is instantiated as a suite of sgcast (resp., mcast) operations, each consisting of 2 rounds: 1 round where Y outputs a message H over n links, 1 round where acks are sent by receivers that have been delivered H. A Ccast to a sink Y by a set of n senders is instantiated as suite of ccast operations, every operation consisting of n ucast of 2 rounds each, 1 round where n senders output their messages, 1 round where Y sends acks to senders which have been heard of by Y. Every expected ack shall be delivered to a sender within some maximum latency. If not the case, a sender starts a new sgcast or mcast or ucast operation. Since acks are messages, they too can be omitted, in which case a sender is led to retransmit its message unnecessarily. With masking protocols, a Xcast is instantiated as a series of consecutive rounds of xcast, no acks required. With centralized masking, only the initial sender performs the xcast rounds. With distributed masking, every participant that hears from the initial sender, directly or indirectly, performs a xcast round (the message received is repeated).

Adversarial choices regarding how to inflict f omissions range between two extreme strategies: f omissions in the first xcast operation, 1 omission every xcast operation. In the former case, termination occurs in 4 rounds (PAR) or 2 rounds (masking), since the adversary has exhausted its omission budget at once – subsequent rounds are failure-free. With the latter strategy, termination occurs in $2(f+1)$ rounds (PAR) or $f+1$ rounds (masking), yielding the worst-case termination

times. Since it is impossible to prohibit its instantiation, the worst-case strategy must be considered when computing guaranteed termination times of Xcast operations.

To summarize, classical models of omissions and related solutions rest on assuming (A1) targeted set of n receivers (membership of $R(Z,F)$ in our case) and exact values of n and f are advance knowledge, (A2) worst-case termination times incurred with $2(f+1)$ or $f+1$ rounds are acceptable, (A3) the bounded failure model is a faithful modeling of reality.

B. Classical Approaches Considered Inapplicable

(A1) With IVNs, it is impossible to predict accurate values of n , hence f . Firstly, by definition, membership of $R(Z,F)$, i.e. n and f , are unknown to Z when Z initiates SGcast. Ditto for $E(Z,F)$, a fortiori. Secondly, n may change at every invocation of SGcast. Thus, *stricto sensu*, neither PAR nor masking protocols can be considered.

(A2) Safety being at stake, one must retain values that shall never be violated at run time. Indeed, if n (thus f) is assigned a value smaller than real n (or f) experienced at run time, then termination times of cyber phases may be higher than those used for safety calculations, leading to hazards or/and catastrophes, inevitably. Consequently, n and f must be assigned “trustable” values, i.e. values possibly significantly higher than strictly necessary most often. This raises a serious concern. With classical solutions, worst-case termination times of a Xcast are linear functions of f , which depends on n . Whenever n and f are assigned high values, acceptable values stipulated for the T_i 's cannot be met.

(A3) If it is possible to experience f omissions during a first xcast operation, why should it be impossible to experience f omissions again in subsequent xcast operations? A round lasts in between a few milliseconds and a few dozens of milliseconds. Due to mobility, vehicle density, and anisotropic radio communications, message losses (cars occasionally disappear behind trucks) and interferences (noise, fading...) may well last longer than hundreds of milliseconds, possibly as much as or beyond the acceptable bounds set for the T_i 's. Moreover, it may well be that the same vehicles experience omission failures over and over again in every xcast round. Proving the opposite is impossible.

Therefore, one is led to consider an adversary much more powerful than embodied in conventional models, i.e. an adversary that can create up to f omissions per round, ad infinitum, possibly impacting the same subset $\{L\}$ of receivers or senders, which we refer to as the constrained unbounded failure model – up to f omissions per round is the constraint. Even with an infinite number of rounds, no vehicle in $\{L\}$ can succeed in sending or receiving the necessary messages – an obvious and serious threat to safety. Classical PAR or masking protocols fall apart in the presence of such an adversary. Given that they may never terminate, it follows that $T_1 = T_2 = T_3 = \infty$ with classical protocols.

Conclusions: assumption (A1) is unrealistic with IVNs; assumption (A2) is invalid with SC scenarios; assumption (A3) is invalid with V2V wireless mobile communications. Novel models and novel solutions are mandatory.

TABLE III. VARIABLES AND NOTATIONS

cs: current number of vehicles in a cohort, upper bound cs^*
r_X : rank of vehicle X in a cohort
$\Psi_t(X)$: X 's situational data current at time t
$\Phi_\theta(X)$: situational data that shall hold true for X at some future time θ
F : type of SC scenario
$M(Z,F)$: message sent by Z , cyber phase 1
$C(X,Z,F)$: candidacy message returned by e-vehicle X , cyber phase 2
$D(Z,F)$: decision message sent by Z , cyber phase 3
n : number of receivers or senders in a multipoint communication round
f : highest number of omissions that may occur in every communication round (a function of n)
n_e : number of eligible vehicles
α : time of channel occupancy for a message in the absence of contention
K_x : worst-case access delay to radio channel in the presence of x contenders
O_x : time of channel occupancy by messages sent by x contenders
T_i : worst-case termination time of cyber phase i , $i = 1, 2, 3$
ρ, ρ_e : dissemination times of a N2N message within a string of vehicles
$\Pi(X)$: proxy set of cohort member X
S : safe inter-cohort spacing
γ : longitudinal location inaccuracy

V. OMISSION FAILURES – THE Ω MODEL AND PROXY SETS

Assume we know n , the exact number of vehicles in $R(Z,F)$. Due to safety concerns, any prediction regarding worst-case f shall be highly trustable, i.e. stated rigorously so as to be carefully scrutinized by safety authorities and other stakeholders in the IVN domain. Trivially, $f = n-1$ is the most pessimistic assumption not leading to impossibility. We need a modeling whereby f could be “tunable” at will, encompassing every possible choice, i.e. $0 < f < n$. Picking up a small (resp., large) f leads to better (resp., poorer) performance, but to smaller (resp., higher) confidence ratios. Since values of n (resp., n_e) may change at every invocation of SGcast (resp., CCast), it is impossible to retain some fixed f (resp., f_e) a priori once forever, as done with classical solutions. Consequently, we need to provide ourselves with a novel modeling for omission failures in order to overcome these difficulties.

A. The Ω Model

Integer function $\xi(n) = \lceil 2n/3 \rceil$ appears to be the most pessimistic, hence safest, function matching the physics of mobile wireless communications, while not leading to impossibility, as shown Table IV. No message is delivered in worst-case runs with $n < 3$. This is not a problem. Indeed, prefixing physical maneuvers with 3 cyber phases is useful only when sensing-based capabilities may not suffice for avoiding hazardous situations from happening, which is notably the case in dense traffic conditions (n is large).

TABLE IV. OMISSIONS AS A FUNCTION OF NUMBER OF VIRTUAL LINKS

	$1 \leq n \leq 2$	$3 \leq n \leq 5$	$6 \leq n \leq 8$	...
$f = \lceil 2n/3 \rceil$	n	$n-1$	$n-2$	...
$n-f$ (actual deliveries)	0	1	2	...

On the contrary, finely tuned insertion of a vehicle on a lane that hosts 1 or 2 vehicles only can be safely performed with, e.g., radars and side-looking sensors. Thus, choosing some integer function $\xi(n) \leq \lceil 2n/3 \rceil$ for defining f is perfectly acceptable. The above reasoning trivially holds with $n = 0$, a case arising when insertion is to be performed in between 2 cohorts, i.e. when the spacing available for insertion is S . It follows that the worst-case conditions we are interested in imply considering a maximally compact cohort of highest size cs^* (corollary MC). This completes the presentation of our omission failures model, denoted Ω , which is based on combining $f = \xi(n)$ with the constrained unbounded model. To the best of our knowledge, Ω appears to be the most extreme omission failure model coined so far, i.e. the closest to models that would lead to impossibility. As a result, worst-case time bounds calculated under the Ω model are endowed with highest confidence ratios.

Imagine now that there is a solution such that Xcast worst-case termination times would not be linear functions of f , hence of n . Problems that result from invalid assumptions (A1) and (A2) would vanish. Such a solution exists, thanks to cohorts.

B. The Power of Cohorts

The power of cohorts lies with the possibility of designing management and coordination protocols based on N2N communications. Since time redundancy cannot be considered with the Ω model, we must turn our attention to space redundancy [15]. Cohorts lend themselves quite well to maintaining multiple copies of “vital” data, as seen with the GD protocol and the dissemination of N2N messages Nm^* . The GD protocol is a particular (simple) instance of the more general protocols based on proxy sets. Recall that every cohort member X periodically generates a N2N “do you hear me” beacon, denoted $Fb(X)$, which is sent to X ’s neighbors, if any. $Fb(X)$ also carries $\psi_t(X)$, which is a lightweight version of $\Psi_t(X)$. Since we are considering members of the same cohort, there is no need to repeat lt, j, r . Thus: $\psi_t(X) = \{id\ X, \text{current time } t, \lambda, v, s\}$. Dissemination of beacons $Fb(X)$ throughout a cohort permits to maintain multiple copies of up-to-date X ’s situational data across some number of cohort neighbors.

Figure 2. Proxy Sets in a Cohort

In order to avoid incurring the overhead that would result from using GD, scope-limited bidirectional dissemination is resorted to for the Fb^* ’s. Scope size must be necessary and sufficient. Proxy sets are endowed with this property.

1) Proxy sets – Definition

Consider SGcast and f “deaf” members in $R(Z, F)$. The idea is to have some “non deaf” member responding to Z on behalf of “deaf” members. Let $\Pi(X)$ stand for the proxy set of X .

Definition: $\Pi(X)$ is a set of z contiguous vehicles such that at least 1 member W of $\Pi(X)$ receives a message $M(Z, F)$ that should have also been received by X , and W is knowledgeable of the message that X would return to Z .

Defining z is complicated by the fact that the positioning of $R(Z, F)$ and the f “deaf” members in a cohort is unknown. Moreover, recall that n , the size of $R(Z, F)$, thus f , may vary at each new instantiation of SGcast. From corollary (MC), we have: $n \leq cs \leq cs^*$. Consequently, $f \leq f^* = \xi(cs^*)$. It follows that $z = \min \{f^*+1, cs\}$. When $z = cs$, a cohort happens to be the proxy set of every member, a particular instance of the following general case. Let u stand for $\lfloor f^*/2 \rfloor$ and d for $\lceil f^*/2 \rceil$. Let r_X denote X ’s rank in its cohort. Let h stand for $\lfloor cs/2 \rfloor$. Recall that every cohort member keeps an up-to-date copy of current cs via the GD protocol. Every vehicle X in lower half of a cohort ($r_X > h$) is associated a downstream proxy set $DS(X)$, of size $\delta(X) = \min \{d, h - r_X\}$, and an upstream proxy set $US(X)$, of size $v(X) = f^* - \delta(X)$. Every vehicle X in upper half of a cohort ($r_X \leq h$) is associated an upstream proxy set $US(X)$, of size $v(X) = \min \{u, r_X - 1\}$, and a downstream proxy set $DS(X)$, of size $\delta(X) = f^* - v(X)$.

Thus: $\Pi(X) = US(X) \cup X \cup DS(X)$. See Fig. 2, where f^* has been assigned a small value for facilitating the illustration.

2) Keeping proxy sets up-to-date

PSD, the simple Proxy Set Dissemination protocol used for bidirectional dissemination of beacons Fb^* is given Fig. 3, where $PH(X)$ stands for the member of $US(X)$ which carries the smallest cohort rank, and $PT(X)$ stands for the member of $DS(X)$ which carries the highest cohort rank. It is straightforward to check that every cohort member if a member of up to f^*+1 proxy sets, its own proxy set included. PSD is a multi-periodic background protocol.

Let us now introduce the Zebra suite, which comprises SGcast, Ccast, Mcast, and the Altruistic protocol. (Name Zebra mirrors the Z shape of the 3-way handshake taking place between vehicles located on 2 different lanes.) This presentation applies to the Z-driven option (Subsection III-B).

PSD protocol run by V member of $\Pi(X)$, $V \neq X$

- case $Fb(X)$ received from successor(V)
if $V \neq PH(X)$ then forward $Fb(X)$ to predecessor(V)
- case $Fb(X)$ received from predecessor(V)
if $V \neq PT(X)$ then forward $Fb(X)$ to successor(V)

Figure 3. Keeping $\Pi(X)$ up-to-date

VI. FAST RELIABLE XCASTING IN THE Ω MODEL

With proxy sets, in the Ω model, at least 1 virtual link of a cyber phase is omission-free. However, vehicles of interest might not be the actual receivers, possibly the case for e-vehicles (resp., actors) with SGcast (resp., Mcast). Moreover, in the Ω model and in worst-case conditions, it takes at least 3 senders for guaranteeing the delivery of some given message to some intended recipient (see Table IV). This is the rationale for the Altruistic protocol.

A. The Zebra Protocols Suite

1) SGcast by Z in 1 V2V round

Z does SGcast\($M(Z,F)$). At least n_f vehicles in $R(Z,F)$ receive $M(Z,F)$. Via the e-test, every such vehicle (say X) is able to infer which vehicles are eligible (details not shown here). X can thus trigger the Altruistic\(M) protocol.

2) Ccast by e-vehicles in y V2V rounds, $1 \leq y \leq n_e$

Every e-vehicle does Ucast\($C(*,Z,F)$), either right after completing the e-test, or prompted by the Altruistic\(M) protocol. Assume X is an e-vehicle. Since every e-vehicle belongs to the proxy set of every other e-vehicle, in addition to $VA(X)$ and $\Psi_t(X)$, message $C(X,Z,F)$ contains the n_e-1 pairs $\{VA(*), \Psi_t(*)\}$ relative to other members of $E(Z,F)$. Since $f_e = \xi(n_e)$, at least n_e-f_e messages will reach Z.

Filtering serves to combat a worst-case adversary strategy regarding timeliness: no omitted message $C(*,Z,F)$ on any virtual link. Since these messages $C(*,Z,F)$ have identical or nearly identical contents, it is useless and time-consuming to let the n_e e-vehicles send each a message $C(*,Z,F)$. Filtering, which is a cross-layer solution involving N2N communications as well as V2V MAC and transport layers, solves this problem. Filtering is not described here due to space limitations.

3) Mcast by Z in 1 V2V round

Z initiates Mcast\($D(Z,F)$) over $E(Z,F)$, without waiting, right after receiving a $C(*,Z,F)$ message. At least n_e-f_e e-vehicles receive $D(Z,F)$. Each activates the Altruistic\(D) protocol. If assigned the actor status, an e-vehicle initiates SC scenario phase 4, followed by phase 5, its physical maneuvers being inferred from being P or Q.

4) The Altruistic protocol

The Altruistic protocol serves to forward a V2V message as an intra-cohort N2N message. Redundant N2N messages can be detected and eliminated. Activated at the end of phase 1, the Altruistic\(M) protocol forwards $M(Z,F)$ towards e-vehicles, which guarantees that Z will be delivered at least 1 message $C(*,Z,F)$, unless the cohort circulating on lane 1 comprises less than 3 vehicles, in which case Z's insertion is safe (see Subsection V-A). Activated at the end of phase 3, the Altruistic\(D) protocol forwards $D(Z,F)$ towards actors. Let p and ρ_e stand for the worst-case forwarding times of N2N messages by the Altruistic protocol, at the end of phase 1, and phase 3, respectively. It follows that "deaf" vehicles which are members of the same cohort are made aware of their status in bounded time, whatever the cohort size. Rather than depending on n , n_e depends on the e-test, i.e. cohort compactness and relative velocities of vehicles. Thus, 4 or 5 are most plausible values for n_e , yielding $f_e \leq 4$ (see Table IV). N2N link failures would

jeopardize the Altruistic protocol. When such a failure is detected, say between neighbors X and Y, Y decelerates (and X accelerates, if possible) until spacing S is created between X and Y. Therefore, no hazard may result from a cohort split. In other words, a partitioning of a N2N (linear) communication network translates into a physical partitioning of a (linear) vehicular network (a cohort).

5) Discussion

The "3 senders at least" rule should be applied whenever possible (not the case with $M(Z,F)$) for transmitting a safety/time critical V2V message. In the above example, vehicles X and Y as well as X's predecessor and Y's successor shall broadcast (Bcast) a SC message typed CS (cohort split) in order to inform other vehicles. X's predecessor and Y's successor are instructed to do so via the Altruistic\(CS) protocol in both (forming) cohorts. This suffices for guaranteeing a timely delivery of message CS. In our ORM example, Z would be alerted by X, and/or Y, and/or their respective immediate neighbors. More generally, in every SC scenario, there is a need for "aborting" quickly a risk-prone maneuver that has been, or is about to be, undertaken. Our results can be used for instantiating fast reliable "abort" Bcast protocols.

B. Performance Analysis

1) The setting

Consider the IEEE 802.11p/DSRC standard. The 6 Mbits/s channels are accessed via a CSMA-CA protocol. SC messages are transmitted on a specific SC channel. Let g stand for the number of SC channel contenders when Z attempts a transmission (SGcast, Mcast), including itself. As for Ccast, without filtering, an e-vehicle has g^* contenders, $g^* = g + n_e - 1$. In general, g (resp. g^*) is a fraction of the number G of vehicles in the vicinity of Z (resp., an e-vehicle). For example, assume an interference radius of 800 m (twice the radius of correct message receptions, which is in the order of 250 m, plus 300 m for accounting for lane curvatures). On a highway of 4 lanes each direction, and under high vehicle density, G would be in the order of 550. Let K_x stand for the worst-case channel access delay incurred due to contention in the presence of x contenders, and O_x stand for the time of channel occupancy by messages sent by x contenders. These variables are computed considering that the x contenders succeed in accessing the SC channel and in transmitting one message each, prior to the vehicle under consideration, at every channel access attempt. W.l.o.g., for the sake of conciseness, we assume that g does not vary while unfolding a SC scenario.

Since CSMA-CA is a stochastic MAC protocol, there are no exact analytic expressions for K_x . Traditional approaches for circumventing this difficulty consist in defining K_x as a random variable, and derive K_x from the first moments (mean, standard deviation...). Unfortunately, doing this does not provide the worst-case bounds sought-after. Rather than retaining (optimistic) stochastic bounds, we have computed strict upper bounds for K_x , derived from considering a deterministic variation of CSMA protocols which is based on deterministic tree searches adapted to wireless mobile communications – see Section VII. A similar observation is in order regarding p and ρ_e (although the time-bounded channel access problem is simpler with N2N linear communications).

2) Analytic expressions of worst-case termination times

One cannot prove safety unless one provides analytical expressions of worst-case time bounds. Simulations are inappropriate in this respect.

Let T stand for the worst-case total duration of the 3 cyber phases. Computing delays are ignored, since they are negligible compared to K_x and O_x . Duration of channel occupancy by a V2V message in the absence of contention is denoted α , α_1 for message $M(Z,F)$, α_2 for message $C(*,Z,F)$, α_3 for message $D(Z,F)$. To be realistic, α includes the time budget needed for framing and IFS. With filtering, T_2 would be smaller than shown in Table V. The longest forwarding chain for reaching the most distant e-vehicle comprises $f+1$ N2N hops, f_e+1 N2N hops for reaching the most distant actor. Thus, $\rho = (f+1)\tau$ and $\rho_e = (f_e+1)\tau$, where τ stands for the 1-hop N2N transmission delay. Observe that with the Zebra suite, T does not depend on n , contrary to classical solutions (n_e does not depend on n). Moreover, T is not a linear function of f , contrary to classical solutions. Only delays ρ and ρ_e are linear functions of f and f_e , respectively. Therefore, for any given g , worst-case termination times T are remarkably stable, i.e. almost fully immune to variations of f , of n . These results derive from the proxy set concept, i.e. from the cohort construct. This stability property is briefly discussed Section VII.

3) Numerical illustration

All time durations given in Table VI are in milliseconds. For the sake of conciseness, we assign the same value α to α_1 , α_2 , and α_3 . Averaged over $M(Z,F)$, $C(*,Z,F)$ and $D(Z,F)$, the body of a SC message is in the order of 600 bits (no encryption). Adding MAC framing and IFS leads to 900 bits, i.e. $\alpha = 0.15$. Values selected for g are 3 (low contention), 20 (moderate contention), and 50 (high contention). Let $n = 10$ and $f = 5$. Let $n_e = 4$ and $f_e = 2$. Assume 600 Kbits/s for N2N link bandwidth. By construction, IFS and framing, as well as contention delays, are negligible on a N2N link. Thus, $\tau = 1.2$, $\rho = 7.2$, and $\rho_e = 3.6$. With trees of 64 nodes, a channel slot time of 0.1, one finds: $K_3 = 4$, $K_6 = 9.2$, $K_{20} = 15.6$, $K_{23} = 16.8$, $K_{50} = 19.6$ and $K_{53} = 20$.

TABLE V. WORST-CASE TERMINATION TIMES IN THE Ω MODEL

SGcast	$T_1 = K_g + O_g + \alpha_1 + \rho$
Ccast	$T_2 = n_e (K_{g*} + O_{g*} + n_e \alpha_2)$
Mcast	$T_3 = K_g + O_g + \alpha_3 + \rho_e$

TABLE VI. WORST-CASE TERMINATION TIMES ILLUSTRATED (MS)

Zebra suite	T_1	T_2	T_3	T
Low contention	11.8	42.8	8.2	62.8
Moderate contention	25.95	83.4	22.35	131.7
High contention	34.45	114.2	30.85	179.5
Classical PAR or masking protocols	∞	∞	∞	∞

VII. PRACTICAL SIGNIFICANCE OF THESE RESULTS

Time has come to investigate issues relative to safety-critical V2V communications more rigorously than done so far. IVNs are complex socio-technical systems, and mobile wireless communications will inevitably play a central role regarding safety. There are many scientific disciplines which have not been harnessed yet in this respect, such as, e.g., systems theory, real-time networking, distributed and dependable computing. More work should be directed at proving, rather than just simulating. Despite their “theoretical” flavor, the results presented in this paper are believed to have deep practical implications with regard to the future of IVNs.

• Stability property and the engineering of IVNs

With the Zebra suite, worst-case termination times T of the 3 cyber phases are almost fully immune to variations of f , of n . Moreover, proxy sets can be dimensioned at will, considering pessimistic values of f^* if so desired, without jeopardizing T . These appealing features shall greatly facilitate the work to be conducted by engineers in charge of deploying real IVNs.

• Multi-hop SC communications

The Zebra protocols have been designed for 1-hop V2V communications. However, they can be used iteratively for instantiating SC multi-hop V2V communications. For example, vehicles located h ($h > 1$) lanes away from a vehicle that initiates a SC scenario would be able to participate safely in this scenario via $h-1$ intermediate and consecutive executions of the Zebra suite performed by vehicles on adjacent lanes, every execution serving to “cross” a lane.

• Highly reliable short-lived SC communications

Figures obtained for T are well within the expected 0.5 second range quoted in the introduction. In our numerical example, under high contention, it takes less than 200 ms, worst-case, for achieving cyber coordination among vehicles involved in a lane change scenario. Assume velocities equal to 108 km/h. While executing the 3 cyber phases, vehicles travel less than 1.88 m, 3.95 m, and 5.39 m, in low, moderate, and high contention conditions, respectively.

• Existing V2V communication protocols are not suitable

Consider multipoint communications first. For the sake of comparison, worst-case termination times with classical centralized masking protocols under the same valuation of parameters would be $T = 932$ ms and $T = 1.74$ second, in moderate and high contention conditions, respectively, i.e. distances travelled during the 3 cyber phases possibly as high as 27.96 m and 52.2 m, respectively, at 108 km/h. With classical PAR protocols, distances that could be travelled are at least twice as high. Besides being obviously unsafe, such distances can be computed only if one postulates a bounded failure model, which is a very risky assumption when considering inter-vehicular multipoint communications.

Consider now MAC level protocols. In agreement with many authors, we have pointed at a severe weakness of CSMA-CA (IEEE 802.11p/DSRC), namely the impossibility of predicting guaranteed upper bounds for channel access delays. More generally, to the best of our knowledge, there is no

published protocol, be it based on CSMA, CDMA, or TDMA, which solves the time-bounded channel access problem under realistic assumptions. Solutions are still lacking for wireless networks [16]. Therefore, this observation is valid a fortiori in the presence of mobility. Various MAC protocols such as location-based or space division based protocols rest on assuming that different vehicles in proximate neighborhood necessarily compute different positioning coordinates, either at the same time or at times approximately equal. This amounts to assuming that inaccuracy γ is negligible. Since safety mandates making the opposite assumption, such protocols cannot be considered for our purposes. It turns out that the cohort construct, not initially devised to that end, is an essential cornerstone for solving the time-bounded MAC problem. In forthcoming papers, we shall present “deterministic” MAC protocols that guarantee the existence of time bounds for channel access delays in the presence of highest traffic density (hence highest contention). Some of them derive from [17], adapted to mobile wireless settings – see [18] for an early example.

VIII. CONCLUSIONS AND PERSPECTIVES

Many problems related to safety/time critical requirements in IVNs remain open. Regarding safety/time-critical multipoint communications, we have argued that existing approaches do not qualify. We have presented a novel solution, notably the Ω unbounded omission failures model, proxy sets, and the Zebra protocols suite. With these results at hands, one wonders: which organization in charge of public safety or safety regulations would clear usage of classical PAR or masking protocols in SC scenarios? Arguing for “best effort” solutions – read “no guaranteed reliability or/and timeliness” – might be acceptable when there is no alternative. Is it still acceptable when other solutions exist, and human life is at stake? We see no difficulty with having different V2V protocols standardized for different needs. Classical PAR or masking protocols, which are appropriate for the handling of non SC communications, would co-exist with protocols specifically designed for the handling of SC communications. This is exactly what we already have for secure and non secure communications [19], and this is common practice in critical cyber-physical systems. Now that we have established analytical results which demonstrate the value of our solution, simulations and experiments can be undertaken.

REFERENCES

- [1] M.L. Sichitiu and M. Kihl, “Inter-vehicle communication systems: a survey”, *IEEE Comm. Surveys & Tutorials*, vol. 10, 2, 2008, pp. 88-105.
- [2] Y. Toor, P. Mühlethaler, A. Laouiti, and A. de La Fortelle, “Vehicle ad hoc networks: applications and related technical issues”, *IEEE Comm. Surveys & Tutorials*, vol. 10, 3, 2008, pp. 74-88.
- [3] C.Y. Koo, V. Bhandari, J. Katz, and N. H. Vaidya, “Reliable broadcast in radio networks: the bounded collision case”, 25th ACM Symposium on Principles of Distributed Computing (PODC), Denver, July 2006, pp. 258-264.
- [4] X. Défago, A. Schiper, and P. Urban, “Total order broadcast and multicast algorithms: taxonomy and survey”, *ACM Computing Surveys*, vol. 36, 4, Dec. 2004, pp. 372-421.
- [5] Q. Xu, T. Mak, J. Ko, and R. Sengupta, “Vehicle-to-vehicle safety messaging in DSRC”, 1st ACM Workshop on Vehicular Ad Hoc Networks (VANET), Philadelphia, Oct. 2004, pp. 19-28.
- [6] M. Raya, P. Papadimitratos, and J-P. Hubaux, “Securing Vehicular Networks”, *IEEE Wireless Communications*, vol. 13, 5, Oct. 2006, pp. 8-15.
- [7] F.J. Ros, P.M. Ruiz, and I. Stojmenovic, “Acknowledgment-based broadcast protocol for reliable and efficient data dissemination in vehicular ad hoc networks”, *IEEE Trans. on Mobile Computing*, vol. 11, 1, Jan. 2012, pp. 33-46.
- [8] J. Rushby, “Formal methods and the certification of critical systems”, *SRI Technical Report CSL-93-7*, Dec. 1993, 313 p.
- [9] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit, “Ad hoc networking with directional antennas: a complete system solution”, *IEEE Journal Selected Areas in Communications*, vol. 23, 3, March 2005, pp. 496-506.
- [10] G. Le Lann, “Cohorts and groups for safe and efficient autonomous driving on highways”, 3rd IEEE Vehicular Networking Conference (VNC), Amsterdam (NL), Nov. 2011, pp. 1-8.
- [11] G. Le Lann, “Integrated safety and efficiency in intelligent vehicular networks: issues and novel constructs”, *European Commission Transport Research Arena (TRA) 2012 Conference*, Athens, April 2012, Elsevier pub., vol. 48, pp. 951-961.
- [12] Y.B. Ko and N.H. Vaidya, “Geocasting in mobile ad hoc networks: location-based multicast algorithms”, 2nd IEEE Workshop on Mobile Computer Systems and Applications, 1999, pp. 101-110.
- [13] V. Annamali, S.K.S. Gupta, and L. Schwiebert, “On tree-based convergecasting in wireless sensor networks”, *IEEE Wireless Communication and Networking Conf.*, vol. 4, 1, 2003, pp. 1942-1947.
- [14] 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, pp. 1912-1951.
- [15] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, “Basic concepts and taxonomy of dependable and secure computing”, *IEEE Trans. Dependable and Secure Computing*, vol. 1, 1, 2004, pp. 11-33.
- [16] P. Suriyachai, U. Roedig, and A. Scott, “A survey of MAC protocols for mission-critical applications in wireless sensor networks”, *IEEE Comm. Surveys & Tutorials*, vol. 14, 2, 2012, pp. 240-264.
- [17] J.-F. Hermant and G. Le Lann, “A protocol and correctness proofs for real-time high-performance broadcast networks”, 18th IEEE Intl. Conference on Distributed Computing Systems (ICDCS), Amsterdam (NL), May 1998, pp. 360-369.
- [18] M. Mohsin, D. Cavin, Y. Sasson, R. Prakash, and A. Schiper, “Reliable broadcast in wireless mobile ad hoc networks”, 39th IEEE Hawai Int. Conf. on Systems Sciences, 2006, pp. 233-242.
- [19] US Department of Transportation – Research and Innovative Technology Administration, “An approach to communications security for a communications data delivery system for V2V/V2I safety: technical description and identification of policy and institutional issues”, White Paper FHWA-JPO-11-130, Nov. 2011, 52 p.