VANETs: Why use Beaconing at All ?

Dario Rossi, ENST Télécom Paris, France (dario.rossi@enst.fr)
Roberta Fracchia, Motorola Labs, France (roberta.fracchia@motorola.com)
Michela Meo, Politecnico di Torino, Italy (michela.meo@polito.it)

Abstract— **We investigate the broadcast problem in intervehicular networks, aiming at assessing a definitive comparison of two antipodean algorithm classes: the first one makes use of *instantaneous information*, while the second one relies on *longer-term knowledge* gained through a beaconing procedure. Using a realistic microscopic model to represent the vehicular traffic flow, we investigate the performance of the above broadcast algorithm classes by simulation. In order to explore a very large algorithmic design space, we devise a Convex Hull framework that allows us to effectively compare and compactly present the boundaries of the solution space for each algorithm class. By the use of such framework we show that the beaconing approach is not justified for broadcast in suburban and highway VANETs, as there is no performance gain that justifies the complexity entailed by the beaconing procedure.**

I. INTRODUCTION

The last years witnessed an evergrowing thirst for connection, which is pushing for the deployment of communication devices along new roads – precisely those where the nodes of vehicular networks are traveling along. As entertainment applications may still favor unicast communication, infrastructured networks may be necessary to cope with such demanding requirements. At the same time, it is likely that many useful services will be successfully built on top of purely ad hoc multi-hop broadcast communication. Such services are targeted to unknown receivers in a given road segment and, rather typically, road-safety services are paradoxically brought as example of “killer application” in vehicular ad hoc network (VANET). Since messages should cease to be relayed outside the road stretch they pertain to, we may define this communication paradigm as a spatially-constrained *broadcast*: given the appropriate context, we coin the *roadcast* neologism and use both terms interchangeably in the following. Roadcast communication may exploit two antipodean distributed paradigms: the *beaconless* approach relies on *instantaneous* information only, whereas the *beaconed* one exploits a *longer-term knowledge*, maintained through the exchange of beacons.

Built on previous knowledge gained in the field, this work explores the benefits of both the approaches in an exhaustive and systematic comparison, whose aim is to find the best paradigm to implement roadcast services in suburban and highway VANETs. Using a realistic microscopic model to represent the vehicular traffic flow, we investigate and compare the performance by simulation. However, as simulation results forcibly reflect several modeling simplifications and assumptions, rather than precisely quantifying the performance of a specific approach, we prefer to define the *boundaries* of the solution space inside which realistic performance figures may lie. To this purpose, we devise a framework based on the Convex Hull concept, which allows to explore a very wide design and parameter space, and is thus very helpful in assisting the design and evaluation of communication algorithms in VANETs.

Considering as performance indices the reliability, the redundancy and the timeliness of the roadcast service, we explore a large number of design and environmental factors (e.g., roadcast algorithm settings, traffic models, vehicular densities, wireless channel, etc.), investigating several thousands parameter combinations. Our main finding is that beaconing is not suitable in suburban and highway VANETs, in the sense that the complexity of the beaconing procedure is not repayed by a significant performance gain: indeed, beaconless approaches are much more reliable, and we will show that their redundancy is the necessary condition for such reliability.

II. SYSTEM ASSUMPTIONS

We evaluate the effectiveness of roadcast services considering a VANET formed along a linear stretch of a suburban road or highway. The assumption of the road linearity bares additional discussion. Assuming that vehicles are equipped with GPS and a navigation system as well, positions can easily be remapped over a digital map, and thus “linearized”. Moreover, road windings can be expected to occupy a relatively small highway portion compared to the roughly straight one. On such road stretch, roadcast propagation may start from any traveling vehicle or stationary road-side unit, and is targeted to all vehicles within a *relevant area*: vehicles outside this area, instead, never relay the message, so that the medium remains available for other services. The message propagates in the opposite direction with respect to the vehicles movement and, hopefully, it should reach all nodes up to the last one in the roadcast area. In order to gather *conservative* results, we adopt a single-lane, single-direction scenario: several parallel lanes indeed increase the network connectivity [1], and a similar effect is induced when vehicles that are traveling in the opposite direction relay the non-pertinent messages *on purpose* [2].

A. Network Algorithms

To compare the beaconed and beaconless paradigms, we have to define representative examples of both broadcast classes. Thus, we first need to individuate the specific *information* at our disposal under either paradigm, and further devise the general means for exploiting such information: in other words, which are the algorithmic *knobs* that can be tuned on the basis of the gathered information.

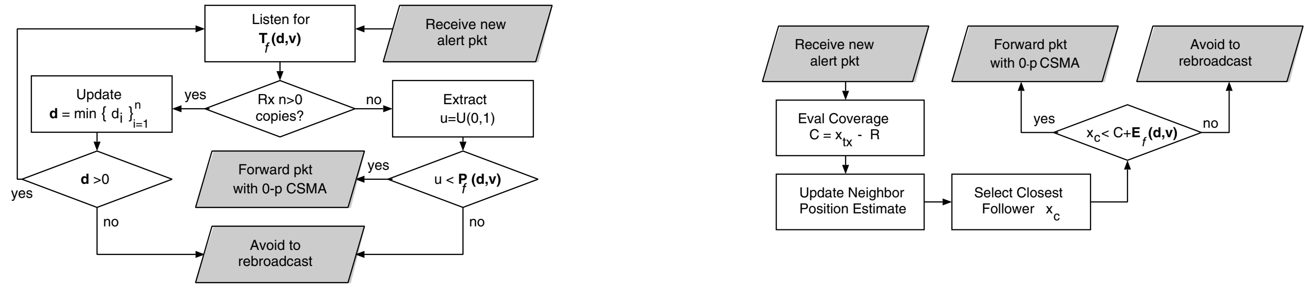

Fig. 1. Beaconless (left) and Beaconed (right) Algorithms Flowchart

1) Beaconless: The beaconless case is the most subtle, as there is a critical need to identify as much useful knowledge as possible in order to take the most effective decision. In all beaconless algorithms, only those pieces of information that are locally available at the receiving node and at the transmitting node (by piggybacking information in the roadcast message) can be used. Following is a list of the instantaneous information that can be used by a beaconless algorithm:

- *Transmission Range*, R, possibly piggybacked in the packet, or inferred by coupling the received signal strength and quality to the transmitter distance D information, or set according to the standard specification given for the adopted technology.
- *Distance from the transmitter*, D, possibly normalized over the transmission range, $d = D/R$. The distance is computed at the receiver through the instantaneous GPS position of transmitter, piggybacked in the message header, and the GPS position of the receiver. The distance d can be used to evaluate the additional coverage that the receiving node would achieve by retransmitting the message: the rational is that the larger the additional coverage, the higher the forwarding probability.
- *Receiver speed*, V, possibly normalized over a maximum vehicular speed, $v = V/V_{\max}$. The speed can be used as an indicator of the local traffic conditions: in particular, small values of speed denote congested traffic (a jam, a platoon of vehicles close to each other and slaken by a slow vehicle, etc.). Under congested conditions, the network is connected and the forwarding probability can be set to small values to reduce the possibility of a broadcast storm.
- *Transmitter speed*, V_T, piggybacked in the roadcast message, normalized as $v_T = V_T/V_{\max}$. We argue that v_T is not likely to be relevant and therefore we disregard it in the following: indeed, as the broadcast propagation occurs in the *opposite* direction with respect to the vehicles movements, the speed of the transmitter barely provides any useful information to the vehicles that follow.

Then, different knobs can exploit the above information, yielding to very different reliability and redundancy performance:

- *Delay of the first transmission attempt*, T_f, is the time after which the packet is handed to the MAC layer. The purpose is to create a short time window that allows the receiver to collect from the neighboring nodes a number of copies of the *same* roadcast message: the number of copies, combined with the information carried by the messages, gives the receiving node some knowledge about its neighborhood (such as position and speed v_T). The larger T_f is, the more accurate the knowledge about the neighborhood is. However, the larger T_f is, the less reactive the message retransmission is.
- *Forwarding probability*, P_f, is the probability that the roadcast message is retransmitted. As already mentioned, the setting of P_f is the core of a beaconless algorithm, and trades off between the opposite needs of having high reliability and low redundancy.
- *Transmission power*, P_T. By properly setting P_T, a desired transmission range R can be selected. In general, this parameter is strongly dependent on the employed technology and regulations cannot be easily changed, therefore we do not consider P_T as a free variable.

The general idea is that whenever one or more copies of the message are received during a waiting phase, the receiver suspend the decision and enters a new waiting phase; otherwise, in case no copies are received, the receiver probabilistically forwards the message on the basis of the information gathered during all the previous waiting phases(s), such as the minimum transmitter distance observed $d = \min\{d_i\}_{i=1}^{n}$ during the waiting phase. Based on this idea, several algorithms can be devised by differently combining the above information and knobs in the left flowchart of Fig. 1: in order to build a fairly large and representative subset of the beaconless class, we select a number of functions, indicated by different IDs in Tab. I, for setting the forwarding delay $T_f(d, v)$ and the forwarding probability $P_f(d, v)$. These functions, that make use of the transmitter distance d and receiver speed v, also depend on the packet transmission duration T_{tx}, the MAC time slot duration T_{slot} and on two free parameters, namely H and K. By abuse of notation, we indicate the functions by their corresponding ID, e.g., P_f^{ID}=1 means that $P_f(d, v)$=$1-(1-d)^K$: a specific algorithm of the considered beaconless class can then be uniquely individuated by a $(T_f^{ID}, P_f^{ID}, K, H) \in \mathbf{N}^4$ tuple. In what follows, we inspect the full cross product of functions $T_f^{ID}, P_f^{ID} \in \{0, 1, 2, 3, 4, 5\}$ and parameters $H, K \in \{1, 2, 4, 8\}$, for a total of $6 \times 6 \times 4 \times 4 = 576$ design tuples per vehicular density ρ as indicated in the right part of Tab. I.

TABLE I

ALGORITHMIC DESIGN SPACE AND SCENARIO PARAMETERS EXPLORED

ID	Beaconless: $T_f(d,v)$	Beaconless: $P_f(d,v)$	Beaconed: $E_f(d,v)$
0	$2T_{tx} + U(0,31)T_{slot}$	P (const)	$E(const)$
1	$T_{tx} + U(0,(1-d)T_x)$	$1-(1-d)^K$	vE
2	$T_{tx} + U(0,(1-v)T_x)$	$1-(1-v)^H$	dE
3	$T_{tx} + U(0,(1-v)(1-d)T_x)$	$1-(1-d)^K(1-v)^H$	-
4	$T_{tx} + U(0,\min(1-v,1-d)T_x)$	$1-(1-d)^{K(v)}$, with $K(v) = 1+(K-1)(1-v)^H$	-
5	$T_{tx} + U(0,(2-v-d)T_x)$	$1-(1-v)^{H(v)}$, with $H(d) = 1+(H-1)(1-d)^K$	-

Beaconed Parameters	Scenario Parameters
$E_f^{ID} \in \{0,1,2\}$	$\rho^{ID} \in$ {Gipps, VDR, TOCA, NaSch }
$B \in \{1,2,3,4,5\}$	
$E \in \{2i\}, i=0..8$	$\rho \in \{5i\}, i=2..8$
Beaconless Parameters	$\rho_c \in \{12..18\}$
$T_f^{ID} \in \{0,1,2,3,4,5\}$	
$P_f^{ID} \in \{0,1,2,3,4,5\}$	$\alpha^{ID} \in \{0,1\}$
$H \in \{1,2,4,8\}$	$\alpha_0 \in \{i/20\}, i=0..5$
$K \in \{1,2,4,8\}$	$\alpha_1 \in \{i/10\}, i=0..5$

2) Beaconing: Besides the information available in the beaconless case, a beaconing service may provide nodes with a more complete knowledge of the neighborhood within a one- or two-hops radius. However, two-hops information is mainly useful in MANETs (where nodes move at relatively slow speed in a bidimensional space), whereas a two-hop knowledge is redundant and outdates at much higher rate in VANETs (due to linear network topology and high speeds). Given this premise, we focus on algorithms that rely on the knowledge of the *position* and *speed* of one-hop neighbors, exchanged through a beaconing procedure: thus, nodes can estimate their neighbor position anytime. Beacons have a small size (about 20 Bytes) and are exchanged periodically every B seconds: to avoid synchronization and reduce beacon collisions, the transmission of each beacon can be jittered as in [4], so that the inter-beacon transmission time is uniformly distributed in the interval $[0.5B,\ 1.5B]$.

Two general distributed approaches can be envisioned at this point, namely *receiver-based* or *transmitter-based* decision: in reason of its robustness in case of packet loss, in the following we focus on the former. Therefore, upon reception of a message, each vehicle decides to relay the message based on the updated *estimate* of its neighbors' positions. Through the knowledge of the transmission range R, vehicles estimate the total distance covered by the message, and the decision boils down to assessing whether the rebroadcast is *necessary* in order to extend the message coverage. In other words, to avoid network connectivity disruption, each vehicle rebroadcasts the message only when it is the further away informed node in the transmission range of the roadcast message.

As the decision is affected by several estimation errors, we argue for the necessity of an error correction margin E in the decision process. Assuming to precisely know the transmission range, the error margin reduces the number of missed forwarding errors due to underestimation of neighbor's distance. The margin E is adapted to the traffic condition through the E_f function reported in Tab. I – although in the beaconed case the type and number of explorable information could actually be much larger (e.g., the error of the transmitter and receiver position estimation, which we leave for further work). The explored beaconed design set (B, E, E_f^{ID}) is the full cross product of the sets $B \times E \times E_f^{ID}$ for a total of $5 \times 9 \times 3 = 145$ tuples per each combination of the scenario parameters that we describe in the next section.

B. Simulation Assumptions

To evaluate the performance of the broadcast communication classes, we developed a custom C discrete event simulator that accurately describes the system behavior from both the network and vehicular dynamics.

Movements of vehicles, as well as distances, are one-dimensional along the direction of the highway, of which we simulate a single-lane 10 km long road stretch: as previously stated, this scenario lead to conservative results, and is thus justified though it may seem rather simplistic at first sight. Vehicular traffic dynamics are described at a *microscopic* level, where each vehicle is individually resolved by its spatial location and speed. Then, a few mathematical rules are used to describe the interaction amongst vehicles, mimiking real drivers' behavior (e.g., such as braking when getting to close to the vehicle ahead, accelerating when no one is in sight, adhering to speed limits, etc.). We implement four different traffic models indicated by ρ^{ID} in Tab. I (namely, NaSch, VDR, TOCA and Gipps), for which description we refer the reader to [5] and references therein. Such models are able to reproduce the most important characteristics of the real traffic dynamics [6], and especially the phase transitions exhibited by the vehicular speed, which is due to the spatial correlation of vehicles. The typical transition happens whenever the vehicular density ρ increases above a critical density ρ_c: the system pass from the *free-flow* regime (where vehicles move nearly at maximum speed without interference from other vehicles) to a *congested* state (where vehicular flow and density are strongly correlated and velocity decreses). We anticipate that the region around the critical density constitute a stiffer scenario [5], where the reliability of safety applications may be compromised. This corresponds to a very hazardous situation, where a high number of vehicles moving at a relatively high speed may be forced to abruptally drop their speed, suddenly braking to avoid an accident. Consequently, as reported in Tab. I, we explore a total of 16 values for the traffic density ρ, with a finer granularity around the critical threshold ρ_c=15.

The roadcast message size is 1000 Bytes, whereas beacon packet size is 20 Bytes, as usually assumed. We assume that messages carry a random identifier chosen by the source, which is cached by every receiver for a small amount of time: in this way, every node is able to forward the same message at most *once*. Also, to limit the broadcast coverage, messages are considered to be relevant only for a L long highway stretch: in other words, messages are not relayed anymore after having traveled L=2 km. We assume moreover that on-board devices transmit such messages at a 2 Mbps rate, with transmission range R=200 m equal for all vehicles. The simulator implements a MAC layer, which accurately models collision and retransmission. We further assume that nodes adopt a 0-persistentCarrier Sense Multiple Access (CSMA) mechanism: in order to avoid collisions, they sense whether the channel is busy before starting a transmission. In case the channel is busy, then the message transmission is delayed, for an amount of time slots uniformly distributed between zero and the contention window size, until the medium is sensed idle. In the CSMA, the contention window is set to 31 time-slots and the time-slot is T_{slot}=20 μs long.

No physical layer is actually implemented, but we take into account varying conditions of the wireless channel by means of two probabilistic packet-level error models α^{ID} (described at lenght in [7] and only briefly overviewed here), and by further varying their settings . Both are high-level models that describe packet loss probability: the first one (i.e., α^{ID}=0) is a simple Bernoulli channel, where each packet is lost with probability p and loss events are independent. In the second model α^{ID}=1, the loss probability depends on the distance from the transmitter: below a given distance packets are never lost, in a "gray zone" of length αR, the loss probability linearly increases and saturates to 1 when the receiver-transmitter distance equals the transmission range. As reported in Tab. I, we consider a loss probability ranging from 0 to 0.25 (in steps of 0.05) for the Bernoulli channel, and a transitional amplitude factor α varying from 0 to 0.5 (in steps of 0.1), in order for the average loss probability to span over the same ranges for both channels.

III. METHODOLOGY AND RESULTS

The methodology we adopt is to perform an extensive set of simulation, based on the guidance of previous knowledge [5], [7], [8], [9], which we briefly summarize here. To clarify our terminology, in the following we refer to simulation *point* as the result of several thousand simulation *runs*, whose stop criterion is that significant variables achieve a confidence level of 99% and a confidence interval of 2%. Each simulation *point* correspond to a full touple of environment settings $(\rho^{ID}, \rho, \alpha^{ID}, alpha)$ and algorithm settings for the beaconed (B, E, E_f^{ID}) and beaconless $(T_f^{ID}, P_f^{ID}, K, H)$ cases. Overall, this would yield to a total of $4 \times 14 \times 2 \times 6 = 672$ combinations of vehicular traffic and wireless errors per single algorithmic setting, which would lead to $672 \times (145 + 576) = 484,512$ different combinations to eplore – a definitively too vaste parameter space.

Luckily, the gained knowledge allow us to build a more efficient investigation path. Specific algorithms falling in the beaconed and beaconless roadcast classes were studied in isolation respectively in [8] and [9]: while the above algorithms clearly do not represent their *whole* respective class, nevertheless their analysis is an helpful starting point. Indeed, both work showed the existence of a sharp transition on the packet reception probability, which is entailed by the phase transition of the vehicular traffic speed. This phenomenon was further explored in [5], which analyzed the impact of different traffic models used in this study from a networking perspective. To summarize the most relevant findings, we may say that, *below* the critical threshold $\rho < \rho_c$, the reception performance is mainly driven by the network connectivity – and thus, that the algorithm class plays a minor role in determining the performance. Conversely, the region *around* the critical density $\rho \simeq \rho_c$ is also critical with respect to reliability – as this performance metric is heavily influenced by both the vehicular traffic model and the specific implementation of the roadcast service. Finally, *above* the threshold $\rho > \rho_c$, the VANET is connected and broadcast communication is very reliable irrespectively of the service implementation – conversely, the number of exchanged messages differs significantly in this zone depending on the broadcast service class.

The above observations have some important consequences. First, the roadcast implementation mainly affects the reliability performance around the critical vehicular density: this region is therefore the most suited to tune the algorithmic setting aiming at maximizing the packet reception probability. Second, as redundancy varies widely at high traffic densities, this region is the most suited to tune the algorithm setting aiming at minimizing the number of transmitted messages. Therefore, in order to reduce the number of simulation to a more manageable amount, we first inspect the algorithmic design space on a smaller set of scenarios (carefully individuated based on the above observations) and select the algorithm settings that yield to the best performance. Then, considering the reduced set of design space touples, we will analyze more deeply the impact of the environmental parameters on the roadcast performance.

A. Algorithmic Design Space Exploration

Let us denote as P the reception probability in the relevant road stretch, and by M the number of relay nodes (or, the number of transmitted messages, beacons excluded): intuitively, these metrics tradeoff, as the redundancy M is the price to be payed in order to achieve a reliability level of P.

In the following, we explore this tradeoff adopting a Convex Hull representation: to put it simply, we plot the (P, M) couples in a bidimensional cartesian space, then find the smallest closed line that encloses all the points. This representation has several advantages: first, it is very compact, and therefore allows us to explore a very wide number of scenarios at once. Second, it is still possible to understand the impact of each parameter in isolation by partitioning the hull over that parameter (i.e., by conditioning the results). Finally, it is possible to build semi-automated processes that help in exploring and trimming the design space.

The convex hull of the whole beaconless and beaconed design space is depicted in the left side of Fig. 2 for $\rho = \rho_c$ and a wireless error-free medium, which corresponding to about 3,000 simulation points – each of which, we recall, is the result of several thousands simulation runs. For the sake of

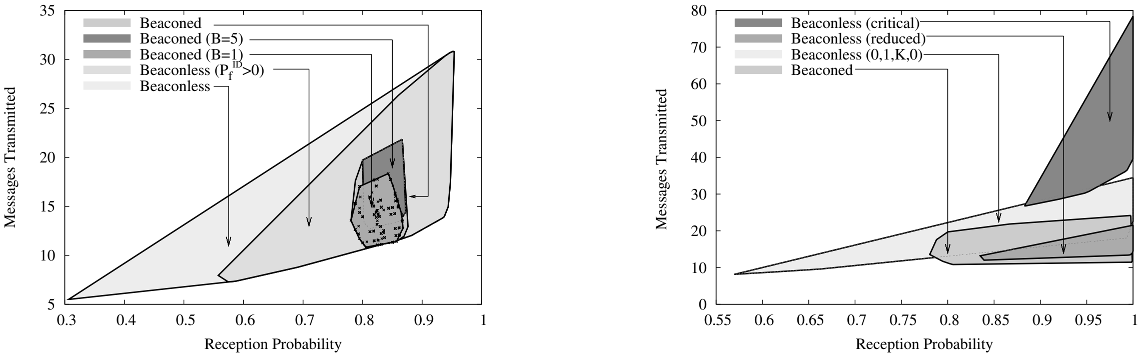

Fig. 2. Convex hull representation of the explored design space around (left) and above (right) the critical density ρ_c

illustration, we show also the 27 individual simulation points corrisponding to the beaconed hull conditioned over B=5. Interestingly, beaconed performance is very clustered and the cloud of points conditioned over the beaconing interval B is very compact: thus, both the beacon interval B the adaptive margin E impact is only marginal, especially with respect to beaconless performance dispersion. Also, it is easy to gather from the picture that beaconless performance is a *superset* of beaconed performance.

Interestingly, it is actually possible to reverse the relationship as shown in the right plot of Fig. 2, i.e., to make a portion of the beaconless solution space a *subset* of the beaconed one – in other words, that it is possible to make design choices such that beaconless algorithms actually outperform beaconed one with respect to both reliability and redundancy. To trim the algorithm design space, we adopt the following approach: considering first the densities around the critical threshold, we lower bound the reliability requirement and discard poor-performing settings; then, we focus on densities above the critical threshold, to further prune the solutions that yield to high redundancy. We stress that performance "requirements" can be specified as *absolute* thresholds (which has the disadvantage of being very sensitive to the threshold value) or *relatively* to the performance of the other settings: following the latter approach, we devise an iterative dichotomic procedure that partition the hull and selects its lower-right portion (corresponding to high reliability P and low overhead M) for the next iteration. The result of the selection process is shown in the right side of Fig. 2, which refers to densities above the critical threshold and an error free wireless medium. Several hulls are reported in the picture, such as the hull corresponding to the whole design space of the beaconed approach, and the *reduced* beaconless group, comprising four $(T_f^{ID}, P_f^{ID}, K, H)$ tuples obtained by the dichotomic procedure. For reference, we also report the hull referring to the $(0, 1, K, 0)$ algorithm proposed in [9]: interestingly, this results show that the exploration of a wider design space enabled us to find more performing settings that could have otherwise be hard to discover. Finally, the picture reports a *critical* beaconless hull, comprising 13 tuples obtained requiring $P \geq 0.85$ for $\rho = \rho_c$ veh/km: this group, obtained with absolute thresholding as opposite to dichotomic hull partitioning, is neverthless very interesting in the context of, e.g., safety applications – where reliability is undouteby the most important performance metric, and where the redundancy is not much of a concern as the propagation of warning messages is not likely to be a continuous nor persistent activity.

B. Sensitivity Analysis to Scenario Parameters

This section presents a sensitivity analysis of beaconless and beaconed algorithm: the aim is to investigate more sistematically the impact of the scenario parameter, so to assess their impact on the performance of the roadcast algorithm. For comparison purposes, we define a beaconed reduced set with the same cardinality (i.e., 4 toupls) of the reduced beaconless group. As the adaptive margin E_f^{ID} function does not yield a clear advantage, and since the beaconed hull is very compact as previously noticed, we select the four extremal points $(B, E, 0)$, where the update interval is $B \in \{1, 5\}$, the error margin is $E \in \{0, 16\}$ and constant $E_f^{ID} = 0$.

Therefore, results presented in this section refer to scenarios generated using four different traffic models, two different channel models each tuned according to six different settings, running two classes of algorithms comprising four different design settings for a total of 384 simulation points per vehicular traffic density. Fig. 3 which depicts, for $\rho \geq 15$, the number of messages versus the reception probability hull (left plot) and the delay at the end of the reception strip versus the reception probability hull (right plot). From either plot, it can be gathered that beaconless approaches are intrinsically more robust to wireless channel losses, due to their distributed flooding based structure. Another important remark is that the *joint* effect of the wireless channel error and the underlying traffic model, so far unexplored to the best of our knowledge, can have dramatic impact on the performance: indeed the reception probability P may drop of a factor of 2 in the beaconless case and a factor of 4 in the beaconed one. Concerning the timeliness, it can be seen that, despite the introduction of a waiting time, vehicles that are far away (i.e., 2 kilometers) at the end of the relevant strip receive the broadcast message after a delay no longer than 200 ms. To this extent, it is worth considering that

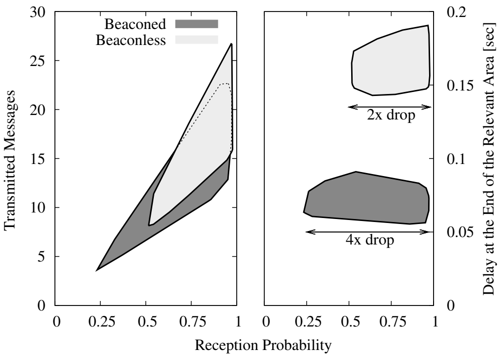

Fig. 3. Sensitivity to wireless channel and vehicular traffic model.

vehicular traffic dynamics happen on a much slower timescale with respect to network dynamics: with a speed of 100 km/hr, in a time interval vehicles could just have moved of about 5 meters, which does not threaten the effectiveness of beaconless algorithms, even for safety applications.

To complete the sensitivity analysis, we define as *impact factor* $I_F(\overline{X}|Y)$ the ratio of the maximum and minimum average values observed for a given performance metric X, achieved by conditioning over a given parameter, but irrespectively of all other parameters. For example, to examine the B impact factor on P, we define 5 sets, one for each value of $B \in \{1..5\}$, and compute the average reception probability $\overline{P}$ over all other parameter values. Finally, we evaluate the impact factor as $I_F(\overline{P}|B) = \max\{\overline{P}|B\}_{B=1}^{5} / \min\{\overline{P}|B\}_{B=1}^{5}$, which intuitively gives a feeling of the imporance of the metric Y on the performance X variability.

Some interesting remarks can be gathered from Tab. II, which report the evaluation of the impact factor for all environmental parameter and design space settings of Tab. I, for both beaconless and beaconed classes. Concerning the traffic model, it is interesting to notice that its impact is more pronouced that the variation of the vehicular density around the critical threshold – which confirms the necessity of carefully taking into account this aspect in order to build a more robust set of results. Then, considering the reception metric, all the environmental parameters (traffic model and density, channel model and loss probability) have greater impact for the beaconed class. Conversely, in the beaconless case, the higher variablity of the redundancy translates into a smaller variability of the recpection metric.

Concerning the design space, the beaconed error margin E is the design parameter that influences performance the most – even more than the beaconing interval B, meaning that a less frequent beaconing coupled to a higher error margin may be a preferable choice. Finally, the most critical design parameter for the beaconless class is, unsurprisingly, the forwarding probability P_f^{ID}; at the same time, the impact factor flattens the importance of H and K, are these free parameters are differently used in different functions – intuitively, their quantitative effect is averaged over possibly different qualitative behaviors.

TABLE II
IMPACT FACTOR OF SCENARIO AND DESIGN PARAMETER

Beaconed

y	$I_F(\overline{P}\|y)$	$I_F(\overline{M}\|y)$
α^{ID}	1.65	1.61
α	2.56	2.80
ρ^{ID}	1.90	1.77
ρ	2.56	2.91
ρ_c	1.63	1.74
E_f^{ID}	1.11	1.48
B	1.39	1.69
E	1.52	2.41

Beaconless

y	$I_F(\overline{P}\|y)$	$I_F(\overline{M}\|y)$
α^{ID}	1.16	1.45
α	1.29	2.34
ρ^{ID}	1.84	3.16
ρ	2.51	4.99
ρ_c	1.48	2.03
P_f^{ID}	2.74	5.64
T_f^{ID}	1.23	1.33
H	1.02	1.08
K	1.12	1.27

C. Why beaconing at all?

In this section we show that the best candidate in the beaconless settings explored outperforms the best beaconed one: we preliminary select the most reliable beaconed service, given by the $(5, 16, 0)$ tuple, with respect to which we normalize the performance of some representative beaconless algorithms.

Left hand side of Fig. 4 depicts, as a function of the vehicular traffic density, the beaconless over beaconed ratios of both i) the number of exchanged messages in the relevant area for each roadcast message, excluding the beaconing overhead and ii) the reception probability in the relevant area as a function of the vehicular traffic density. The picture not only shows that, as expected, beaconless is largely superior to beaconed for what concerns the reception probability metric (especially around critical densities), but also shows that redundancy can be limited. More precisely, the critical group's tuple $(3, 3, 2, 2)$ yields to the upper bound of the reception probability at the expense of, at most, a factor 3 in the number of exchanged messages with respect to the beaconed case. The tuple $(0, 1, 8, 0)$ corresponds to the original algorithm of [9], which limits the overhead to slightly more than a factor of 2.

Finally, the $(1, 4, 8, 2)$ tuple confirms that when vehicles take into account their *own speed*, they may effectively reduce the amount of transmitted messages without compromising at all the service reliability: indeed, the additional number of messages transmitted around the critical density is offset by the significant gain on the reception probability and, moreover, the $(1, 4, 8, 2)$ overhead decreases for densities above the critical threshold. In other words, the redundancy increase (25% at ρ_c) is actually *needed* in order to ameliorate the reliability peformance (14% at ρ_c).

IV. RELATED WORK

During the last few years, several beaconed and beaconless algorithms have been proposed and analyzed in the VANET context, some of which are briefly overviewed here. Generally speaking, the beaconing information is used either in a *position* based fashion or a *topology* based fashion, although the latter approach is not suitable for VANETs as pointed out in [10]. Position-based routing can be mainly divided in *geographic forwarding* [4], where each node forwards incoming packets to exactly *one* of its neighbors, and *restricted flooding* [11], where flooding is restricted to a certain region of the network:

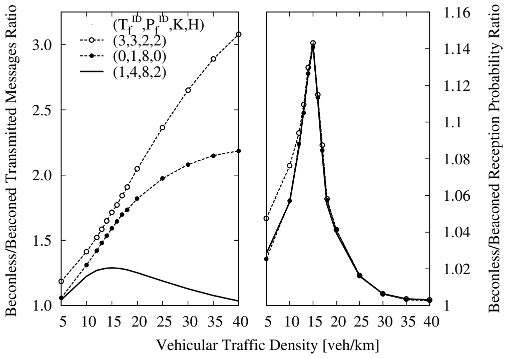

Fig. 4. Beaconless vs beaconed redundancy (left) and reliability (right) ratios

the beaconed algorithms considered in this work belong to the latter family. Beacon-less algorithms are the object of [12], [13], [14], [15], [16]. The idea of letting the decision process depend on the distance is already present in [12], although the decision process is threshold-based rather than probabilistic. A uniform distributed waiting time is present in [13], whereas in [14], [15] the waiting time depends on the position of the previous transmitter and on the destination node. Also, position-based unicast beaconless routing has been studied in [16]. More recently, approaches started to appear that rely on DSRC [17] or IEEE 802.11b [18].

To the best of our knowledge, the idea of adapting the rebroadcast probability to the vehicular speed –which account for significant performance gain– has not been used in the beaconless context so far, whereas [11] proposed to adapt the rate of the beaconing procedure to the vehicular speed. However, we want to stress that the aim of this work is broader than the proposal of a specific algorithm and, to the best of our knowledge, this is the first work that explores the beaconless versus beaconed comparison in the VANET context in a systematic and exhaustive way. Under this light, the work closest to ours [3], [13] compares beacon-less strategies versus beacon-based approaches in the MANET context: as such, authors adopt rather simple mobility models that neglect the existence of correlation among the mobile nodes, which, as we have shown, can play an important role in determining the achievable performance.

V. CONCLUSION

Multi-hop ad hoc broadcast services may be implemented using either *beaconed* or *beaconless* strategies: by means of simulation, we explored a fairly large algorithmic *design* space, taking into account the most important information that can be used to build efficient broadcast algorithms. Using a realistic microscopic model to represent the vehicular traffic flow, we investigated the performance of the above broadcast classes on a broad *parameter* set.

Our main achieviengs are the following. First, simulation results lead us to the rather strong conclusion that the beaconing approach is not justified in suburban and highway VANETs – as there is no evident performance gain to justify the complexity entailed by the beaconing procedure. In other words, the additional amount of messages exchanged by beaconless algorithms is actually *needed* in order to ameliorate the broadcast service reliability in critical situations. Second, we shown that in the beaconless case, taking into account vehicular speed can very effectively assist the rebroadcast decision – both to increase the reception probability at low vehicular densities and to decrease the overhead at high densities. Third, results were gathered and presented through the use of a convex hull framework, which is interesting per-se. Indeed, the framework is very powerful in that it allows to compactly compare any performance metric of interest, and explore a wealth of design and environmental parameter. In a sense, we may say that the convex hull representation is less "misleading" than a single curve, in that is individuate different "regions" in which realistic performance may fall – rather than aiming at precisely quantifying real-world performance by means of simulation, which forcibly reflect a number of simplifying assumptions.

REFERENCES

[1] M.M. Artimy, W. Robertson, W.J. Phillips, "Connectivity in Inter-Vehicle Ad Hoc Networks," *Proc. of IEEE CCECE'04,* May 2004.
[2] Füßler, M. Mauve, H. Hartenstein, D. Vollmer and M. Käsemann, "Location based routing for vehicular ad-hoc network," Proc. MOBICOM'02, Sep. 2002
[3] B. Williams, T. Camp, ' 'Comparison of Broadcast Techniques for Mobile Ad Hoc Networks," *Proc. of ACM MOBIHOC'02*, Jun. 2002
[4] B. Karp, H.T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks," *Proc. ofACM MOBICOM'00*, Aug. 2000
[5] R. Fracchia, M. Meo and D. Rossi, "On the Impact of Traffic Models on Inter-vehicular Broadcast Communication" *Proc. of MedHocNet'06*, Jun. 2006
[6] B.S. Kerner, "Experimental features of the emergence of moving jams in free traffic flow," *J. Phys. A 33,*, pp. 221-228, 2000.
[7] R. Fracchia, M. Meo, D. Rossi, "VANETs: To Beacon or Not To Beacon ?", *Proc. of the Autonet'06 at IEEE Globecom'06*, Nov. 2006
[8] R. Fracchia, M. Meo, and D. Rossi, "Knowing Vehicle Location HELPs Avoiding Broadcast Packets Storm," *Proc. of MP2P'06 at IEEE PerComm'06*, Mar. 2006
[9] R. Fracchia, M. Meo, and D. Rossi, "Avoiding broadcast storms in inter-vehicular warning delivery services," *Proc. of ASWN'06*, Apr. 2006
[10] M. Mauve, A. Widmer, H. Hartenstein, "A survey on position-based routing in mobile ad hoc networks," *IEEE Network Magazine*, Vol. 15, Iss. 6, Nov.2001, pp. 30-39.
[11] S. Basagni, I. Chlamtac, V.R. Syrotiuk, B.A. Woodward, "A distance routing effect algorithm for mobility (DREAM)," *Proc. of ACM MOBICOM'98*, Oct. 1998.
[12] S.Y. Ni, Y.C. Tseng, Y.S. Chen, J.P. Sheu, "The broadcast storm problem in a mobile ad hoc network," *Proc. of ACM MOBICOM'99*, 1999.
[13] Y.C. Tseng, S.Y. Ni, E.Y. Shih, "Adaptive approaches to relieving broadcast storms in a wireless mulihop mobile ad hoc network," *IEEE Transactions on Computers*, May 2003, Vol. 52, No. 5, pp.545-557.
[14] L.B. Briesemeister, L. Shafers, G. Hommel, "Disseminating Messages among Higly Mobile Hosts based on Inter-Vehicle Communications," *IEEE Intelligent Vehicles Symposyum*, Oct. 2000.
[15] M. Heissenbüttel, T. Braun, T. Bernoulli, M. Wälchli "BLR: beaconless routing algorithm for mobile ad hoc networks," *Elsevier Computer Communications*, Vol. 27, No. 11 , Jul. 2004
[16] C. Lochert, M. Mauve, H. Füßler, H. Hartenstein "Geographic Routing in City Scenarios," *ACM SIGMOBILE Mobile Computing and Communications Review*, Vol. 9, No. 1, Jan. 2005
[17] T. ElBatt, S. K. Goel, G. Holland, H. Krishnan, J. Parikh, "Cooperative collision warning using dedicated short range wireless communications," *Proc. ACM VANET'06*, Sep. 2006
[18] Q. Chen, D. Jiang, V. Taliwal, L. Delgrossi, "IEEE 802.11 based vehicular communication simulation design for NS-2," *Proc. ACM VANET'06*, Sep. 2006