VANETs: To Beacon or not to Beacon?

Roberta Fracchia, Michela Meo, Dario Rossi
Politecnico di Torino, Italy
email: {name.surname}@polito.it

***Abstract*— We address the broadcast problem in intervehicular networks by considering two antipodean approaches: the first one makes use of *instantaneous* information on vehicles position, while the other one relies on its *longer-term knowledge*, gained through a beaconing procedure. Using a realistic microscopic model to represent the vehicular traffic flow, we investigate and compare the performance of the proposed algorithms through simulation. We show that a significant performance tradeoff exists: indeed, though the approach based on longer-term knowledge proves to be more efficient in terms of the channel utilization, the instantaneous strategy is intrinsically more robust to errors due to the wireless channel. Finally, we show that this fundamental tradeoff can be turned into an advantage, and we investigate the effectiveness of an *hybrid* solution that combines the diversity of the above approaches.**

I. Introduction

In the last years, wireless communications have enjoyed an amazing expansion in many different directions – from bringing connectivity in many under-development areas at lower costs with respect to cabling, to offering high-revenues Internet connectivity services on-board of inter-continental flights. This evergrowing thirst for connection is pushing for the deployment of communication devices in a new area, i.e., the so-called vehicular ad hoc networks (VANETs).

Typically, given the intolerable number of deaths and injuries caused by car accidents, road-safety services are brought as a paradoxical example of "killer application" in VANETs: in this vision, vehicles sport on-board wireless communication facilities so that, when dangerous situations are detected (either by specific on board sensor devices or by drivers initiative), a *warning message* can be automatically propagated to vehicles that follow by adopting ad hoc networking capabilities. Wireless technologies, such as Dedicated Short Range Communication (DSRC) [1], could be thus used to enhance the perceptive capabilities of the driver in a specular way to what the rear-view mirror does, by providing a "networked spyglass" that allows to foresee dangers beyond the driver cognitive horizon.

Inter-vehicular research [2]–[4] is also fueled by the appeal of an entirely new market segment, which includes geographically-contextualized advertisement, entertainment applications and services aimed at improving passengers traveling comfort. In many cases, the nature of the services and the unknown and varying identity of the users will favor the use of *broadcast*, rather than unicast, communication [5]: for example, road works could advertise their presence through a wireless device beside the usual traffic signals and flags; sensors on board of vehicles (or even directly deployed in the concrete) may detect icy road pavement; under heavy fog, the network itself could reveal the presence of otherwise invisible neighbors. Finally, suburban tollbooths or intelligent traffic-lights in urban environments could receive and relay traffic or road conditions information coming from infrastructured networks.

In this context, adopting a realistic microscopic model to represent the vehicular traffic, we investigate the performance of two distributed algorithms for broadcast communication in VANETs. The approaches we consider are *orthogonal* in the sense that they are based on opposite mechanisms: the Beaconless approach relies on *instantaneous* information on vehicles position only, whereas the Beaconed one exploits its *longer-term knowledge*, maintained through the exchange of beacon packets. Let us anticipate that an interesting performance tradeoff exists: though the Beaconed approach is close to the optimum in terms of the channel utilization, the Beaconless one is intrinsically more robust to wireless channel errors. This tradeoff suggests that no definite and unique answer exists to whether it is better to implement and exploit beaconing or not: different applications or services may find more favorable either strategy depending on its specific requirements. Luckily though, we show that this fundamental tradeoff can be also turned into an advantage, and we investigate the effectiveness of an *hybrid* solution that combines the diversity of the above approaches.

II. Related Work

The problem of routing in ad-hoc networks has been extensively studied in the context of Mobile Ad hoc Networks (MANET), whose results have been the natural starting point for VANET research as well. However, due to the intrinsic differences of these networks, when applied to VANETs, the most promising routing strategies are not as effective as for the MANET context. Moreover, although unicast routing service may be required, it is well recognized that most applications often resort to broadcast communications [5]. Broadcast algorithms can be mainly divided into two groups: namely, *beacon-less* strategies versus *beacon-based* approaches. The performance of these algorithms has been extensively studied and compared in the MANET context [6], [7], adopting rather simple mobility models that neglect the existence of correlation among the mobile nodes: therefore, broadcast communication needs further investigations in the VANET context – that feature very constrained string-type topologies, on the one hand, but strongly correlated movements and very high speeds, on the other hand.

The beacon-less algorithms closest to ours can be found in [7] and [8], where several (either beacon-based or beacon-less) approaches designed for MANETs are compared. More specifically, [7] introduces a uniform distributed waiting time, while our approach, instead, actually *increases* the lowest waiting time in order to ensure that every node gathers sufficient information of the system status. Moreover, the idea of letting the decision process depend on the distance is already present in [8], but the decision process is *threshold-based* rather than probabilistic. Among the beacon-less algorithms specifically designed for inter-vehicle communication, we may cite [9]–[11]. In [9] only one forwarding node is selected by dividing the road portion within the transmission range into segments and by choosing the vehicle in the further non-empty segment; however, this broadcast algorithm requires MAC layer modifications to the IEEE 802.11 standard and does not consider delay constraints. Finally, [10] introduces a waiting time, proportional to the distance from the transmitter, decremented by one unit at each idle slot, while in [11] the waiting time depends on the progress that the node can provide towards the destination.

The use of *beacons* to discover and maintain neighbor relationships is gaining popularity in both unicast and broadcast inter-vehicular communications. In principle, the beaconing information could be used either in a *topology*- or a *position*-based fashion. Recent work [12] though, highlighted that topology-based approaches, while very appealing for MANETs, are not suitable for VANETs: indeed, the high vehicle mobility makes the topology discover and maintenance task very expensive, introducing an excessive amount of control message overhead. Therefore, approaches that rely on the knowledge of the *geographical position* of nodes in the network, gathered through Global Positioning System (GPS), are growing consensus in both academic and corporate research: the basic idea is to exploit more precise information (i.e., position, direction and speed) while avoiding to maintain complex global topological information. This class can be mainly divided in *geographic forwarding* [13], where each node forwards incoming packets to exactly *one* of its neighbors, and *restricted flooding* [14], where flooding is restricted to a given region. The broadcast beacon-based algorithms designed for VANETs and closest to ours are [15], [16]. A routing algorithm for urban-like environments is presented in [15]: every node decides to forward the message only if it estimates to be closer to the *trajectory* toward the destination with respect to its neighbors. In [16] the rebroadcast decision is delayed of an amount of time proportional to the distance from the transmitter. As opposite to traditional beaconing in vehicular communications, where each node advertises only its speed and position, the approach requires each vehicle to advertise its *complete* neighbor list, generating thus a remarkable amount of network traffic: our approach, instead, is to seek for a mechanism that can be easily implemented, while providing at the same time satisfactory performance even under realistic traffic.

III. The Network Scenario

A. Scenario Assumptions

In order to evaluate the effectiveness of the broadcast communication, we consider a broadcast message propagation in a highway-like scenario. Broadcast propagation may start from any traveling vehicle or from stationary road-side units, 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 possible communication services.

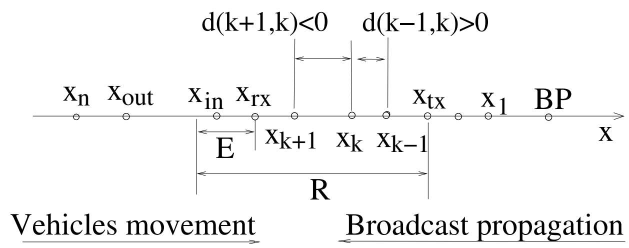

Fig. 1. Schematic representation of the broadcast propagation

In order to describe the algorithms, we adopt the notation sketched in Figure 1. Let the road be represented by an x-axis in the direction of vehicles movement: the relevant area starts in the Broadcast Point (BP), and comprises n vehicles positioned along the x-axis at $x_1 > x_2 > \cdots > x_n$. The broadcast message is forwarded over the relevant area exploiting multi-hop ad hoc communications: the message propagates from $x = BP$ in the opposite direction with respect to the vehicles movement (i.e., toward decreasing values of x) and, hopefully, it should reach all nodes up to the n-th one that is the last one in the relevant area.

Let the distance between nodes i and k be denoted by $d(i, k) = x_i - x_k$, where, clearly, the distance is positive if node i is closer to BP than node k. The assumption of the road linearity bares additional discussion: indeed, it could be argued that roads are actually quite convoluted and, in fact, the geographical distance $d(k-1, k+1)$ can be smaller than $d(k-1, k)$ for some k. However, we point out that, being vehicles equipped with GPS, they are also likely equipped with a navigation system as well: therefore, a digital map will be available at the receiver. In this case, the position advertised by transmitters in the broadcast packets can be re-mapped to a "linearized" road portion, and thus the actual road-distance, rather than the air-distance, could be easily considered. Moreover, such road winding can be expected to occupy a relatively small highway portion compared to the roughly straight one. Finally, the study of a linear road stretch is a preliminary but necessary step, before more realistic but complex scenarios (e.g., involving intersection, motorway overpasses and highway junctions) could be taken into account.

Another important remark is that, although broadcast packets likely refer to a single traffic direction, they are nevertheless received by both directions of the traffic flow. However, for the sake of simplicity, in the following we consider the road stretch in the broadcast propagation direction only. First, we point out that this is basically equivalent to assume that vehicles traveling in the opposite direction can discriminate, via the GPS, if the received message is pertinent to the direction of lanes they are traveling along (which can be done simply by testing whether the direction of the broadcast message propagation is the opposite with respect to their traveling direction). Second, we stress that vehicles traveling in the opposite direction could relay the non-pertinent messages *on purpose*: indeed, in low density networks, these vehicles could helpfully extend the network connectivity and the message coverage – although this may be challenging due to high relative speeds among vehicles, it may as well represent an advantage to exploit.

B. Communication Assumptions

Let us now elaborate on nodes communication capabilities. We do not focus on a specific technology, but we assume that the on-board device transmits at a data rate of 2 Mbps on a R=200 m transmission range, where the low data-rate has been chosen by virtue of its better resilience to the wireless channel errors. Also, we do not investigate the content of the broadcast message but we assume that i) the position of the broadcast initiator, and ii) the position of the last relay vehicle are reported, so that message rebroadcast can be easily confined in the relevant area. Besides, we assume that packet header contains the iii) broadcast initiator identifier as well as iv) a randomly chosen packet identifier, assigned once by the original source: all nodes are required to cache this information on a soft-state table and, at every forwarding hop in the network, each node performs a table lookup, avoiding to rebroadcast a message carrying the same identifiers more than once.

For what concerns the MAC protocol, nodes adopt a 0-persistent Carrier Sense Multiple Access (CSMA) mechanism: in order to avoid collisions, before starting a transmission, they sense whether the channel is busy. In the latter case, the message transmission is delayed, for an amount of time slots $T_{slot} = 20\,\mu$s uniformly distributed between zero and the contention window size $CW = 31$, until the medium is sensed idle.

Finally, in the case where a position-based routing layer is implemented by the communication device, we adopt the common beacon format and the standard beaconing procedure described in the literature. To be more precise, beacons are usually about 20 Bytes long and carry information regarding i) the vehicle identifier, ii) its geographical position and iii) its speed. Every vehicle caches the beacon information along with the time of its reception, which will be later used to estimate the neighbor's position. The inter-beacon transmission interval B is usually a *fixed* value between 1 s and 5 s. Moreover, in order to avoid synchronization and beacon collisions, we jitter the transmission of each beacon as in [13], so that the inter-beacon transmission time is uniformly distributed in [0.5B, 1.5B]. We point out that the beacons transmission requires a very low channel utilization. Indeed, considering 20-Bytes beacons and B equal to 1 s, the beaconing task of every vehicle consumes 160 bps of physical layer bandwidth: at the highest density we consider, namely 50 veh/Km, the 20 vehicles lying in the same transmission range would thus use 3.2 Kbps. Given a physical layer channel capacity of 2 Mbps, the channel utilization would be thus as low as 0.16%; similarly, we can derive a channel utilization of 0.03% when $B = 5$ s.

C. Vehicular Traffic Models

The performance of wireless communication algorithms in mobile ad hoc networks strongly depends on the adopted mobility model, especially when, as in our case, a high mobility scenario is analyzed. In [17], by considering a number of popular traffic models, we showed that, although network connectivity remains largely unaffected by the specific used model, the traffic model plays a critical role in determining the performance of inter-vehicular communication algorithms under study. In this paper we adopt a realistic *microscopic* traffic model that falls in the category of Coupled Maps (CM) models. These models have coarse-grained discrete time, while space is continue, and, as testified by independent empirical traffic measurements [18], [19], they display properties similar to the real traffic dynamics (considering both individual vehicle movements and the correlation between the behavior of neighboring vehicles). The most popular CM model is due to Gipps [20] but, in reason of its simpler notation, we report here the formulation of Krauß's [21]. Each vehicle is *individually* represented by a state vector (x, v, d) describing its spatial location, speed and distance from the vehicle ahead, and, at each time slot, the state vector of each vehicle is updated according to the following set of rules:

$$\begin{cases} \textbf{Velocity-update} \\ \quad v_{safe} = v + \frac{d-v}{\overline{v}/b} \\ \quad v_{des} = min\{v+a, v_{safe}, v_{max}\} \\ \quad v = max\{0, v_{des} - a\epsilon\eta\} \\ \textbf{Motion-update} \\ \quad x \leftarrow x + v \end{cases}$$

where a denotes the maximum vehicle acceleration, b the maximum deceleration, ϵ is the noise amplitude and η is a random number in $[0, 1]$. Moreover, if $\tilde{v}$ is the velocity of the car ahead at time t, we further indicate with $\overline{v} = (v+\tilde{v})/2$ the average of the vehicles velocities used to determine the safe speed bound v_{safe}. The first rules describe deterministic car-following behavior: drivers try to accelerate except when the gap from the vehicle ahead is too small or when the maximum speed is reached. Vehicles speed is limited not only by the maximum velocity v_{max} but also by the desired acceleration a and by the safe speed bound v_{safe}, which depends on the local conditions of the traffic. The speed is also subject to a random perturbation: with a probability η, a vehicle ends up being slower than what calculated deterministically: this parameter simultaneously models effects of i) speed fluctuations at free driving, ii) over-reactions at braking and car-following, and iii) randomness during acceleration periods. In our simulations, we set v_{max} to 135 km/hr and the time-step granularity is $t =$ 1 s, even if these units are assumed implicitly and left out of the equations. The parameters are set to $(a, b, \epsilon) = (0.2, 0.6, 1.0)$, values that are standardly used in the literature [21].

IV. Broadcast Algorithms

In this section we describe the two algorithms that we use to compare BEACONED and BEACONLESS broadcast communications. For comparison purposes, in what follows we will also consider the *optimum centralized* solution that selects the

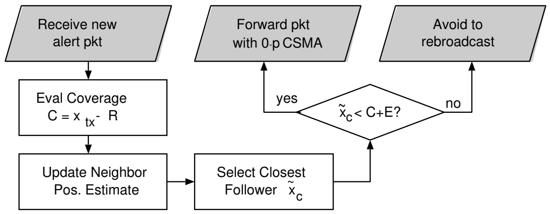

Fig. 2. Flow-charts of the BEACONED algorithms

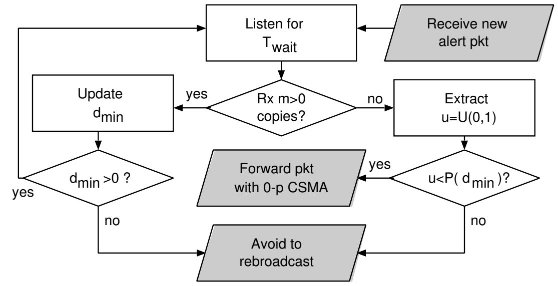

Fig. 3. Flow-charts of the BEACONLESS algorithms

minimum rely nodes set that maximizes the message coverage, as well as the very simple *p–flooding* strategy, ruling that, upon the reception of a new message, each node forwards the packet with independent probability p.

A. Beaconed Algorithm

The basic idea of the BEACONED algorithm is that, whenever some information that has to be broadcasted is received, each node estimates the position of its neighbors by exploiting the information previously exchanged by routing layer through the use of beacon packets: the neighbors' position estimate is then used to decide whether a message should be forwarded or not. The forwarding decision aims at trading-off the number of transmitted messages and the probability to inform the vehicles: the optimal trade-off is derived when only the furthest away informed vehicle forwards the message.

We assume that the routing layer maintains neighborhood state informations, called the neighbors set $\mathcal{S}$ recording for each neighbor i the tuple $(i, \hat{t}_i, \hat{x}_i, \hat{v}_i)$ which stores the node identifier i, the time of the beacon transmission $\hat{t}_i$, the position $\hat{x}_i$ and the speed $\hat{v}_i$ of the node i at the beacon transmission time. Moreover, neighbor table entries are maintained based on their age using a timer T, which can be considered as an "age threshold": below the threshold T, the routing-layer information is considered to be up-to-date, whereas any member of the $\mathcal{S}$ set older than T is discarded. By preliminary simulations, we verified that, as the channel utilization due to the beaconing is very low, the use of a jittered beacon transmission keeps collisions of beacon messages negligible. As a consequence, when the wireless channel is error-free, in every beaconing interval B, nodes gather a *complete* information of their neighborhood. However, in order to increase the robustness in case of lossy wireless media, we select the age threshold to be $T = 2B$, thus considering valid the information received during the last or the second-last beaconing cycle.

Considering the nodes of Figure 1, let us assume that at time t node rx receives a message from node tx. The re-broadcasting decision process at rx, sketched in the flow-chart of Figure 2, initially evaluates the broadcast message coverage area as the zone from x_{tx} to $x_{tx} - R$; note that this information is *precise*, since the *actual* transmitter position x_{tx} is piggybacked in the broadcast packet header. Then, rx updates the *estimate* of its neighbors position, by using the information stored on its neighbor set $\mathcal{S}_{rx}$ as $\tilde{x}_i = \hat{x}_i + (t - \hat{t}_i)\hat{v}_i,\ \forall i \in \mathcal{S}_{rx}$. The receiver then individuates the the *closest follower* c among its neighbors: c is the closest neighbor in the message propagation direction, thus $c = \arg\min_{i\in\mathcal{S}_i}(x_{rx} - \tilde{x}_i > 0)$. Finally, rx decides to rebroadcast the message when either i) c is *outside* the coverage range (i.e., $c = out$ in Figure 1), or ii) c is *inside* the coverage range but within a distance E from the coverage border (i.e., $c = in$). The above two rules, separately described for the sake of clarity, can be expressed in a single threshold-based decision: thus, the message is forwarded iff, $\tilde{x}_c < x_{tx} - R + E$.

Observe that errors on the neighbors position estimation have negative impact on the performance; in particular, the *sign* of the error leads to opposite consequences: over-estimating the distance of a neighbor raises the number of unnecessary relay nodes, whereas distance under-estimation reduces the number of reached vehicles, thus increasing the probability of a premature interruption of the broadcast propagation. Therefore, the safety margin E is included in order to reduce the "missed forwarding" errors (i.e., nodes that erroneously decide to avoid rebroadcasting) in the decision process, in order to account for: i) under-estimation of the neighbor position and ii) possibly missing neighborhood information due to beacon loss.

B. Beaconless Algorithm

The second approach we propose, orthogonal to the previous one, does not rely on any topological information of the network: every node, being unaware of other vehicles' position, should decide independently, and thus probabilistically. The BEACONLESS approach is based on two main ideas: i) to base the forwarding decision on the distance from the closest neighboring relay node, and ii) to introduce a short waiting time before the message forwarding. The rationale behind the idea of letting the decision depend on the distance is the following: when a node hears the message for the first time and its distance from the transmitting node is small, then the additional coverage that can be achieved by re-broadcasting the message is also small[1]: thus, the decision to forward the message should be taken with low probability. The role of the waiting time is to allow nodes to listen for new copies of the broadcast message: this yields to a better estimate of

[1]This assumption is particularly relevant if vehicles have the same transmission range and in the case of unidirectional propagation.

the smaller relay distance, which is then used to tune the forwarding decision process.

The BEACONLESS scheme, whose flow-chart is presented in Figure 3, works as follows. As soon as a node rx has received for the first time a broadcast message from node tx, the node sets the variable d_{min} to the distance $d(tx, rx)$ and starts sensing the channel for a time T_{wait}, during which the node checks if other copies of the same message are received. The waiting time T_{wait} is set proportional to the time T_{tx} necessary to transmit a full-siwe message augmented by a factor f plus a random amount uniformly distributed in $[0, CW]$ of T_{slot} long time-slots:

$$T_{wait} = fT_{tx} + U[0, CW]T_{slot} \quad (1)$$

Notice that T_{wait} is composed by a fixed and a variable component. The role of the fixed component fT_{tx} is to ensure that node rx can acknowledge if other vehicles are forwarding the message, and thus decide to refrain from rebroadcast[2]. The role of the variable component $U[0, CW]T_{slot}$ is twofold: first, it avoids the synchronization of retransmissions from nodes that decided to rebroadcast the message; second, and most important, it allows the *unordered* retransmission of nodes belonging to the same transmission range.

Assume that, after time T_{wait}, node rx has collected m copies of the message from m nodes $n_1, \ldots, n_m$. If at least one node is further away from the BP, i.e., it exists a node n_i such that $d(n_i, rx) < 0$, then rx can safely avoid to forward the message: indeed, the message has already covered an area which is outside the transmission range of rx and a re-broadcast from rx would be useless. Otherwise, if all the transmitting nodes are closer to BP than rx, i.e., if $\forall i,\ d(n_i, rx) \geq 0$, then rx computes the minimum estimated distance $d_{min} \leftarrow \min_i\{d_{min}, d(n_i, rx)\}$ and enters a new waiting phase.

Conversely, when during the sensing period no copies of the message are received, the message is then forwarded with a probability $P(d_{min})$, increasing with the minimum estimated distance d_{min} from the relaying hosts. Thus, in the case where the node decides to rebroadcast, the broadcast message is delivered to the MAC layer. Among the several function families that can be used to set the probability $P(d_{min})$, we choose:

$$P(d_{min}) = 1 - (1 - d_{min}/R)^K, \quad K > 1 \quad (2)$$

Note that in the support $d_{min} \in [0, R]$, the curves are monotonically increasing in both d_{min} and K, and that the function degenerates into d_{min}/R when $K = 1$. The impact of K in (2) is thoroughly discussed in the next section.

V. PERFORMANCE EVALUATION

Results reported in this section are obtained with a discrete event simulator that accurately describes the scenario presented in Section III-B. In particular, the simulator properly describes networking dynamics like collisions, propagation

[2]Since we verified by simulation that the performance results are only marginally affected by any value of $f > 1$, in the following we set $f = 2$.

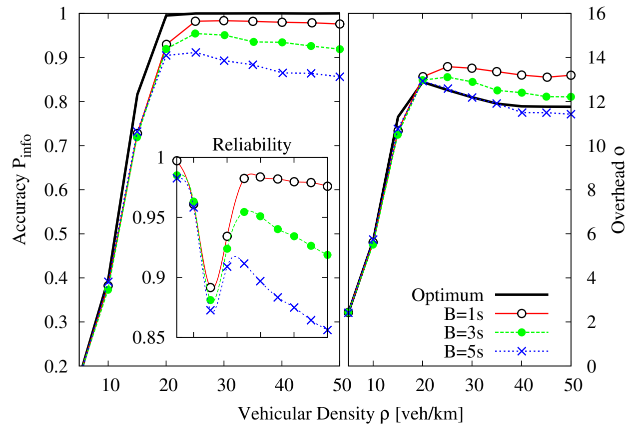

Fig. 4. BEACONED system performance: beaconing interval B impact

delays, transmission times, CSMA delays and implements the vehicular traffic model described in Section III-C.

The performance metrics we consider are: i) the *accuracy* of the message propagation, defined as the percentage of vehicles reached by the broadcast message, that is also the probability that a vehicle is informed; ii) the *reliability*, defined as the ratio of the accuracy with respect to a reference (e.g., optimum) value; iii) the propagation *overhead*, expressed in terms of the number of transmitted messages, iv) the *redundancy*, defined as the ratio of the overhead with respect to a reference (e.g., optimum) value, and v) the *timeliness* expressed as the delay by which the last informed vehicle receives the message. The performance indices are obtained with a confidence level of 99% and a confidence interval of 2%, considering a 2 km long relevant area out of a 8 km long simulated road stretch.

In the following, when we talk about optimum strategy we refer to the case in which the only node that relays a message is the furthest node receiving the message. This strategy, when the wireless channel is error-free, is the optimal one, since it jointly minimizes the number of transmitted messages and maximizes the number of informed vehicles. As a results, on a linear topology, also the number of collided messages is minimized.

A. Ideal Wireless Channel

In this section we evaluate the performance of the proposed algorithms, compared to that of the optimal strategy, when the wireless medium is error free.

Let us first consider the BEACONED approach. Figure 4 depicts the algorithm performance as a function of the density ρ and for different values of the beaconing interval $B \in \{1, 3, 5\}$ s; initially, we do not consider any error margin on the position estimate, thus we set $E = 0$ m. The very left plot of Figure 4 reports the algorithm accuracy P_{info} , that is the percentage of informed vehicles. The optimum centralized strategy is also reported as a reference. It is interesting to notice that when vehicles are in free-flow ($\rho < 20$ veh/km), BEACONED performance is very close to the best case –which is mainly driven by the network connectivity– for any of the considered values of B: indeed, since the road is uncongested, vehicles unlikely modify their speed in 5 seconds or less and

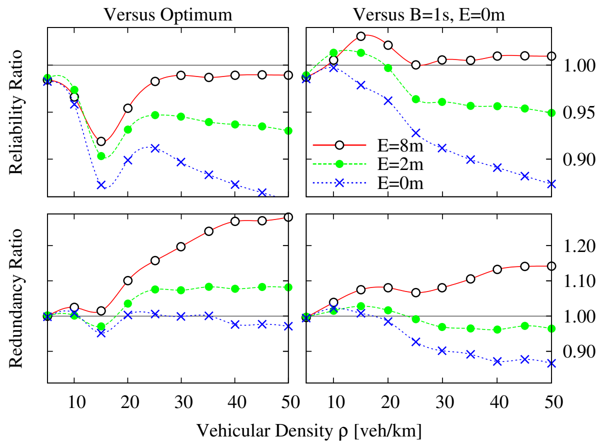

Fig. 5. BEACONED system performance: error margin E impact

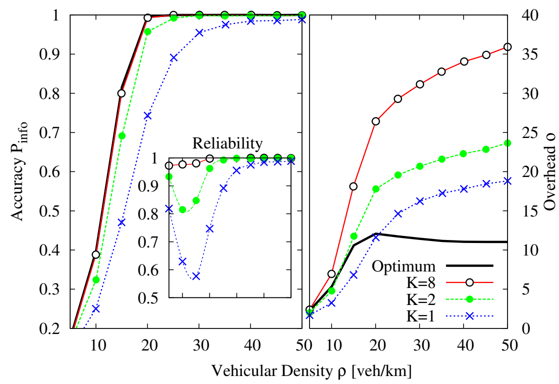

Fig. 6. BEACONLESS system performance: exponent K impact

the position estimate is thus very accurate. On the contrary, when traffic is jammed, vehicles speed may rapidly change and the accuracy of the position estimate is possibly significantly threatened as B increases.

The inset plot depicts the relative system reliability, that is the relative accuracy with respect to the optimum strategy. Interestingly, reliability shows a negative peak around the critical density irrespectively of the beaconing frequency. However, for higher densities, a smaller beaconing interval translates into higher system reliability, whereas for large values of B the reliability decreases as ρ increases above the critical threshold.

Finally, the right plot of Figure 4 reports the algorithm overhead O, expressed as the amount of traffic generated in the relevant area in terms of the number of transmitted messages. Note that we do *not* include the beaconing messages overhead, reported in Section III-B in terms of channel utilization, thus directly investigating the overhead due to the broadcast message propagation. Notice that O also corresponds to the number of *relay* nodes, since every vehicle forwards the message at most one time: it can be seen that, when $B = 1\,\text{s}$ or $B = 3\,\text{s}$, the number of relay nodes running this distributed algorithm only slightly exceeds the optimum one obtained by a centralized strategy. When $B = 5\,\text{s}$, a systematic distance under-estimation leads nodes to refrain from rebroadcasting (i.e., a lower overhead) which in turns reduces the algorithm effectiveness (i.e., a lower accuracy).

In order to reduce the beaconing overhead and ameliorate the system reliability as well, let us now introduce an error margin E on the position estimate: intuitively, the higher the adopted error margin is, the higher the accuracy is. Considering the worst case $B = 5\,\text{s}$, we analyze the benefits of introducing the error margin E, by considering $E \in \{0, 2, 8\}\,\text{m}$. Figure 5 depicts the reliability (top) and redundancy (bottom) ratios with respect to the optimum strategy (left) and with respect to the case with $B = 1\,\text{s}$ and $E = 0\,\text{m}$ (right). Observing the top right plot, it can be gathered that the introduction of an error margin $E = 8\,\text{m}$ makes the $B = 5\,\text{s}$ system more reliable with respect to systems with higher beaconing frequency and without error margin (i.e., $B = 1\,\text{s}$, $E = 0\,\text{m}$). Also, considering the bottom right plot, these benefits come at the expense of a modest 10% increase on the relay redundancy: in other words, only about two additional messages per broadcast cycle are needed, but at the same time the periodic beaconing overhead is reduced by a factor of 5. Observing the left plots, it can be seen that the introduction of an error margin $E = 8\,\text{m}$ brings the $B = 5\,\text{s}$ system performance very close to the optimum in terms of reliability, at the expense of a redundancy increase lower than 25%.

Let us now consider the performance of the BEACONLESS approach when the wireless medium is error free. Figure 6 depicts the results for different values of K in (2) as function of the vehicular traffic density. More on details, the figure reports the accuracy (left plot), the reliability (inset plot), and the overhead (right plot), achieved by the BEACONLESS approach and by the optimum centralized strategy, used again as a reference. First of all, notice that, for high values of K, reliability can be made arbitrarily close to the one of the optimum strategy, even around the critical density region where the BEACONED algorithm has significant performance problems. Moreover, unlike in the BEACONED case, reliability increases with vehicular density, because the probability that at least one node forwards the message increases. Finally, the overhead, that also increases as the vehicular density increases, in the $K = 8$ case is at most three times the one of the optimum case.

For what concerns the delay, shown in Table I, it is worth noticing that, as the overall distance covered by the broadcast message increases ($\rho \leq 20$ veh/Km), the average delay incurred by the message increases as well. At higher densities, the BEACONLESS delay further increases because the number of probabilistically chosen relay increases and thus the average time to access the medium. Conversely, the BEACONED delay, accordingly with the optimum one, slightly reduces, since the number of necessary relay nodes, that are the furthest vehicles in the coverage area, reduces as the density arises.

B. Non-ideal Wireless Channel

The previous section has highlighted some important differences between the BEACONED and the BEACONLESS approach: the former is the least *redundant* and the most performing in terms of delivering time, while the latter is the most *reliable*. However, these differences are marginal,

TABLE I

OPTIMUM, BEACONLESS AND BEACONED AVERAGE DELAY [MS]

Density [veh/Km]	Optimum	BEACONED		BEACONLESS	
		E=0m	E=8m	K=1	K=8
5	39	41	38	84	131
20	58	59	60	214	303
50	47	54	50	271	419

since by tuning the algorithm parameters, the BEACONED approach can be made more reliable (e.g., increasing E) and the BEACONLESS less redundant (e.g., decreasing K). In order to complete our comparative analysis of the two approaches we now introduce more realistic wireless channel models, so that we can also assess the *robustness* of the two algorithms to errors over the wireless channel. If not otherwise specified, we consider the BEACONED strategy with $B = 5\,\mathrm{s}$ and $E = 8\,\mathrm{m}$ and the BEACONLESS with $K = 8$.

Let us now introduce two wireless channel models: in the first one, the probability that a message is not correctly received is Bernoulli with probability β, independently from the receiver distance from the transmitter. Conversely, in the second model the loss probability depends on the distance between transmitter and receiver: indeed, it is well known that the electromagnetic signal degrades with the distance, and empirical studies [22], [23] show that a *transitional* region exists. We denote by αR the amplitude of the transitional region and we assume that: i) below $(1-\alpha)R$ communication is lossless, ii) within $(1-\alpha)R$ and R the packet loss probability sharply increases and iii) above R, that is outside the transmission range, communication will certainly fail. In the following, we usually refer to the average channel loss probability, which is equal to β and $\alpha R/2$ for the Bernoulli and transitional channel respectively. Summarizing,

$$P_{Bern}(d) = \begin{cases} \beta & \text{if } d \le R \\ 1 & \text{if } d > R \end{cases}$$

$$P_{Tran}(d) = \begin{cases} 0 & \text{if } d \le (1-\alpha)R \\ \frac{d-(1-\alpha)R}{\alpha R} & \text{if } (1-\alpha)R < d \le R \\ 1 & \text{if } d > R \end{cases}$$

The above models, despite their simplicity, allow a first grade inspection of the system behavior in presence of non-ideal wireless channels; besides, an advantage of this level of abstraction is that such models are not tied to a peculiar physical modulation, nor to a specific hardware platform, nor to a given environment, but are applicable to a more general extent. Also, we point out that wireless channel models such as Gilbert's [24], which introduce time-correlation in the wireless channel, are not suited for our setup: indeed, since the message propagation happens only once along a given channel, the effect of the time-correlation is likely to be negligible. On the contrary, we expect the effect of the space-correlation introduced by the transitional model to play a remarkable role.

Figure 7 depicts, as a function of the average loss probability, the accuracy of the BEACONLESS and the BEACONED algorithms, for two different values of the vehicular density, namely 20 veh/km (left plot) and 40 veh/km (right plot). First,

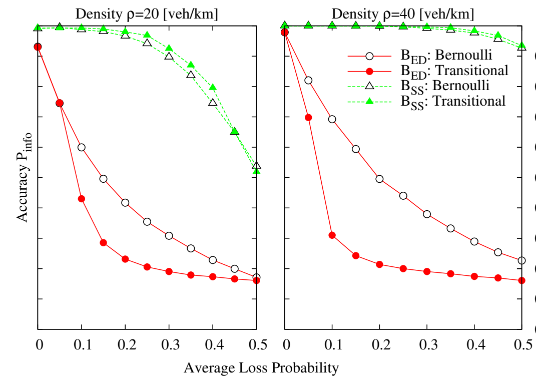

Fig. 7. BEACONLESS (B_{SS}) and BEACONED (B_{ED}) accuracy as a function of wireless channel errors, for Transitional and Bernoulli channel models, at low (left) and high (right) vehicular density

notice that the BEACONLESS algorithm is far more robust with respect to the BEACONED service irrespectively of the channel model. Indeed, even in the worst case, when half of the packets are lost due to wireless channel errors, more than half of the vehicles are effectively warned by the BEACONLESS algorithm, whereas this percentage drops to less than one fifth when the BEACONED strategy is considered. The reason is that for the BEACONLESS strategy packet losses have the effect of virtually reducing the density of the relay nodes, and, thanks to the high redundancy, the effect of errors is visible only at low density (e.g., 20 veh/Km) and high error probabilities. The forwarding decision taken by BEACONED approach is instead based on nodes position estimate so as to reduce redundancy; therefore, if error probability is large it can easily translate in the interruption of the message propagation. Notice also that we consider the possibility that beaconing messages too may be lost due to radio errors, leading to a more imprecise neighborhood knowledge. Moreover, the BEACONED strategy is not only affected by the loss amount, but it is also sensitive to their "placement": indeed, since the algorithm is very effective in selecting only the farther relay nodes, and since the transitional channel yields higher loss probabilities at longer distances, the performance degradation is more remarkable under the transitional channel model. On the contrary, the performance of the BEACONLESS strategy is almost the same for the two channel models and it is driven only by the amount of losses.

Given the non-ideality of the wireless channel, it is no longer possible to devise an optimum centralized solution to the problem. We use as a reference the simple probabilistic p–flooding strategy, regulating the access to the channel with the 0-p CSMA: when $p = 1$ the system reliability is enforced, but at the price of a broadcast packet storm; conversely, by lowering p, to e.g. 1/2, the redundancy is halved at the expense of lower reliability. Figure 8 depicts, for a vehicular density of $\rho = 40\,\mathrm{veh/km}$, the reliability and redundancy of the BEACONED strategy ($B = 5\,\mathrm{s}$ with $E = 8\,\mathrm{m}$ and $B = 1\,\mathrm{s}$ with $E = 0\,\mathrm{m}$), the BEACONLESS and 1/2–flooding approaches normalized versus the 1–flooding performance for both the considered wireless channel models. Note that, although the BEACONLESS algorithm reduces the redundancy of about

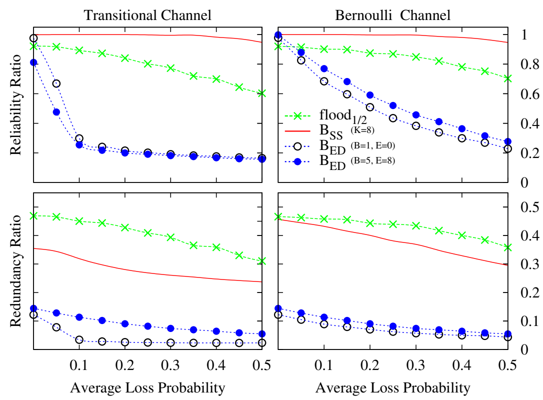

Fig. 8. BEACONLESS (B_{SS}) and BEACONED (B_{ED}) reliability (top) and redundancy (bottom) ratio over the 1-flooding algorithm, for Transitional (left) and Bernoulli (right) channel models

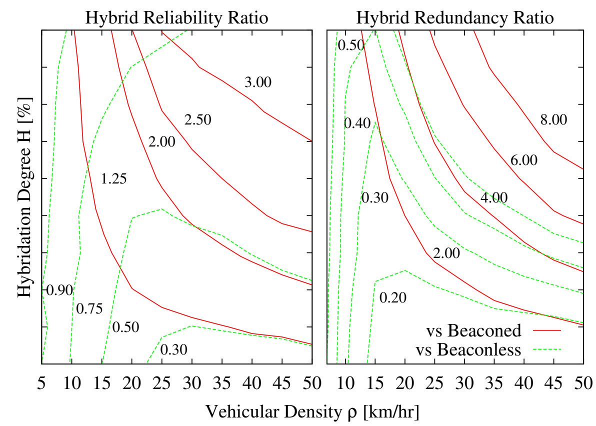

Fig. 9. Hybridating the BEACONLESS and BEACONED approaches: Reliability, redundancy contour plot as a function of the hybridation degree H and the vehicular density ρ

60%–70%, its reliability is practically indistinguishable from 1–flooding. For instance, comparing the BEACONLESS to the 1/2–flooding strategy, observe that for the Bernoulli channel with $\beta = 0.5$, the BEACONLESS approach successfully alerts 28% more vehicles using 8% less messages. Conversely, the BEACONED approach reduces of about 95% the number of exchanged messages, but the system reliability drops significantly (slightly more than 20% in the worst case), and even the introduction of a safety error margin does not mitigate the effect of the wireless channel errors.

C. To Beacon or Not To Beacon?

The previous sections highlighted that several tradeoffs exist between BEACONLESS and BEACONED approaches. However, it is not possible to give a definite and unique answer to whether it is worth introducing a beaconing procedure or not in the VANET context: indeed, different applications or services may prefer one solution or the other, depending on their specific requirements. For example, a road-safety service should implement a BEACONLESS algorithm, since its intrinsic redundancy guarantees high reliability and robustness to wireless channel error. On the contrary, when the safety of the drivers in no longer endangered, the smaller overhead of the BEACONED strategy may be preferred.

Another option consists in proposing *hybrid* solutions that trade-off among reliability, robustness and overhead. In hybrid solutions, the decision of whether to relay the packet is taken by combining the decisions that the BEACONED and BEACONLESS approaches would take. Let $B_{ED} \in \{0, 1\}$ and $B_{SS} \in \{0, 1\}$ be the boolean random variables representing the output of the decision processes of the BEACONED and BEACONLESS algorithms, respectively. B_{SS} is set according to the function $P(d_{min})$ in (2). In the hybrid solution, the two decisions are combined in the following way: the packet is forwarded according to the boolean variable $B_{ED} \vee [(\eta < H) \wedge B_{SS}]$, where η is a uniform random number in $[0, 1]$. In other words, the probability that the packet is forwarded in the BEACONED strategy is increased by the probability to be forwarded in the BEACONLESS strategy, and the last term is weighted by a factor H that represents the *hybridation degree*.

We run several simulations, considering different settings of the BEACONED and BEACONLESS algorithms, channel error models and probabilities, and different hybridation degrees as well. In order to quantify the effectiveness of the hybrid solution, we report results for the Bernoulli channel with packet loss probability $\beta = 0.5$. We consider $B = 5$ s and $E = 8$ m for the BEACONED approach, and $K = 8$ for the probabilistic rebroadcast function (2). Figure 9 depicts the contour plot for the redundancy and reliability ratio of the hybrid approach over both the pure-BEACONED and the pure-BEACONLESS algorithms, as a function of the vehicular density ρ and of the hybridation degree H. For example, let us focus on reliability (left plot). As expected, the hybrid approach allows to greatly enhance reliability of BEACONED (of even a factor of 3), in a way which is far more effective than what was provided by the introduction of the error margin. When $\rho > 30$, it is possible to double the BEACONED effectiveness (line labeled with 2.0 for the BEACONED) while achieving half of the BEACONLESS effectiveness (line labeled with 0.5 for the BEACONLESS). The curves diverge for $\rho < 30$, where a higher hybridation degree is required to keep a factor 2 gain over the BEACONED approach. The hybridation required to achieve half of BEACONLESS reliability, instead, sharply decreases as the vehicular density decreases, because for sparser networks connectivity and wireless errors dominate the performance of the algorithms. Figure 9 also suggests that the hybridation degree H should not be kept constant over all vehicular densities, rather it should be adaptively adjusted to the traffic conditions. As an ongoing work, we are exploring the possibility to estimate the vehicle density through the beaconing procedure so as to adaptively set the values of the hybrid strategy parameters: intuitively, whenever the estimated vehicular density is low, the hybridation degree should be higher and vice-versa.

Finally, Figure 10 depicts the delay at the farmost informed node as a function of the vehicular density for the Bernoulli channel with $\beta = 0.5$ for different hybridation percentages.

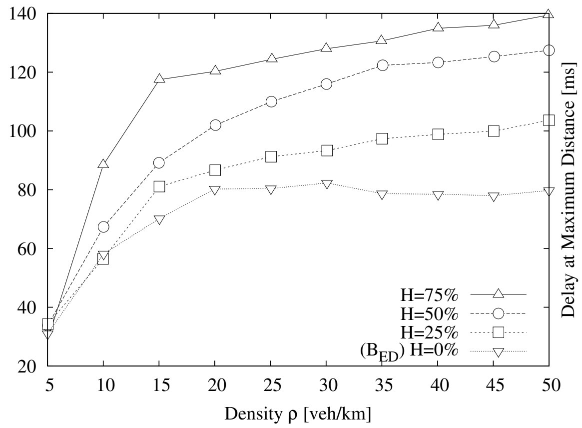

Fig. 10. Hybridating the BEACONLESS and BEACONED approaches: Delay at maximum distance, for different hybridation degrees H, as a function of the vehicular density ρ

The picture not only confirms the small BEACONED delay, but it also shows that hybridation provides a mean for *smoothly* tuning performance between the two orthogonal algorithms. Also, it should be kept in mind that the delay performance is very effective, since vehicles at the end of the relevant area -that is 2 Km away from the danger- are informed in a few tenths of second, when they thus had the chance of moving of a few meters only.

VI. CONCLUSION

In this paper we explored the design space of broadcast communication services for VANETs, considering two kinds of algorithms: the BEACONLESS one, that relies solely on *instantaneous* information on vehicles position, and the BEACONED one, that maintains and exploits *longer-term knowledge* of the vehicular network topology.

By extensive simulation, we compared the different approaches under a realistic microscopic traffic model. Results show that, in the case of ideal wireless channel, by properly tuning the system parameters both algorithms *accurately* propagate a broadcast message. Moreover, if the beaconing overhead is neglected, the BEACONED approach is more *efficient* in terms of the channel utilization. Conversely, the redundancy connatural to the probabilistic decision makes the BEACONLESS algorithm intrinsically more *robust* with respect to errors due to the wireless channel.

Two interesting conclusive considerations can be drawn. First, it is not possible to clearly state a winner of the contest: indeed, different applications or services may prefer to use one strategy or the other depending on their specific requirements. For example, by being more accurate and robust, the BEACONLESS approach is suitable for road safety applications. Second, and most important, we showed that *hybrid* approaches are a promising alternative, as they provide the freedom of finely tuning the performance tradeoffs based on the application needs.

REFERENCES

[1] Dedicated Short Range Communications (DSRC), http://www.leearmstrong.com/DSRC/DSRCHomeset.htm
[2] eSafety Forum, http://europa.eu.int/information_society/programmes/esafety/text_en.htm
[3] CarTALK2000, http://www.cartalk2000.net
[4] Fleetnet, http://www.fleetnet.de
[5] Q. Xu, T. Mak, J. Ko, R. Sengupta, "Vehicle-to-Vehicle Safety Messaging in DSRC," *Proc. of ACM VANET'04*, Page(s): 19-29, 2004
[6] B. Williams and T. Camp, "Comparison of Broadcast Techniques for Mobile Ad Hoc Networks," ' *Proc. of ACM MOBIHOC'02*, Jun. 2002
[7] Y.C. Tseng, S.Y. Ni, E.Y. Shih, "Adaptive Approaches to Relieving Broadcast Storms in a Wireless Multihop Mobile Ad Hoc Network," *IEEE Transactions on Computers*, May 2003, Vol. 52, No. 5
[8] 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/IEEE MOBICOM'99*, 1999.
[9] G. Korkmaz, E. Ekici, F. Ozguner, U. Ozguner, "Urban Multi-Hop Broadcast Protocol for Inter-Vehicle Communication Systems," *Proc. of ACM VANET'04*, Oct. 2004.
[10] E. Fasolo, R. Furiato, A. Zanella, "Smart Broadcast Algorithm for Inter-vehicular Communications," *Proc. of WPMC'05*, Sep. 2005.
[11] H. Fuessler, H. Hartenstein, J. Widmer, M. Mauve, W. Effelsberg, Contention-Based Forwarding for Street Scenarios, *Proc. of WIT'04*, Mar. 2004.
[12] M. Mauve, A. Widmer, H. Hartenstein, "A survey on position-based routing in mobile ad hoc networks," *IEEE NetworkMagazine*, Vol. 15, Iss. 6, pp. 30-39, Nov.-Dec. 2001.
[13] B. Karp and H.T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks," *Proc. of IEEE/ACM MOBICOM'00*, Aug. 2000.
[14] S. Basagni, I. Chlamtac, V.R. Syrotiuk, B.A. Woodward, "A distance routing effect algorithm for mobility (DREAM)," *Proc. of IEEE/ACM MOBICOM'98*, Oct. 1998.
[15] J. LeBrun, C. Chuah, D. Ghosal, M. Zhang, "Knowledge-Based Opportunistic Forwarding in Vehicular Wireless Ad Hoc Networks," *Proc. of IEEE VTC'05*, Jun. 2005.
[16] H. Alshaer and E. Horlait, "An Optimized Adaptive Broadcast Scheme for Inter-vehicle Communication," *Proc. of IEEE VTC'05*, Jun. 2005.
[17] R. Fracchia, M. Meo and D. Rossi, "On the Impact of Traffic Models on Inter-vehicular Broadcast Communications" *Proc. of IEEE MedHoc-Net'06*, Jun. 2006
[18] S. Panwai and H. Dia "Comparative Evaluation of Microscopic Car-Following Behavior" *IEEE Transaction on Intelligent Transportation Systems,* Vol. 6, No. 3, Sep. 2005
[19] B. S. Kerner, "Experimental features of the emergence of moving jams in free traffic flow," *J. Phys. A 33,*, pp. 221-228, 2000.
[20] P.G. Gipps, "A behavioural car following model for computer simulation," *Transportation Res. B*, (15)1981, pp. 105–111.
[21] S. Krauß, P. Wagner. C. Gawron, "Metastable states in a microscopic model of traffic flow,", *Phy. Review E*, (55)1998
[22] G. Anastasi, E. Borgia, M. Conti, E. Gregori and A. Passarella, "Understanding the real behavior of Mote and 802.11 ad hoc networks: an experimental approach," *Pervasive and Mobile Computing (Elsevier)*, vol. 1, Jul. 2005, pp. 237–256.
[23] M. Zuniga and B. Krishnamachari, "Analyzing the Transitional Region in Low Power Wireless Links, " *Proc. of IEEE SECON'04*, Oct. 2004.
[24] E.N. Gilbert, "Capacity of a burst-noise channel," *Bell Sysr. Tech. J.*, p. 1253, Sept. 1960.