![Consiglio Nazionale delle Ricerche logo]

Consiglio Nazionale delle Ricerche

Beaconing Performance in IEEE 802.11p Vehicular Networks: the Effect of Radio Channel Congestion

F. Librino, E. Renda, P. Santi

Technical report

![IIT logo]

Istituto di Informatica e Telematica

Beaconing Performance in IEEE 802.11p Vehicular Networks: the Effect of Radio Channel Congestion

Federico Librino, **M. Elena Renda**, **Paolo Santi**

E–mail: {`name.surname`}`@iit.cnr.it`

Abstract—In this paper, we study the performance of the beaconing mechanism underlying active safety vehicular applications in presence of different levels of channel congestion. The importance of this study lies in the fact that channel congestion is considered a major factor influencing communication performance in vehicular networks, and that ours is the first investigation of the effects of congestion based on extensive, real-world measurements.

The results of our study reveal that congestion has a profound impact on the most important beaconing performance metric, namely, packet (beacon) inter reception time, influencing not only the average value, but also the shape of the distribution. Congestion also considerably increases the frequency of potentially dangerous situation-awareness blackouts, with a likely negative impact on the effectiveness of active safety applications. Our study also reveals that multi-hop propagation of beaconing information can be used as an effective means of lessening the negative impact of congestion on beaconing performance.

I. INTRODUCTION

Active safety applications are a class of vehicular applications aimed at improving a driver's situation awareness, and road safety conditions in general. Applications belonging to this class build on top of a low-level beaconing mechanism, according to which each vehicle periodically broadcasts its status information (position, kinematic data, etc.) to surrounding vehicles. The effectiveness of the beaconing mechanism, which can be informally understood as the extent to which a vehicle is *promptly* informed of the status of *all* surrounding vehicles, is the pre-requisite for any active safety application to meet its often strict design requirements. The importance of beaconing is well understood within the vehicular networking community, which has devoted substantial efforts to characterizing its performance initially by means of analysis/simulation [3], [8], [14] and, more recently, based on real-world measurements [1], [12].

Early studies on beaconing performance [8], [14] have revealed that the beaconing mechanism itself, if not adequately designed, can cause channel congestion in presence of medium to dense vehicular traffic conditions. In turn, channel congestion can severely degrade beaconing effectiveness, potentially impairing the strict communication requirements imposed by active safety applications running on top of beaconing. For this reason, congestion control protocols have been proposed in the literature, with the goal of tuning communication parameters (such as transmission power, beaconing rate, etc.) in order to confine channel congestion below a certain threshold perceived as corresponding to "acceptable" channel conditions. For instance, ETSI indicates a Channel Busy Time (CBT) $\leq 25\%$ as an "acceptable" congestion level for active safety applications [4]. For an overview of the congestion control problem and relative solutions, the interested reader is referred to [13].

Despite the considerable body of literature devoted to congestion control in vehicular networks, very little is known to date on the effects of congestion level on beaconing performance in a real-world scenario. The likely reason of this is that the IEEE 802.11p standard for vehicular communications, and corresponding vehicular communication platforms, have been released only in 2010, and measurement studies have been extensively published in the literature in the last 2-3 years only – see Section II.

In this paper, we fill this gap and present, to our best knowledge, the first measurement-based study of beaconing performance in presence of different levels of channel congestion. Experiments were performed with a five vehicles platoon driven in a typical car following configuration in normal vehicular traffic conditions. The vehicles at the head and tail of the platoon generated a constant traffic floor emulating congestion caused by a large number of vehicles exchanging beacons, while the three central vehicles exchanged beacons. By monitoring beaconing performance on the three vehicles in presence of increasing background traffic, we were able to characterize beaconing performance in presence of congestion. Beaconing performance was mostly evaluated in terms of packet inter-reception (PIR) time, which is identified in [3], [12] as the most important beaconing metric.

The results of our study reveals several interesting insights. First, we discovered that congestion not only severely impact the average PIR time, but also the shape of its distribution: while the PIR time distribution behaves like a power law in a congestion-free channel [12], it behaves like a power law with exponential cutoff in presence of congestion. This exponential cutoff is the consequence of a

much heavier head of the distribution caused by congestion, while the tail of the distribution remains heavier than in the case of congestion-free channel. A second important observation concerns the frequency of situation-awareness black-out events, which is increased of 3 to 4 times when CBT is increased from 18% to 29%. Interestingly, the results also revealed that multi-hop propagation of beaconing information can be used as an effective means of lessening the negative impact of congestion on beaconing performance: when CBT is 29%, multi-hop beaconing reduces the frequency of black-out events of a factor 2.5 with respect to the case of single-hop beaconing. Finally, another major contribution of this study is presenting two Markov-chain based models to accurately predict the average black-out frequency observed on a vehicular link.

II. RELATED WORK

The problem of estimating beaconing performance in presence of channel congestion [8], [14], and of designing congestion control protocols [7], [10], [13], [15], have been extensively studied in the literature. To our best knowledge, though, all existing studies are based on analysis and/or simulation.

Measurement-based investigation of beaconing performance has been addressed in a few recent papers. In [6], the authors consider an intersection collision warning application, and evaluate PDR and RSSI as a function of the distance of the two vehicles from the intersection. In a similar study [11], Mangel et al. evaluates how NLOS conditions impact PDR and RSSI as vehicles approach an intersection. The effects of visibility conditions on channel quality are studied also in [5], where the authors focus on vehicle-to-infrastructure communications.

In [1], the authors present an extensive analysis of PDR in different scenarios for what concerns propagation environment, data rate, etc. The authors also analyze temporal, spatial, and symmetric correlation of PDR values, and conclude that, while temporal and spatial correlation are weak, symmetric correlation is instead quite strong.

In a previous work [12], some of the authors of this paper have characterized beaconing performance not only in terms of PDR, but also in terms of the most relevant PIR metric. In [12], it is shown that PDR and PIR are loosely correlated metrics, thus indicating that PDR cannot be used as a representative metric of situation-awareness. However, all measurements reported in [12] were performed with two vehicles communicating in a congested-free channel. Thus, characterizing real-world beaconing performance in presence of a congested channel is a problem remained open so far.

III. EXPERIMENT SETUP

A. Hardware configuration

Experiments were performed with a setup similar to that described in [12], except for the number of used vehicles which was five instead of two. For vehicular communications we used IEEE 802.11p compliant NEC LinkBird-MX units. Each one was deployed on a different vehicle, together with an omnidirectional WiMo antenna ($108\,mm$ long, $5\,dBi$ gain), a laptop and a GPS receiver. Channel 180 at $5.9\,Ghz$ was selected for radio communication among all the vehicles, being it the one recommended for safety applications (control channel). The transmission power was fixed to $20\,dBm$, with a PHY layer data rate of $3\,Mbps$ and a $10\,Mhz$ channel bandwidth.

We used three compact cars as beaconing vehicles, with the antennas installed at the centre of the roofs as recommended in [9], [11]. Two mini-vans were instead employed as background traffic sources; the antennas on their roofs were instead shifted in the direction of the three beaconing vehicles, as detailed below, in order to increase the impact of congestion on the beaconing process.

B. Multihop beaconing application

The 802.11p protocol stack, together with the multi-hop beaconing application, has been implemented through a Java application running on each vehicle. The multi-hop propagation of data inherently requires two main elements. On the one hand, the usage of an efficient data structure is necessary to keep the information about surrounding vehicles updated; on the other hand, an effective way of including this information (or a portion of it) in the beacons has to be designed, without violating the constraint on the beacon size, which is fixed to $100\,B$. As regards the first issue, we remind that, in a real vehicular environment, several vehicles are encountered during a single trip. Clearly, it is useless to keep the information sent by every single vehicle, since only the beacons sent by the closest ones are likely to be relevant for active safety applications. If A is the considered vehicle, we define as $\mathcal{R}(A)$ the subset of vehicles whose beacon packets carry relevant information for A. The elements in $\mathcal{R}(A)$, as well as its cardinality, are not fixed, but change over time. Therefore, a flexible and dynamic data structure has to be employed. Our choice was to use the `HashMap` data structure, already available in Java. This structure stores `<key,value>` pairs, which can be easily inserted and deleted when necessary. In addition, it is possible to set the *initial size* and the *load factor*, indicating the fraction of occupancy which triggers an automatic size increase. A *refresh* function is also available to remove elements not refreshed for a predefined time interval. As regards the data stored, in our implementation we use the vehicle ID as `key`, and the corresponding *situational information* as the `value`. The former is uniquely associated to
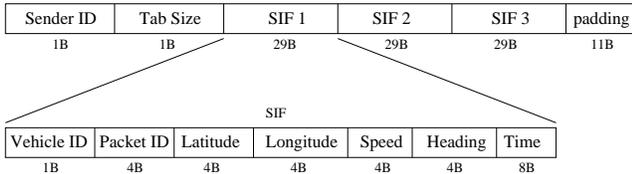
Fig. 1. Beacon format.

a given vehicle, while the latter consists of the following fields: ID of the last received packet ($4\,B$), *latitude* ($4\,B$), *longitude* ($4\,B$), *speed* ($4\,B$), *heading* ($4\,B$) and *GPStime* ($8\,B$). In general, the number of entries in the `HashMap` may vary. However, in our experiments we fix it to 3, since we know in advance that only three out of the five vehicles are sending beacons.

In accordance with recommendations from standardization bodies [2], [16] and previous studies [12], the beacon size is fixed to $100\,B$, not including security overhead. It follows that, in general, information regarding a limited number of vehicles can be included in it. While this can be a problem in dense traffic situations, in our experiments the `HashMap` contains only the information of the three beaconing vehicles, which can easily fit into a single beacon payload. As a result, the beacon contains:

- the sender vehicle ID ($1\,B$);
- the number of vehicles (up to three in our experiments) whose *situational information fields (SIF)* are included in the beacon ($1\,B$);
- the SIFs of each beaconing vehicle, including the sender.

Each SIF contains the ID of the vehicle and the corresponding situational information, as defined above, requiring $29\,B$ overall. The resulting beacon payload is $89\,B$, which is finally padded to $100\,B$, as shown in Figure 1. We remark that in a general scenario with more than three beaconing vehicles, more refined schemes are needed to define the beacon payload. It could be possible to dynamically select a subset of vehicles (belonging to $\mathcal{R}(A)$) whose situational information is broadcast through the beacon. Such a selection could be done according to some priority requirements, or in a round robin fashion. Alternatively, the data to be included in the beacon could be compressed, with the aim of reducing the SIF size. The design of a more sophisticated multihop beaconing strategy is outside the scope of this paper.

The beaconing application running on vehicle A consists of a `send` and a `receive` thread. The former triggers the transmission of a new beacon every $100\,ms$. Before preparing the beacon, it updates in the `HashMap` the situational information of A through a GPS reading. Subsequently, it prepares the beacon payload with the data in the `HashMap`, as described above. The `receive` thread on vehicle A is in charge of updating the `HashMap` upon reception of a beacon from vehicle B, if needed. It first checks the ID of the sender vehicle at the beginning of the beacon payload. If B is not one of the three nodes performing beaconing, the entire packet is discarded, and no additional operations are required; otherwise, the thread parses the rest of the payload. For each SIF, it is first checked whether the corresponding *vehicle ID* equals A. If this is the case, the information in the SIF are ignored. If instead the *vehicle ID* is different, say C, then the $packetID$ in the SIF is compared with the one already recorded in the `HashMap`: if it is larger, it means that the information contained in the beacon is fresher than the stored one, and therefore the `value` in the `HashMap` corresponding to the `key` C is updated with the information contained in the beacon. If the $packetID$ is smaller, an outdated information has been received, which is then ignored.

It can be observed that both the `send` and `receive` threads have access to the `HashMap` data structure. In order to avoid simultaneous operations on the structure, which may impair its consistency, we used the `Semaphore` class in Java, which is designed to ensure exclusive access to the `HashMap`.

Both the threads produce measurement logs in two different text files, recording the content of the `HashMap` at every send or receive event, together with the corresponding system time. In the log file of the `receive` thread, also the ID of the sender vehicle contained in the received beacon is reported.

C. CBR application

In the experiments, two of the vehicles are devoted to background traffic generation. Rather than implementing an entirely new application, we preferred to slightly modify the beaconing application to obtain a constant bit rate (CBR) transmission of interfering packets, whose rate can be tuned at the beginning of the experiment. While the `receive` thread does not require any modification, some details of the `send` thread have to be changed, namely the packet size (and payload content) and the interval T_i between the transmissions of two subsequent packets.

The size of the background packets payload is fixed to $500\,B$. Since this payload is never parsed in our experiments, only the first byte must necessarily contain the sender vehicle ID. In fact, as explained above, this byte is checked at the receiver to verify whether the received packet is a valid beacon or a background packet.

Once the desired background traffic rate R is known, the number of packets to be sent per second can be also derived, as well as the interval T_i. The `send` threads then sends a packet every T_i milliseconds, including in the payload only the situational information of the sender vehicle itself, which is useful for data post-processing. The remaining $471\,B$ are then padded to reach the required payload length.

Clearly, the effective rate is not exactly equal to R, due to two reasons. First, we are not taking into account the additional bytes of the headers; secondly, the Java application can only send an integer number of packets per

second. Nevertheless, we have verified by preliminary tests performed in the lab that both these factors have only a negligible impact on the effective resulting bitrate.

IV. MEASUREMENT CAMPAIGN

The aim of our experiments was to shed light on the impact of wireless channel congestion on the beaconing effectiveness. More precisely, we wanted to investigate the extent to which background traffic affects:

- the Packet (beacon) Delivery Rate (PDR), defined as the ratio between the number of successfully received beacons and the number of sent beacons;
- the Packet Inter-arrival Time (PIR), defined as the interval between two subsequent beaconing receptions.

We performed four different experiments, conducted during four Pisa-Florence-Lucca trips. Similarly to [12], each trip was about $160\,km$ long, and consisted of two parts: the former was on a freeway from Pisa to Florence, with speed limit of $90\,km/h$ and two lanes per direction, while the latter was on a highway from Florence to Lucca, with speed limit of $130\,km/h$ and two/three lanes per direction.

Three experiments were performed with a five vehicle configuration – see Figure 2 – and constant R background traffic, with CBR adjusted to have a value of R of $500\,kbps$, $650\,kbps$ and $800\,kbps$ during the three experiments, corresponding to channel busy time values of 18%, 24% and 29%, respectively. A fourth experiment was performed with the three beaconing vehicles only (vehicles V_0, V_1, V_2 in Figure 2) and no background traffic, to estimate the baseline beaconing performance in absence of channel congestion. While we could have used data from previous measurement campaigns [12] to estimate baseline beaconing performance, we decided to undertake a more rigorous approach using exactly the same vehicle configuration and hardware devices used during the five vehicle experiments, so to rule out possible effects of vehicle heights and inhomogeneous hardware performance.

Referring back to Figure 2, notice that vehicles I_0 and I_1 generating background traffic – called *interfering vehicles* in the following – where placed at the head and tail of the platoon. Notice also that I_0 and I_1, with an approximate height of $1.75\,m$ and $2.07\,m$, were significantly taller than the beaconing vehicles V_0, V_1 and V_2, whose heights were $1.45\,m$, $1.46\,m$ and $1.54\,m$, respectively. Using relatively taller vehicles in head and tail position was a design choice made to generate as a uniform traffic floor as possible. Vehicles I_0 and I_1 were intended to mimic a background traffic which is likely to be present in a wider vehicular network, where also surrounding vehicles actively transmit and receive data packets. Since performing experiments with a large number of vehicles is very expensive and logistically challenging, we adopted this feasible solution instead. Finally, observe that since we performed the experiments
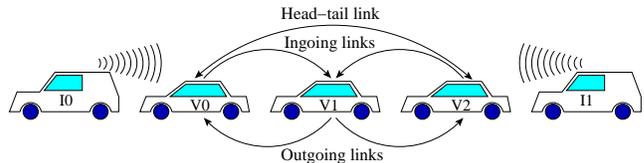

Fig. 2. Reference vehicles configuration for experiments with background traffic.

mostly over 2-lane roads, all the vehicles were allowed to change lane, when possible; this implies that a line of sight was often available also between non adjacent vehicles.

Since beacons are broadcast, we can identify 6 different beaconing links, between any pair of beaconing vehicles. We call *outgoing links* the ones from V_1 towards V_0 and towards V_2. We call *ingoing links* the ones from V_0 and V_2 towards V_1. Finally, direct communications between V_0 and V_2 are also possible, although the presence of V_1 often blocks the line of sight between them. We call the (bidirectional) link between V_0 and V_2 as *head-tail link*.

Each link (V_i, V_j) described above, with $i, j \in \{0, 1, 2\}$, is a physical link, and its performance can be measured by analyzing the beaconing packets transmitted by V_i and received by V_j. However, the presence of a third beaconing node V_k makes it possible to define also a corresponding multi-hop link $(V_i, V_j)_{MH}$, which takes into account also the packets sent by V_i which are not directly received by V_j, but instead delivered via relaying through the path $V_i - V_k - V_j$. Although multi-hop is more relevant for the *head-tail* link, the augmented diversity offered by relaying can be beneficial also for the other links, as will be shown later. From a practical point of view, the measurement of each link performance is done by observing the output files of the beaconing application. Since the `receive` thread writes the content of the `HashMap` immediately after its update every time a new packet is correctly received (together with the sender vehicle ID), the evaluation of the physical link (V_i, V_j) is obtained by considering only the records relative to packets received from V_i. When multi-hop is considered, on the contrary, all the records are taken into account, thus including also the `HashMap` updates due to beaconing packets received from V_k.

V. IMPACT OF WIRELESS CHANNEL CONGESTION

We first investigate the effect of different wireless channel congestion levels on beaconing performance. Beaconing performance is measured in terms of PIR time, which is identified in [3], [12] as the most relevant beaconing metric.

In Figure 3, the complementary cumulative distribution function (ccdf) of the PIR over the head-tail links is shown. A first observation concerns the shape of the distribution, which is strongly influenced by radio channel congestion: while *the PIR time ccdf behaves like a power law without congestion* (confirming what found in [12]), *the distribution behaves like a power law with exponential cutoff in*
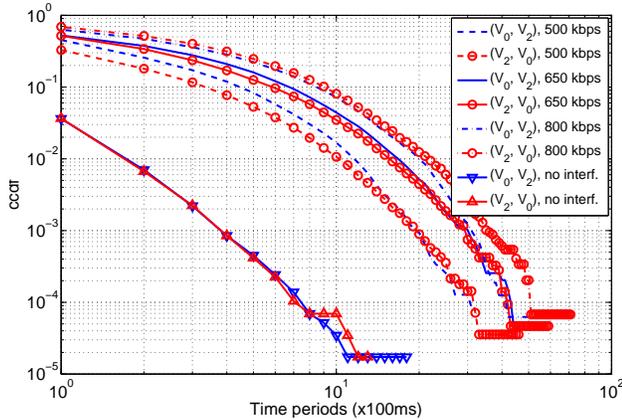
Fig. 3. PIR time ccdf of the head-tail links for different wireless channel congestion levels.
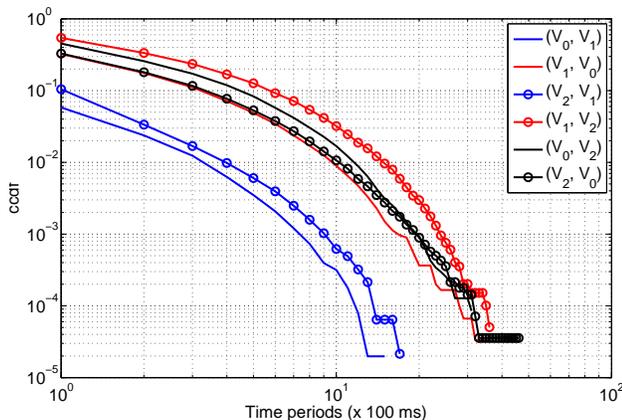

Fig. 4. PIR time ccdf of the six vehicular links, when the congestion level is fixed to $R = 500 \, kbps$.

presence of congestion. While we defer a more accurate characterization of the PIR time distribution in presence of congestion to Section VII, we would like to emphasize here that, while the exponential cutoff of the distribution tail caused by congestion might hint to a thinner tail of the PIR time ccdfs with congestion as compared to the congestion-free distribution, this is not actually the case. As reported in Figure 3, channel congestion severely impacts the *head* of the PIR distribution, implying that the tail of the PIR distribution with congestion, despite its exponential trend, is still fatter than the tail of the congestion-free distribution.

From Figure 3, we also observe that beaconing performance does not depend on the link direction. A small difference between the link (V_0, V_2) and the link (V_2, V_0) is observed only for the congestion level $R = 500 \, kbps$, with a slightly worse performance for the former.

To gain a better understanding of the effect of link directionality at intermediate congestion levels, we analyze the PIR time ccdf of the six links when congestion level is $R = 500 \, kbps$ – see Figure 4. We first observe that *incoming links ((V_0, V_1) and (V_2, V_1)) consistently experience a better*

performance than the other links. This is due to the fact that these links are relatively short, and that they are unlikely to experience the hidden terminal problem. In fact, transmitter vehicles (either V_0 of V_2) are relatively close to the source of background traffic, implying that, when gaining access to the channel, the respective interferer vehicle (I_0 for V_0 and I_1 for V_2) is likely to sense the ongoing beaconing transmission, refraining from transmitting background traffic. Conversely, hidden terminal is likely to be the cause of the bad performance experienced by outgoing links ((V_1, V_0) and (V_1, V_2)). In this case, transmitter vehicle V_1 is relatively far from the two interferers; hence, its beacon transmission is likely not to be sensed by I_0 and/or I_1. Thus, interferers might not refrain from transmitting background traffic during V_1's transmission, causing interference at the receiver vehicles V_0 and/or V_2. It is interesting to observe that hidden terminal effect slightly outweighs also the effect of NLOS, since performance of link (V_1, V_2) is slightly worse than that of the longer, and likely NLOS, link (V_0, V_2).

Focusing on the ingoing links, we also notice that the performance on (V_0, V_1) is better than that on (V_2, V_1). This performance difference might be due to the average link length, as reported in Figure 5. As seen from the figure, although the average length of the (V_0, V_1) link is always longer, this was particularly pronounced during the experiment with $R = 500 \, kbps$. Furthermore, also the location of the interferers play a significant role. Figures 6 and 7 report the joint pdf of the lengths of the two links (V_0, V_1) and (I_0, V_0), and of links (V_2, V_1) and (V_2, I_1), respectively. Looking at the x-axis, we confirm that the average link length of (V_0, V_1) is smaller than that of (V_2, V_1), as already assessed by Figure 5. In addition, looking at the y-axis, we also notice that the average link length of (I_0, V_0) is larger than that of (V_2, I_1). In other words, not only is V_2 farther from V_1, but its strongest interferer is also closer to it, thus fully explaining the performance difference between link (V_0, V_1) and (V_2, V_1). A similar explanation can be used also to justify the performance gap of the two outgoing links, namely (V_1, V_0) and (V_2, V_1).

The Packet Delivery Rate is also influenced by congestion, as reported in Figure 8, where we plot the PDR as a function of the link length for the ingoing and the outgoing links. We focus on the scenario with the highest congestion level ($R = 800 \, kbps$) in order to highlight its impact. As noticed above, the performance of the two ingoing links are quite similar to each other, as well as the one of the two outgoing links. However, the trend of PDR vs. distance is very different. *The PDR on the outgoing links tends to decrease with the link length, while the opposite is true for the ingoing links*. From the joint pdfs of the link lengths in the scenario with $R = 800 \, kbps$, which are not reported here due to lack of space but are similar to those in figures 6 and 7, we observe that even when the lengths of the links (V_0, V_1) and (V_1, V_2) increase, the lengths of the links
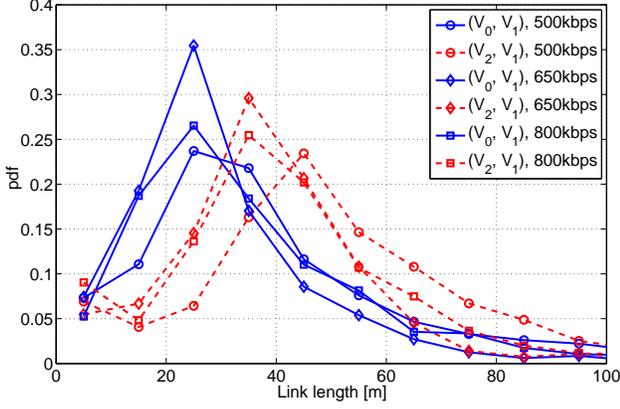
Fig. 5. Probability distribution functions of the length of the ingoing links, for different congestion levels.
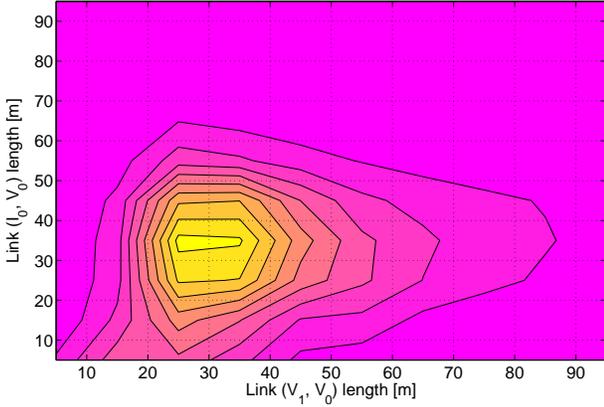

Fig. 6. Joint probability distribution function of the lengths of the two links (V_0, V_1) and (I_0, V_0), when the congestion level is set to $R = 500\,kbps$.
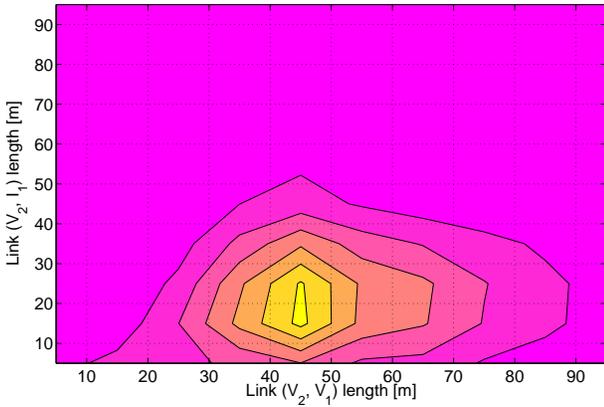

Fig. 7. Joint probability distribution function of the lengths of the two links (V_2, V_1) and (V_2, I_1), when the congestion level is set to $R = 500\,kbps$.
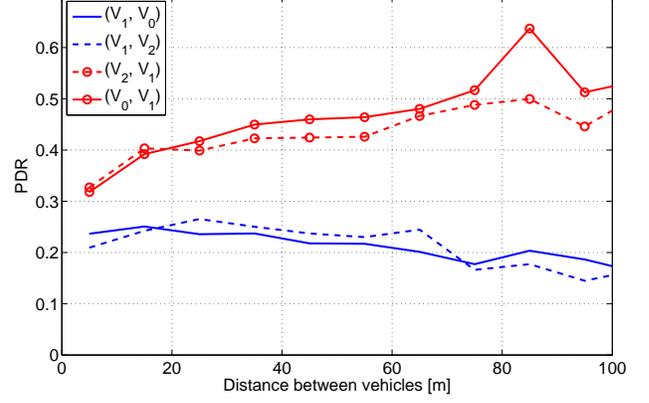

Fig. 8. Packet delivery rate as a function of the link length, when the congestion level is set to $R = 800\,kbps$.

(I_0, V_0) and (V_2, I_1) do not change very much. Focusing on the outgoing links, this means that as long as the distance from the transmitter increases, the receiver remains quite close to the interferer, and the likelihood of an interfering transmission from, say, I_1 corrupting beacon reception at V_2 slowly increases. On the contrary, if we look at the ingoing links, as long as the distance to the transmitter increases, so do the distance to the interferer, due to the fact that vehicle configuration was not changed during the experiments. Hence, the impact of a possible interfering transmission at vehicle V_1 tends to decrease with distance, with a beneficial effect on the experienced PDR.

VI. IMPACT OF MULTIHOP BEACONING

In a real vehicular network, especially in presence of medium to heavy traffic, several vehicles are likely to communicate within the same area. Having a large number of communicating vehicles might cause channel congestion problems which, as shown in the previous section, can severely impact beaconing performance. Spatial diversity can be exploited as a way of counteract this effect. When multi-hop beaconing is employed, each node includes, in its beaconing packets, information about surrounding vehicles, thus actually relaying the beacons sent from other nodes. As a result, the same information can be delivered to the same vehicle through different paths, and a substantial performance improvement can be achieved if the direct link is weak.

As described in Section III, our beaconing application includes in the beacon the information of all the three beaconing vehicles. In case multi-hop performance is evaluated, we derived the PIR on link $(V_i, V_j)_{MH}$ considering all the entries of the output file generated by the `receive` thread of the beaconing application running on V_j. By doing this, we took into account all the updates regarding the situational information of V_i in the `HashMap` on V_j, no matter whether the beaconing packet which triggered the
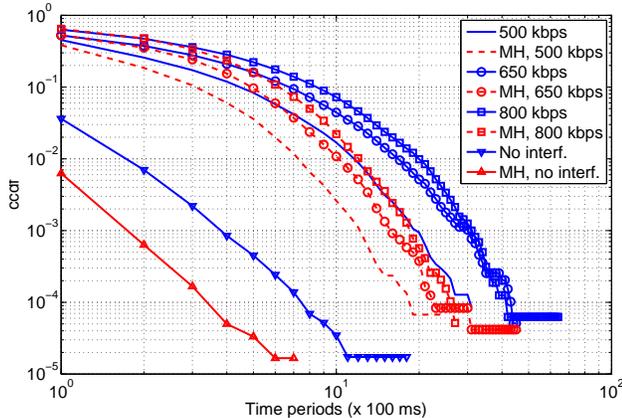
Fig. 9. PIR time ccdf on the (V_0, V_2) link, with and without multi-hop beaconing, for different levels of congestion.

Fig. 10. PIR time ccdf of the links from and towards V_2 when $R = 650\,kbps$, with and without multi-hop beaconing.

update was generated from V_i. On the contrary, when the performance of simple beaconing is evaluated[1], we filtered the output file written by the `receive` thread on V_j to contain only the lines corresponding to updates in the `HashMap` due to beacons received from V_i.

Figure 9 depicts the ccdf of the PIR on the link (V_0, V_2), with and without multi-hop beaconing. The same performance has been observed for the (V_2, V_0) link. As expected, multi-hop beaconing offers a relevant improvement in all the considered scenarios. In particular, we observe that when congestion level is low ($R = 500\,kbps$), or even absent, the entire ccdf curve is shifted downwards. In this situation, the impact of congestion on beaconing performance is not overwhelming, and the effect of the link length and of the more frequent NLOS conditions on (V_0, V_2) is still relevant. When multi-hop is used, shorter links are involved, where in addition LOS is often available, thus highly reducing the average PIR, especially for the scenario with no interferers.

When congestion level is higher ($R = 650\,kbps$ and $R = 800\,kbps$), its effect on beaconing performance becomes predominant. In this case, the multi-hop curves show a considerable improvement only in the tails. In fact, even if we can assume that link (V_0, V_1) is more reliable, due to vehicle configuration, the two links (V_0, V_2) and (V_1, V_2) are hampered by the same interferer node (I_1). Therefore, each of them is likely to experience relatively long PIRs. However, assuming that the fading coefficients of the two channels are almost uncorrelated, the average time necessary to decode a packet over either one of the two links is statistically lower. This explains why the spatial diversity benefit then becomes more evident in the ccdf tails. In absolute terms, *multihop beaconing reduces the probability of experiencing a black-out event ($PIR > 10$) of approximately one order of magnitude*, with a relatively smaller reduction in black-out probability for relatively

---

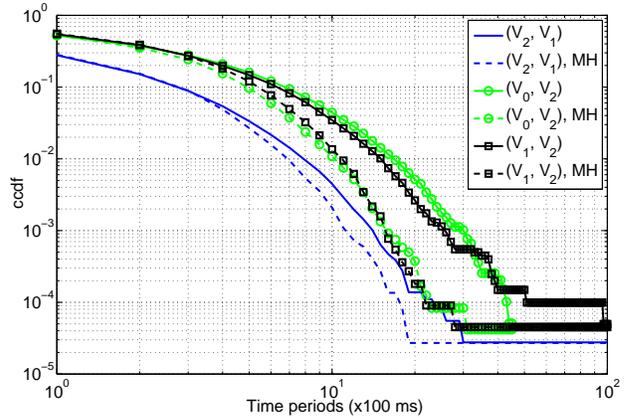[1]The data reported in Section V refers to simple beaconing.

higher congestion level.

In scenarios with a high interference level, as observed, relaying can be beneficial. It is worth noting that this beneficial effect may not be limited to the *head-tail* links. In Figure 10, the ccdf curves of the PIR over the links from and towards V_2 are reported, when $R = 650\,kbps$.

We have already commented on the benefits on multihop beaconing on link (V_0, V_2). However, measurements show that an improvement of beaconing performance with multihop is clearly visible also for links (V_1, V_2) and (V_2, V_1). Interestingly, the beneficial effect of multihop beaconing is more visible on the problematic (V_1, V_2), showing that *spatial diversity can be effectively used to lessen the impact of the hidden terminal problem.*

VII. MODELING CONGESTED LINKS

As observed in Section V, the shape of the PIR time distribution is clearly influenced by congestion: while, as already observed in [12], it behaves like a power law in absence of congestion, it instead follows a power law with exponential cutoff in presence of congestion. In this section, we assess the latter statement, showing how well the measurement data can be fitted by the following power law with exponential cutoff:

$$\mathbb{P}[PIR > k] = ak^{-b}e^{-ck} \tag{1}$$

where a, b and c are parameters to be determined. More precisely, a and b jointly influences the ccdf behavior for low PIR values, which is close to that of a power law, while c determines the shape of the distribution tail. On the grounds of the measured data, we performed an iterative procedure based on Gauss-Newton algorithm to derive the three parameters of the distributions which best approximate the curves of interest.

As an example, we report in Figures 11 and 12 the ccdf of the PIR in some of the analyzed cases. In Figure 11 we plot
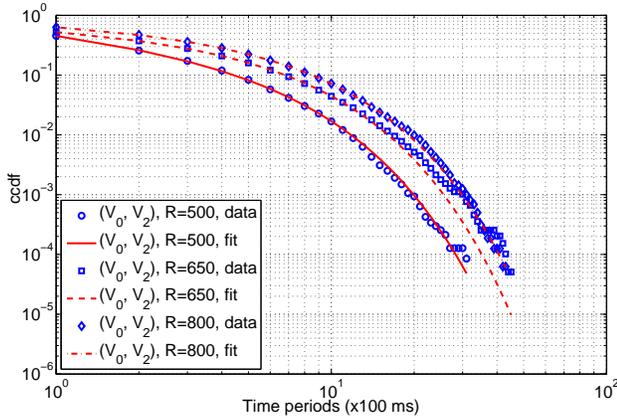
Fig. 11. PIR time ccdf of link (V_0, V_2) for different congestion levels: measured data and fitting curves.
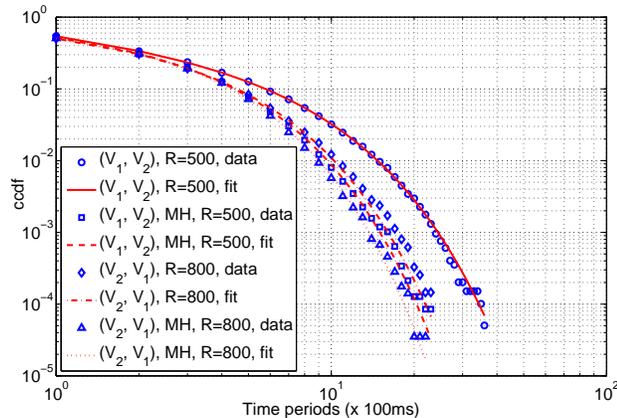

Fig. 12. PIR time ccdf of link (V_1, V_2) and (V_2, V_1) for different congestion levels, with and without multi-hop beaconing: measured data and fitting curves.

the ccdf of the PIR on link (V_0, V_2), for different levels of congestion. The curves obtained through proper selections of the parameters a, b and c are always quite accurate, which confirms the good approximation offered by a power law with exponential decay. The same can be observed for the PIR ccdf of ingoing and outgoing links, with and without multi-hop beaconing, as illustrated in Figure 12.

After deriving the optimal parameters for all the curves, which are reported in Table I, we can make the following observations:

- when increasing the congestion level, the value of a also quickly increases, while b becomes lower and c is unaffected or does not show a clear trend. This is consistent with the fact that congestion mainly impacts on low PIR values, being large PIR values also due to link length or NLOS conditions;
- when enabling multi-hop beaconing, both a and b show very small variations (a slightly increases, b slightly decreases), while c become much higher (often more than doubled). This confirms that multi-hop is

particularly effective in reducing the distribution tail, where spatial diversity can offer the highest benefits, as explained in Section VI;
- the values of a are relatively larger for outgoing links than for the ingoing links, while the opposite holds for the values of b, as expected, since the outgoing links are relatively more impaired by congestion;
- when congestion is not present, the ccdf curves are approximated by a power law, which can be obtained from the general expression in (1) by setting $c = 0$. This is consistent with the results assessed in previous work [12].

VIII. MODELING BLACK-OUT EVENTS

In this section, we present Markov chain-based models for modeling black-out events. Black-out events are defined as occurrence of a PIR time exceeding a threshold, empirically set to $1sec$, corresponding to loosing at least $\ell = 10$ consecutive beacons. As explained in [12], black-out events severely impair onboard situation-awareness, which motivates our interest in modeling these events.

In the following, we analyze the black-out statistics collected during the experiments, and compare these with predictions obtained from different models. The first considered black-out model is the one proposed in [12], according to which the average inter-black out time T_{bo}, defined as the average interval between two successive black-outs, can be estimated as follows:

$$T_{bo} = \mathbb{E}[PIR] \cdot \frac{1}{p_{bo}} \ , \tag{2}$$

where $p_{bo} = \mathbb{P}[PIR > k]$ is the black-out probability[2] and $\mathbb{E}[PIR]$ is the expected PIR time, both derived from the measured PIR time distribution. In fact, a new PIR value is expected to be observed every $\mathbb{E}[PIR]$ seconds and, assuming independence of black-out events, we can model them as Bernoulli trials.

Indeed, equation (2) is inaccurate, due to the fact that, as observed in [12], black-out events *are not* independent. In fact, black-outs are typically caused by bad channel conditions, which usually show strong temporal correlation. Hence, the expression in (2) can be considered an upper bound to the actual inter-blackout time, which is tighter in scenarios with low channel congestion. In fact, in such scenarios black-outs occur with relatively low probability, T_{bo} is larger, and the temporal correlation between two subsequent black-outs is lower. On the contrary, when interference is high, channel conditions are worse, black-outs are much more frequent, and also their time correlation increases. As a result, the expression in (2) becomes too

[2]In the following, we assume that the PIR can take only values in $\mathbb{Z}^+(T_B)$, that is, multiples of the inter-beaconing period $T_B = 100\,ms$. Therefore, we omit the term T_B and consider $\mathbb{P}[PIR > k]$ as the probability that PIR is larger than kT_B.

	0 kbps (c=0)		500 kbps			650 kbps			800 kbps		
	a	b	a	b	c	a	b	c	a	b	c
(V_0, V_1)	0.008	3.1917	0.089	0.6648	0.4261	0.3563	0.3394	0.3552	0.7375	0.1112	0.4335
(V_1, V_0)	0.0089	3.3109	0.4289	0.4696	0.2813	0.7060	0.2477	0.3017	0.9576	0.0844	0.2316
(V_2, V_1)	0.0579	2.5031	0.1239	1.3697	0.1711	0.4146	0.3110	0.3991	0.7371	0.1715	0.3814
(V_1, V_2)	0.0686	2.2716	0.6741	0.3701	0.2183	0.7192	0.1241	0.2757	0.8863	0.1105	0.2157
(V_0, V_2)	0.0362	2.4589	0.5837	0.4232	0.2565	0.6636	0.1571	0.2340	0.7788	0.1301	0.2086
(V_2, V_0)	0.0366	2.5139	0.4235	0.4872	0.2564	0.6471	0.3151	0.2208	0.8462	0.1372	0.2028
$(V_0, V_1)_{MH}$	0.0029	4.6356	0.0983	0.5190	0.5178	0.4239	0.0980	0.5076	0.8275	-0.0665	0.5445
$(V_1, V_0)_{MH}$	0.0034	5.6724	0.4979	0.2544	0.4126	0.8089	0.0259	0.4346	1.063	-0.1126	0.3466
$(V_2, V_1)_{MH}$	0.0256	4.9623	0.1395	1.1761	0.3126	0.4862	0.0827	0.5428	0.8283	-0.0121	0.4928
$(V_1, V_2)_{MH}$	0.0064	3.4325	0.7953	0.1335	0.4203	0.8621	-0.1437	0.4427	0.9921	-0.0879	0.3274
$(V_0, V_2)_{MH}$	0.0063	3.3136	0.5886	0.4148	0.4295	0.8444	-0.1308	0.4736	0.979	-0.1654	0.4158
$(V_2, V_0)_{MH}$	0.0055	3.3025	0.4032	0.4855	0.4576	0.8039	0.034	0.4596	1.0278	-0.1682	0.3960

TABLE I

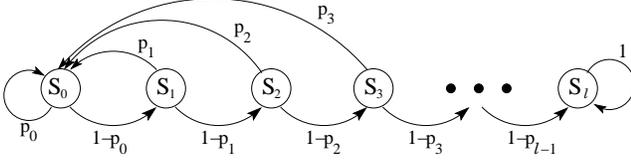VALUES OF THE PARAMETERS FOR THE PIR CCDF FITTING CURVES OF THE ANALYZED LINKS.

Fig. 13. Graphical representation of the Markov Chain modeling process $\mathcal{H}$.

optimistic in presence of congestion. The results presented in the remainder of this section confirm this intuition.

A different approach to model black-out events, which partially takes into account the channel time-correlation, is based on Markov chains. We can define a Markov process $\mathcal{H}$ to model beaconing packet reception as follows. Consider a process with states S_i, $0 \le i \le \ell$. A state change occurs every time a new beacon is sent from the transmitter to the receiver. The process is in state S_i if the last i beacons have been lost. Therefore, the corresponding Markov chain is represented in Figure 13.

We observe that state S_ℓ is an absorbing state, and correponds to a black-out event. As a result, the average time between two subsequent black-outs can be approximated by the expected absorbing time of the Markov chain, starting from state S_0.

In order to do this, the transition probabilities p_i, with $0 \le i \le \ell-1$ are to be known. These probabilities can be derived from the PIR time cdf, already analyzed in the previous sections. In fact, p_i, with $1 \le i \le \ell - 1$, is the probability that the PIR is no larger than $i+1$ beaconing intervals, given that it is larger than i beaconing intervals (missing i beacons corresponds to a PIR time of $i+1$ beaconing interval). Thus we can write:

$$1 - p_i = \mathbb{P}[PIR > i+1|PIR > i] , \qquad (3)$$

which implies

$$p_i = 1 - \frac{\mathbb{P}[PIR > i+1]}{\mathbb{P}[PIR > i]} \qquad (4)$$

The case $i = 0$ corresponds to the probability that two consecutive beacons are received, i.e.:

$$p_0 = 1 - \mathbb{P}[PIR > 1] \qquad (5)$$

The corresponding transition matrix is:

$$\mathbf{M} = \begin{bmatrix} p_0 & 1-p_0 & 0 & 0 & \dots & 0 & 0 \\ p_1 & 0 & 1-p_1 & 0 & \dots & 0 & 0 \\ p_2 & 0 & 0 & 1-p_2 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ p_{\ell-1} & 0 & 0 & 0 & \dots & 1-p_{\ell-1} & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & 1 \end{bmatrix} \qquad (6)$$

The expected absorbing time is computed through matrix $\mathbf{Q}$, which is obtained from matrix $\mathbf{M}$ by deleting the rows and columns corresponding to the absorbing states (in our case, therefore, by deleting the last row and the last column). The $\ell \times 1$ column vector $\mathbf{T}$ contains, as the i–th element, the expected absorbing time starting from state S_{i-1}. It is derived as $\mathbf{T} = (\mathbf{I}_\ell - Q)^{-1}\mathbf{1}_\ell$, where $\mathbf{I}_\ell$ is the identity matrix of size ℓ and $\mathbf{1}_\ell$ is a column vector of length ℓ with all elements equal to 1. Since we are interested in computing the average time T_{bo}^H between two subsequent black-out events, we have that $T_{bo}^H = (\mathbf{T}(1)-\ell)T_B$, where we need to subtract ℓ to identify the instant before the black-out beginning, while T_B is the beaconing interval.

Given the particular structure of $\mathbf{M}$, a closed form expression also exists:

$$T_{bo}^H = \left(\sum_{i=0}^{\ell-1} \prod_{j=\ell-1-i}^{\ell-1} \frac{1}{1-p_j} - \ell \right) T_B \qquad (7)$$

The inter-black out time estimate derived from this model can be considered an upper bound to the actual one. In fact, the channel time correlation is significant over several beaconing times. In our model, this means that the transition probabilities among the states closer to the absorbing state S_ℓ are more accurate. In fact, the probability of going from S_i to S_{i+1} implicitly takes into account the entire Markov process evolution in the previous i states. Conversely, this is not true for the states closer to S_0, and especially for S_0 itself. In fact, when the process enters S_0, it completely loses the memory about the previous state which, however, is likely to have a strong influence on the transition to the next state. Intuition suggests that if several beacons have
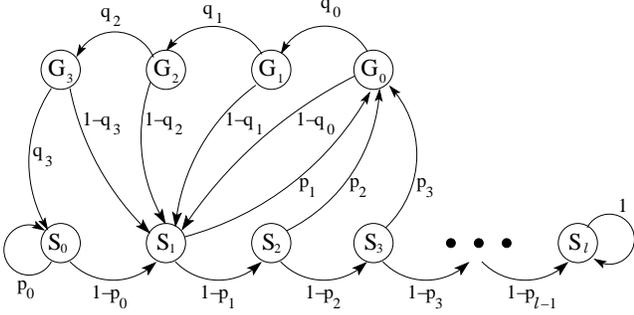
Fig. 14. Graphical representation of the Markov Chain modeling process $\mathcal{L}$, with $k = 3$.

Fig. 15. Average inter-black out time on the link (V_1, V_0), for different values of the congestion level R (expressed as a CBT value).

been lost before the last reception, the probability of losing the next one is higher than in the case of having successfully received also the previous beacons. When computing the expected value, it results higher than in reality, since as S_0 is reached, the model forgets the consecutive losses of the previous beacons. As confirmed by the results reported in the following, T_{bo}^H is an upper bound of the measured inter-black out time, which becomes tighter when the channel conditions are bad (e.g., relatively high contention), and state S_0 is visited fewer times.

A more refined Markov chain model can be defined, by keeping memory of past states also when beacons are successfully received. A simple way to do this is to introduce some states G_i, with $0 \leq i \leq k$. The system is in state G_i when the previous i beacons have been correctly received. The number of additional states k can be tuned (it is set to 9 in the following), and we call $\mathcal{L}$ the new Markov process, which is reported in Figure 14. Whenever a beacon is lost, process $\mathcal{L}$ behaves as $\mathcal{H}$; however, upon reception of a beacon, it moves to state G_0, rather than S_0. The process then proceeds along the G_i states as long as beacons are received, finally arriving in S_0. If instead a beacon is lost, the process moves to S_1, as before. This second model, which better represents the channel behaviour after a packet loss, requires in turn additional information. If we call N_b the number of correctly received consecutive beacons between two packet losses, the probabilites q_i's can be found based on the ccdf of the distribution of N_b:

$$q_0 = \mathbb{P}[N_b > 1] \qquad (8)$$

$$q_i = \frac{\mathbb{P}[N_b > i]}{\mathbb{P}[N_b > i - 1]} \qquad (9)$$

The derivation of the corresponding transition matrix $\mathbf{M}_L$ is straightforward. By calling $\mathbf{Q}_L$ the matrix obtained by removing the row and the column corresponding to S_ℓ, the vector of the average absorbing times of the process is:

$$\mathbf{T}_L = (\mathbf{I}_{\ell+k+1} - \mathbf{Q}_L)^{-1} \mathbf{1}_{\ell+k+1} \qquad (10)$$

and the expected time between two consecutive blackouts is

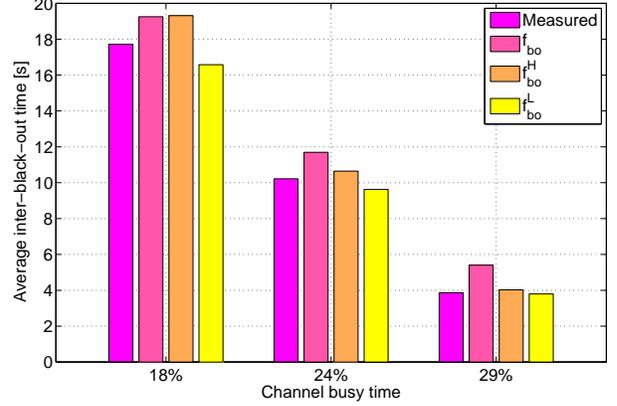$$T_{bo}^L = (\mathbf{T}_L(1) - \ell)T_B \qquad (11)$$

We first notice that $T_{bo}^L < T_{bo}^H$. In fact, the two processes follow the same behaviour along states S_i's. Nonetheless, when a beacon is received, process $\mathcal{L}$ moves to state S_0, while process $\mathcal{H}$ moves to state G_0. Now, it can be shown that $q_0 < p_0$. In fact, both represent the probability of receiving the following beacon and are obtained from the measurements collected during the experiments; however, while p_0 is measured over all the received beacons, q_0 is measured only over the beacons received immediately after a packet loss. Since the channel is time correlated, it follows that the probability of receiving two consecutive beacons after a packet loss is lower than the unconditioned probability of receiving two consecutive beacons.

Also the values obtained through process $\mathcal{L}$ are approximated. In fact, memory is lost whenever the process enters state G_0 and state S_1. This is more evident for sparse beacon losses, since from S_1 the only available transition is to G_0. In this state, the packet loss probability is higher than in reality, since it is averaged over all the beacons received after a loss. It follows that f_{bo}^L is often even lower than the actual expected time interval between black-out events, as confirmed by the results.

We report in Figures 15, 16 and 17 the values of the inter-black out time computed with the three models, compared with the measured value. The plots refer to different links and different congestion levels. The results confirm that T_{bo}, assuming independence between black-out events, always gives an optimistic estimate of the true inter-black out time, independently of the congestion level. Conversely, both T_{bo}^H and T_{bo}^L, which take into account channel correlation, predict inter black-out times increasingly closer to the measured values as congestion level increases. This is due to the fact that correlation is better modeled during periods with no successful beacon reception, which happen more frequently when congestion is high.

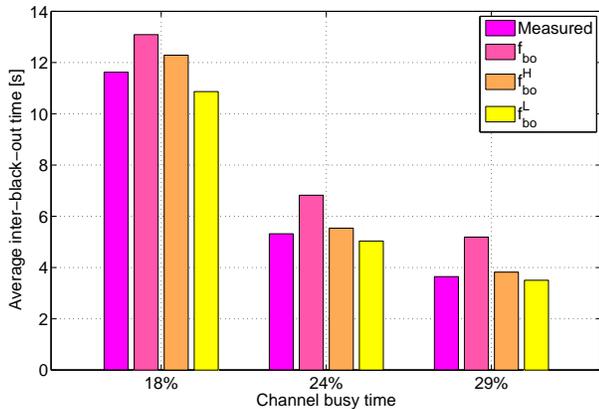As explained above, we always have $T_{bo}^H > T_{bo}^L$; while the

Fig. 16. Average inter-black out time on the link (V_0, V_2), for different values of the congestion level R (expressed as a CBT value).
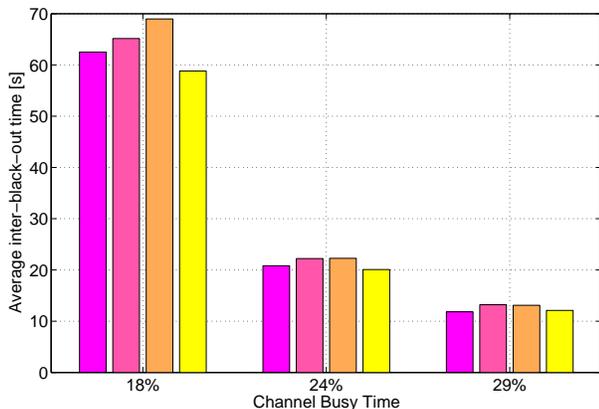

Fig. 17. Average inter-black out time on the link $(V_0, V_2)_{MH}$, when multi-hop beaconing is enabled, for different values of the Congestion level R (expressed as a CBT value).

former is also higher than the effective black-out frequency, the latter is instead always below. This last inequality is not necessarily true, but in most cases T_{bo}^H and T_{bo}^L can be effectively used to identify an upper and a lower bound to the actual inter-black out time, respectively.

Before ending this section, we want to comment on the notable effect of congestion on inter-black out time: when CBT is increased from about 18% to about 29%, the inter-black out time is reduced of a factor 4 in the (V_1, V_0) link, and of a factor 3 in the (V_0, V_2) link. Thus, an *increase of about 10% of the observed CBT causes a 3- to 4-fold increase of black-out frequency*. This finding experimentally confirms the necessity of keeping the congestion level under strict control, in accordance with recommendations from standardization bodies: for instance, ETSI recommends a CBT $\leq 25\%$ [4]. Interestingly, *multi-hop beaconing can be exploited to substantially reduce occurrence of black-out events, also in presence of congestion*: comparing the values reported in figures 16 and 17, we observe that the

inter-black out time at CBT 29% is increased of a factor 2.5 when multi-hop beaconing is used on link (V_0, V_2), with a corresponding decrease of black-out frequency.

IX. CONCLUSIONS

In this paper, we have presented the first measurement-based study of beaconing performance in presence of channel congestion. Our study has revealed the profound impact of congestion on beaconing performance, thus strongly motivating further research on congestion control algorithms. Another major finding of this study is showing the effectiveness of multi-hop beaconing in lessening the negative impact of congested channel conditions on beaconing performance. While we have proven the effectiveness of multi-hop beaconing, it is important to observe that such benefits were achieved in a scenario in which beacons reported information on *all* surrounding vehicles. Reporting such complete information within limited size beacons is likely to be unfeasible in presence of dense vehicular traffic conditions, which are exactly those causing channel congestion. For this reason, implementing "smart" multi-hop beaconing protocols that include only partial information in the beacon while substantially benefiting situation awareness is an interesting direction for future research.

REFERENCES

[1] F. Bai, D.D. Stancil, H. Krishnan, "Toward Understanding Characteristics of Dedicated Short Range Communications (DSRC) From a Perspective of Vehicular Network Engineers", *Proc. ACM Mobicom*, pp. 329–340, 2010.
[2] "Standard Specification for Telecommunications and Information Exchange Between Roadside and Vehicle Systems - 5Ghz Band Dedicated Short Range Communications (DSRC)", *ASTM E2212-03*, 2003.
[3] T. ElBatt, S.K. Goel, G. Holland, H. Krishnan, J. Parikh, "Cooperative Collision Warning Using Dedicated Short Range Wireless Communications", *Proc. ACM VANET*, 2006.
[4] ETSI, "ITS: Decentralized Congestion Control Mechanisms for ITS Operating in the 5GHz Range", ETSI TS 102 687, 2011.
[5] J. Gozalves, M. Sepulcre, R. Bauza, "IEEE 802.11p Vehicle to Infrastructure Communications in Urban Environments", *IEEE Communications Magazine*, Vol. 50, n. 5, pp. 176–183, 2012.
[6] K. Hong, D. Xing, V. Rai, J. Kenney, "Characterization of DSRC Performance as a Function of Transmit Power", *Proc. ACM VANET*, pp. 63–68, 2009.
[7] C.-L. Huang, Y.P. Fallah, R. Sengupta, H. Kirshnan, "Adaptive Intervehicle Communication Control for Cooperative Safety Systems", *IEEE Network*, Vol. 24, n. 1, pp. 6–13, 2010.
[8] D. Jiang, Q. Chen, L. Delgrossi, "Optimal Data Rate Selection for Vehicle Safety Communications", *Proc. ACM VANET*, pp. 30–38, 2008.
[9] S. Kaul, K. Ramachandran, P. Shankar, S. Oh, M. Gruteser, I. Seskar, T. Nadeem, "Effect of Antenna Placement and Diversity on Vehicular Network Communications", *Proc. IEEE Secon*, pp. 112–121, 2007.

[10] J.B. Kenney, G. Bansal, C.E. Rohrs,"LIMERIC: a Linear Message Rate Control Algorithm for Vehicular DSRC Systems", *Proc. ACM VANET*, pp. 21–30, 2011.

[11] T. Mangel, M. Michl, O. Klemp, H. Hartenstein, "Real-World Measurements of Non-Line-Of-Sight Reception Quality for 5.9Ghz IEEE 802.11p at Intersections", *Proc. Nets4Cars*, pp. 189–202, 2011.

[12] F. Martelli, M.E. Renda, G. Resta, P. Santi, "A Measurement-based Study of Beaconing Performance in IEEE 802.11p Vehicular Networks", *Proc. IEEE Infocom*, 2012.

[13] M. Sepulcre, J. Mittag, P. Santi, H. Hartenstein, J. Gozalvez, "Congestion and Awareness Control in Cooperative Vehicular Systems", *Proceedings of the IEEE*, Vol.99, pp. 1260-1279.

[14] M. Torrent-Moreno, D. Jiang, H. Hartenstein, "Broadcast Reception Rates and Effects of Priority Access in 802.11-based Vehicular Ad Hoc Networks", *Proc. ACM VANET*, pp. 11–18, 2004.

[15] M. Torrent-Moreno, J. Mittag, P. Santi, H. Hartenstein, "Vehicle-to-Vehicle Communication: Fair Transmit Power Control for Safety-Critical Information", *IEEE Trans. on Vehicular Technology*, Vol. 58, n. 7, pp. 3684–3703, 2009.

[16] VSC Consortium, "Vehicle Safety Communications Project Task 3 – Final Report: Identify Intelligent Vehicle Safety Applications Enables by DSRC, *DOT HS 809 859*, 2005.