

Contention Window Analysis for Beacons in VANETs

René Reinders, Martijn van Eenennaam, Georgios Karagiannis and Geert Heijenk

Department of Computer Science, University of Twente, The Netherlands

r.reinders-1@student.utwente.nl,

{e.m.vaneennaam, g.karagiannis, geert.heijenk}@utwente.nl

Abstract—The recently standardised IEEE 802.11p Wireless Access in Vehicular Environment (WAVE) protocol is envisioned to be used for Vehicle-to-Vehicle and Vehicle-to-Infrastructure communications. We provide an analysis of 802.11p’s performance when used for the exchange of small status messages known as beacons. These beacons enable vehicles to establish a Cooperative Awareness from which many applications can draw their inputs. Applications have been proposed to increase safety, efficiency and comfort of future road traffic.

This paper investigates the impact of adapting settings for the Contention Window (CW) based on the traffic density in order to increase the efficiency of the 802.11p broadcasts w.r.t. reception probability and delay with a focus on real-time vehicle control. In contrast to the literature available on adapting CW settings in unicast, we find that the initial CW value of 15 proposed in the IEEE 802.11p standard performs significantly well under almost all studied circumstances.

I. INTRODUCTION

A Vehicular Ad-Hoc Network, or VANET, is a wireless ad-hoc network that supports the communication (1) amongst vehicles and (2) between vehicles and Road Side Units (RSUs). The main goal of VANETs is to support Intelligent Transportation System (ITS) applications which aim at providing safety and comfort to passengers and increasing traffic efficiency.

A traffic efficiency application that can be used to increase passenger comfort and reduce traffic jams is the Cooperative Adaptive Cruise Control (CACC) [1], [2]. Using the VANET, the vehicles (and if required RSUs) are cooperating to obtain the necessary lead vehicle dynamics information (such as speed, position, and acceleration) and a general view of traffic ahead [3]. This information is used to enhance the performance of the currently available Adaptive Cruise Control (ACC) system. In particular, the system controls the accelerator, engine powertrain and vehicle brakes to maintain a desired time-gap or time headway to the vehicle ahead [4]. CACC enables the shock waves, i.e. travelling disturbances in the traffic flow on a highway [5], to be filtered out of the main traffic flow, thus greatly improving throughput [6].

CACC uses a Cooperative Awareness that can be visualised as a radar screen on which all nearby vehicles are placed. To build Cooperative Awareness, each vehicle periodically sends a short status message which contains information such as position, speed and acceleration [7], a practice called beaconing. CACC relies on accurate and timely situational awareness to perform its task. It is worth mentioning that the

Cooperative Awareness constructed by beaconing is not only useful for CACC, but virtually any ITS application.

For communications in the VANET, the IEEE 802.11p protocol is used [8]. IEEE 802.11p uses a Media Access Control (MAC) protocol that is based on a Carrier Sense Multiple Access protocol with Collision Avoidance (CSMA/CA). Due to the fact that beacons may be sent several times per second [9] and vehicular density can vary greatly (and become large during traffic jams), it is expected that the channel may become congested [3], resulting in a deterioration of the Cooperative Awareness and hence the performance of the ITS applications.

When a node tries to access the medium and finds the channel busy, it chooses a random backoff time from the interval $[0, CW]$ and delays the medium access for the duration of backoff. CW represents the size of the Contention Window. If no acknowledgement is received (e.g. a collision occurs) the CW size is doubled and the process starts over. However, because of the lack of acknowledgments when performing a beacon broadcast, CW is never increased for beaconing [8].

The probability that two nodes will try to access the medium at the same time is small, but when there are e.g. 100 nodes the probability that two nodes will choose the same time increases. Increasing the CW might be a way to deal with this but making the CW too large will increase delay. A solution may lie in adapting the CW in response to the number of nodes. If there are many nodes, CW could increase and if there are fewer nodes CW could decrease.

The contributions of this paper are as follows:

- Analyse the effects of various CW sizes on beacon reception probability in VANETs
- Provide an analysis of both delay and inter-arrival times, which is of particular interest in VANETs because the received information can be fed into real-time control systems. Understanding of these metrics is an important input in the process of developing CACC systems.

Section II provides an overview of the background and related work. Section III outlines the experiments and Section IV reports on the findings. Section V discusses the impact of the findings and VI concludes this paper.

II. BACKGROUND AND RELATED WORK

In a VANET, beacons are sent on a regular interval, i.e. the beaconing interval τ_g . Many vehicular applications assume a beaconing interval of 100ms or a generation rate λ_g of

10Hz [10], [9]. Beacons contain a message defined by the European ITS VANET Protocol (EIVP) Cooperative Awareness Message (CAM) as defined in [11]. As such, a beacon is approximately 400 bytes long (including security fields). Beacons are transmitted in the one-hop communication range, typically 250m. As identified in [3] the load on the channel can increase rapidly as traffic density increases and 802.11 broadcast performance deteriorates as the load on the network increases [12], [13], [14]. Though the topic of performance deterioration in a channel where broadcasts dominate has been extensively studied in the area of flooding and network-layer broadcasts [14], beaconing does not generate a Broadcast Storm because (in most implementations) beacon generations are independent. But similar to flooding and network layer broadcast, deteriorated performance at the MAC layer results in a reduced probability of successful reception and in increased delay. For beaconing, this may be particularly detrimental to the performance of real-time controllers as used in CACC, so it is paramount to optimise the performance of beaconing. In contrast to the beaconing performed by access points, timing synchronisation by 802.11 nodes in DCF mode or some MANET routing protocols, beaconing in VANETs does not support a service, it *is* the service [3].

Several methods can be applied to improve the performance of beaconing. The optimisation criterion for improving beaconing performance has to be built upon the concept of *fairness*, as stressed in [15]. We add to this the concept of *relevance*; a node should only transmit information which may be relevant to others, especially if the wireless channel is congested. Furthermore, we note in [3] that a beaconing system must exhibit the property of *gracefull degradation*. In addition, a node must have a notion of the quality of the received information, so automated control may be switched off if reception rates or delay requirements are violated and, for instance, a CACC system cannot operate. Ideally such cases would be rare but are not unthinkable.

A. Generation Rate

The generation rate (λ_g) is the rate at which beacons are sent to the MAC for transmission. Since they are used to create a Cooperative Awareness, λ_g should be in the order of several beacons per second to provide the system with accurate information about the close surroundings [15], [16], [10]. Though some works consider a fixed λ_g of 10Hz [9], [17], we motivate in [3] that generation rate adaptation as a network-layer mechanism is one of the instruments to make beaconing more scalable. Increasing the rate results in more beacons being sent and a higher temporal resolution. But this comes at the price of an increase in collision probability, especially in dense traffic. Decreasing the rate in such scenarios results in fewer beacons and hence fewer collisions. A drawback of this approach is that if the rate is too low, it also affects the Cooperative Awareness as it does not get updated often enough [18], but in case of severe channel congestion it is better to at least have some information, than no information at all. In [3] we suggest adapting λ_g based on channel estimations and road

traffic dynamics as collected in the Cooperative Awareness. Fairness is important here, as no node should suffer *starvation*; where it is not able to transmit its beacons while others still can. Controlling λ_g is possible for several reasons [12]:

- 1) Inherent unreliability of the wireless channel: the application must already cope with this by applying filtering (e.g. interpolation, Kalman estimates). Due to the nature of wireless communication, upper layers already need to deal with the probability a message does not arrive. The lack of acknowledgements and reservation by means of RTS/CTS contributes to the unreliability.
- 2) More load leads to more loss. Under certain conditions, reception rate λ_r may be raised by lowering λ_g . This is a well-known phenomenon in networks, where raising the supplied load beyond a certain point will lead to a reduction in throughput due to losses. In the case of 802.11 broadcast, there is no reaction to load by means of contention window increase.
- 3) As traffic dynamics decrease (e.g. in traffic congestion, speeds are significantly lower than in free-flow), a high λ_g is not always necessary for the application to operate. This is important in the light of *relevance*: if a node has not moved it does not need to notify its neighbours.

B. Transmission Power

The transmission power affects how many nodes will receive the beacon, but also how many will suffer interference from the transmission. With low transmission power, only the closest neighbour may receive the beacon, a more remote node might not. With high transmission power, a significant number of nodes might receive the beacon, but the collision probability is also higher [15] and more nodes will receive interference. The goal of transmission power control is to increase spatial frequency reuse. The power control method must be fair: a higher transmission power of a sender should not be selected at the expense of preventing other vehicles to send/receive their beacons. Torrent-Moreno *et al.* [15] and Yang *et al.* [13] provide adaptive Transmission Power Control solutions.

C. Contention Window Size

Several works evaluate the impact of different CW settings on the performance of 802.11p unicast. CW is initialised to CW_{min} and doubled after a failed transmission (e.g. the sender has not received an acknowledgement within a certain timeout) up to CW_{max} [8]. As a result, a transmission can be attempted multiple times up to a retry limit, after which the transmitting station gives up.

[19] and [20] study the performance of 802.11p unicast transmissions for the Access Classes defined by 802.11p's Enhanced Distributed Channel Access (EDCA). [19] finds that the AC for high priority traffic (small CW values) suffers from a drop in performance as the number of nodes increases. The work in [20] studies two CW adaptation mechanisms in a V2I setting after identifying that static CW settings do not perform optimal under the dynamically changing vehicular environment. Note that both studies only evaluated a single

Parameter	Values
Nodes per lane	1,2,5,10,20,30,40,50
Generation rate (λ_g)	1,5,10,15,20,25,30
Contention Window (CW)	3,7,15,63,255,1023
Beacon size	3200 bits
Datarate	3Mbit/s
MAC buffer	14 frames

TABLE I
SIMULATION PARAMETERS

AC at a time. In a sense, a priority scheme (like EDCA) only makes sense if there is other traffic to give priority over, as such researching high priority settings in isolation makes little sense. However, these works do provide valuable insight in the operation of various CW values in a unicast setting.

This paper expands this to broadcast traffic, where because of the lack of acknowledgments, CW is never increased. Beaconing performance has been studied by means of analytical models in [18], [3], and shows declining reception probability with increasing number of nodes. Furthermore, [18] also provides a delay analysis which shows an increasing delay as the number of vehicles in range increases. In these works, an increase in the number of nodes results in an increase in load on the network because the supplied load per node remains constant. In [12] it is suggested that mimicking the CW increase in response to load in broadcast as it is done in unicast could improve beaconing performance.

III. SIMULATIONS EXPERIMENTS

For the simulation experiments the discrete event simulation environment OMNeT++ [21] is used in conjunction with the MiXiM framework [22]. This is a mobility simulation framework for wireless and mobile networks. A beaconing model using IEEE 802.11p was implemented.

Furthermore, propagation loss was disabled in order to isolate collision loss. This can be justified by noting that Bianchi [23] and others [24], [18] also exclude propagation loss in their analytical approach. We are interested in 802.11p's ability to coordinate transmissions without causing collisions. By disabling propagation loss, all loss can be attributed to collisions. The collision model is SiNR-based using BER curves as standard in MiXiM. Because of these assumptions, the results provide valuable insight but cannot be directly translated to reality.

Vehicles are uniformly placed on a stretch of 4-lane highway of 250m. The data rate is set to 3Mbit/s and all nodes are within the transmission range of all other nodes. This yields a scenario without Hidden Terminals. The beacon message length was set to 3200 bits. All the PHY and MAC properties used in the IEEE 802.11p simulation model conform to [8]. The nodes remain stationary during each simulation run.

During simulation, each node transmits 100 beacons after a small random initial delay. These beacons are sent with λ_g , which is varied in the different simulation experiments. The number of nodes per lane and CW also vary per simulation experiment. An overview of the different values used is shown

in Table I, though not all results are presented in this paper due to space limitations. The value of nodes per lane ranges from very sparse traffic: 1 per 250m per lane to congested: 50 per 250m, which translates to 200 per km per lane, a bumper-to-bumper scenario. The remaining parameters are set according to the default values used by the MiXiM framework [22]. Simulation experiments were performed using the method of independent replications as described by Jain in [25]. Each simulation scenario was performed with enough replications to get statistical significant results, and all confidence intervals are well below 5% of the metric.

IV. SIMULATION RESULTS

Several performance metrics can be used to study the performance of beaconing in VANETs. These performance metrics are reception rate, reception probability, delay and inter-arrival times of beacons. The last may be of particular interest for performance evaluation in the light of the CACC application. We will now briefly list the metrics.

Beacon Reception Rate (λ_r) - the number of beacons per second node i receives from node j .

Beacon Reception Probability (P_s) - the probability that a beacon sent by node j is received by node i .

Delay (D) - the time needed to send a data frame from node j to i . After generation, a beacon is handed over to the MAC layer. The 802.11p MAC performs backoff and, once it has acquired access to the channel, transmits the message. This measures the freshness of the beacon upon reception and accounts queueing-, backoff- and propagation delay. This is only calculated for received beacons.

Inter-arrival Time (τ_r) - the duration of time between two successive receptions of a beacon from node j . This is important from the point of view of a real-time CACC controller which may be tracking node j as its lead vehicle. Ideally, the inter-arrival time would equal the inter-generation time (e.g. 100ms for $\lambda_g=10\text{Hz}$).

A. Beacon Reception Probability

Figs. 1a–1c show P_s for each of the CW values, in relationship with the number of nodes per lane on the 4-lane simulated highway. Fig. 1a shows beaconing at 5Hz. Up to 10 nodes per lane there is no difference between the different CW values. As the traffic density increases further, the higher CW sizes suffer a greater decline than the smaller CW sizes. The decline sets in at 40 nodes in the system. The beacons are approximately 1ms long, resulting in 200ms of transmissions to be performed per second, or a 20% utilisation.

The scenario with $\lambda_g=10\text{Hz}$ is shown in Fig. 1b. In this case, P_s already declines at 5 nodes per lane. Here too, we find a 20% utilisation. This 20% also holds for Fig. 1c, where $\lambda_g = 25\text{Hz}$. When looking at the point where P_s drops below 50%, we can perform a similar calculation and observe this occurs at approximately 33% utilisation.

The acceptable reception probability depends on the type of application supported by the VANET. As specified by Ma

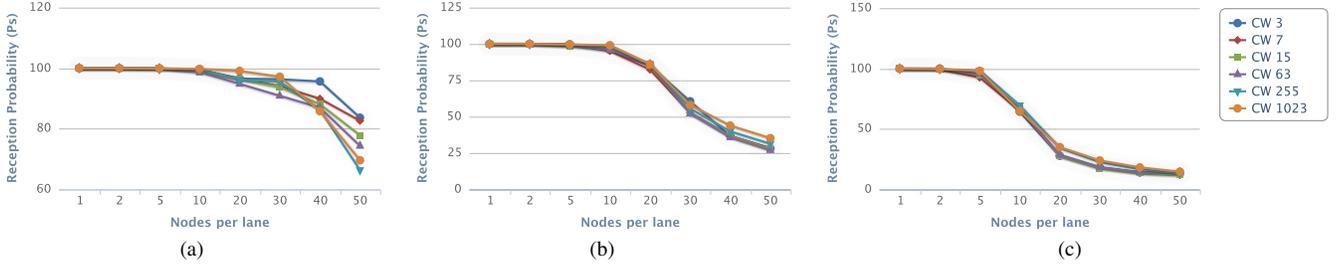


Fig. 1. Beacon Reception Probability P_s for: (a) $\lambda_g=5\text{Hz}$ (b) $\lambda_g=10\text{Hz}$ (c) $\lambda_g=25\text{Hz}$

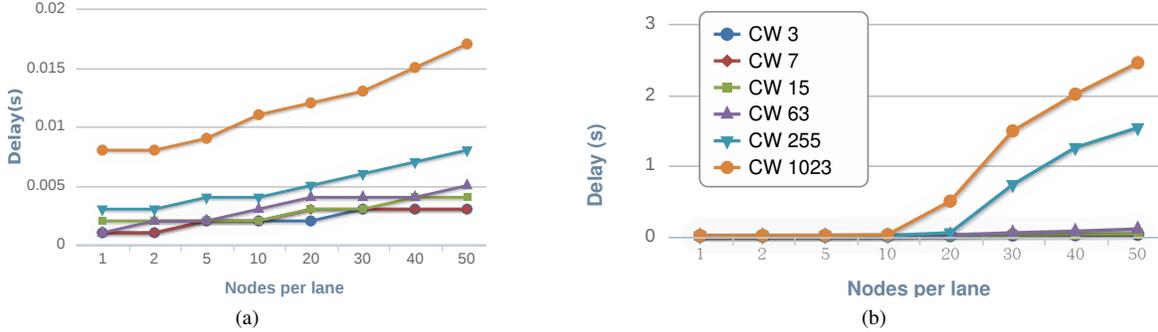


Fig. 2. Average Delay increases with traffic density: (a) $\lambda_g=1\text{Hz}$ (b) $\lambda_g=10\text{Hz}$

and Chen in [26], P_s should not drop below 0.99 for safety-critical applications. As evident in the results this requirement is readily violated as traffic density increases. For other types of vehicular applications, such as CACC, the reception probability can be much lower than 99%, depending on the loss tolerance of the CACC controller [10] by means of e.g. Kalman filtering.

The simulation results show that a P_s of 99% is only achieved if the number of nodes per lane is low, or if λ_g is low. If there are 5 nodes per lane, P_s is around 99%, for all λ_g values and for all CW values up to moderate traffic densities. The P_s is nearly the same for all studied CW values. Hence, we can conclude that increasing the CW size when the number of nodes increases does not raise the beacon reception probability. This is further elaborated upon in Sec. V.

These simulation experiments show that P_s can be increased by decreasing the beacon generation rate when the number of vehicles using the same VANET radio channel is high, a relation also found analytically in [3].

B. Delay Analysis

Fig. 2 shows the delay experienced by beacons under various traffic densities for the considered CW values at $\lambda_g=1\text{Hz}$ and 10Hz . Average delay is larger for higher CW value (see Fig. 2a) and increases as the number of nodes increases. The delays found coincide with the findings of Vinel *et al.* in [18], approximately 8ms with up to 80 vehicles in the system in case $\lambda_g = 10\text{Hz}$. Work on a CACC system in [2] by Naus *et al.* models a communication delay of 10ms. From our results, the assumed communication delay in the CACC system is, on average, not violated, unless very large CW values are used

or the system enters the saturated regime (i.e. utilisation goes toward 100%). This is what happens to the large CW values in Fig. 2b. What is actually happening here is the effect of queue-buildup: beacons are generated faster than they can be transmitted. A beacon now not only has a contention delay, but also accumulates queueing delay. Our default 802.11p implementation with a buffer size of 14 shows this only for a CW of 256 and 1023 for $\lambda_g = 10\text{Hz}$, but as λ_g is raised further the other CW values also start to rise. The queueing delay increases when the queue size becomes larger. As a result, a node is only transmitting old information, even though newer information is available in the buffer.

A histogram can be used to estimate the probability density function of a variable. Fig. 3 provides a histogram showing the distribution of delay on a sparse road. Note that a node with a larger CW (and hence larger backoff counter) has more opportunities of being interrupted while performing backoff countdown. As a result, the delay will be larger. The delay distribution plotted in Fig. 3a shows a typical long-tailed distribution. Most received beacons have a delay below 1ms. These are cases where a station performs clear channel assessment, finds the channel idle and immediately transmits. The tail contains transmissions which did not directly find the channel idle, and backoff was performed. This backoff may have been interrupted by other transmissions, a node will then freeze its backoff counter. The result is a (potentially) very long delay.

Fig. 3b shows the first 10ms in more detail. The bars from left to right indicate the CW values in ascending order. We can observe that the smaller CW values have a greater proportion of their transmissions in the first interval. This is due to the

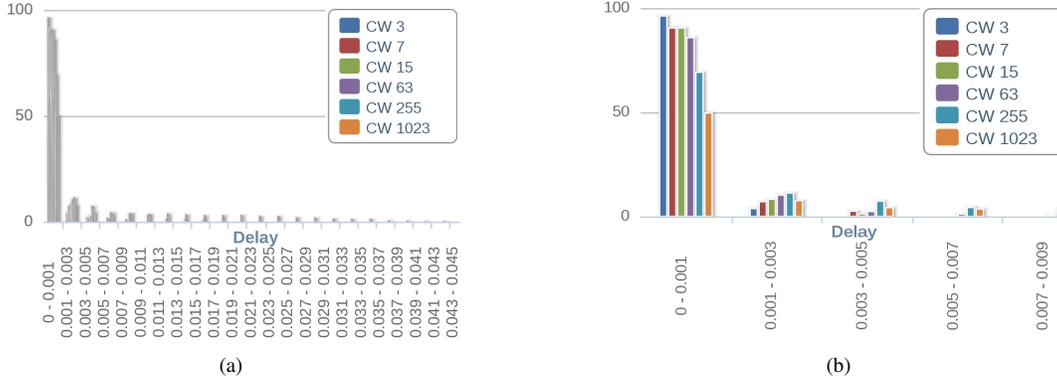


Fig. 3. Delay distribution on a sparse road ($n=2$, $\lambda_g=10\text{Hz}$): (a) long tail (b) details of the first few values. The bars from left to right in each bin correspond to the legend entries from top to bottom.

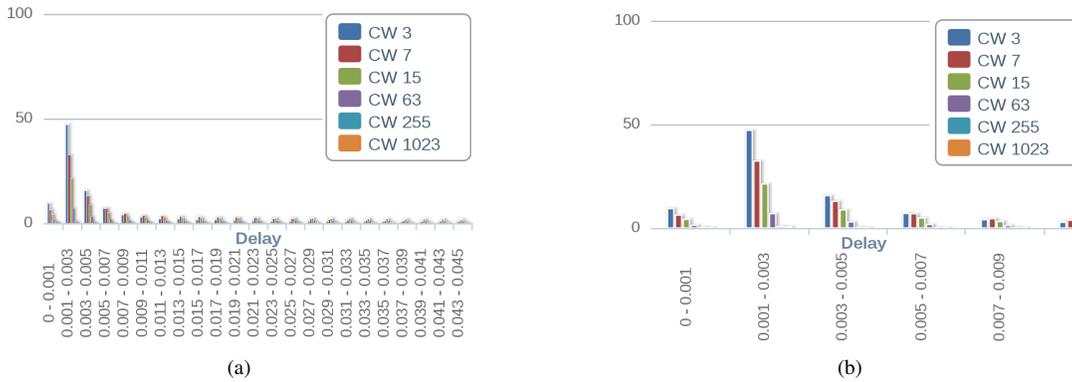


Fig. 4. Delay distribution on a dense road ($n=40$, $\lambda_g=10\text{Hz}$): (a) a long tail similar to Fig. 3a, but now the peak has moved to the right (b) details of the first few values. Note that CW 255 and 1023 hardly show up because of their large delays (see Fig. 2b).

fact that a bin size of 1ms (used to generate these plots) can contain CW3's initial attempt, plus its later attempt (when its backoff counter reaches zero). In the case of larger CW values, more attempts fall into later bins and hence the percentage in the first bin is smaller.

Fig. 4 shows the delay distribution in a dense scenario (40 vehicles per lane, 160 nodes in total). What can be observed is that a node's initial attempt to access the channel often finds the channel busy. Hence the peak present near the origin in Fig. 3 moves to the right. The long tail, however, becomes heavier: larger delays occur more often. In this scenario delays in the order of seconds were observed (not visible in the plots). Fig. 4b shows the plot in more detail, and shows CW values 255 and 1023 have a very low percentage of short delays, indicating the heavy weight of the tail.

An important observation, also made in [18], is that delay requirements are easier to satisfy than P_s requirements. Delay requirements for safety-critical information are specified as 500ms [26], which seems relaxed when compared with the 200ms requirements of a real-time CACC system [10]. However, even 200ms is not a big problem in single-hop broadcasts. But not only the delay of a single beacon is important; the time between reception of two consecutive beacons from the same node can easily surpass 200ms as shown in the next section.

C. Inter-arrival Time

In a sparse network, the load on the channel is not high. As such, we observe a large success probability P_s . This implies that the probability of node i missing a beacon from vehicle j is small. However, as the traffic density increases P_s diminishes, with a detrimental effect on the inter-arrival time τ_r of beacons, as shown in Fig. 5 and 6. Inter-arrival times are greater for larger CW values. As can be seen in Fig. 6, some inter-arrival times are as large as one second. Clearly, this will have a detrimental effect on CACC performance.

An explanation for the observed increase in inter-arrival time is as follows. A beacon is transmitted at time t_0 , but it is lost in collision. At time t_1 that node will generate the next beacon, and so on. In a sense, this periodicity can be seen as an implicit retransmission (albeit with updated information). This property makes it possible to operate real-time systems which are able to operate on deliberately delayed samples which are replayed with a known delay. Missing samples can then be interpolated. If the distributions are as in Fig. 5, samples could be buffered and replayed with a constant delay of 200ms, and the controller would still use most of the samples, resulting in good operation. If the distributions of Fig. 6 are used, many beacons would arrive late, resulting in poor performance.

V. DISCUSSION

The CW values studied in this work show little impact on P_s , however large CW sizes tend to cause a very long delay. Therefore, we can deduce that the currently used default initial CW size (i.e. 15) for beaconing is an adequate value. Contention Window increase in unicast works under the assumption that load on the channel is not constant, but of a bursty nature. If a node finds the channel busy, it will defer increasingly longer so the burst of transmissions by other nodes is at some time over. After that, there is a larger probability of finding the channel idle. In beaconing, however, the load is relatively constant. If a node is to perform channel sensing at arbitrary moments, it would find the channel idle with the same probability. What a backoff counter of 1023 actually does, is allow (or force) a node to wait 1023 slottimes (and potentially, equally many transmissions by other nodes) and then transmit with the same probability of success as when the backoff counter would have been 0.

As suggested in [12], a combination of beacon generation rate control and power control may be a more adequate way to deal with an increase in traffic density than Contention Window increase alone.

P_s is generally high, until the number of vehicles increases beyond a certain value. Or more specific, P_s declines as the supplied load increases. In contrast to the suggestion in [12], adaptation of the CW did not yield the expected beneficial effects on P_s . However, it should be noted that in this research *all* nodes in a scenario used the same CW settings. This situation differs from unicast, because in unicast nodes will not all have the same CW at the same time, as such the different CW values give one node an advantage over the other. This is also the case with 802.11e's EDCA which is also part of 802.11p. As such, in order to properly mimick CW increase as done in unicast a node could decide to increase the CW based on its own situation: it should be able to estimate the P_s of its transmissions and increase its CW if P_s drops below a certain threshold. This would then function like the occurrence of a timeout when waiting for an acknowledgement in response to a collision in unicast.

On average, delay is short. However, in case of a congested channel, a node can spend a long time in contention. This time can even be so long (values in the order of seconds have been observed, confirming findings in [27]) that a new beacon has been generated. This beacon has been added to the MAC queue and is waiting for transmission. In this case, from an application perspective, the beacon currently in contention is obsolete by the new one in the queue. The result is excessive delays and queue drops - a newly generated beacon finds the queue full and is dropped. The magnitude of this problem is subject of future work, but it may be addressed in several ways: 1) make the queue of unity size - a new beacon will overwrite the old one or 2) make a preemptive backoff procedure or 3) only generate a new beacon if the old one has left the MAC queue. In turn, a solution to this problem would also lower the supplied load on the channel, because the node does not have

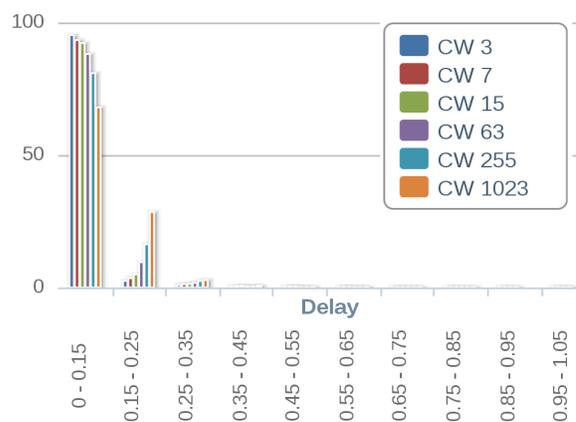


Fig. 5. Distribution of Inter-arrival times τ_r for $n=20$ and $\lambda_g=10\text{Hz}$.

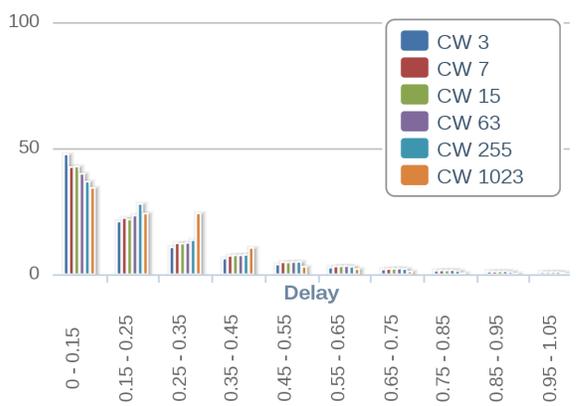


Fig. 6. Distribution of Inter-arrival times τ_r for $n=40$ and $\lambda_g=10\text{Hz}$.

to perform contention for both the new and the old beacon. In the simulation experiments the default First-Come-First-Serve MAC queue reinforces the performance drop with respect to delay and packet loss.

Delay can be large, but this only holds for a small percentage of the beacons (most of the distribution lies well within τ_g). Note however, that under an increase in traffic density the average delay will rise. Whereas for sparse traffic the delay distribution approximates an exponential distribution, the distribution for dense traffic is of a different kind. This is important to keep in mind when analytically modelling delay of a beaconing system.

This work considered 802.11p transmissions *without* propagation effects and hidden terminals, focussing only on collision loss. It should hence be noted that realistic P_s will be lower, and as a result the inter-arrival times can also be expected to increase.

VI. CONCLUSION AND FUTURE WORK

This paper provided an analysis of the effects of the Contention Window size on the performance of beaconing in VANETs. It was expected that increasing the Contention Window in reaction to an increased traffic density could improve performance. The effects on Beacon Reception Probability,

Delay and Beacon Inter-arrival Time were researched in a simulation study. The conclusion is that Contention Window increase as performed in the experiments does not improve beaconing performance.

The analysis of the beacon reception probability shows that 802.11p broadcast performance drops below 50% when the channel utilisation exceeds 33%. This occurs when either λ_g or the number of vehicles sharing the same collision domain exceeds a threshold. The beacon delay analysis shows that, if a beacon is successfully received, it is on average less than 10ms old. An increase in traffic density results in an increase in delay. The beacon inter-arrival time analysis shows that, under high vehicle densities, a vehicle can miss several successive beacons from a vehicle it is tracking.

The periodicity of beaconing implicitly retransmits data, and as such a real-time system using samples derived from beacons should be able to operate under a wide variety of traffic densities if it can replay samples with a fixed delay and interpolate missing samples.

Future work includes adaptively controlling beacon generation rate and transmission power in order to arrive at a beaconing system which is able to operate under a wide range of vehicular densities. The MAC buffer for beaconing is also a topic of interest, as the queueing strategy for beacons may be different from the one used in unicast. Furthermore, the impact of the delay and inter-arrival times of received beacons on the performance of a CACC system is a very relevant topic.

ACKNOWLEDGMENT

This work is supported by the Dutch NL Agency/HTAS (High Tech Automotive Systems) Project Connect&Drive, Project no. HTASD08002.

REFERENCES

- [1] A. Bodhankar, "Cooperative adaptive cruise control (CACC)," *KRESIT, IIT Bombay*, Nov 2004.
- [2] G. Naus, R. Vugts, J. Ploeg, M. van de Molengraft, and M. Steinbuch, "Cooperative adaptive cruise control, design and experiments," *American Control Conference; Baltimore, MD, United States, accepted*, 2010.
- [3] E. M. van Eenennaam, W. Klein Wolterink, G. Karagiannis, and G. J. Heijenk, "Exploring the Solution Space of Beaconing in VANETs," *First IEEE Vehicular Networking Conference, VNC2009, Tokyo, Japan*, Oct 2009.
- [4] G. Marsden, M. McDonald, and M. Brackstone, "Towards an understanding of adaptive cruise control," *Transportation Research Part C* 9, pp. 33–51, 2001.
- [5] A. Ghazy and T. Ozkul, "Design and simulation of an artificially intelligent VANET for solving traffic congestion," *Mechatronics and its Applications, ISMA09, 6th international symposium*, pp. 1–6, 2009.
- [6] R. Pueboobpaphan and B. van Arem, "Understanding the Relation between Driver/Vehicle Characteristics and Platoon/Traffic Flow Stability for the Design and Assessment of Cooperative Adaptive Cruise Control," *Proc. Transportation Research Board 89th Annual Meeting, Washington DC, USA*, 2010.
- [7] T. Mak, K. Laberteaux, and R. Sengupta, "A MultiChannel VANET Providing Concurrent Safety and Commercial Services," in *VANET05, Cologne, Germany*, September 2005, pp. 1–9.
- [8] Task Group P, "IEEE draft amendment to standard for information technology - telecommunications and information exchange between systems - local and metropolitan networks - specific requirements - part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications: Amendment: Wireless access in vehicular environments," *IEEE 802.11p Draft 9.0*, Sept 2009.
- [9] "ETSI TR 102 638 v1.1.1: Intelligent transport systems (ITS); vehicular communications; basic set of applications; definitions," European Telecommunications Standards Institute: Technical Committee Intelligent Transport System (ITS), Tech. Rep., 6 2009.
- [10] Connect & Drive project, Project no. HTASD08002, "Connect & Drive WP-1 requirements document," Tech. Rep., April 2009.
- [11] COMESafety, "D31 European ITS Communication Architecture," Informarion Society Technologies, Tech. Rep., 2008.
- [12] E. M. van Eenennaam, W. Klein Wolterink, G. Karagiannis, and G. J. Heijenk, "Towards Scalable Beaconing in VANETs," *Fourth ERCIM Workshop on e-Mobility*, pp. 103–108, May 2010.
- [13] L. Yang, J. Guo, and Y. Wu, "Channel adaptive one hop broadcasting for vanets," in *IEEE ITSC 2008*, Oct. 2008, pp. 369–374.
- [14] Y.-C. Tseng, S.-Y. Ni, Y.-S. Chen, and J.-P. Sheu, "The broadcast storm problem in a mobile ad hoc network," *Wirel. Netw.*, vol. 8, no. 2/3, pp. 153–167, 2002.
- [15] M. Torrent-Moreno, P. Santi, and H. Hartenstein, "Distributed fair transmit power adjustment for vehicular ad hoc networks," in *SECON '06*, vol. 2, Sept. 2006.
- [16] Q. Xu, T. Mak, T. Ko, and R. Sengupta, "Vehicle-to-vehicle safety messaging in dsrc," 2004, pp. 19–28.
- [17] *IEEE P1609.4/D9, August 2010, title=IEEE Draft Standard for Wireless Access in Vehicular Environments (WAVE) - Multi-Channel Operation*, pp. 1 –61, 2010.
- [18] A. Vinel, Y. Koucheryavy, S. Andreev, and D. Staehle, "Estimation of a successful beacon reception probability in vehicular ad-hoc networks," *IWCMC*, June 2009.
- [19] S. Eichler, "Performance evaluation of the IEEE 802.11p WAVE communication standard," *Vehicular Technology Conference, 2007. VTC-2007 Fall. 2007 IEEE 66th*, pp. 2199–2203, 30 2007-Oct. 3 2007.
- [20] Y. Wang, A. Ahmed, B. Krishnamachari, and K. Psounis, "IEEE 802.11p Performance Evaluation and Protocol Enhancement," *IEEE intl. conf. on Vehicular Electronics and safety*, Sept 2008.
- [21] OMNeT++ discrete event simulation framework, <http://www.omnetpp.org>.
- [22] MiXiM - Simulation Framework for Wireless and Mobile Networks, <http://mixim.sourceforge.net>.
- [23] G. Bianchi, "Performance analysis of the IEEE 802.11 Distributed Coordination Function," *Selected Areas in Communications, IEEE Journal on*, vol. 18, no. 3, pp. 535–547, Mar 2000.
- [24] P. Engelstad and O. Østerbø, "Non-saturation and saturation analysis of IEEE 802.11e EDCA with starvation prediction," in *MSWiM '05: Proceedings of the 8th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*. New York, NY, USA: ACM, 2005, pp. 224–233.
- [25] R. Jain, *The art of computer systems performance analysis: Techniques for experimental design, measurement, simulation and modeling*. John Wiley & Sons, Inc., 1991.
- [26] X. Ma and X. Chen, "Delay and broadcast reception rates of highway safety applications in vehicular ad hoc networks," in *Mobile Networking for Vehicular Environments*, 2007, pp. 85–90.
- [27] K. H. Suleiman, T. Javidi, M. Liu, and S. Kittipiyakul, "The Impact of MAC Buffer Size on the Throughput Performance of IEEE 802.11," Jan 2008.