On the potential of multicast in millimeter wave vehicular communications

Gereon Mendler
Dept. of Computer Science
University of Twente
Enschede, The Netherlands
g.t.mendler@student.utwente.nl

Geert Heijenk
Dept. of Computer Science
University of Twente
Enschede, The Netherlands
geert.heijenk@utwente.nl

Abstract—In this paper we explore the potential of doing multicast transmissions in millimeter wave vehicular communications using beamwidth-aware scheduling. Using a dataset of measured vehicle positions in a highway scenario, we show that often a considerable number of neighbouring vehicles can be reached with a single millimeter wave transmission. Furthermore, we propose a fully distributed beamwidth-aware scheduler using sub-6GHz control information exchange that opportunistically exploits multicast opportunities. We show that this scheduler can provide a significant delay reduction for millimeter wave vehicular communications.

Index Terms—millimeter wave, multicast, V2X

I. INTRODUCTION

Intelligent Transportation Systems (ITS) require increasing amounts of data to be exchanged between vehicles. Beyond the basic exchange of status messages (Cooperative Awareness Messages (CAM) in Europe, Basic Service Messages (BSM) in the US), modern vehicles generate large amounts of sensor information that is of potential benefit to surrounding vehicles. This could amount to several Gbit/s [1].

For V2X communications, spectrum in the 5.9 GHz band has been reserved, and communication standards have been defined for exchanging information for various ITS applications in this band, based on a variant of the IEEE 802.11 standard, often referred to as IEEE 802.11p. Furthermore, 3GPP has developed a series of standards to make cellular systems more suitable for automotive applications, referred to as C-V2X and LTE-V. However, these standards will not provide sufficiently high data rates to fulfil the requirements of all applications. For that reason millimeter wave (mmWave) communications is being considered as a complement to sub-6Ghz V2X technologies [2], [3]. Next to to high available bandwidth, mmWave communications is characterized by a limited communication range, and the need for beam-forming to transfer a sufficient amount of energy from transmitter to receiver.

[2] presents a proposal for a sub-6GHz assisted mmWave MAC, using position and scheduling information exchanged omnidirectionally using CAM/BSM messages on sub-6GHz to control beamforming and scheduling of the mmWave transmissions for V2X communication. One of the challenges identified, but not solved, in [2] is that of beamwidth-aware scheduling, where potentially packets can be transmitted to multiple

receivers in a single transmission, with a larger beamwidth. Indeed, for many ITS applications multiple receivers could be interested in the same information. In vehicular networks, geocasting over multiple hops is seen as an important mode of communications, but already within the direct transmission range of a vehicle (single hop) broadcasting or, more generally, multicasting is an important mode of communications, either to increase efficiency, to reduce the delay experienced, or because the intended receivers are not explicitly known.

The key research question addressed in this paper is to what extent multicasting can be beneficial for mmWave vehicular communications. We will first look into the trade-off between range and beamwidth and the resulting number of vehicles that can be reached using a single transmission. We will do so by using an empirical dataset regarding vehicle positions on a highway, combined with an accepted propagation model for mmWave communications. We will then extend the scheduling approach proposed in [2] to evaluate to what extent the potential of multicasting can be exploited in a real system.

The contributions of this paper are as follows: (i) We explore the potential of multicasting for mmWave communications; (ii) we introduce a scheduling mechanism that tries as much as possible to transfer messages to multiple nodes in a single transmission using a wider beamwidth; and (iii) we evaluate the performance of multicasting for mmWave using the proposed scheduling mechanism and compare it to unicasting.

The remainder of this paper is organized as follows. Related work is discussed in Section II. Section III explores the potential of multicast for mmWave communication by analyzing the beamwidth - range trade-off. A beamwidth-aware scheduling algorithm for mmWave multicast is introduced in Section IV, and subsequently evaluated in Section V. In Section VI, we draw conclusions, and identify potential future work.

II. RELATED WORK

The current paper builds on the research presented in [2]. The authors of that paper propose a system where mmWave MAC control functions are offloaded to omnidirectional sub-6GHz communications. They propose and evaluate a distributed scheduling algorithm for mmWave communications for which the control information is exchanged by piggybagging on CAMs/BSMs broadcasted periodically anyway (bea-
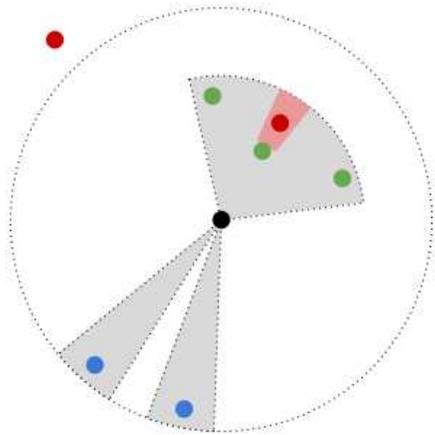
Fig. 1. Beamwidth - range trade-off (blue: in unicast range; green: in multicast range; red: out-of-range)

Parameter	Value
Carrier frequency	60.48 GHz
Transmission power	10 dBm
Receiver sensitivity	-66 dBm
Transmitter antenna gain	0 to 11.5 dBi
Receiver antenna gain	11.5 dBi
Transmitter efficiency	100%
Values for A; C	1.77; 70 (Line of sight) 1.71; 78.6 (1 obstruction)

cons). Basically, a node wishing to do mmWave transmissions to other nodes will indicate that wish, including the length of the transmission and the list of recipients in a so-called RTS (Request-to-send) message, piggybagged to a beacon. Intended recipients, receiving such an RTS message will plan the mmWave transmission, and announce the scheduled mmWave transmission from the transmitter to this receiver in a CTS (Clear-to-send) message piggybagged to their next beacon message. In this scheduling, it will take into account already scheduled mmWave transmissions from the same or other transmitters conflicting with the new transmission. MmWave transmissions may conflict due to either the transmitter or the receiver already being engaged in another transmission, or due to interference caused by other transmit beams within the receive beam. In the evaluation, the proposed mechanism shows superior delay performance, compared to a standard mmWave MAC (i.e., IEEE 802.11ad). The paper mentions the challenge of multicasting through beamwidth-aware scheduling, the topic of this paper, but does not provide any analysis or solutions in that direction.

Other papers addressing mmWave communications for vehicular networking include [3]–[7], but none of these papers addresses multicasting or addressing multiple receivers at the same time. Many papers do study mmWave communications to multiple users simultaneously in general. The authors of [8] for example, show that for a purely 802.11ad-based control, the control overhead increases linearly with the number of users, whereas the throughput gain is decreasing with the number of users. Contrary to the current paper, these studies are done in a very general setting, without taking into account the possibilities and limitations of vehicular networks, and without considering the exchange of control information on sub-6GHz. To the best of our knowledge, multicasting for mmWave communications, using sub-6GHz control or in a vehicular network setting has not been addressed before.

III. EXPLORATION OF MMWAVE MULTICAST POTENTIAL

Reaching multiple receivers with a single mmWave transmission requires the beam width to be wide enough to cover the receive antennas of all receivers. This comes at the cost of reduced antenna gain and as a result a reduced transmission range, as illustrated in Figure 1. In order to get insight into the multicasting potential, i.e., in the trade-off between beam width and communication range, we consider real vehicles locations in a highway setting, and explore to what extent other vehicles are in or out of range, within multicast range, or blocked by other vehicles. To that end, we combine an empirical model of mmWave communications with a dataset of vehicle positions in a highway scenario ([9]).

A. mmWave Communications Model

The communication and propagation model used adheres to the parameters used in [2] where possible. The listed parameters only refer to the data plane based on IEEE 802.11ad; it is assumed that the IEEE 802.11p based control plane performs sufficiently well within the relevant range and is therefore not explicitly taken into account. The data plane uses a variable beam-width antenna with a minimal beam width of 25.7 degrees, which corresponds to a 14 sector antenna. Moreover, a transmission power of 10 dBm is used at 100% efficiency. Consequently the transmitting antenna gain can vary between 0 and 11.5 dBi. The communication itself utilizes the second channel of 802.11ad, which is part of the unlicensed spectrum for all major regions. The channel is 2.16 GHz wide and centered around 60.48 GHz. A receiver sensitivity of -66 dBm allows for a data rate of 693 Mbps in correspondence with the Modulation and Coding Scheme (MCS) 13. Usually this coding scheme is cited as requiring a receiver sensitivity of -78 dBm, which however refers to the control channel requirements for 802.11ad. As previously mentioned, this aspect is offloaded to the control plane (802.11p) for this study. Additionally the antenna gain on the receiving end is assumed to be 11.5 dBi, the maximum achievable with a 14-sector antenna, and therefore requires beam alignment for successful communication between receiver and transmitter.

The propagation loss is modeled using the empirical path loss model presented in [10] explicitly meant for 60GHz Vehicular Ad-hoc Networks (VANETs). The model includes measurements under Line-Of-Sight (LOS) conditions as well
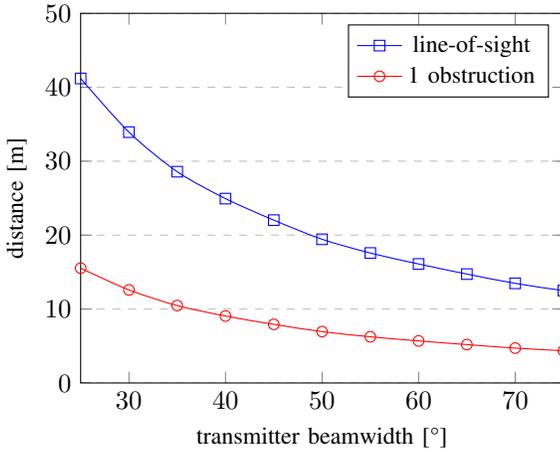
Fig. 2. Transmission ranges for mmWave communication in VANETs

Fig. 3. Average number of neighbours per vehicle compared to traffic density

as communication paths obstructed by one or more other vehicles. The path loss in dB is given as:

$$PL(dB) = A * 10 * log10(d) + C + 15 * d/1000, \quad (1)$$

where d is the distance between transmitter and receiver (in meters), and A and C are constants dependent on the number of blocking obstacles and visibility conditions. The values for A resp. C are taken from [4] and given in Table I, together with the other parameters of the model. Figure 2 shows the ranges for various beam widths and number of obstacles resulting from the model. The maximum transmission range for line-of-sight varies from 41.2m for a 25.7° beamwidth (unicast) to below 10m for omnidirectional transmissions, whereas with a single obstacle, these ranges reduce to 15.5m for a 25.7° beamwidth and a few meters for omnidirectional. Since 2 obstacles would reduce the range to a few centimeters, we do not further consider this case.

B. Vehicle Mobility Data

This study is based on the empirical US Highway 101 dataset, captured in June 2005 by Next Generation SIMulation (NGSIM) researchers ([9]). This data was originally intended for driver behaviour analysis and provides information about vehicle types, dimensions, locations and trajectories. The study area measures 640 meters in length and consists of the 5 lanes southbound US Highway 101. The data was collected every 100 milliseconds over three 15 minute periods; in this study we will solely focus on snapshots between 7:50am and 8:05am on June 15th, 2005. Due to the capturing approach, only 7500 snapshots (corresponding to a 12.5 minute interval) track all vehicles present in the observation area. The vehicle density during this period varies between 175 and 275, and averages at 196 vehicles per kilometer.

Vehicular obstructions in this dataset were discovered by means of geometric analysis using shapely, a BSD-licensed Python package facilitating the manipulation and analysis of planar geometric objects. The combination of these data with the mmWave communications model presented in Section
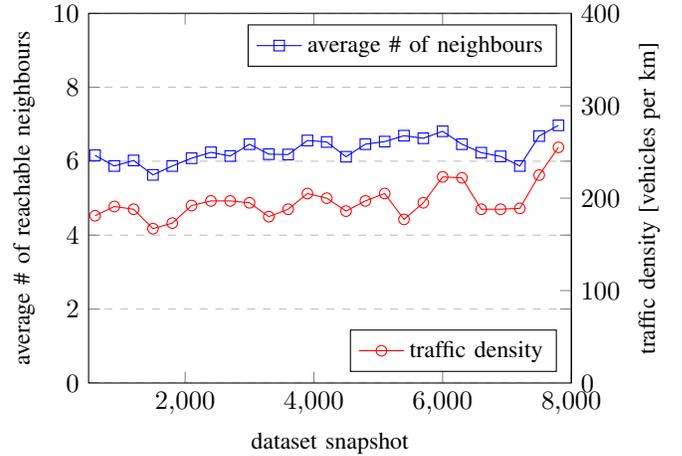
III-A provides insight into the reachable neighbours of every vehicle in the dataset. Fig. 3 shows the aggregation of this information, assuming a 25.7° antenna beam at both transmitter and receiver, i.e., unicast communications. Overall, a vehicle is in communication range with an average of 6.3 neighbours. Between the frames of the dataset, the average varies between 5.6 and 7.0 in correlation to the traffic density. Basically, increasingly dense traffic leads to more reachable neighbours, which intuitively makes sense since the average distance between vehicles would decrease accordingly. There are no isolated vehicles in any evaluated snapshot, in fact at least two other vehicles are reachable in all cases. Finally, we observed that vehicles can have up to 11 neighbours within communication range.

C. Initial Results

In most cases multicasting, sending a single transmission to multiple receivers, involves using a wider beam at the transmitter. However, using a wider beam negatively affects the communication range. A transmitter can transmit to multiple receivers at the same time (multicast), provided all are within its transmit antenna beam, and within range, given the beam width. Combining the communication model with the information about neighbours and obstructions extracted from the data provides a first insight into the multicast potential of mmWave vehicular communications. Here, we determine the number of multicast opportunities for every vehicle in every snapshot. An opportunity refers to the potential of transmitting to a distinct set of multiple receivers simultaneously. Note that the multicast opportunities are potentially overlapping, such that neighbours may be included in multiple ones. Selecting a proper set of multicast opportunities, so that all required neighbours receive a message, is one of the tasks of the scheduling function, which we will address in the next section.

On average, a vehicle has 4.96 multicast opportunities. Across all frames, only 1.78% of vehicles do not have any such opportunities. The average vehicle has the potential to multicast to 4.54 distinct sets of size two, and 0.41 groups
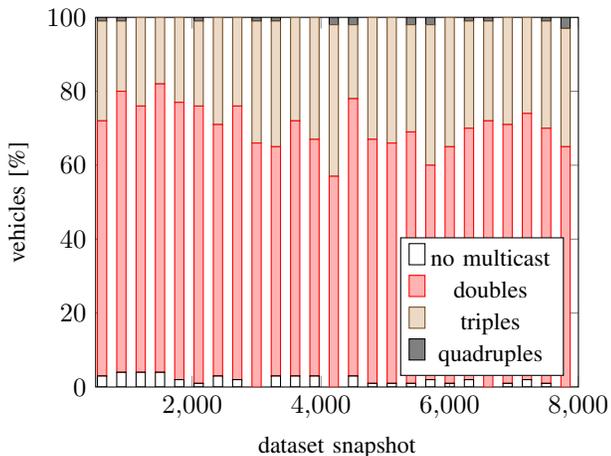
Fig. 4. Best multicast chances per vehicle, aggregated for each snapshot

of three vehicles. Opportunities for quadruple multicast are exceedingly rare, roughly 1 in 155 vehicles. Fig. 4 shows the breakdown per snapshot for the best opportunities per vehicle. With an average of 69%, and varying between 57.8% and 78.5%, the majority of vehicles can reach at most two of their neighbours simultaneously. For triplets, the respective values are 28.6% for the average and 17.8% and 40.6% for the lowest and highest snapshot. Reaching four vehicles at once is highly unlikely due to the small propagation range at larger angles and the size of vehicles, on average only 0.6% of vehicles are able to do so.

IV. PROPOSED BEAMWIDTH-AWARE SCHEDULING

MmWave communications is subject to several constraints with respect to when communication between a transmitter - receiver pair can take place. Coordinating transmission times and antenna directions for transmitters and receivers in such a way that conflicts at the receiver and severe interference are avoided is a task of the scheduling function. Ideally, scheduling exploits parallel transmissions and multicast opportunities as much as possible. In this section, we first give our assumptions with respect to scheduling, and then propose a beamwidth-aware scheduling mechanism.

A. Assumptions

We assume that for a node transmitting and receiving cannot take place simultaneously, that receivers can only receive from a single transmitter at a time, and that receiving a transmission may be subject to interference if an interferer within the receive beam of the receive antenna transmits in the direction of the receiver. Furthermore, we assume that a transmitter can transmit to multiple receivers at the same time (multicast), provided all are within its transmit antenna beam, and within range, given the beam width.

We define the proposed multicast scheduling mechanism using similar assumptions as in [2], as described in Section II. Through the exchange of CAMS/BSMs broadcasted

periodically on sub-6GHz, nodes have up-to-date information regarding positions of other nodes in their environment. Because the communication range at sub-6GHz is considerably larger than in mmWave, we assume this information is available up to at least 2 times the mmWave transmission range. The periodically broadcasted CAMs/BSMs also allow for piggybagging messages for the scheduling mechanism ensuring sequential consistency. We assume that within the area relevant for mmWave transmissions, the sequential exchange of CAMs/BSMs is taken care of by the access mechanism at sub-6GHz, and nodes receive all relevant scheduling messages.

B. Scheduling Algorithm

The proposed scheduling algorithm is based on the use of an RTS/CTS mechanism. Similarly to the MAC protocol presented in [2], the RTS and CTS messages are piggybagged on CAMs/BSMs, also called beacons in this context. A single beacon can contain one or more RTS messages, each indicating a request to have a transmission opportunity to a set of one or more neighbours, which is feasible in a single multicast transmission. A transmitter, wishing to send to a certain set of intended receivers, typically will issue an RTS for every feasible multicast or unicast opportunity to any subset of the set of intended receivers. Efficient coding of this list of RTSs is possible but beyond the scope of this paper. Further, a beacon can contain at most one CTS message, sent by one of the intended receivers indicated in one of the earlier sent RTSs, and including an identifier of that RTS and the intended starting time of a 50 milliseconds long communication window. As a consequence, the scheduling is ultimately determined by the receiving parties. The transmitter, after receiving the CTS, will start the transmission specified in the RTS at the indicated starting time. The receiver sending the CTS and all other intended receivers, after overhearing the CTS message, will start receiving the transmission at the indicated time, aiming their receive beam in the direction of the scheduled transmitter. The sequential consistency of the beacons ensures that receivers schedule mmWave transmissions sequentially, taking already scheduled transmissions into account.

For the receiver-controlled scheduling, each receiver decides between several RTS messages received from its neighbours according to multiple factors. From overheard CTS messages, it has learnt which transmissions have already been scheduled, possibly also including itself as a multicast destination. The latter obsoletes the need for scheduling certain RTSs. The receiver will remove all RTSs where it is not a receiver, or where the set of receivers intersects with that of an already scheduled RTS from its list of RTSs. For all the remaining RTSs, the receiver determines the earliest possible transmission time without conflicts. These conflicts can be incurred by transmissions already scheduled by other nodes, causing the transmitter and / or some of the intended receivers to be already engaged in a transmission for some time interval. Among all candidate RTSs, it will select the earliest multicast transmission (RTS), unless the earliest unicast transmission (RTS) can be scheduled at least a communication window

of 50 ms earlier, in which case the unicast transmission is scheduled. If a choice has to be made among multiple multicast transmissions that can be scheduled at the same time, the one with most receivers is chosen. This is to maximise the number of receivers of a transmission. If a choice has to be made among multiple unicast transmissions the one of which the transmitter has fewest neighbours is chosen. This is to keep the probability of scheduling a multicast transmission at a later moment as high as possible, and because transmitters with many neigbours have a higher chance of being selected for transmission anyway. Clearly, this scheduling mechanism favours multicast transmission over unicast transmission. However, in case selection of a multicast transmission would cause enough idle time for a unicast transmission to be transmitted first, such a unicast transmission is preferred, and scheduling of a multicast transmission is postponed. Ultimately, the scheduled transmission is announced by this receiver in a CTS message, piggybagged on the next transmitted beacon, as described above.

The scheduling algorithm at the receiver is described in pseudocode in Algorithm 1. In this algorithm, r_u^* is the best unicast RTS and r_m^* the best multicast RTS considered for scheduling by this node so far. R_ν is the set of all received RTS messages, r with this node, ν, as a receiver for which the set of receivers, I_ν, is disjoint with the receivers from RTSs already scheduled by other nodes. So, RTSs for which the transmission to some receivers has already been scheduled are not considered anymore. N_r denotes the set of reachable neighbours of the transmitter of r. Further, $t(r)$ denotes the earliest time at which r can be scheduled without any conflict with already scheduled transmission. The earliest possible multicast RTS is scheduled, unless there is none or the best unicast RTS can be scheduled at least ψs earlier, where ψ is set to the duration of the communication window.

As stated above, upon hearing a CTS message, all intended receivers for that transmission (as indicated in the relevant RTS) will also try to receive the scheduled transmission. In their next submitted beacon, all receivers of the transmission will include acknowledgements for the received mmWave transmission, so that the transmitter can issue RTS messages for transmissions to the remaining receivers if not all intended receivers did successfully receive the transmission.

V. EVALUATION

In order to get insight into the impact of scheduling constraints on mmWave multicast performance, and in order to assess the proposed scheduling algorithm, we evaluate the performance of the proposed algorithm. As much as possible, we use the evaluation scenarios of [2], combined with real mobility data, as described in Section III-B.

A. Scenario and Assumptions

We use a Monte Carlo type of simulation, considering all vehicles positions from the mobility traces described in III-B at a sequence of time instances (snapshots) 30 seconds apart, so that they can be considered independent. At those time

---

Algorithm 1 Multicast Scheduling per Vehicle ν

$r_u^* \leftarrow null$
$r_m^* \leftarrow null$
for all r **in** R_ν **do**
  if $|I_r| = 1$ **then**
    if $t(r) < t(r_u^*)$ **or** $(t(r) = t(r_u^*)$ **and** $|N_r| < |N_{r_u^*}|)$
    then
      $r_u^* \leftarrow r$
    end if
  else
    if $t(r) < t(r_m^*)$ **or** $(t(r) = t(r_m^*)$ **and** $|I_r| > |I_{r_u^*}|)$
    then
      $r_m^* \leftarrow r$
    end if
  end if
end for
if $r_m^* = null$ **or** $t(r_u^*) + \psi < t(r_m^*)$ **then**
  $schedule(r_u^*)$
else
  $schedule(r_m^*)$
end if

---

instances, we consider a communication need for a randomly selected subset of [5, 10, 15, 20, 25, 30, 35, 40] percent of the vehicles (similar to [2]). The selected transmitters want to transmit to all their neighbours for a communication window of 50 milliseconds long, containing 600 * 1600 bytes at a data rate of 693 Mbps (MCS 13). Neighbours are those nodes that can be reached using mmWave unicast transmissions. Because vehicle mobility during a sequence of these communication windows is very small, we assume vehicles to be stationary during a snapshot. We again assume the mmWave communications model described in Section III-A. For unicast, a 25.7° antenna beam at both transmitter and receiver is used. For multicast, a transmit antenna beam of a multiple of 25.7° is used, with a 25.7° antenna beam at the receiver. For the scheduling, perfect knowledge across nodes is assumed. Basically, any node is aware of all ongoing and planned transmissions, as well as the multicast opportunities and reachable neighbours of all other nodes. This is based on an underlying assumption that the beacons carrying RTS/CTS messages are transmitted every 50ms on sub-6GHz with a range that is at least double the mmWave communication range, which seems reasonable given typical ranges of vehicular communications in those frequency bands. This allows vehicles to listen to the scheduling decisions of all its neighbours' neighbours and plan opportunistic multicast transmissions accordingly. For both mmWave and sub-6GHz communication, failed transmissions and interference are not explicitly modelled. We argue that for mmWave communications, the relatively narrow beams at both transmitter and receiver and the hence required beam alignment, combined with the short ranges, contribute to less prevalent interference compared to scenarios without directional antennas.
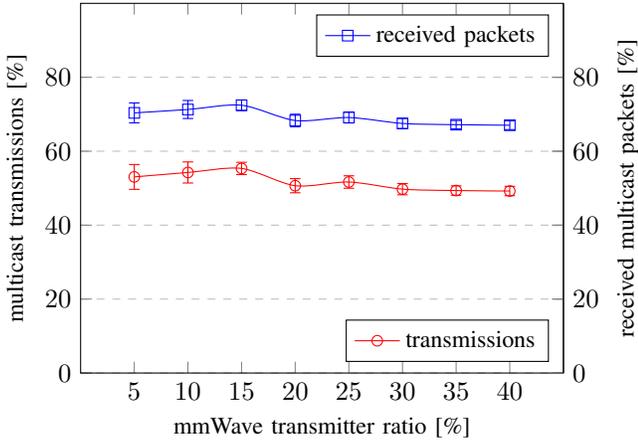
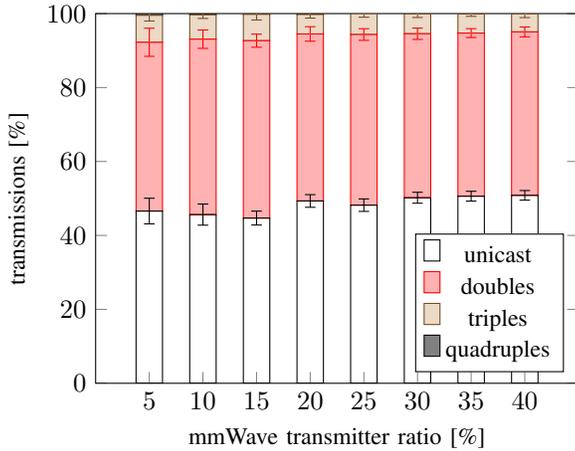Fig. 5. Portion of multicast transmissions and packets

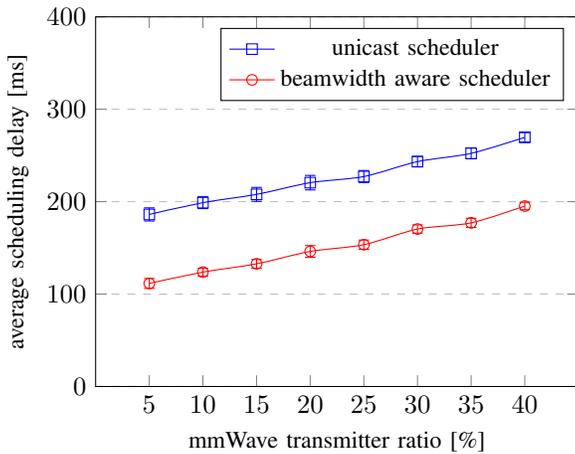Fig. 6. Number of receivers per transmission

Fig. 7. Average scheduling delay

B. Results

The performance of the proposed scheduling mechanism is evaluated by comparing it to a similar mechanism without multicast exploitation. We are interested in the number of receivers per transmission; i.e., is it unicast or multicast, and in case of the latter, with how many receivers? Furthermore, we are interested in the resulting scheduling delay for received transmissions, measured as the time between the generation of the data at the transmitter (the emergence of the communication need) and the arrival of the first bit of data at the receiver. This delay is averaged over all received packets in a snapshot, so a multicast packet is counted multiple times. Subsequently, this is averaged over all snapshots. Results are displayed with 95% confidence intervals

In Figure 5, it can be observed that roughly 50% of the transmissions done in the system are multicast transmissions. On average, a multicast transmission reaches slightly over 2 neighbours simultaneously, resulting in around 70% of the received packets originating from multicast transmission, as can also be observed in the figure. These numbers decrease only slightly with increasing traffic load, i.e., with a higher ratio of mmWave transmitters. This slight decrease can be attributed to the phenomena that at higher load, multicast transmissions have a higher probability of being delayed due to scheduling conflicts at some receiver. As a result, an intended receiver may give preference to a unicast transmission, to reduce idle time. If this unicast receiver could also be reached with a multicast transmission, the number of recipients for that multicast transmission is reduced, possibly to 1, effectively resulting in another unicast transmission.

To get more insight in the behaviour of the scheduler, Figure 6 displays the number of receivers per transmission. Please note that this figure is different from Figure 4, which is based on the best multicast opportunity per node (with the highest number of receivers). This figure is based on all transmission of a node. In a scenario without any scheduling conflicts that would be the best multicast opportunity with possibly some additional multicast and/or unicast transmission to reach all neighbours. From Figure 6 we can observe that around 5% of the transmissions have at least 3 receivers (triple), whereas the number of transmissions with even 4 receivers (quadruple) is, hardly visible, below 1%. We can observe that with increasing load fewer triples are scheduled, which we attribute to more frequent scheduling conflicts, leading to scheduling of transmissions with fewer receivers. Surprisingly, the number of doubles scheduled is not really affected by the load. Until 15% transmitter ratio, this number is even growing, which may be explained by the reduction of scheduling conflicts due to triples.

In Figure 7, we compare the average scheduling delay of the proposed beamwidth-aware scheduler and a scheduler that solely schedules unicast packets. This last unicast scheduler is similar to the one used in [2]. The figure shows that the average scheduling delay is consistently lower for beamwidth aware scheduling than for unicast scheduling. The delay increases
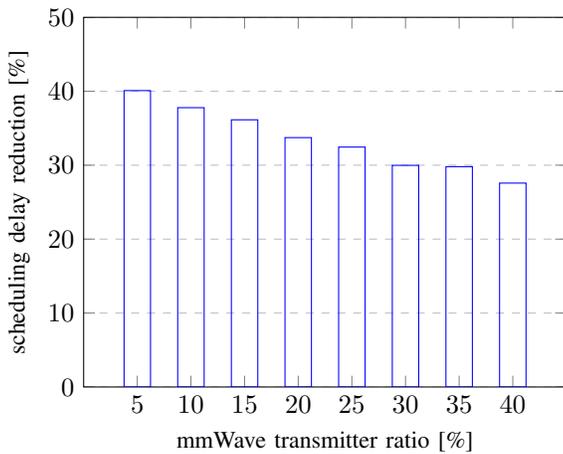
Fig. 8. Average delay reduction

linearly with the transmitter ratio for both scheduling methods, since the increase in transmitters corresponds with an increase in the number of transmissions and hence scheduling conflicts. We have indications that these linear trends continue for higher loads, up to 100% transmitter ratio. However, to analyse the scheduler in capacity-limited scenarios, e.g., with much more frequent communication demand, the current snapshot-based simulation, assuming stationary users during a snapshot and no correlation between snapshots is not suitable. Our expectation is that unicast scheduling will reach the capacity limits, resulting in excessive delay, at considerably lower loads than multicast scheduling, due to the less efficient resource usage. Overall, we observe that the proposed multicast scheduler provides a significant reduction of scheduling delay by exploiting multicast transmission opportunities. This can also be seen from Figure 8, which shows a scheduling delay reduction between 25 and 40 percent for beamwidth-aware scheduling.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have explored the potential of multicast and beamwidth-aware scheduling in millimeter wave vehicular communications. Using a dataset of real vehicle locations combined with an empirical mmWave communications model we have observed that mmWave transmitters could reach often multiple, up to four, receivers with a single transmission by transmitting with a wider antenna beam. To exploit this potential, we have proposed a fully distributed beamwidth-aware scheduler, relying on coordination using beacons transmitted on sub-6GHz. From the simulation model presented in this paper, we conclude that our proposed scheduler indeed significantly reduces the number of transmissions by doing multicast transmissions. As a result, we obtain a significant delay gain compared to unicast scheduling.

We expect that at high loads, this way of scheduling will also provide a significant throughput gain. Since the current simulator considers independent snapshots of vehicle positions, it is less suitable for high-load situations in which correlation of snapshots and vehicle mobility have to be taken into account. In future work, we would like to further investigate this, also studying the impact of interference and packet loss, both in mmWave communications and in sub-6GHz beacon transmissions. Other future work is to extend this fully distributed beamwidth-aware scheduling using sub-6GHz communications to multihop communication scenarios using relaying, e.g., to support mmWave geocasting.

REFERENCES

[1] G. Heijenk, Wireless networks hit the road: towards cooperative automated driving, Enschede: University of Twente, 2019, 54 p.

[2] B. Coll-Perales, J. Gozalves, M. Gruteser, "Sub-6GHz assisted MAC for milimeter wave vehicular communications", IEEE Communications Magazine, March 2019, pp. 125 - 131.

[3] G. Naik, B. Choudhury and J. Park, "IEEE 802.11bd & 5G NR V2X: Evolution of Radio Access Technologies for V2X Communications," in IEEE Access, vol. 7, pp. 70169-70184, 2019, doi: 10.1109/ACCESS.2019.2919489.

[4] B. Coll-Perales, M. Gruteser and J. Gozalvez, "Evaluation of IEEE 802.11ad for mmWave V2V communications," 2018 IEEE Wireless Communications and Networking Conference Workshops (WCNCW), Barcelona, 2018, pp. 290-295, doi: 10.1109/WCNCW.2018.8369031.

[5] A. Molina-Galan, B. Coll-Perales and J. Gozalvez, "C-V2X Assisted mmWave V2V Scheduling," 2019 IEEE 2nd Connected and Automated Vehicles Symposium (CAVS), Honolulu, HI, USA, 2019, pp. 1-5, doi: 10.1109/CAVS.2019.8887840.

[6] J. Choi, et al., "Millimeter-Wave Vehicular Communication to Support Massive Automotive Sensing", IEEE Communications Magazine, vol. 54, no. 12, pp. 160-167, December 2016.

[7] H. Assasa and J. Widmer, "Implementation and Evaluation of a WLAN IEEE 802.11ad Model in ns-3", Proceedings of the Workshop on ns-3 (WNS3), pp. 57-64, Jun. 2016.

[8] W. Wu, Q. Shen, M. Wang and X. Shen, "Performance analysis of IEEE 802.11.ad downlink hybrid beamforming," 2017 IEEE International Conference on Communications (ICC), Paris, 2017, pp. 1-6, doi: 10.1109/ICC.2017.7996975.

[9] U.S. Department of Transportation Federal Highway Administration, (2016), Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data, Provided by ITS DataHub, http://doi.org/10.21949/1504477, 2016.

[10] A. Yamamoto, et al., "Path-Loss Prediction Models for Intervehicle Communication at 60 GHz", IEEE Transactions on Vehicular Technology, vol. 57, no. 1, pp. 65-78, Jan. 2008.