

Optimal Geocast Scheduling under Multicasts and Relaying in mmWave Vehicular Networks

Thijs Havinga*, Suzan Bayhan†, and Geert Heijenk†

* Ghent University, Ghent, Belgium, Thijs.Havinga@UGent.be

†University of Twente, Enschede, The Netherlands, {s.bayhan, geert.heijenk}@utwente.nl

Abstract—Due to the increasing volume of data generated by sensors in modern cars, the need for high data rate links rises in vehicular networking. A promising way to achieve this is using mmWave communications, although beamforming is needed to overcome the high propagation losses at these frequencies. In addition, relaying might be needed to extend the coverage to larger distances for the delivery of a message in a specific geographical area, called *geocasting*. Moreover, reaching multiple receivers at once via *multicasting* can be achieved by using a wider antenna beamwidth, which comes at the cost of transmission range, while spatial sharing can be exploited using narrow beams. This paper investigates if using multicasts is beneficial for routing and scheduling of mmWave geocasts that need to be delivered before a timeout. We consider a non-time-slotted system with realistic antenna model and multiple data rates, for which we seek an optimal solution by modeling it as a mixed-integer linear program. Our numerical evaluations show that using multicasts is especially advantageous in scenarios with multiple highway lanes. Furthermore, we devise a heuristic algorithm that efficiently finds a route and creates a transmission schedule for a geocast. Several methods to include multicast links are evaluated, of which some consistently outperform the unicast-only method.

Index Terms—mmWaves communications, vehicular networks, geocast, multicast.

I. INTRODUCTION

Vehicles equipped with diverse sensors perceiving the environment promise many benefits including improved traffic safety, road use efficiency, and reduced fuel consumption [1]. However, in a dynamic and crowded environment, a vehicle may not obtain all the needed information in time and needs to communicate with other vehicles for a more complete view of its operating environment. While typically sub-6GHz bands are used for such vehicle-to-vehicle (V2V) communications, with an increasing volume of data from (camera) sensors, mmWave communications with plenty of available bandwidth presents a more promising alternative [2]. However, the downside of this technology is its severe path loss and sensitivity to blockage, resulting in very limited transmission range.

These shortcomings of mmWave bands can be partly overcome by directional transmissions which let the transmitter and receiver point their antennas towards each other using a narrow beam. Yet, the process of beamforming leads to significant overhead in a dynamic environment due to the need for several training frames to obtain the channel quality so that the transmitter can determine which antenna sectors and data rate to use [3]. As presented in [2] and [4], position information can be acquired via sub-6GHz communication to overcome this

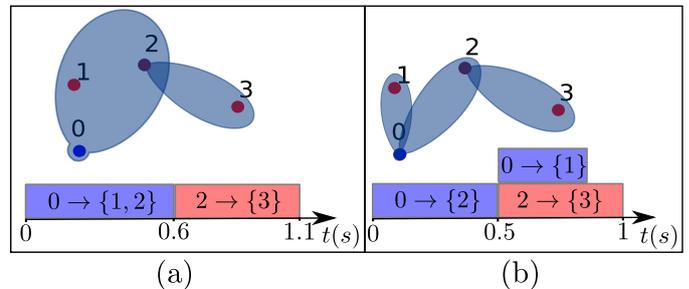


Fig. 1. Topology and transmission schedule for node 0 to reach nodes 1, 2 and 3 in case of using a multicast (a) or using unicasts only (b).

training challenge. From the position information, link quality can be estimated based on antenna and propagation models, so that the right links can be chosen. Also, offloading scheduling decisions to sub-6GHz band reduces control overhead.

While leveraging sub-6GHz bands mitigate the challenge of beam alignment, deciding on the optimal direction, the beamwidth, and the data rate of transmissions is not straightforward for various reasons. First, given that vehicular settings involve geocasts, i.e., delivering a message in a certain geographical area, multicasts can be exploited to transmit a message to multiple destinations simultaneously. On the other hand, unicasts might be more favorable as concurrent transmissions between different pairs of nodes can be exploited more frequently. Additionally, the multicast data rate needs to be determined considering the user with the weakest signal in a multicast group, resulting in a trade-off between number of nodes reached in a single transmission and the achieved throughput. For example, consider the scenario in Fig. 1, in which node 0 wants to reach nodes 1, 2 and 3. In this case, the delay for reaching the last receiver in case of using a multicast (a), is larger than using unicasts only (b), because the wider beam of the multicast requires a lower data rate to cover the distance to node 2. However, if node 1 would also need to transmit to another node, using a multicast is likely to be favorable, since in that case the node will be idle sooner. In addition to the introduced multicast vs. unicast decision, a transmitter might need to decide whether to transmit its message in a single hop or using one or more relays.

Since communications in vehicular scenarios are typically safety-critical, messages must be delivered before a certain timeout. For example, a vehicle may inform other vehicles that it executes an emergency break, or that it wants to merge into a

certain lane. If these messages are delivered when the vehicles can no longer act on it, the content is not useful anymore. Therefore, it is paramount to consider aforementioned aspects of spatial sharing and using various beamwidths as well as data rates such that the number of intended receivers which get the message before the deadline is maximized.

In this paper, we seek answers to the following questions: can multicasts and relays be exploited to improve the delivery ratio of messages for vehicular scenarios? How can a mmWave geocast be efficiently routed and scheduled using a lightweight algorithm given the possibility of using different beamwidths, data rates, and multicast receiver groups? Prior works such as [5], [6] have shown the potential of relays and multicasts for mmWave vehicular scenarios, but mostly separately or under the assumption of a (synchronous) time-slotted system and with simplistic antenna models. Therefore, these works disregard the flexibility of IEEE 802.11-based systems, and do not incorporate the implications of realistic beams. Moreover, these works generally consider all nodes to be intended receivers, whereas in vehicular networks only a select group of nodes (in the geocast area) should be reached, while others may be used for relaying. In particular, our contributions in this paper are twofold:

- We consider multicast, relaying and spatial sharing to find the optimal routing and scheduling for multiple mmWave geocasts in a vehicular network under a realistic antenna model and in a non-time-slotted system. Solving the formulated problem optimally using linear programming gives insight into the theoretical performance of the system, either when only unicasts are allowed or when multicasts may also be used.
- Given the high complexity of the optimal solution, we design a low-complexity heuristic algorithm for a single geocast. Numerical evaluations show a reduction in computational complexity, which comes at the cost of limited performance degradation. Two methods for leveraging multicasts improve the schedule consistently.

II. RELATED WORK

mmWave communications for vehicular networks: Due to the abundant spectrum in these bands (i.e., 30 to 300 GHz), mmWave communications have attracted a lot of interest for their potential use in vehicular scenarios. However, especially in vehicular environments, neighbor discovery, beamforming and scheduling might lead to significant overhead [2]. As shown in [4], using vehicle position information can outperform traditional beamforming approaches of IEEE 802.11ad in terms of average network throughput, up to a mean position error of 3m. Instead of using sector level sweeps, antenna and propagation models are needed to estimate the received power. Position information should be received out-of-band, for example using sub-6GHz technology, to achieve overhead-free beamforming. The authors of [2] address the use of sub-6GHz V2V technologies for the scheduling of beamformed mmWave transmissions. Suggestions for future research are to exploit the possibility to schedule multiple receivers at the

same time and relaying to reach vehicles at larger distances and to obtain spatial sharing.

Multicast scheduling in mmWave networks: There exist some studies on routing and scheduling with directional antennas to provide either multicast, relaying, spatial sharing, or a combination of these. Among these studies, [7] and [8] consider only multicast grouping and both describe a method in which for one transmitter the optimal sets of multicast receivers are determined. In [9], the potential of using multicasts in mmWave vehicular networks is shown for a single data rate and time-slotted system. The authors of [10] address both multicast and relaying, a problem that they prove to be NP-complete. They propose a time-slotted scheme in which the maximum achievable rate for each frame is obtained. The algorithm performs better than using only multicast or relaying. Both relaying and spatial sharing is considered in [11], but they consider unicasts only. Their algorithm maximizes the number of transmissions in a time slot. In [5], a heuristic algorithm is presented that offers the combination of multicast, relaying and spatial sharing, which is proven to be NP-hard. In this method, the closest transmitter is selected for each receiver. Then, each transmitter determines to which of the selected receivers it will send in the upcoming time slot by maximizing the sum throughput. The results show the importance of using relaying and spatial sharing. A routing and scheduling protocol for broadcast messages in multi-rate wireless networks with directional antennas is presented in [6]. They consider multicast and relaying as well as spatial sharing. Their results show that the use of multiple data rates and directional antennas leads to shorter transmission delay compared to the systems which do not include those. *Different from these works, our work aims at designing a scheduling and routing scheme in a non-time-slotted system exploiting multicast, relaying, and spatial sharing jointly. Additionally, we aim at geocasts, which also means that not all nodes need to receive the message, but they may be used for relaying. In contrast to the aforementioned works, we consider a realistic antenna model which is essential to determine the supported data rates under a beam configuration.*

III. SYSTEM MODEL

Let us consider a network of N nodes placed in a 2D-plane, capable of sending and receiving mmWave messages. We denote this set by $\mathcal{N} = \{0, 1, \dots, N-1\}$. We denote by $\mathcal{M} = \{m_0, \dots, m_j, \dots, m_{M-1}\}$ the set of mmWave messages generated during the considered time of a scenario from t_{\min} to t_{\max} . A mmWave message m_j is generated at time $t_{m_j}^{\text{gen}}$, which is triggered by an application on a higher level by node $n_{m_j} \in \mathcal{N}$, called the origin node. We assume the nodes are static for the duration of a scenario.

The set of intended receivers \mathcal{R}_{m_j} of message m_j are nodes other than the origin node, which are positioned in a specific geographical destination area, or are individually addressed. An intended receiver should receive the message before the timeout O , determined by the origin node. Per scenario, we use a fixed timeout for all messages, which all

have a fixed size S (in Mbits). Initially, only the origin node can send the mmWave message. After another node receives this message, it can transmit it to others. We assume that the mmWave communication is half-duplex, meaning that nodes cannot transmit and receive simultaneously. Furthermore, a node can only transmit or receive one message at a time.

To model the propagation of the mmWave communication, we use an empirical path loss model derived for vehicular communications at 60 GHz from [12]. This model includes measurements under line-of-sight (LOS) and non-LOS, when the link is obstructed by one or more vehicles. The path loss $\Gamma_{T \rightarrow R}$ in dB between a transmitter T and receiver R is [12]:

$$\Gamma_{T \rightarrow R}(d_{T \rightarrow R}) = A \cdot 10 \cdot \log_{10}(d_{T \rightarrow R}) + C + 15 \cdot \frac{d_{T \rightarrow R}}{1000}, \quad (1)$$

where $d_{T \rightarrow R}$ is the distance between the antenna of the transmitter and receiver (in m), A is the path loss exponent and C is a constant. Both A and C depend on the number of vehicles obstructing the link. To determine this number, we use the number of nodes that are crossed when a straight line from the transmitter towards the receiver is outlined, using a fixed car width (2 m) and length (5 m).

Each node transmits with a fixed power P_T (in dBm). Both the transmitter and receiver use a beamforming antenna. The transmitter can choose to transmit using a beamwidth $w \in \mathcal{W}$ and a transmission direction angle $a \in \mathcal{A}$. The set of beamwidths is determined by the number of sectors of the antenna. To calculate the directional antenna gain, we use a symmetrical antenna model using an average side-lobe level and Gaussian shape main-lobe from [13]. This model is used in the mmWave-based IEEE 802.15.3c standard and is considered simple, yet realistic. The directional antenna gain G of a node transmitting using half-power beamwidth w (in degrees), observed at an angle $\alpha \in [-180^\circ, 180^\circ]$ relative to the direction of transmission, is given by:

$$G(w, \alpha) = \begin{cases} G_0 - 3.01 \cdot \left(\frac{2\alpha}{w}\right)^2, & 0 \leq |\alpha| \leq w_{ml}/2 \\ G_{sl}, & w_{ml}/2 \leq |\alpha| \leq 180^\circ \end{cases}$$

$$G_0 = 20 \log_{10} \left(\frac{1.6162}{\sin(w/2)} \right)$$

$$G_{sl} = -0.4111 \cdot \ln(w) - 10.579$$

where $w_{ml} = 2.58 \cdot w$ is the main lobe width (in degrees), which may range from 15° to 60° . G_0 is the maximum antenna gain and G_{sl} is the side lobe gain. In this paper, we allow beamwidths of 15° , 30° , 45° , 60° and 360° (where the latter provides 0 dB directional antenna gain). The transmitter can point its antenna towards a discrete number of angles in the range $[0^\circ, 360^\circ]$, for which we consider a granularity of 1° .

If a node has been addressed as a receiver, it will point its antenna towards the transmitter using the narrowest possible beamwidth w_{\min} , resulting in an antenna gain equal to $G(w_{\min}, 0) = G_0(w_{\min})$. The power from transmitter T using beamwidth w as received at an angle α at the receiver R is then given by:

$$P_{T \rightarrow R} = P_T + G(w, \alpha) + G_0(w_{\min}) - \Gamma_{T \rightarrow R}(d_{T \rightarrow R}). \quad (3)$$

A mmWave transmission is possible if the power from transmitter to receiver is higher than a certain threshold, the receiver sensitivity σ (in dBm), such that the signal can still be decoded. The receiver sensitivity is defined by the used Modulation and Coding Scheme (MCS), based on a maximum allowed Packet Error Rate (PER). We use the directional multi-gigabit PHY specification of the IEEE 802.11 standard [14], comprising of 12 different MCSs for Orthogonal Frequency Division Multiplexing (OFDM). The available data rates comprise the set \mathcal{D} and the corresponding receiver sensitivity is given by $\sigma(r)$.

While it is not possible for a node to transmit or receive multiple messages at a time, concurrent transmissions between sets of nodes that do not overlap are possible. For the sake of simplicity, we ignore the interference that might occur due to concurrent transmissions. Due to the antenna directivity, especially the receiving beam, the influence of interference will be limited in realistic scenarios. Moreover, as moving vehicles keep distance to each other, their antennas will not be very close to each other, such that it is unlikely that they pick up significant power from unwanted signals. Although movement is not taken into account, in realistic scenarios the location of a vehicle in the future (at the time of transmission) can be estimated based on speed, acceleration and heading information. To account for errors, beam calculations can be made more conservative. Nevertheless, relative movement between vehicles with the same heading will be limited. Albeit our method provides a centralized solution on a per-scenario basis, it can be applied to a distributed system based on beacons, as we show in [15]. In this way, the schedule is only fixed up until the next beacon, meaning that the problem is divided into smaller pieces, such that each time the input is updated with fresh information on positions.

IV. PROBLEM STATEMENT

Now, we can formulate routing and scheduling of mmWave transmissions as an optimization problem. Let us first assume that all knowledge is available to an oracle, e.g., the generation time of the messages, the intended receivers for all messages, the channel gains and vehicle locations. The oracle is given an unconstrained number of transmission opportunities, called frames ($f_i \in \mathcal{F}$). It will then determine for each frame the message to be transmitted (m_j), the transmitter (T), the receiver set (\mathcal{R}_i) and the starting time of the transmission ($t_{f_i}^{start}$), which is not restricted to a slot. Table I lists the symbols including our decision and helper variables in the optimization problem. All decision variables are binary, except the start time and duration of a frame, which are continuous.

We consider two objectives with different priority as seen in (4) and (5). First, the program maximizes the number of intended receivers reached within the timeout for all messages. Then, for all messages, the program minimizes in (5) the time between message generation and the time at which the last frame ends.

TABLE I

SYMBOLS AND DEFINITIONS FOR THE PARAMETERS, DECISION AND HELPER VARIABLES USED IN THE MILP.

$$1) \max \sum_{f_i \in \mathcal{F}} \sum_{m_j \in \mathcal{M}} \sum_{R \in \mathcal{N}} \rho_{f_i, m_j, R}^s \quad (4)$$

$$2) \min \sum_{m_j \in \mathcal{M}} \max_{f_i \in \mathcal{F}} \{s_{f_i, m_j} \cdot (t_{f_i}^{start} + \delta_{f_i}^f - t_{m_j}^{gen})\} \quad (5)$$

subject to

$$\sum_{m_j \in \mathcal{M}} s_{f_i, m_j} = q_{f_i}, \quad \forall f_i \quad (6)$$

$$\sum_{T \in \mathcal{N}} u_{f_i, T} = q_{f_i}, \quad \forall f_i \quad (7)$$

$$\sum_{\mathcal{R}_l \in \overline{\mathcal{R}}} w_{f_i, \mathcal{R}_l} = q_{f_i}, \quad \forall f_i \quad (8)$$

$$2 \cdot w_{f_i, \mathcal{R}_l} \leq u_{f_i, T} + y_{\mathcal{R}_l, T}, \quad \forall f_i, \forall \mathcal{R}_l, \forall T \quad (9)$$

$$2 \cdot v_{f_i, R} \leq w_{f_i, \mathcal{R}_l} + x_{\mathcal{R}_l, R}, \quad \forall f_i, \forall \mathcal{R}_l, \forall R \quad (10)$$

$$\delta_{f_i}^f = \sum_{\mathcal{R}_l \in \overline{\mathcal{R}}} w_{f_i, \mathcal{R}_l} \cdot \delta_{\mathcal{R}_l}^s, \quad \forall f_i \quad (11)$$

$$t_{f_i}^{start} \geq s_{f_i, m_j} \cdot t_{m_j}^{gen}, \quad \forall f_i, \forall m_j \quad (12)$$

$$\phi_{f_i, f_{i'}}^< = 1 \iff t_{f_i}^{start} + \delta_{f_i}^f \leq t_{f_{i'}}^{start}, \quad \forall f_i, \forall f_{i'}, f_i \neq f_{i'} \quad (13)$$

$$u_{f_i, n} + v_{f_{i'}, n} \leq \phi_{f_i, f_{i'}}^< + \phi_{f_{i'}, f_i}^< + 1, \quad \forall f_i, \forall f_{i'}, f_i \neq f_{i'}, \forall n \quad (14)$$

$$u_{f_i, T} + u_{f_{i'}, T} \leq \phi_{f_i, f_{i'}}^< + \phi_{f_{i'}, f_i}^< + 1, \quad \forall f_i, \forall f_{i'}, f_i \neq f_{i'}, \forall T \quad (15)$$

$$v_{f_i, R} + v_{f_{i'}, R} \leq \phi_{f_i, f_{i'}}^< + \phi_{f_{i'}, f_i}^< + 1, \quad \forall f_i, \forall f_{i'}, f_i \neq f_{i'}, \forall R \quad (16)$$

$$3 \cdot l_{f_i, f_{i'}, m_j, n} \leq \phi_{f_i, f_{i'}}^< + s_{f_i, m_j} + v_{f_{i'}, n}, \quad \forall f_i, \forall f_{i'}, f_i \neq f_{i'}, \forall m_j, \forall n \quad (17)$$

$$s_{f_i, m_j} + u_{f_{i'}, n} \leq o_{m_j, n} + \sum_{f_{i'} \in \mathcal{F} \setminus f_i} l_{f_i, f_{i'}, m_j, n} + 1, \quad \forall f_i, \forall m_j, \forall n \quad (18)$$

$$s_{f_i, m_j}^< \leq s_{f_i, m_j}, \quad \forall f_i, \forall m_j \quad (19)$$

$$s_{f_i, m_j}^< = 1 \implies t_{f_i}^{start} + \delta_{f_i}^f \leq t_{m_j}^{gen} + O, \quad \forall f_i, \forall m_j \quad (20)$$

$$3 \cdot \rho_{f_i, m_j, R}^s \leq \rho_{m_j, R} + v_{f_i, R} + s_{f_i, m_j}^<, \quad \forall f_i, \forall m_j, \forall R \quad (21)$$

$$\rho_{f_i, m_j, R}^s \leq 1 - \sum_{f_{i'} \in \mathcal{F} \setminus f_i} \rho_{f_{i'}, m_j, R}^s, \quad \forall f_i, \forall m_j, \forall R \quad (22)$$

The constraints can be briefly explained as follows. Consts. 6-8 ensure that if a frame is used, there is exactly one message included, one transmitter assigned and one receiver set selected for it, respectively. Following, Const. 9 specifies that if a receiver set is assigned to a frame, the transmitter corresponding to this set should send the frame. Furthermore, if a receiver is addressed in a frame, the specific set that contains this receiver should be chosen for the frame, as modeled by Const.10. Const.11 ensures that the duration of a frame is equal to what is given for the receiver set that is selected for the frame. Next, if a message is included in a frame, the frame should start at a time equal to or greater than the generation time of that message, which is ensured by Const.12. Const.13 enforces the helper variable $\phi_{f_i, f_{i'}}^<$ to be 1 if and only if frame $f_{i'}$ starts not earlier than the end of another frame f_i , i.e., these frames do not overlap. Const.14 ensures that if a node is the transmitter of frame f_i and a receiver of frame $f_{i'}$, these should be sent at non-overlapping times. Furthermore, we ensure by Const.15 that

Symbol	Definition
Parameters	
$\mathcal{N} = \{\dots, n/T/R, \dots\}$	Nodes
$\mathcal{F} = \{\dots, f_i, \dots\}$	MmWave frames
$\mathcal{M} = \{\dots, m_j, \dots\}$	MmWave messages
$\overline{\mathcal{R}} = \{\dots, \mathcal{R}_l, \dots\}$	Set of all possible receiver sets
$t_{m_j}^{gen}$	Generation time (in s) for message m_j
O	Timeout of a mmWave message (in s)
$\rho_{m_j, n}$	$1 \iff n$ is intended receiver for m_j
$o_{m_j, n}$	$1 \iff n$ is origin node for m_j
$x_{\mathcal{R}_l, R}$	$1 \iff$ receiver R is in set \mathcal{R}_l
$y_{\mathcal{R}_l, T}$	$1 \iff T$ is transmitter for set \mathcal{R}_l
$\delta_{\mathcal{R}_l}^s$	Duration (in s) for transmission to \mathcal{R}_l
Decision variables	
q_{f_i}	$1 \iff$ frame f_i is used
s_{f_i, m_j}	$1 \iff$ message m_j is included in f_i
$u_{f_i, T}$	$1 \iff$ transmitter T sends f_i
w_{f_i, \mathcal{R}_l}	$1 \iff$ receiver set \mathcal{R}_l is included in f_i
$t_{f_i}^{start}$	Start time of frame f_i
Helper variables	
$\delta_{f_i}^f$	Duration (in s) of frame f_i
$v_{f_i, R}$	$1 \iff$ receiver R is addressed in frame f_i
$\phi_{f_i, f_{i'}}^<$	$1 \iff$ frame f_i ends before start of frame $f_{i'}$
$s_{f_i, m_j}^<$	$1 \iff$ frame f_i containing m_j ends before timeout
$\rho_{f_i, m_j, R}^s$	$1 \iff$ intended receiver R for message m_j is reached in time in frame f_i
$l_{f_i, f_{i'}, m_j, n}$	$1 \iff$ node n received message m_j in frame $f_{i'}$ before f_i

frames may not overlap if a node is the transmitter for both and by Const.16 if a node is a receiver of distinct frames. We formulate Const.17 and 18 to allow a node to relay a message in a frame if it has received the message in a frame before. Finally, for a receiver R to be successfully reached with message m_j in a frame f_i , the following three constraints must hold: the receiver should be an intended receiver for the specific message, the receiver should be addressed in the frame and the frame including this message is received before the timeout. This is ensured by Consts.19-21. Moreover, once the receiver successfully received the message in a frame, delivery of the same message in another frame does not count as a successful transmission anymore, as Const.22 specifies.

To solve this mixed integer linear problem (MILP), first we pre-process the given scenario for obtaining the input parameters to this program. We determine all possible receiver sets for all nodes, meaning the nodes that a transmitter can reach (i.e., the power from transmitter to receiver is higher than the receiver sensitivity), by an exhaustive search through all combinations of beamwidths, data rates, and transmission direction angles. For each receiver set, we find the highest data rate such that all nodes can be reached. This determines the transmission duration when choosing this receiver set. If multiple beamwidths are possible, our algorithm prefers the narrowest beam to ensure higher gain and lower interference, whereas the transmission direction angle is chosen arbitrarily if multiple are valid.

Solving the formulated MILP by Gurobi [16], we observe that for larger scenarios, the number of input variables and also the computation time of the solver increases rapidly. Hence, we will design a lower-complexity heuristic for the problem. We will consider only a single message, while ensuring the algorithm is suitable to apply in a distributed system in which multiple concurrent geocasts need to be scheduled, as we elaborate on in [15].

V. HEURISTIC ROUTING AND SCHEDULING

Our proposal to find a route and schedule for a **single** geocast consists of the following three steps: i) link assessment, ii) graph reduction, and iii) transmission tree generation. Before we present each step in details in the following sections, let us first provide a high-level overview of our proposal.

First, the heuristic algorithm pre-processes the scenario in a non-exhaustive way to find which receiver sets a transmitter can reach and which beamwidth, data rate, and transmission direction angle it should use for each receiver set. Due to the shape of the realistic antenna model, this is not straightforward. We will call hereafter this step *link assessment*.

The links that are obtained in the previous step can be used to set up a graph, in this case a *directed hypergraph*. A directed hyperedge represents which receivers a transmitter can reach using a certain beamwidth, data rate, and transmission direction angle. In the problem defined before, the origin node needs to reach the intended receivers only, but can use a path via other nodes by relaying. Thus, the problem is to find an optimal interconnect from a transmitter (root vertex) to a given set of vertices (terminals) and under a certain objective function. This is generally referred to as the directed Steiner tree problem or the Steiner arborescence problem. The complexity of the directed Steiner tree problem is shown to be NP-hard [17]. In this specific case, it is the Steiner arborescence problem in hypergraphs [18]. For an increasing number of nodes, the number of options to reach all terminals grows rapidly. If there exists a link between each node and every distinct receiver set, the number of links in the hypergraph equals $N \cdot (2^{N-1} - 1)$. The properties of a scenario can be used to reduce the graph such that only the vertices and edges that are likely to contribute to the result will be included. This is the *graph reduction* step.

Once a reduced graph is obtained, a sophisticated method is needed to find a minimum Steiner tree for which several heuristics exist in the literature. A different problem arises when we aim at maximizing the number of intended receivers reached before the timeout, instead of minimizing the overall transmission duration. When taking a delay bound into account, finding a minimum Steiner tree is called the *constrained Steiner tree problem* [19]. Moreover, since we consider directional antennas, multiple outgoing edges from the same node should be considered subsequently, as a node can only transmit or receive one message at a time. In this case, the delay depends on the order at which edges are taken and thus scheduling comes into play. Eventually, we will find

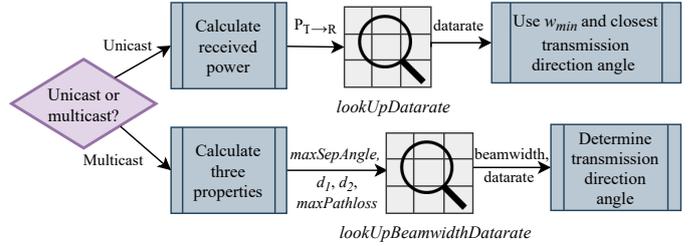


Fig. 2. Schematic summary of the link assessment step.

what we call a *transmission tree* which specifies the route taken and the corresponding schedule of the transmissions.

VI. STEP-1: LINK ASSESSMENT

In this step, summarized in Fig. 2, for a unicast transmission the transmitter will always use the narrowest beamwidth w_{\min} to maximize the data rate. Furthermore, the transmission direction angle will be the one closest to the angle between transmitter and receiver, as calculated from their positions (e.g., acquired from sub-6GHz beacons). In the rest of the paper, we assume perfect alignment for unicast transmissions. The data rate r^* to be used is:

$$\max_{r \in \mathcal{D}} \{r \mid P_T + G_0(w_{\min}) + G_0(w_{\min}) - \Gamma_{T \rightarrow R}(d_{T \rightarrow R}) \geq \sigma(r)\}.$$

The value for r^* will be the output of a one-dimensional sorted look-up table listing the minimum received power needed for all data rates as its entries. This lookup table is referred to as `lookUpDatarate` further on. In case of a multicast transmission, we propose to use three properties of the receiver set to select the beamwidth and data rate. These properties are chosen such that they represent the topology of the receiver set best while minimizing the computational costs. Using these properties, we will create a three-dimensional lookup table, that gives the beamwidth and data rate to use for the receiver set of interest, which will be called `lookUpBeamwidthDatarate`.

The first property is the maximum separation angle (*maxSepAngle*) of all pairs of receivers as seen from the transmitter. For example, consider the situation of Fig.3, which illustrates the coverage of the optimal beam from transmitter 0 to receivers 1, 2, 3 and 4. The pair with maximum separation angle as seen from the node 0 consists of receiver 1 and 4.

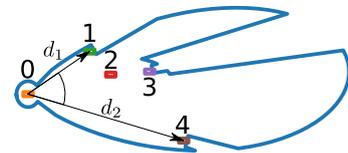


Fig. 3. Situation with transmitter 0 and receivers 1, 2, 3 and 4, along with the coverage of the optimal beam for 0 to reach all receivers.

The second property incorporates the distances between the transmitter and each receiver of the aforementioned pair, displayed as d_1 and d_2 in Fig.3. Incorporating both distances ensures that the varying horizontal coverage due to the Gaussian shape is considered on both sides of the beam.

The last property is the maximum path loss ($maxPathloss$) from the transmitter to any of the receivers in the set. The receiver with the highest loss needs to be considered in beam configuration as otherwise the beam may not cover that receiver, i.e., its received power is too low to decode the signal successfully. Note that the pair of maximum separation may not contain the receiver with the highest path loss. For instance, in Fig.3, the receiver with the minimum received power is node 3 due to the obstacle formed by node 2. Also, the receivers in the pair of maximum separation might be blocked by obstacles causing that the beam is deformed and just incorporating the distance to the transmitter is not sufficient for determining whether the beam covers both receivers.

We generate the lookup table by creating multiple topologies with four nodes; one transmitter and three receivers. We place the receivers in such a way that all possible values (with a certain granularity) for the aforementioned properties are evaluated. For the results presented later in this paper, the parameter $maxSepAngle$ ranges from 0 to 180°, with a step size of 3°. Parameter d_1/d_2 ranges from 0 to 340 m, in steps of 5 m. Lastly, $maxPathloss$ ranges from 0 to 126 dB, evaluated using Eq. 1 with $d_{T \rightarrow R}$ in steps of 1 m. During the link assessment step, for a receiver set of interest, the transmitter calculates the three properties, and rounds towards the granularity used. Next, the transmitter retrieves the beamwidth and datarate to use from the table using these properties as input. The transmission direction angle is determined by rotating the beam such that one receiver is just within it, creating the largest coverage for the other receivers.

VII. STEP-2: GRAPH REDUCTION

Using the *link assessment* step, all possible unicast and multicast links can be identified. However, including all links might lead to a large graph with many irrelevant links. Therefore, as a first step, we will generate a graph with only those nodes that have a lower path loss to any intended receiver than the origin node. This is done as it is unlikely that a better path exists via a node that has a higher path loss, since the total duration of transmissions needed to reach the intended receiver will likely be higher. The nodes that have a lower path loss to any intended receiver than the origin node are called *multipoint relays* (MPRs), as in OLSR [20]. Furthermore, the origin node and all intended receivers are MPRs themselves.

In addition, we introduce the notion of *MPR destinations* (MPRD). The MPRDs of a node are the intended receivers which can be reached by using this node as MPR. Furthermore, intended receivers are MPRDs for themselves and the origin node has all intended receivers as its MPRDs. With the MPRs, we identify the nodes that will be in the graph. Now, let us introduce how we identify the links that will be in the reduced graph. Only if certain conditions apply, a link between an MPRs and a (set of) MPR(s) is evaluated by applying the link assessment step. In this way, fewer receiver sets will need to be evaluated and also fewer links will appear in the hypergraph. The conditions for evaluating a link are different for a unicast and a multicast link as explained next.

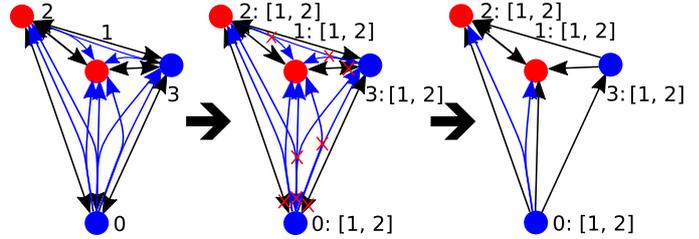


Fig. 4. Scenario displaying all possible links (left), deletion of irrelevant links by graph reduction (middle) and the resulting reduced graph (right).

A unicast link between two MPRs, specifically from μ_T to μ_R , is evaluated only if they have an MPRD in common. As the origin node has all intended receivers as MPRDs and the intended receivers are an MPRD for themselves, links between them will be allowed. Furthermore, relaying nodes might form a path with lower cost towards an intended receiver in this way. However, an additional condition is that the path loss between μ_R and at least one of the shared MPRDs should be lower than or equal to the path loss between μ_T and this specific MPRD. If this does not hold, μ_R will never have a lower transmission duration to one of the MPRDs of μ_T (denoted by $\Delta(\mu_T)$), and thus this link can be neglected. Lastly, the path loss from μ_T to μ_R should be smaller than or equal to the path loss from μ_T to the specific shared MPRD. Otherwise, it would be favorable to let μ_T directly transmit to the specific MPRD. The condition for a link from μ_T to μ_R is formally given by:

$$(\exists \Delta \in \Delta(\mu_T) \mid \Delta \in \Delta(\mu_R) \quad \wedge \quad \Gamma_{\mu_R \rightarrow \Delta} \leq \Gamma_{\mu_T \rightarrow \Delta} \\ \wedge \quad \Gamma_{\mu_T \rightarrow \mu_R} \leq \Gamma_{\mu_T \rightarrow \Delta}).$$

Next, we consider a multicast link from μ_T to a set of MPRs, \mathfrak{M}_R . If an MPR in \mathfrak{M}_R is not in the MPRD set of μ_T , it should have a **unique** MPRD among all MPRDs of the receivers in this link, that is in the MPRD set of μ_T . The idea behind this is that every MPR should have an MPRD that is not handled by other receivers of the multicast link. The condition for a multicast link from μ_T to \mathfrak{M}_R is formally:

$$(\mu_{R_1} \in \Delta(\mu_T) \vee (\exists \Delta \in \Delta(\mu_{R_1}) \mid \Delta \in \Delta(\mu_T) \wedge \Delta \notin \Delta(\mu_{R_2}), \\ \forall \mu_{R_2} \in \mathfrak{M}_R, \mu_{R_2} \neq \mu_{R_1}), \forall \mu_{R_1} \in \mathfrak{M}_R).$$

For example, on the left in Fig. 4, we show which links may exist after the link assessment step for a toy scenario with three nodes. In this case, node 0 is the origin node and it has nodes 1 and 2 as intended receivers for its message. By applying the conditions, in the middle we show which links are removed and the final reduced graph is shown on the right.

VIII. STEP-3: TRANSMISSION TREE GENERATION

Now that the reduced graph has been set up, a tree should be created which determines the transmitter and receiver sets to choose to reach the intended receivers. This tree will be called the *transmission tree* \mathcal{T} from now on. For this, we will modify the heuristic for the constrained Steiner tree problem of [19] to support directional antennas and multicast transmissions.

In short, the heuristic from [19] works as follows. First, the shortest paths between all pairs of nodes in the graph $\mathcal{G}(\mathcal{MPR}, \mathcal{E})$ are calculated, provided that the delay of this shortest path is less than the delay bound O . The shortest path includes which nodes need to transmit to reach R from T and is given by $\mathcal{P}(T \rightarrow R)$. The cost of the shortest path from T to R is given by $\mathcal{P}_C(T \rightarrow R)$, which is equal to the sum of edge costs (transmission durations) on this path. After the shortest paths are determined, a tree \mathcal{T} will be built greedily by determining which nodes to visit. In the beginning, only one node (the origin node) is part of the visited nodes \mathcal{V} . The edges and vertices of the shortest path from a visited node to a terminal (intended receiver) for which the cost function f_{cost} has the lowest value are added to the tree. In this way, the number of visited nodes increases, whereas the number of nodes still to find decreases step by step. Adding new paths to the tree is done until all terminal nodes are covered.

Two cost functions are provided in [19]. The output of both cost functions is infinite if the sum of the path delay from the origin node to the visited node and the additional delay of the shortest path is higher than the delay bound. However, in case of directional antennas, the order of transmissions by the same node is important. Hence, the order in which transmissions are added to the tree sets the schedule that has to be followed. We will now present our modifications that are specific for these type of systems. We will refer to the resulting algorithm as the *Directional constrained Steiner tree (DCST)* heuristic. Alg.1 presents the pseudocode of *DCST* for **unicast** transmissions.

Several cost functions can be used to determine the order in which intended receivers will be handled and by which visited node. Instead of the two presented in [19], we introduce a cost function that fits better to the systems with directional antennas as it exploits simultaneous transmissions. Namely, instead of the path delay from a transmitter to a receiver, the delay of that node in the transmission tree that is set up so far is used. The cumulative delay of node n in the transmission tree is given by $D_{\mathcal{T}}(n)$ and changes while building the tree. For a specific edge to be included in the tree, only the delay of the nodes participating in the transmission which is represented by this edge is updated. Specifically, the updated delay is equal to the current delay of the transmitter, summed with the duration of the transmission. The updated delays for a transmission from T to R having transmission tree \mathcal{T} are given by:

$$D'_{\mathcal{T}}(T) = D'_{\mathcal{T}}(R) = D_{\mathcal{T}}(T) + E_C(T \rightarrow R). \quad (23)$$

Next to that, the cost function ensures that the tree grows first towards the direction where more MPRDs can be handled, such that the chance is higher that more receivers will be reached within the timeout. This is done by dividing the updated delay by the number of MPRDs of the receiver that still need to be handled after this transmission, $|\mathcal{MPRD}_{s_r}(R)|$. The cost function is then given by:

$$f_{cost} = \begin{cases} \frac{D'_{\mathcal{T}}(R)}{|\mathcal{MPRD}_{s_r}(R)|+1} & \text{if } D'_{\mathcal{T}}(R) \leq O \\ \infty & \text{otherwise,} \end{cases} \quad (24)$$

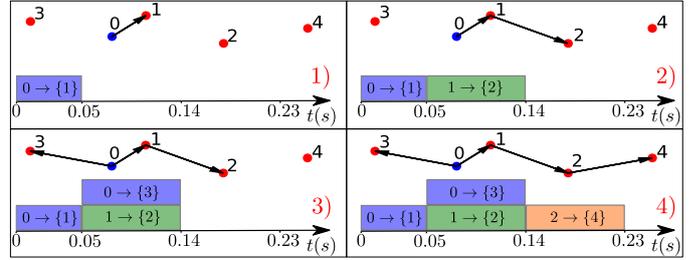


Fig. 5. Steps of the Directional constrained Steiner tree algorithm.

where the addition of 1 in the denominator is to eliminate division by zero. As an example, Fig. 5 shows how the transmission tree and corresponding schedule is generated for a scenario in which origin node 0 wants to reach nodes 1, 2, 3 and 4.

The shortest path to a single node as calculated in lines 1-3 of Alg.1 will usually not include multicast transmissions, since the cost of a multicast is at least as high as a unicast link due to the unicast having a higher data rate. To enable multicast transmissions, their benefit should be assessed in a different way. A multicast transmission is only favorable when reaching multiple receivers in one transmission results in lower delay than reaching them via unicast transmissions. To address this in the most accurate way, a condition we use is that a multicast link should have a lower cost (transmission duration) than the cost of the minimum spanning tree formed by unicast links between the transmitter and receivers of the multicast link. We refer to the multicasts for which this requirement holds as *relevant multicasts*. We introduce the following three methods to evaluate whether a multicast should be included in the transmission tree: *stepwise*, *average cost* and *post-processing*.

1) *Stepwise*: At each step in Alg. 1 (lines 9-25), first the unicast choice will be determined. Afterwards, the algorithm evaluates if there is a relevant multicast transmission that includes this unicast receiver. If there exists one, it will be used instead of the unicast link. If there are multiple relevant multicasts available, the algorithm chooses one that covers the most receivers. When there are still multiple available, the path with the lowest cost (using f_{DM}) is chosen. For the value of $|\mathcal{MPRD}_{s_r}(R)|$ in (24), only the distinct MPRDs of the complete receiver set are counted. For the remainder of the steps in the algorithm, multicast links that include at least one already visited node are not considered further on.

2) *Average cost*: Instead of only single nodes, the receiver sets of relevant multicasts are also included in the set of nodes still to visit (\mathcal{U}). The algorithm operates in the same way except that to let multicasts compete with unicasts, inside the cost function the duration of a multicast transmission is divided by the number of receivers addressed. In this way, an average cost is used, similar to the average broadcast time used in [6]. If a multicast transmission is included in this way, all links that include one of its receivers are now excluded from \mathcal{U} . This rather simple approach ensures that multicasts are favored less often if they have a relatively high cost compared to a unicast.

Algorithm 1: Directional constrained Steiner tree

Input : $\mathcal{G}(\mathcal{MPR}, \mathcal{E}), n_{m_j}, \mathcal{R}_{m_j}, O$ **Output:** Transmission tree \mathcal{T} .

```
1 foreach  $T, R \in \mathcal{MPR}$  do
2   Compute the shortest path within the timeout.
3    $\mathcal{P}(T \rightarrow R) \leftarrow$  Nodes along shortest path.
4    $E_C(T \rightarrow R) \leftarrow$  Transmission duration.
5  $\mathcal{V} \leftarrow \{n_{m_j}\}, \mathcal{U} \leftarrow \mathcal{R}_{m_j}, \mathcal{T} \leftarrow \emptyset$ 
6 foreach  $n \in \mathcal{N}$  do
7    $D_{\mathcal{T}}(n) \leftarrow 0$ 
8 while  $\mathcal{U} \neq \emptyset$  do
9    $minCost = \infty$ 
10  foreach  $T \in \mathcal{V}$  do
11    foreach  $R \in \mathcal{U}$  do
12      if  $\mathcal{P}(T \rightarrow R)$  exists then
13        foreach  $n_1, n_2 \in \mathcal{P}(T \rightarrow R)$  do
14           $D'_{\mathcal{T}}(n_1), D'_{\mathcal{T}}(n_2) \leftarrow$ 
15             $D_{\mathcal{T}}(n_1) + E_C(n_1 \rightarrow n_2)$ 
16          if  $f_{cost} < minCost$  then
17             $nextPath \leftarrow \mathcal{P}(T \rightarrow R)$ 
18             $minCost \leftarrow f_{cost}$ 
19            Save updated delays  $D'_{\mathcal{T}}$  for nodes
20            in  $nextPath$ .
21  if  $minCost == \infty$  then /* No paths */
22    return  $\mathcal{T}$ 
23   $\mathcal{T} \leftarrow \mathcal{T} + nextPath$ 
24   $\mathcal{V} \leftarrow \mathcal{V} \cup \{n \mid n \in nextPath\}$ 
25   $\mathcal{U} \leftarrow \mathcal{U} \setminus \{R \mid R \in nextPath\}$ 
26  foreach  $n \in nextPath$  do
27     $D_{\mathcal{T}}(n) \leftarrow D'_{\mathcal{T}}(n)$ 
28 return  $\mathcal{T}$ 
```

However, with this approach the full potential of multicasts is not exploited, since even if the average cost is not lower than a unicast, using a multicast transmission might be beneficial.

3) *Post-processing*: Here, the DCST heuristic is first evaluated with unicast transmissions only. Afterwards, unicast transmissions can be replaced by multicast transmissions if the multicast receivers are used in the unicast tree. Yet, it should be evaluated first whether at least one receiver receives the message earlier than in the unicast alternative. Only then, the multicast is potentially beneficial for the tree. Besides, transmissions that get delayed due to the multicast transmission should not violate the timeout. This is checked by inserting the multicast transmission in the already created unicast schedule. If a transmitter can send multiple multicast transmissions for which these two conditions hold, the one with the largest receiver set is chosen. When there are still multiple available, the one with the lowest cost is chosen.

IX. PERFORMANCE EVALUATION

Now, we assess the performance of our proposals using a system-level Python simulator. While our key goal is to address the research questions in Sec.I, we also investigate the decrease in computation time and performance introduced by our proposals over the optimal or exhaustive approach.

Performance of the heuristic

First, we evaluate the link assessment step by generating 5,000 scenarios with a receiver set of two to four nodes, placed randomly in the vicinity of the transmitter. For each scenario, we derive the beamwidth, datarate and transmission angle to use using the lookup tables with the granularity as specified in Sec. VI. We use a Python implementation, run on a laptop with an Intel i7-6700HQ processor with 16 GB RAM. Our analysis shows that the computation of the link assessment step takes on average 0.84 ms, which is only 1.5% of the time needed by an exhaustive search. To evaluate the performance degradation of using this step compared to the optimal solution, we derived topologies from snapshots of the five-lane Highway 101, captured in June 2005, as presented in [21]. This dataset includes, amongst others, the positions of vehicles at every 100 ms during three 15 min intervals. The considered segment is 640 m long in total, but we consider only the first part (starting North), such that a fixed number of vehicles is included in a scenario. From 1000 snapshots we took five nodes to generate a scenario in which the northernmost node was assigned the transmitter, which needs to deliver a mmWave message with a timeout of 150 ms to three of the four remaining nodes. We derive the optimal solution by solving the MILP using Gurobi, which uses the links determined by an exhaustive search. From the links selected by the optimal solution, 16% of the links are not detected by the link assessment step, which gives some insight in the performance degradation of this step.

Next, we evaluate the graph reduction step for the same scenario. Our analysis shows that 72% of the links determined by an exhaustive search are excluded, reflecting the degree of saving in computation overhead. Meanwhile, 95% of the links used by the optimal solution are still included despite the graph reduction step, which implies the accuracy of our reduction approach. To examine the effects of both steps, we compare the computation time of the transmission tree generation step when using all possible links (exhaustively determined) to when using only those resulting from the link assessment and graph reduction step. We evaluate 1000 snapshots with either five nodes, of which three were intended receivers, or seven nodes, of which four were intended receivers. The average computation time for the transmission generation tree step (implemented in Python on the same hardware) using the method *stepwise* with reduced links is 0.96 ms and 3.1 ms, in the smaller and larger scenario, respectively. This is 32% and 11% of the time needed when using all links. For the method *average cost* the numbers are 1.1 ms (31%) and 4.1 ms (11%), respectively. As *post-processing* is more complex, the effect is slightly less. The computation time is 1.7 ms and 4.7 ms for the smaller and larger scenario, respectively, which is 44%

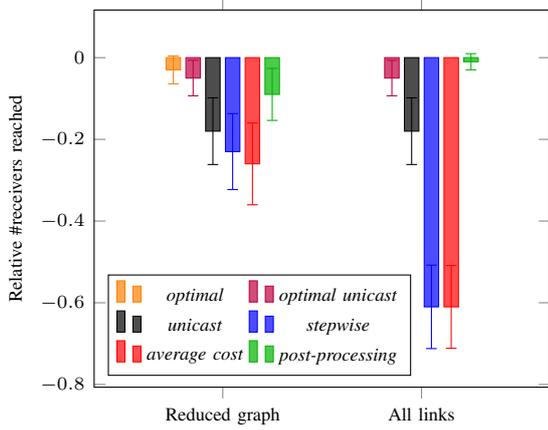


Fig. 6. Average relative number of intended receivers reached before the timeout as compared to *optimal*, either using the reduced graph or all links.

and 17% of the time needed using all links. Thus, indeed these steps facilitate a significant reduction in the computation time, which has a larger influence on more complicated scenarios.

Let us now examine the performance of the transmission tree generation step either with reduced graph or all links as input, as compared to the optimal solution. We use 100 snapshots with 20 s separation from [21] with 5 nodes. We assign the northernmost node as origin node. From the remaining nodes, we select three nodes randomly as intended receivers. The timeout is set to 100 ms. For each scenario, we derive the optimal solution from the Gurobi solver (referred to as *optimal*) and our heuristic methods. We also investigate the performance of the optimal solution when only unicasts may be used, by allowing only unicast links while pre-processing the scenario. We refer to this method as *optimal unicast*.

Fig. 6 shows the average number of receivers reached in time (including 95% confidence intervals) as compared to *optimal* with all links as input. Using *unicast* with reduced graph, in 81% of the scenarios all intended receivers were reached in time. First, looking at *optimal unicast*, the difference as compared to optimal is limited for this small scenario, both using the reduced graph and all links (remember that *optimal* with all links is the baseline). Moreover, *unicast* outperforms *stepwise* and *average cost* in this small scenario. The difference is even more clear when all links are used, as this negatively influences *stepwise* and *average cost*. This is due to the fact that if more multicast links are available, it is more likely that these methods choose to use them, while it is not guaranteed that those eventually lead to performance improvement. Thus, the filtering effect of the graph reduction step is useful. The method *post-processing* outperforms *unicast*. Using all links, it comes very close to the optimal solution. However, this method is still sub-optimal as it is restricted to use the nodes leveraged in the unicast tree.

Multicast usage

To understand the benefit of using multicasts, we compare the performance of the three options for selecting multicast links using *DCST* against that of *DCST* with unicast links only,

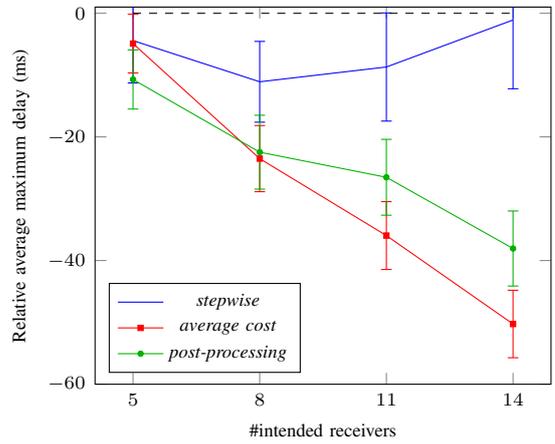


Fig. 7. Average maximum delay as compared to *unicast* for different number of intended receivers in a geocast.

referred to as *unicast*. For each method (*unicast*, *stepwise*, *average cost* and *post-processing*), we generate the reduced graph (using link assessment) and create a transmission tree using Alg.1. We use 15 nodes for all snapshots, in which the origin node is the northernmost node. Either 5, 8, 11 or 14 nodes (all but the origin node) are assigned as intended receiver. To compare the scenarios, we set the timeout arbitrarily high (10 s), such that all multicast methods reach the same number of intended receivers as *unicast*. As metric, we report the delay of the last receiver and show it relative to *unicast*.

Fig. 7 depicts the average maximum delay including 95% confidence intervals for these three methods for increasing number of intended receivers. From the figure, we observe that the method *stepwise* does not show a significant improvement as compared to *unicast* and its performance varies. But, for *average cost*, the relative delay reduction ranges from 2.5% to 19.3% with increasing number of intended receivers. For *post-processing*, the delay decrease ranges from 5.4% to 14.6%. This shows that these methods provide a consistent improvement as compared to *unicast*. However, these two methods differ in nature, as follows from the use of relaying nodes. The percentage of intended receivers among all receivers reached for 5 to 11 intended receivers ranges from 73% to 88% for *average cost*, whereas this is 89% to 98% for *post-processing*. This means that the latter approach does not yet fully exploit the relaying nodes, as it closely follows the unicast tree. On the other hand, not all non-intended receivers reached using *average cost* might contribute to relaying.

Lastly, solving the problem by the mathematical solver gives insight into multicast usage under various scenarios. Using wider beams, multiple receivers can be reached at once using a high data rate, if these receivers are located at limited vertical distance from the transmitter. It seems therefore that multicasts are more likely to be used in scenarios with multiple highway lanes, in which vehicles are more horizontally spread. To examine whether this is true, we consider two scenarios; one with five nodes and three randomly assigned intended receivers

TABLE II
MULTICAST USAGE FOR SCENARIOS WITH A DIFFERENT NUMBER OF
HIGHWAY LANES AND NODES.

#Lanes	#Nodes	Vehicles per km per lane	Fraction multicast (%)
1	5	199.8	9.3
	7	181.8	3.5
3	5	136.4	57.3
	7	144.1	55.9
5	5	107.5	43.8
	7	111.2	56.9

and one with seven nodes and four intended receivers. We evaluate 100 snapshots from the dataset of [21] with 20 s separation in which either vehicles of the first lane, the first three lanes or all five are included. We assign the northernmost node as the origin node and select the intended receivers randomly from the remaining nodes. We provide the reduced graph as input to the Gurobi solver, to limit its computation time. Table II shows the fraction of receivers that are reached via a multicast for each scenario. As the traffic varies per lane, the average vehicular density per lane differs. For reference, we provide this value in the third column. We observe that multicasts are less frequently used in a scenario with only one lane. For three lanes, the fraction is higher than for five lanes in the smaller scenario, and it is similar in the larger scenario, which can be caused by the higher vehicular density per lane. Also, the limited width of the antenna beams might lead to this. However, we observe that multicasts are heavily used, since up to 57% of the receivers are reached via multicast.

X. CONCLUSIONS

While mmWave bands with plenty of spectrum present many opportunities for vehicular communications, it is not straightforward how these bands should be exploited to meet the delivery deadlines imposed by the safety-critical applications. In particular, the optimal beamwidth and data rate for transmission, whether a node should transmit in multicast or unicast mode, or whether it should use relays are some open questions. To address these questions, we provided a linear program and a heuristic method for routing and scheduling mmWave messages considering multicast, relaying, and spatial sharing. Different from the literature, we consider a non-time-slotted system and a realistic antenna model. Due to the high computation-cost of the optimal solution, we have developed a 3-step heuristic consisting of link-assessment, graph reduction, and transmission tree generation. We have shown a significant reduction in computational complexity, while the heuristic still shows reasonable performance as compared to the optimal solution. Three methods for selecting multicasts are proposed, of which two show consistent improvement as compared to using unicasts only. In relatively small scenarios for which an optimal solution could be found, our results show that using multicast is beneficial especially for multi-lane highways, arising from the fact that up to 57% of the receivers are reached via a multicast.

Future work includes a realistic implementation to schedule **multiple** mmWave geocasts using sub-6GHz beacons, for

which we give a basis in [15]. Also, it is interesting to examine the influence of interference on multicast usage to create an interference-aware routing and scheduling protocol.

REFERENCES

- [1] C. Campolo, A. Molinaro, and R. Scopigno, "From today's VANETs to tomorrow's planning and the bets for the day after," *Vehicular Communications*, vol. 2, no. 3, pp. 158–171, 2015.
- [2] B. Coll-Perales, J. Gozalvez, and M. Gruteser, "Sub-6GHz Assisted MAC for Millimeter Wave Vehicular Communications," *IEEE Comms. Magazine*, vol. 57, no. 3, pp. 125–131, 2019.
- [3] B. Coll-Perales, M. Gruteser, and J. Gozalvez, "Evaluation of ieee 802.11ad for mmwave v2v communications," in *2018 IEEE Wireless Comms. and Nw. Conf. Workshops (WCNCW)*, 2018, pp. 290–295.
- [4] I. Mavromatis, A. Tassi, R. J. Piechocki, and A. Nix, "Beam alignment for millimetre wave links with motion prediction of Autonomous Vehicles," in *Antennas, Propagation RF Technology for Transport and Autonomous Platforms 2017*, 2017, pp. 1–8.
- [5] G. H. Sim, M. Mousavi, L. Wang, A. Klein, and M. Hollick, "Joint Relaying and Spatial Sharing Multicast Scheduling for mmWave Networks," in *IEEE International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 2020, pp. 127–136.
- [6] Y. Chang, Q. Liu, X. Jia, and K. Zhou, "Routing and transmission scheduling for minimizing broadcast delay in multirate wireless mesh networks using directional antennas," *Wirel. Commun. Mob. Comput.*, vol. 15, no. 1, p. 87–99, Jan. 2015.
- [7] H. Park, S. Park, T. Song, and S. Pack, "An Incremental Multicast Grouping Scheme for mmWave Networks with Directional Antennas," *IEEE Communications Letters*, vol. 17, no. 3, pp. 616–619, 2013.
- [8] D. Gong, Y. Yang, and H. Li, "Link-layer multicast in smart antenna based 802.11n wireless lans," in *2013 IEEE 10th International Conference on Mobile Ad-Hoc and Sensor Systems*, 2013, pp. 484–492.
- [9] G. Mendler and G. Heijenk, "On the potential of multicast in millimeter wave vehicular communications," in *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*, 2021, pp. 1–7.
- [10] Y. Niu, L. Yu, Y. Li, Z. Zhong, and B. Ai, "Device-to-Device Communications Enabled Multicast Scheduling for mmWave Small Cells Using Multi-Level Codebooks," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 3, pp. 2724–2738, 2019.
- [11] A. Taya, T. Nishio, M. Morikura, and K. Yamamoto, "Concurrent Transmission Scheduling for Perceptual Data Sharing in mmWave Vehicular Networks," *CoRR*, vol. abs/1903.09402, 2019.
- [12] A. Yamamoto, K. Ogawa, T. Horimatsu, A. Kato, and M. Fujise, "Path-Loss Prediction Models for Intervehicle Communication at 60 GHz," *IEEE Trans. on Vehicular Technology*, vol. 57, no. 1, pp. 65–78, 2008.
- [13] I. Toyoda and T. Seki, "Antenna Model and Its Application to System Design in the Millimeter-wave Wireless Personal Area Networks Standard," 2011.
- [14] IEEE Std 802.11-2016 (Revision of IEEE Std 802.11-2012), "Telecommunications and information exchange between systems Local and metropolitan area networks — Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE, Standard, 2016.
- [15] T. Havinga, "MmWave Beam Control for Geocasting in Vehicular Networks," Master's thesis, University of Twente, 2021.
- [16] *Gurobi Optimizer Reference Manual*, Gurobi Optimization LLC, 2021. [Online]. Available: <http://www.gurobi.com>
- [17] M. Jones, D. Lokshtanov, M. S. Ramanujan, S. Saurabh, and O. Suchý, "Parameterized Complexity of Directed Steiner Tree on Sparse Graphs," *CoRR*, vol. abs/1210.0260, 2012.
- [18] M. Brazil and M. Zachariasen, *Optimal interconnection trees in the plane: theory, algorithms and applications*, ser. Algorithms and Combinatorics. Springer, 2015.
- [19] V. P. Kompella, J. C. Pasquale, and G. C. Polyzos, "Multicast routing for multimedia communication," *IEEE/ACM Transactions on Networking*, vol. 1, no. 3, pp. 286–292, 1993.
- [20] T. Clausen and P. Jacquet, "RFC3626: Optimized Link State Routing Protocol (OLSR)," USA, 2003.
- [21] "Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data," U.S. Department of Transportation Federal Highway Administration, ITS DataHub, 2016.