Feedback-Based Hidden-Terminal Mitigation for Distributed Scheduling in Cellular V2X

Philip Wendland, Guenter Schaefer
Telematics / Computer Networks Research Group
Technische Universität Ilmenau, Germany
{philip.wendland, guenter.schaefer}@tu-ilmenau.de

***Abstract*—Hidden terminal situations, in which two transmitters may interfere at a single receiver, but the transmitters cannot notice each other's interference, can have severe effects. Traditional access technologies use RTS/CTS to remedy this. This is infeasible in vehicular scenarios, because the broadcast communication would lead to many RTS/CTS dialogues and every potential receiver might not be known to the transmitter. Hence, there are no mitigations in IEEE 802.11p, ITS-G5 or LTE-V2X mode 4.**

In this paper, we describe a method to leverage the periodicity of the semi-persistent scheduling in LTE-V2X mode 4 to proactively consider and reactively detect hidden terminal problems. Our approach uses the padding bits of periodic messages, e.g., CAMs, to indicate which resource has been successfully decoded by the terminal. This information can be used by terminals to detect whether its own transmissions have not been decodable by others, and react to this situation by selecting new resources. Additionally, we present an extended candidate resource selection procedure that considers these acknowledgements to extend the station's sensing range. Due to semi-persistent scheduling, already acknowledged resources are likely to result in hidden terminal situations in future. Our approach significantly outperforms standard LTE-V2X mode 4 in our evaluation.

***Index Terms*—connected vehicles, vehicle-to-everything, radio access technologies**

I. INTRODUCTION

In future, vehicles are envisaged to periodically broadcast status update messages (called *cooperative awareness message* (CAM) in Europe), which contain the current position, speed, heading, etc., likely reducing the risk of collisions or enabling cooperative driving maneuvers. There are two competing access technologies: IEEE 80211.p (ITS-G5 in Europe) and *long term evolution vehicle-to-everything* (LTE-V2X). The hidden terminal problem, explained in Figure 1, is generally ignored in all *vehicle-to-vehicle* (V2V) communication technologies. While *request-to-send/clear-to-send* (RTS/CTS) is usually used as mitigation for unicast traffic, this option is not feasible for broadcast traffic because of scaling and other issues, which will be described in more detail in Section II.

In this paper, we present a heuristic approach to tackle this problem for LTE-V2X mode 4. We leverage the *semi-persistent scheduling* (SPS) of LTE-V2X to mitigate this issue: as resources are being periodically reused, it is still worth reacting to hidden terminal problems even if the first collision already occurred. As a terminal cannot detect hidden terminal effects by itself, we rely on explicit feedback by other terminals. We make dual use of this feedback and use it as acknowledgement mechanism that might trigger terminals to select new resources ahead of time, and as input for the candidate resource selection procedure to prevent hidden terminal effects from occurring in the first place. Specifically, we make the following contributions:

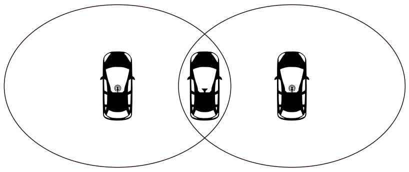

Fig. 1: Visualization of the hidden terminal effect. The receiving vehicle (middle) is unable to decode packets from either transmitting vehicles. Still, the transmitting vehicles (left and right) cannot detect each others interference.

1) We present a heuristic feedback mechanism that is used to detect and react to radio frame collisions.
2) We propose to make secondary use of this feedback to extend the sensing range of terminals and reduce the number of hidden terminal situations.
3) We evaluate our proposed modifications using a large-scale system-level simulation along with application-oriented metrics.

II. RELATED WORK

Due to the repeated reuse of radio resources as dictated by SPS in LTE-V2X mode 4, any issues that are a result of incorrect resource allocation are very likely to reoccur [1]. In general, several publications target the evaluation and improvement of LTE-V2X mode 4. The focus is often to reduce the number of repeated collisions by introducing more entropy into the resource allocations [1]–[3]. These schemes either do not or only marginally improve the *packet delivery ratio* (PDR). While our approach could be combined with these ideas, the reactive component also achieves a reduction of the number of repeating collisions. Moreover, we target a root cause of collisions: the hidden terminal problem. For broadcast-type communication as in most V2V scenarios, there are typically no mitigations for the hidden terminal problem due to various incompatibilities of RTS/CTS: First, when

---

This research has been supported by the Free State of Thuringia and the European Social Fund under grant 2016 FGR 0039.

addressing a packet to multiple receivers, RTS/CTS would need to be performed with all of them. Second, the assumption of RTS/CTS is that the RTS/CTS messages are significantly smaller than the actual data that is to be transmitted, which might not be true in V2V scenarios. Third, the transmitter might not know all potential receivers beforehand.

There are few publications that specifically address the hidden terminal problem for cellular V2X. In [4], a distributed, geographically assisted scheduling scheme is proposed. The selection of resources is based on the geographical position and ordering of the vehicle on the road. This protocol is based on the assumption that interference is only subject to distance, which might not be true in reality. The same authors previously presented a similar idea that was optimized for intersection scenarios [5]. This approach divides the radio resource in blocks and assigns a block to each street arriving at the intersection. This might be problematic in heterogeneous scenarios, e.g., a small road crossing a large one. Additionally, significant interference might arise from other intersections.

In this paper, we present a heuristic approach to limit hidden terminal effects that leverages SPS. It has no or low overhead, is designed to be flexible, and does not rely on central infrastructure or other additional technologies.

III. LTE-V2X MODE 4

LTE-V2X supports vehicular communication for the down-, up-, and sidelink. The down- and uplink mean communication between the base station and terminals, and vice versa. The sidelink denotes direct communication between terminals. There are two different scheduling modes for the sidelink and vehicular use case: mode 3, in which the base station controls the resource allocation, and mode 4, which does not rely on the presence of a base station. Mode 4 denotes the baseline performance of LTE-V2X and might be the only deployed mode. In mode 4, which is the focus of this paper, the vehicles have to coordinate the access to the wireless channel by themselves in a distributed way.

A. Physical Layer

LTE-V2X uses *single-carrier frequency division multiple access* (SC-FDMA) for the sidelink. This means that the channel is subdivided into multiple subchannels (five in Europe [6]) and the transmission is slotted in the time domain [7], requiring a tight clock synchronization between terminals. A transmission is always one subframe, i.e., millisecond, long, but can consist of multiple, adjacent subchannels [8]. Aside from the payload data, which is transmitted in a transport block using the *physical sidelink shared channel* (PSSCH), additional scheduling information (the *sidelink control information* (SCI)) is transmitted using the *physical sidelink control channel* (PSCCH). The SCI contains explicit reservations in case the terminals intends to reuse its radio resource in a fixed time interval (e.g., $100\,\mathrm{ms}$) [7]. Terminals operate in half-duplex fashion, which also applies to different subchannels, so that a terminal is unable to receive messages during the time it transmits itself.

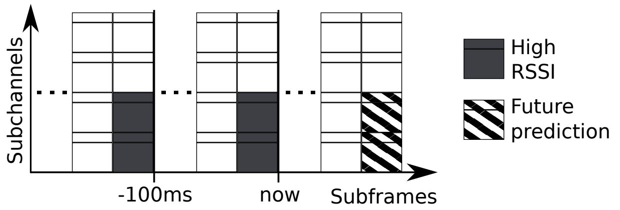

Fig. 2: Visualization of LTE resource grid and RSSI prediction

B. Semi-Persistent Scheduling and Resource Selection

Due to the aforementioned differences at the physical layer, the *media access control* (MAC) protocol for mode 4 is inevitably different to the *carrier sense multiple access/collision avoidance* (CSMA/CA) protocol used in IEEE 802.11p or ITS-G5. In *long term evolution* (LTE), SPS requires terminals to reuse radio resources for a given number of repetitions. This reuse happens without re-evaluating the current channel conditions. The number of repetitions is determined by the reselection counter, which is a random number between 5 and 15, drawn each time a *new* resource has been selected, or at the end of the reuse interval with probability *probResourceKeep*, in which case the same radio resources will be further reused [9].

In the other case, new resources have to be selected. The predictability of other transmitters due to the resource reuse (SPS) is leveraged. The set of possible resources is a two-dimensional matrix across all subchannels (frequency domain) and subframes (time domain), so that the latency requirement of the message ($100\,\mathrm{ms}$ for CAMs) to be transmitted can be fulfilled. From this matrix of resources, $80\,\%$ of the candidate resources with the highest collision risk will be excluded. Afterwards, the resource will be selected randomly from the remaining $20\,\%$. There are three conditions for excluding resources from the set. The first condition is a precaution related to the half-duplex problem. If the terminal has been using a portion of a subframe for transmission in the past, it could not have monitored it or received reservations by other terminals. Therefore, future resources affected by incomplete monitoring due to the half-duplex operation will be excluded from the set of candidate resources. For the remaining resources, the explicit reservations in the SCI can be respected. Whether an explicit reservation will be respected or ignored depends on whether the *reference signals received power* (RSRP) of the corresponding resource in the past is above a certain threshold, which depends on priority of the received frame and the priority of the packet to be transmitted. This RSRP threshold is called $\mathrm{Th}_{a,b}$. If there are many reserved resources and $\mathrm{Th}_{a,b}$ is so lenient that less than $20\,\%$ of the original resources remain in the set, it will be increased by $3\,\mathrm{dB}$ as often as necessary. If less than $80\,\%$ of the resources have been excluded, the RSSI is being averaged over the last ten repetitions, i.e., the last second, and only the resources with the lowest average RSSI will be left in the candidate ($20\,\%$) set [7]. This is visualized in Figure 2.

IV. FEEDBACK-BASED HIDDEN-TERMINAL MITIGATION

There are two different parts to our proposed hidden terminal problem mitigation: First, a mechanism that enables terminals to transmit explicit feedback about the current channel conditions and acknowledge radio resources that have been successfully decoded, subsequently called *acknowledgement feedback*. With this feedback, other terminals can notice if they selected inadequate resources and can react by selecting new resources even if the reselection counter did not yet reach zero. Second, a newly designed candidate resource selector that considers this feedback.

A. Acknowledgement Feedback

The repeated use of radio resources due to SPS leads to repeating collisions [1], but also makes the resource allocation more predictable. Our proposal also leverages the periodic reuse and tries to detect these repeating collisions.

Terminals transmit a matrix that indicates which *radio resources* have previously been successfully decoded. The main difference to traditional acknowledgements is that radio resources instead of packets are acknowledged. This means that there is no need to address the feedback to individual transmitters, and the information can be aggregated and then sent via local broadcast, piggyback to existing messages in the padding bits if necessary. When using five subchannels, five bits suffice to report a single subframe. One hundred bits are enough to report $100 \div 5 = 20$ subframes and so forth. LTE-V2X can use large amounts of padding due to the fixed transmission duration and low granularity of the subchannel allocation scheme. The number of padding bits depends on the message size, the number of required subchannels and the *modulation and coding scheme* (MCS) index, but can even be higher than 500, which is enough to encode the last 100 subframes. If the number of padding bits does not suffice to encode the complete repetition interval, only the latest radio resources will be reported. This gives an incomplete picture, but only from this particular terminal. There are many terminals that periodically transmit with resources that are likely evenly distributed across time, so that it is likely that enough feedback information is available to terminals that select new resources. In some cases, it might be advantageous to use additional subchannels or a different MCS index, but we leave the evaluation of this compromise as future work.

The received acknowledgement feedback is aggregated by every terminal, which then has the number and ratio of reported successful and unsuccessful receptions for each resource. This information can be proactively used when selecting candidate resources, which will be described in the next section. Additionally, it has a reactive component to it: If a terminal remembers which resources it has used in the past, it can assume that there is a correlation between the acknowledgement ratio and the delivery ratio of the transmitted packet. A wireless channel is reciprocal, i.e., if a transmitter can reach a receiver, the receiver will be able to reach the transmitter. In our case, the initial transmitter and receiver, i.e., the acknowledging terminal, will transmit at different points in time, i.e., from slightly different positions. Additionally, acknowledgements could be addressed to a different transmitter, which used the same resource. This problem occurs if the reporting terminal was able to capture a stronger transmission despite interference. However, it is very likely that there are more receivers close to both transmitters, which would report non-acknowledgements back to them. Our approach is a heuristic due to this and the potentially incomplete reports caused by insufficient padding bits.

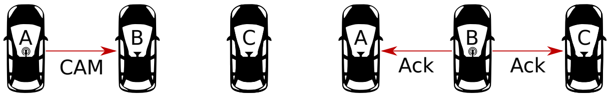

(a) Transmission (b) Acknowledgement

Fig. 3: Visualization of the proactive hidden-terminal mitigation. At first, vehicle A transmits a CAM that can only be received at vehicle B, which acknowledges the successful reception of this particular radio resource. Vehicle C will avoid this resource because it now knows that it will likely be occupied in the future (due to SPS) by a transmitter that it cannot detect itself.

If radio resources used for transmission by a terminal resulted in low acknowledgement ratios, they will react dropping their current resources and selecting new ones ahead of time, i.e., before the reselection counter reached zero. The minimum required acknowledgement ratio, below which terminals have to select new resources, needs to be configured. Additionally, there should be a hysteresis period that allows terminals to drop resources ahead of time only once in a given time interval. In our evaluation, a conservative minimum acknowledgement ratio of 60 % and a hysteresis period of three seconds was empirically determined and configured.

B. Resource Selection Procedure

The idea of the proposed resource selection procedure is that the acknowledgements and periodicity of the transmitter behavior can be used to *extend the sensing range*. As terminals will report acknowledgements from transmitters further away that could be invisible to the terminal that is selecting resources, they should be regarded as occupied, at least with a certain weight. This is illustrated in Figure 3.

The standard-compliant candidate resource selector is already very complex because it has interdependencies between the reasons for exclusion of candidate resources. Additionally, the RSRP thresholds depend on priorities of foreign packets compared to the packet to be transmitted and might need to be adjusted multiple times depending on the current number of reserved resources. While this serves as *quality of service* (QoS) aspect, we think that QoS aspects are better implemented in *distributed congestion control* (DCC) components or other ways that are outside the scope of this paper.

Our proposed candidate resource selector additionally considers acknowledgements, but eliminates the aforementioned interdependencies. For each candidate resource, we build a rating. A higher rating means a lower predicted collision probability. The actual resource will be chosen from the $20\,\%$ of the candidate resources with the highest rating.

Various conditions can decrease the rating of a candidate resource. The rating function is described as pseudo-code in Algorithm 1.

Algorithm 1 Calculate candidate resource rating

```
Require: candidate resource cr
  rating = 1000 − rssi_weight × cr.rssi_factor
  if cr.reservation_present then
    rating -= res_penalty + res_weight × cr.rsrp_factor
  end if
  if cr.ack_present then
    rating -= ack_penalty + ack_weight × cr.ack_ratio
  end if
```

The RSSI and RSRP are reception energies, determined like in standard LTE-V2X. Depending on the transmission requirements, a candidate resource can consist of multiple consecutive subchannels. Only the subchannel with the highest RSSI or RSRP is considered in our approach. The corresponding factors (cr.rssi_factor and cr.rsrp_factor) are calculated by a linear interpolation (in the logarithmic domain) between -110 and $0\,\mathrm{dBm}$, i.e., $(p \div 110\,\mathrm{dBm}) + 1$ when $-110\,\mathrm{dBm} \leq p \leq 0\,\mathrm{dBm}$. The lower bound of $-110\,\mathrm{dBm}$ is close to the theoretical noise floor per subchannel of $-111\,\mathrm{dBm}$, assuming $-174\,\mathrm{dBm/Hz}$ of thermal noise over a subchannel that occupies $2\,\mathrm{MHz}$.

As it can be seen, if no reservation is present, there will be no penalty for reservations, and if no acknowledgement is present, there will be none for acknowledgements. There is a static penalty for the presence of reservations (res_penalty) and acknowledgements (ack_penalty). Additionally, there is a weight factor for the average RSSI (rssi_weight), the RSRP (res_weight) and the acknowledgement ratio (ack_weight). These parameters should add up to one thousand, so that the rating of a candidate resource is between zero and one thousand. In our simulations, we empirically chose $\mathrm{res\_penalty} = 250$, $\mathrm{res\_weight} = 250$, $\mathrm{ack\_penalty} = 125$, $\mathrm{ack\_weight} = 125$, and $\mathrm{rssi\_weight} = 250$. In total, half of the weight is used for explicit reservations, and one fourth each is used for the average RSSI and acknowledgements.

V. EVALUATION

We use a system-level simulation approach in order to validate our approach. The simulation framework uses a combination of the *objective modular network testbed in C++* (OMNeT++), the INET module, Artery [10], Veins [11] and *simulation of urban mobility* (SUMO) [12]. The model for LTE-V2X mode 4 is a custom implementation.

A. Simulation Parameters and Scenario

For this simulation, we assumed the following parameters for LTE-V2X mode 4: A transmission power of $23\,\mathrm{dBm}$, five subchannels per subframe, and a CAM rate of $10\,\mathrm{Hz}$. In practice, the size of CAMs can vary depending in the presence and length of digital signatures or a *lowFrequencyContainer*, but the SPS is not capable of correctly considering varying packet sizes. Therefore, we assumed that every CAM has a size of $200\,\mathrm{Bytes}$, including overhead. We assumed a MCS index of six, which is in between the lower and upper bound (zero to eleven) specified for Europe [6]. This results in a transport capacity of $1032\,\mathrm{Bits}$ when using one subchannel and $2088\,\mathrm{Bits}$ when using two subchannels [7]. Hence, two subchannels are necessary to carry a CAM message. Furthermore, $488\,\mathrm{Bits}$ of padding are necessary to fill the transport blocks, which is sufficient to report acknowledgements for the last 97 subframes.

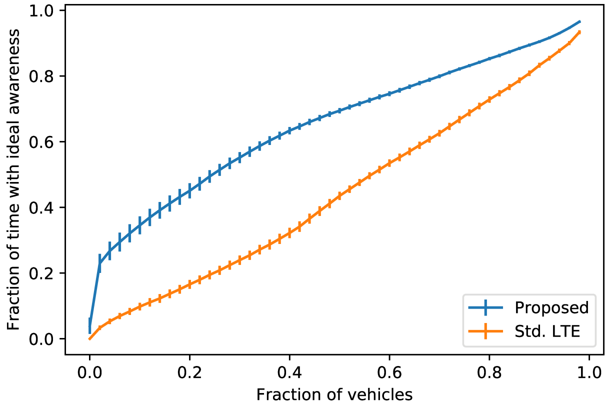

Fig. 4: Quantile function plot: fraction of time with ideal awareness for the range of $0-300$ m. Vertical lines indicate $95\,\%$ confidence intervals. Horizontal connecting lines are for visualization purposes only.

The vehicles use a half-wave dipole antenna in our simulation. We adopted the error probabilities from a study by Huawei and HiSilicon for our receiver model [13]. Unlike in previous publications [1], [2], we used a dual-slope model developed by Cheng et al. [14] as our wireless channel model. It is tailored for $5.9\,\mathrm{GHz}$ and urban V2V scenarios. We use the Luxembourg traffic scenario [15] for SUMO at 08:00, which results in a large traffic load and hence, network load. Due to the large size of it, we only use a part located in the city center, which can be seen in previous publications [1], [2]. The simulation scenario contains around one thousand vehicles, includes small and large roads and simulates traffic lights. We made sure that the transient phase was long enough, performed 32 runs with different seeds for SUMO in order to determine confidence intervals, and excluded statistics from vehicles at the edges of the scenario, which would not be affected by interferers further away. Excluding the transient phase, the simulation scenario is sixty seconds long. We did not enable DCC because it leads to performance decreases or no significant increases [2], [16]. We set *probResourceKeep* to zero, as it results in a large drop of the cooperative awareness [1].

B. Results and Discussion

Table I shows a summary of the simulation results. First, we have to define the *awareness* for a given vehicle under

TABLE I: Simulation Results (0 to $300\,\mathrm{m}$)

	PDR	Mean awareness	Mean percentage of time with ideal awareness	Mean percentage of time with at least 95 % awareness
Proposed	89.21 (±0.05)	99.48 (±0.00)	65.69 (±0.28)	99.87 (±0.01)
Std. LTE	82.41 (±0.08)	98.80 (±0.01)	44.59 (±0.37)	98.62 (±0.03)

consideration. Let A be the set of vehicles that it has received CAMs from. The lifetime of a CAM is one second and older messages will be excluded from the set. Let B be the set of vehicles in the surroundings of a given vehicle for a given range, e.g., zero to $300\,\mathrm{m}$. Then, the *awareness* of the vehicle under consideration is $\alpha = |A \cap B| \div |B|$. In vehicular traffic, accidents are relatively rare, but the consequences are often severe. In such a safety-critical environment, it is desirable to have ideal awareness ($\alpha = 100\,\%$) *all the time*. When only using the mean awareness as metric, potentially dangerous situations can still be hidden in the data. Therefore, we visualize the *fraction of time with ideal awareness* as quantile function plot in Figure 4. It shows the fraction of time with ideal awareness as distribution over all the vehicles, and is also explained in previous publications [1], [2]. With our proposed modifications, the awareness significantly improves with this and all other metrics, which can also be seen in Table I.

Due to the stochastic receiver model and the stochastic part of the channel model, there is no distinct distance at which it is physically impossible to communicate due to low reception power. Therefore, and in order to exclude those physical layer effects, our evaluation focusses on the range between zero and $300\,\mathrm{m}$. Table I also contains the mean PDR across all vehicles for that range. It can be seen that the PDR substantially increases from $82.41\,\%$ to $89.21\,\%$ with our proposed acknowledgement feedback and candidate resource selector.

VI. CONCLUSION AND FUTURE WORK

In this paper, we presented a method to better deal with hidden terminal effects for broadcast-type traffic with low or no overhead. This method includes statistical acknowledgements that acknowledge used radio resources instead of individual packets, and in conjunction with semi-persistent scheduling, can be used to extend the sensing range beyond the capabilities of a normal terminal. Additionally, the feedback can be used by transmitters to detect that their chosen radio resource performs poorly. As a reaction, they can drop these resources and select new ones to prevent further collisions. We based our implementation on LTE-V2X mode 4, and evaluated it using a large-scale, system-level simulation approach. Our approach outperforms standard LTE in all used metrics and is a useful method to improve the reliability of V2V communication when using periodic resource allocations.

As future work, a DCC component for this mechanism should be developed. Additionally, the compromise between the MCS index, number of used subchannels, and performance improvement by the acknowledgements should be further investigated.

REFERENCES

[1] P. Wendland, G. Schaefer, and R. Thomä, “An application-oriented evaluation of LTE-V’s mode 4 for V2V communication,” in *ACM SAC 19, 34th ACM/SIGAPP Symposium On Applied Computing*, Limassol, Cyprus, Apr 2019.

[2] ——, “LTE-V2X mode 4: Increasing robustness and DCC compatibility with reservation splitting,” in *2019 IEEE ICCVE - International Conference on Connected Vehicles and Expo*, Graz, Austria, November 2019.

[3] S. Jung, H. Cheon, and J. Kim, “Reducing consecutive collisions in sensing based semi persistent scheduling for cellular-V2X,” in *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*, Sep. 2019, pp. 1–5.

[4] R. Molina-Masegosa, M. Sepulcre, and J. Gozalvez, “Geo-based scheduling for C-V2X networks,” *IEEE Transactions on Vehicular Technology*, vol. 68, no. 9, pp. 8397–8407, Sep. 2019.

[5] R. Molina-Masegosa, J. Gozalvez, and M. Sepulcre, “An IEEE 802.11p-assisted LTE-V scheduling for reliable multi-link V2X communications,” in *2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)*, Aug 2018, pp. 1–5.

[6] ETSI, “Intelligent transport systems (ITS); access layer specification for intelligent transport systems using LTE vehicle to everything communication in the 5,9 GHz frequency band,” European Telecommunications Standards Institute (ETSI), Technical Specification (TS) 103 613, 2018, version 1.1.1.

[7] ——, “LTE; evolved universal terrestrial radio access (E-UTRA); physical layer procedures,” European Telecommunications Standards Institute (ETSI), Technical Specification (TS) 136 213, May 2019, version 14.9.0.

[8] ——, “LTE; evolved universal terrestrial radio access (E-UTRA); physical channels and modulation,” European Telecommunications Standards Institute (ETSI), Technical Specification (TS) 136 211, May 2019, version 14.9.0.

[9] ——, “LTE; evolved universal terrestrial radio access (E-UTRA); medium access control (MAC) protocol specification,” European Telecommunications Standards Institute (ETSI), Technical Specification (TS) 136 321, May 2019, version 14.10.0.

[10] R. Riebl, H. Günther, C. Facchi, and L. Wolf, “Artery: Extending veins for VANET applications,” in *2015 International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, June 2015, pp. 450–456.

[11] C. Sommer, R. German, and F. Dressler, “Bidirectionally coupled network and road traffic simulation for improved IVC analysis,” *IEEE Transactions on Mobile Computing*, vol. 10, no. 1, pp. 3–15, January 2011.

[12] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz, “SUMO–simulation of urban mobility,” in *The Third International Conference on Advances in System Simulation (SIMUL 2011), Barcelona, Spain*, vol. 42, 2011.

[13] Huawei and HiSilicon, “DMRS enhancement of V2V,” 3GPP Working Group Study R1-160284, 2016. [Online]. Available: https://www.3gpp.org/ftp/tsg_ran/WG1_RL1/TSGR1_84/Docs/R1-160284.zip

[14] L. Cheng, B. E. Henty, D. D. Stancil, F. Bai, and P. Mudalige, “Mobile vehicle-to-vehicle narrow-band channel measurement and characterization of the 5.9 GHz dedicated short range communication (DSRC) frequency band,” *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 8, pp. 1501–1516, Oct 2007.

[15] L. Codecá, R. Frank, S. Faye, and T. Engel, “Luxembourg SUMO traffic (LuST) scenario: Traffic demand evaluation,” *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 2, pp. 52–63, 2017.

[16] A. Mansouri, V. Martinez, and J. Härri, “A first investigation of congestion control for LTE-V2X mode 4,” in *WONS 2019, 15th IEEE Wireless On-demand Network systems and Services Conference, 22-24 January 2019, Wengen, Switzerland*, Jan 2019. [Online]. Available: http://www.eurecom.fr/publication/5791