# A Visible Light Vehicle-to-Vehicle Communication System Using Modulated Taillights

Michael Plattner
Department Mobility and Energy
University of Applied Sciences Upper Austria
Hagenberg im Mühlkreis, Austria
Email: `Michael.Plattner@fh-hagenberg.at`

Gerald Ostermayer
Research Group Networks and Mobility
University of Applied Sciences Upper Austria
Hagenberg im Mühlkreis, Austria
Email: `Gerald.Ostermayer@fh-hagenberg.at`

***Abstract*—In this paper, we propose a visible light vehicle-to-vehicle communication system by modulating the taillights of a car and receiving the signal with a camera. Safety critical communication in applications like platooning requires a fast and secure wireless connection. Such a connection can be established by using our optical communication as an out-of-band channel to transmit a public key to another car following on the road. We are able to transmit 60 bit/s via the optical channel with an average BER (Bit Error Rate) of 3.46% and it takes about 5 seconds on average to receive the transmitted code word containing a 128-bit key in a non-synchronized system. Such an optical channel is very hard to manipulate for a third party and hence the transmitted public key can be used to verify the identity of the communication partner and man-in-the-middle attacks are made more difficult.**

***Keywords*–Automotive applications; Connected vehicles; Vehicle safety; Visible light communication; Differential phase shift keying**

## I. INTRODUCTION

Advanced Driver Assistance Systems (ADAS) and semi-autonomous driving technologies are already built into the newest luxury class cars. However, some future functionalities require cars to communicate with each other, e.g., platooning. Platooning means to group vehicles on the road into platoons and decrease the distance between these vehicles. An electronic coupling of the participating cars allows to accelerate and decelerate simultaneously. This enables the cars inside the platoon to drive in the slipstream of the cars ahead. Due to the smaller amount of drag, fuel can be saved and emissions are reduced. Additionally, this is a method of increasing the capacity of the roads and hence traffic jams can be prevented.

In a platoon, the car ahead needs to control the car following it remotely [1]. Thus, it is crucial to establish a secure connection between the participants. In this paper, we propose a method to transmit a public key via an optical out-of-band channel using the taillights of a car. The following car receives the message using a camera pointing in direction of driving. The public key can then be used to establish a secure encrypted communication via, e.g., 802.11p representing the main channel between two vehicles, where the identity of the car in front is verified via the out-of-band channel. For an attacker, it would be very difficult to fake such a transmission, as the true identity of the sender can be verified using the camera image.

Due to various environmental conditions on the road, caused by daytime, weather, shadows, car model, other light sources, etc., the system needs to adapt to those conditions to ensure a low transmission error rate. Therefore, the system uses convolutional neural networks (CNN) for detecting a transmitting car and for classifying the states of its taillights. By using a broad spectrum of training data in different environments, the resulting system is able to adapt to various lighting conditions and transmitting car models.

In section II, we give an overview of other projects working on VLC using cameras, especially for vehicle-to-vehicle communication. Section III describes our concrete approach for vehicular VLC using modulated taillights. The results of our proof-of-concept are then evaluated in section IV. In section V, we then conclude our findings and give an outlook into future work.

## II. RELATED WORK

Visible Light Communication (VLC) refers to an optical wireless communication system that uses the modulation of light in the visible spectrum (400–700 nm) that is principally used for illumination [2]. The information is encoded on top of the illumination light. A precondition for most cases of VLC is that the modulation of the visible light is not perceived by the human eye. Low modulation frequencies lead to noticeable flickering. Therefore, the base frequency of the illumination light must be higher than the Critical Flicker Frequency (CFF). The CFF is defined as the frequency at which an intermittent light stimulus appears to be completely steady to a human observer [3]. It depends on various factors, e.g., the age of a person, but on average the human eye is able to notice flickering of visible light if the frequency is below 35-50Hz [4].

Viriyasitavat et al. [5] used an off-the-shelf scooter taillight and a photo diode for a VLC system and developed a channel model for vehicle-to-vehicle (V2V) visible light communication. They limited the maximum distance between scooter taillight and photo diode to 10m, for higher distances highly directed light sources would be needed, e.g., lasers (VCSEL) like used by Lu et al. [6]. In contrast to these works, we used a camera instead of a photo diode for receiving the signal. This way we can distinguish between multiple transmitters by analyzing the image using computer vision algorithms, but of course the modulation frequency is limited in comparison to the cut-off frequency of a photo diode.

Using a camera on the receiving side also enables computer vision algorithms to crop the regions of interest of the image that show the modulated taillights of a preceding car. However, assuming the system is able to recognize the correct information of the transmitter in each frame, we are

ISBN: 978-1-61208-781-8

still limited by Shannon's sampling theorem. Luo et al. [7] suggest a modulation variant called Undersampled Phase Shift On-Off Keying (UPSOOK), which overcomes this problem by utilizing the rolling shutter effect of CMOS cameras. This method chooses a modulation frequency between the critical flicker frequency and the cut-off frequency of a camera, which depends on the exposure time. This way the information encoded onto the modulated light source can be recognized by a camera with a very short exposure time, but the human eye does not perceive flickering. Other works from Liu et al. [8] or Lee et al. [9] also propose multiple variations of the utilization of the rolling shutter effect. Liu et al. [8] extend the modulation method by encoding the message into phase shifts between frames and use two modulated light sources to detect errors caused by slight synchronization offsets. Lee et al. [9] propose a variation where the modulated light is visible in the whole camera image and hence are able to encode information into the width of every single pulse of the modulation signal. However, information encoded into pulses that are not visible in the camera image is lost. For this reason, this is not applicable for our system, as the taillights of a car only cover a very small portion of the image.

## III. APPROACH

This paper proposes an out-of-band channel for vehicle-to-vehicle communication using the taillights of a car. A 128-bit public key is transmitted from a car to its follower on the road. This key can then be used to establish a wireless, secure and encrypted connection in the main communication channel over, e.g., the 802.11p standard. The taillights are modulated using Undersampled Differential Phase Shift On-Off Keying (UDPSOOK) [8]. This enables a CMOS camera with a very short exposure time to receive the signal utilizing the rolling shutter effect, while there is no flickering perceivable for the human eye.

### A. Rolling shutter effect

CMOS cameras, which are widely used in digital cameras, Digital Single-Lens Reflex (DSLR) cameras and smartphones use a rolling shutter. This means the image is sampled line by line and hence, e.g., the top of the image is sampled earlier than the bottom of the image. If the object in front of the camera is moving or changing while the image is captured, weird patterns occur like shown in Figure 1. In contrast, a global shutter camera captures all the pixels of the image at once, and hence quick changes or movements while capturing an image with short exposure time have no effect.

### B. Modulation of taillights

The taillights used in the system are state-of-the-art LED taillights, which means they have very low latency when turning on or off (single-digit nanoseconds) compared to conventional halogen taillight bulbs. The chosen modulation frequency must be higher than the critical flickering frequency of approx. 50 Hz. In order to modulate information onto the signal, we need a consistent state throughout the communication. This is achieved by setting the modulation frequency to an exact multiple of the receiving camera's frame rate, which is 30 FPS in our system. However, the modulation frequency should be chosen as low as possible, to detect the states of the taillights easier in the camera images, but high enough to not cause perceivable flickering. Figure 2 compares close-up images of a modulated LED with different modulation frequencies. The stripe pattern occurs due to the rolling shutter effect, the width of the stripes depends on the modulation frequency.

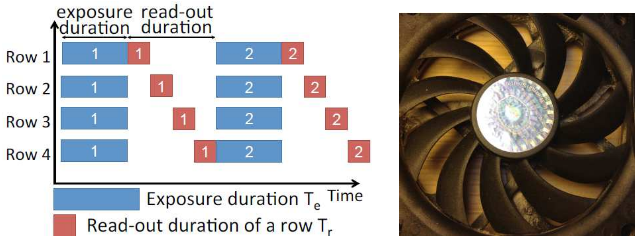

(a) Global shutter operation (b) Global shutter image

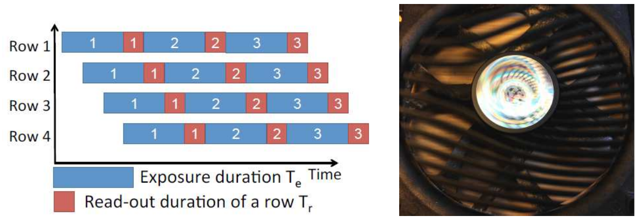

(c) Rolling shutter operation (d) Rolling shutter image

Figure 1. Comparison between global shutter and rolling shutter [9]

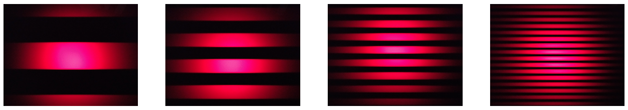

(a) 60 Hz (b) 120 Hz (c) 240 Hz (d) 480 Hz

Figure 2. Close-up pictures of a modulated LED with different frequencies

As depicted in Figure 3, if the camera is farther away from the light source, the area covered by it gets smaller. However, the stripe pattern remains the same, therefore, the state of the taillights in a specific frame depends on the position of the car. Thus, the taillight's state in a single frame is not sufficient to transmit information. The UDPSOOK [8] modulation method hence uses two consecutive frames to encode the information in the phase shift between them. For this reason, the modulation frequency must be an exact multiple of the camera's frame rate. Thus, the strip pattern and the state of the taillights respectively, stay the same in two successive frames if there was no phase shift between them. If there was a phase shift, the stripe pattern changes and the state of the light source changes, which can be detected by the receiver. Figure 4 depicts the idea how this modulation method encodes the information and the receiver samples the signal at a lower frequency, independent of the offset. Precondition for this method is that the modulation frequency of the light source is an exact multiple of the camera's frame rate and that there are no striking movements of the transmitting LED between two frames.

With this modulation method, the information is encoded into the phase shifts between frames. The phase switches between 0 and $\pi$, which is just an inversion of the signal.

Figure 3. Modulated taillights in different distances from close-up to farther away

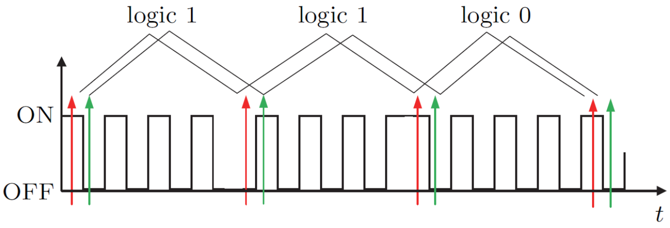

Figure 4. Different sampling timing for UDPSOOK [8]

If the signal is sampled by the receiving camera and the phase changed in comparison to the previous frame, a logical 1 was transmitted. Otherwise, if the phase is the same, a logical 0 was transmitted. Thus, we can transmit one bit per frame and per modulated light source. In our system, we utilize two taillights, that can be modulated separately. So, we can transmit two bits per frame, which results in a transmission rate of 60 bit/s when having a 30 FPS camera at the receiving car. The center high mount stop light (CHMSEL) of cars is not suitable to transmit information by modulation, because a modulated light source is perceived as half-on by the human eye. A stop light must only be turned on, if the car is braking, otherwise it must be turned off. So, in our system, we can transmit data while the car is driving and the modulated taillights are perceived as normal taillights by other drivers. When braking, the transmission is stopped and the taillights are continuously on. For other drivers this looks like the light is brighter and the modulated taillights can additionally be used as brake lights. An alternative for this limitation of the data transmission would be to only adjust the brightness of the LED to differentiate between normal brightness and braking brightness of the taillights. Thus, the transmission would not be interrupted, however this was not in the scope of this work.

This modulation method is prone to errors in some special cases, where the light source does not show one distinct state of the stripe pattern, but the transition between the ON and OFF state. An example for this is shown in Figure 5. In this case, it is hard to detect the correct state and changes of the phase between frames. So, the only way here is to reduce the probability of such a situation to occur by using a stripe pattern, where the stripes are as wide as possible, which means to use a very small modulation frequency. However, we need to make sure to never cause perceivable flickering, even if a phase shift is applied, like shown in Figure 6, and therefore, the frequency is halved for one single pulse. For these reasons we chose a modulation frequency of 120Hz, with a 30 FPS camera at the receiving side in our system.

Figure 5. Modulated taillights showing transitions between ON and OFF, depending on the vertical position in the image

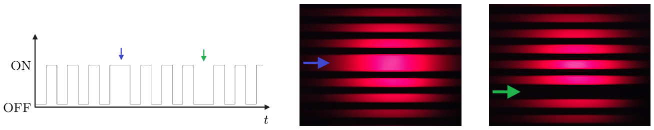

(a) Phase shifts in signal (b) Two consecutive ON states (c) Two consecutive OFF states

Figure 6. Phase shifts shown in signal and resulting stripe pattern

### C. Demodulation of the signal

This system will be used to establish a secure wireless connection between two cars driving in succession on the motorway to enable platooning. This means our system needs to transmit a public key from one car to another car behind, before the platooning process is started. Thus, the two cars are still driving with a safety distance of usually between 30 m and 50 m. This is the main operating range of the optical out-of-band communication channel.

For receiving the signal, we use a common CMOS camera with a rolling shutter and a fixed exposure time of 1 ms. Additionally, the ISO value and gain of the camera is set to its maximum, to still have a bright enough image and to be able to detect the car inside the frame of the camera. The receiving process of a single bit consists of the three following steps:

*1) Vehicle Detection:* The first step of the receiving process is to detect the bounding box of the vehicle in front. We decided to use the YOLO framework proposed by Redmon et al. [10], [11], [12] for detecting the car, Junsheng Fu [13] created a vehicle detection pipeline using the YOLO framework in Python. This framework is able to detect different types of cars and car models and is even able to detect miniature cars printed onto cardboard like used in our first prototype without the context of a real road. However, the detection process is not fast enough to detect the car in real-time in every frame of the camera stream. So, we only detect the car in every 20th frame of a video. This might result in errors if the transmitting car is moving too much inside the camera frame while using a deprecated position of the car for receiving the data. Concrete consequences are covered in the evaluation section. However, faster vehicle detection using a better detection algorithm or ASICs (application-specific integrated circuits) would prevent errors caused by movements of transmitter or receiver.

*2) Taillights ROI Estimation:* When the position of the car in the image is detected, we estimate the regions of interest (ROI) for the taillights using a static calculation depending on the bounding box of the detected car. This is a rather simple

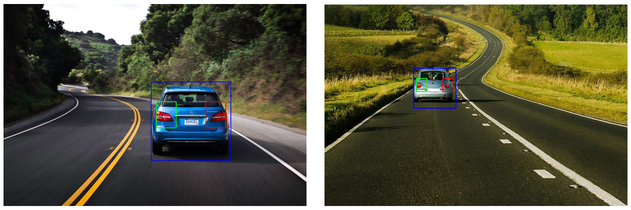

(a) Mercedes-Benz B-Class (b) VW Beetle

Figure 7. Taillight detection on example images of cars on a road

approach, but is sufficient when the rear end of the car was detected accordingly. For this calculation the positions of the car $C$ and the taillights $TL$ and $TR$ are defined by $l$, $r$, $t$ and $b$ representing the left, right, top and bottom border of the bounding box, respectively.

$$l_{TL} = l_C + \frac{r_C - l_C}{16} \tag{1}$$

$$r_{TL} = l_C + \frac{5 \cdot (r_C - l_C)}{16} \tag{2}$$

$$l_{TR} = r_C - \frac{5 \cdot (r_C - l_C)}{16} \tag{3}$$

$$r_{TR} = r_C - \frac{r_C - l_C}{16} \tag{4}$$

$$t_{TL} = t_{TR} = t_C + \frac{b_C - t_C}{4} \tag{5}$$

$$b_{TL} = b_{TR} = t_C + \frac{7 \cdot (b_C - t_C)}{12} \tag{6}$$

Figure 7 shows two examples of cars on a road where we detected the car using the YOLO framework and based on that estimated the positions of the taillights. The cars are marked with a blue rectangle and the positions of their taillights are highlighted red for the right and green for the left one. As we are using a rather big ROI for the taillights, it fits for the majority of car models. Additionally, we do not need the whole taillight inside the ROI, as the transmitted data can be reconstructed using just a fraction of the light source. However, it might be helpful for future works to use a more sophisticated detection approach, as the static calculation relies on an accurate detection of the car's rear.

*3) Taillight State Recognition:* The detected ROI containing the taillights of the car, are reshaped to match the size of 28x28 RGB pixels and then passed into a simple convolutional neural network to classify the current state of the taillight. The network is using two convolutional layers with 16 and 64 channels with two max-pooling layers and a single hidden dense layer with 64 neurons. All three layers use the ReLu activation function. For the two neurons for the states ”on” and ”off” in the output layer the softmax activation function is used.

This network was implemented using Keras with the TensorFlow backend. For training the network, a dataset of approx. 4450 labeled images was used. The training images show the taillights of sending prototypes in different environment and lighting conditions, to get an adaptive network, performing well in different scenarios with different car models. In order to prevent overfitting on the training data, a dropout probability of 50% was implemented during the training. After only 6 epochs of training, the network reached an accuracy of more than 98% in cross-validation. The final network scored an accuracy of 99.4% on an unseen evaluation dataset.

The proposed network design might be changed in future work, if it is not capable of classifying states of real car taillights appropriately. However, the performance of the network is sufficient for the current use case.

The recognized taillight state is then compared with the previous state of the taillight to demodulate the sent message. If the recognized state is the same as in the last frame, a logical 0 was received, otherwise a logical 1.

### *D. Channel Coding*

Wireless communication, especially visible light communication with cameras, is fragile. As already mentioned, some potential error causes cannot be prevented. To get a stable connection, even if errors occur, we used Reed-Solomon [14] channel coding. In our system we want to transmit a 128 bit public key, using the optical out-of-band channel built by the taillights of a car. The transmitted message effectively is always the same, the transmitter just waits until somebody receives the message and starts connecting via the main wireless channel using the transmitted key. We use a RS(24,16) channel coding with 8-bit or 1-byte symbols. This means we use code words with a length of 24 symbols, where 16 symbols carry the message and the remaining 8 symbols are used for error detection and error correction. With 8 error correction symbols we are able to detect and correct 4 erroneous symbols in a code word.

Due to the lack of a synchronization signal, we need a starting sequence of 8 bits to indicate the start of a new code word. Including the starting sequence, the code word to send has a length of 200 bits. After sending 200 bits, the transmitter restarts to send the message again. With two modulated taillights we need 100 frames to transmit 200 bits. As we use a 30 FPS camera, it takes 3.33 seconds to transmit a 128 bit public key via the optical out-of-band channel.

Figure 8 shows a block diagram of the previously explained parts of our system. As depicted, we use two modulated light sources to send a message, where both of them are captured by a single camera.

## IV. EVALUATION

For the evaluation we used various videos recorded with the Canon EOS 1100D DSLR camera. We used different cardboard car models in scale 1:24 for the transmitter in different settings. The distance between the transmitter and the camera was approx. 1.5 m, which represents a distance of 36 m in the real world. We recorded test videos in dark indoor and bright outdoor environments, with and without movement of the sending model inside the video frame.

Figure 9 shows the BER for the optical data transmission for different subsets of test videos. In total, the average BER for the evaluated videos is 6.81% with a standard deviation 5.18%. If we divide the videos into two subsets with videos where the transmitter is and is not moving inside the video frame, we see that movement of the transmitter causes many bit errors. This obviously is due to the fact that we only detect the transmitter every 20th frame to be able to receive the message in almost real-time. Without movement of the

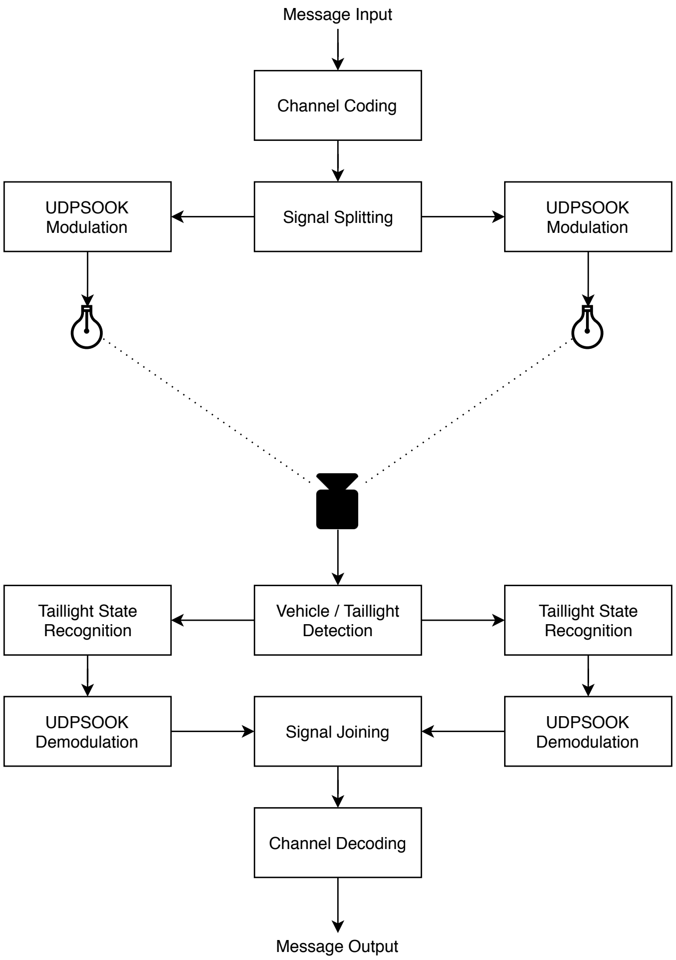

Figure 8. Optical data transmission block diagram

transmitter the BER is at 3.46% on average with a standard deviation of 1.94%, with a moving transmitter it is significantly worse in our setup. Splitting the set of videos with not moving transmitters further down to videos with bright and dark environments we see something interesting. The mean BER of both sets is pretty much the same with 3.6% for dark and 3.2% for bright videos, but the standard deviation of 2.44% is much bigger for dark environments compared to 0.84% for bright ones. This is caused by two different factors. In bright environments, the vehicle detection works very well, but the taillight state recognition is much harder compared to videos with dark light settings. In dark videos, the hard part is to detect the transmitter. However, if the detection was successful, it is pretty easy to recognize the state of the taillight because of the good contrast in the dark background. In some of the dark videos, the vehicle detection works fine, so the BER is very low, but if the car is not detected properly, a burst of bit errors occurs and increases the BER massively. On the other hand, in bright light setting the detection of the transmitter is easier, but there are more single bit errors caused by the taillight state recognition. The difference between dark and bright lighting in the message error rates depicted in Figure 10 encourages this assumption. Error bursts in the received bit strings cause corrupted messages that cannot be corrected. Single bit errors can be corrected using the error correction code of the channel coding, and hence the message error rate of 13.41% on average for bright videos is better compared to dark videos with a mean message error rate of 22.63%.

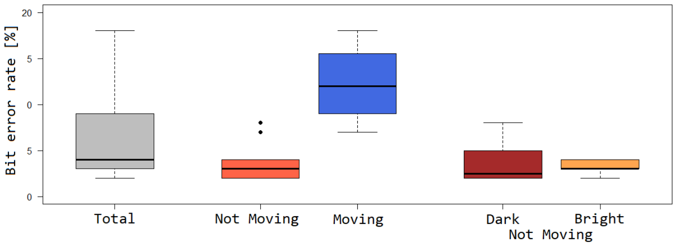

Figure 9. Boxplots of BER for different settings

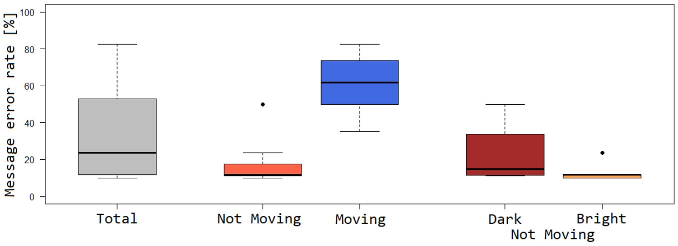

Figure 10. Boxplots of Message Error Rates for different settings

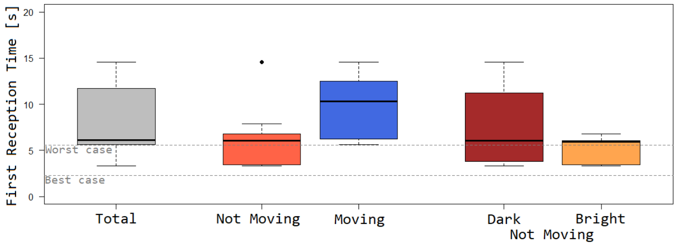

Figure 11. Boxplots of First Reception Time for different settings

What we can also see in this chart is that the message error rate is never below 10% in our experimental setup. These errors are caused by the offset in synchronization between transmitter and receiving camera. Approximately every 20th message the stripe pattern that appears due to the UDPSOOK modulation of the taillights is moving over the area of the taillight inside the camera frame and causes ambiguous taillight states that cannot be decoded properly. Those transitions of stripes were already mentioned as know error causes. They cause error bursts that usually affect two consecutive messages and hence about 10% of the sent messages get corrupted. In our setup, we want to send the messages as quick as possible, therefore, we decided to not use interleaving approaches in our channel coding. Interleaving might correct messages that got corrupted by error bursts, but the time until a single message is transmitted would be significantly higher.

Another interesting metric to evaluate is the first reception time of the sent message in the test videos. Boxplots for this are shown in Figure 11. In our test scenario, we send code words with a length of 200 bits, this means we need at least 100 frames with 2 bits per frame to transmit the whole code word. With a frame rate of 30 FPS, this takes 3.33 seconds. On average, it takes 5.244 seconds until the message was successfully transmitted for the first time, however, the minimum reception time in our test videos is 3.30 seconds and hence lower than the theoretically possible time to send the whole code word. This is possible because of the channel coding in our system. Assuming we have an error free transmission of the code word, we are able to decode the message even before the whole code word was sent. In our case, if we received the starting sequence and the first 20 bytes of the 24-byte code word correctly, we can set the last 4 symbols of the code word to any value and the channel coding enables us to decode the message correctly. This can be done because we can detect and correct 4 erroneous symbols in the code word, using the 8 error correction symbols. In the case of an error free transmission, those 4 errors are the 4 symbols that were not transmitted yet. This means we just need the 8-bit starting sequence and 20 bytes of the code word for a correct reception. These are 136 bits, where we need 68 frames to transmit them, which takes 2.267 seconds. This means if the test video starts exactly when the first bit of the starting sequence is sent, the first reception of the message is possible after 2.267 seconds. The worst case with error free transmission would be that the video starts after the first bits of the starting sequence were sent. In this case, we would not be able to detect that the first message is sent and therefore, the first reception would be possible after 5.567 seconds. Those two times for the best and the worst case with error free connection are also shown as dashed horizontal lines in Figure 11.

## V. CONCLUSION

We can conclude that our proof-of-concept experiment of an optical out-of-band channel for vehicle-to-vehicle communication using modulated taillights was successful. We managed to build prototypes of different car models in a scale of 1:24 with LEDs representing the taillights. Those LEDs where modulated using the UDPSOOK modulation method, where a camera with a very short exposure time of 1 ms is able to capture distinct states of the light source, but the human eye is not able to perceive any flickering. On the receiving side we used a Canon EOS 1100D DLSR camera to receive the signal and record evaluation videos, but actually any other CMOS camera with rolling shutter can be used. The only precondition is that videos can be recorded while the shutter speed is set manually.

The results showed that we were able to transmit messages with an average BER of 3.46% with a standard deviation of 1.94% in videos where no striking movements of the transmitter inside the camera frame occur. As our system is designed for vehicle-to-vehicle communication on the highway with two cars driving in succession, we can assume that the relative position of the car in front is quite stable and only changes slowly. Of course, the total BER is improved by the Reed-Solomon channel coding. However, only 90% of the sent 200-bit code words can be decoded correctly, due to error bursts when the transitions of UDPSOOK modulation pulses are aligned with the taillights of the sending car. Such error bursts can only be corrected with interleaving methods in the channel coding, which would make the time until a code word is received for the first time significantly longer. In our evaluation, it took 5.244 seconds on average to receive the correct code word, in a platooning application, this would mean that after just a few seconds we can establish a secure and encrypted connection between two cars in the main wireless communication channel, e.g., using 802.11p. This connection then might be running for dozens of minutes or even multiple hours, if the two cars have a similar path to their destination.

A future goal for this work is to port this proof-of-concept system using a 1:24 prototype into a full-sized car and testing the performance of the communication system on the road. In the real world, there will be much more interfering factors like other cars that are not sending any information using their taillights or other light sources like traffic lights, lamp posts or the sun, which might interfere. Another idea is to not just use a single camera for receiving the signal, but to have two or multiple cameras that add redundancy and therefore, a better BER, or they could be used to filter specular reflections from other light sources by merging the camera images like suggested by Plattner and Ostermayer [15].

## REFERENCES

[1] D. Swaroop, "String stability of interconnected systems: An application to platooning in automated highway systems," Ph. D. Dissertation, University of California, 1997.

[2] S. Arnon, Visible light communication. Cambridge University Press, 2015.

[3] H. Luczak, Arbeitswissenschaft. 2, vollständig bearbeitete Auflage. Springer-Verlag Berlin und Heidelberg, 1998.

[4] G. Kircheis, M. Wettstein, L. Timmermann, A. Schnitzler, and D. Häussinger, "Critical flicker frequency for quantification of low-grade hepatic encephalopathy," Hepatology, vol. 35, no. 2, 2002, pp. 357–366.

[5] W. Viriyasitavat, S.-H. Yu, and H.-M. Tsai, "Short paper: Channel model for visible light communications using off-the-shelf scooter taillight," in Vehicular Networking Conference (VNC), 2013 IEEE. IEEE, 2013, pp. 170–173.

[6] H.-H. Lu, Y.-P. Lin, P.-Y. Wu, C.-Y. Chen, M.-C. Chen, and T.-W. Jhang, "A multiple-input-multiple-output visible light communication system based on vcsels and spatial light modulators," Optics Express, vol. 22, no. 3, 2014, pp. 3468–3474.

[7] P. Luo, Z. Ghassemlooy, H. Le Minh, X. Tang, and H.-M. Tsai, "Undersampled phase shift on-off keying for camera communication," in Wireless Communications and Signal Processing (WCSP), 2014 Sixth International Conference on. IEEE, 2014, pp. 1–6.

[8] N. Liu, J. Cheng, and J. F. Holzman, "Undersampled differential phase shift on-off keying for optical camera communications," Journal of Communications and Information Networks, vol. 2, no. 4, 2017, pp. 47–56.

[9] H.-Y. Lee, H.-M. Lin, Y.-L. Wei, H.-I. Wu, H.-M. Tsai, and K. C.-J. Lin, "Rollinglight: Enabling line-of-sight light-to-camera communications," in Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services. ACM, 2015, pp. 167–180.

[10] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," arXiv, 2015.

[11] J. Redmon and A. Farhadi, "Yolo9000: Better, faster, stronger," arXiv preprint arXiv:1612.08242, 2016.

[12] ——, "Yolov3: An incremental improvement," arXiv, 2018.

[13] J. Fu. Vehicle detection for autonomous driving. Visited on 2020-06-02. [Online]. Available: https://github.com/JunshengFu/vehicle-detection

[14] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," Journal of the society for industrial and applied mathematics, vol. 8, no. 2, 1960, pp. 300–304.

[15] M. Plattner and G. Ostermayer, "Filtering specular reflections by merging stereo images," in Scandinavian Conference on Image Analysis. Springer, 2019, pp. 164–172.