A Systematic Study on the Impact of Noise and OFDM Interference on IEEE 802.11p

Bastian Bloessl*, Florian Klingler†, Fabian Missbrenner† and Christoph Sommer†
*CONNECT Center, Trinity College Dublin, Ireland
†Heinz Nixdorf Institute and Dept. of Computer Science, Paderborn University, Germany
{bloessl,klingler,fabian.missbrenner,sommer}@ccs-labs.org

***Abstract*—To design, test, and evaluate applications for Vehicular Ad Hoc Networks (VANETs), researchers rely heavily on network simulations. These allow conducting experiments in a fast, cheap, and reproducible manner. In general, the accuracy of simulation results depends to a large degree on the quality of the simulation models. Here, the model of the physical layer is particularly crucial for the realism of the results. Given its relevance, it is unfortunate that there is a dispute within the community on how interference should be modeled. To fill this gap, we conduct a systematic study of the IEEE 802.11p physical layer in which we cross-validate results from simulations, off-the-shelf devices, and lab equipment. The results of these experiments are all coherent and indicate that intra-technology interference, i.e., interference from other IEEE 802.11p devices, has a similar impact than noise. Treating interference like noise is, therefore, not just a simplification that is adopted by many network simulators, but accurately captures reality.**

I. INTRODUCTION AND RELATED WORK

Future cars will be part of Vehicular Ad Hoc Networks (VANETs) and employ wireless communication to exchange information directly with each other, enabling a wide range of applications – from cooperative awareness to cooperative mobility. To design, test, and evaluate these applications, researchers rely heavily on computer simulation, as this allows experimenting with protocol parameters in a fast, cheap, and reproducible manner [1]. The realism and accuracy of these simulators depends to a large degree on the quality of their simulation models. Especially the simulation model of the physical layer is crucial for the fidelity of the results, as its main task is to decide whether a given combination of signal, interference, and noise would allow a frame to be decoded [2].

Yet, there is still a dispute within the community on how interference should be modeled – or if it can be modeled accurately at all [3], [4]. Popular network simulators, like *ns-3* and *Veins*, assume that interference can be treated as being similar to noise. The decision to treat interference like noise is easy to understand, if we consider that the WLAN simulation model was adapted from readily available models operating partly on the license-exempt ISM band in the 2.4 GHz band. On this band, we can find a large range of interference sources, like microwave ovens, cordless telephones, and ZigBee transceivers. It is, therefore, hard to characterize interference in detail [5]. Considering different interference sources, or even the combination of different interference sources, would extremely complicate simulation models. Given this complexity, the most practical solution was often to treat all interference in the same way. However, with that regard, IEEE 802.11p is special as it uses dedicated spectrum in the 5.9 GHz band. Interference is, therefore, limited to intra-technology interference, i.e., interference from other IEEE 802.11p devices, which might cast doubt on earlier assumptions.

Some works argue that there is, in fact, a large difference in the impact of interference and noise. A popular example is the work of Fuxjaeger and Ruehrup [6]. Experimenting with frame capturing, the authors noticed that the Device Under Test (DUT), an off-the-shelf IEEE 802.11p card, could cope with interference better than expected. The authors compared their results to the NIST error model, a state-of-the-art model used by popular simulators such as *Veins* and *ns-3*. This model is empirically validated with off-the-shelf WiFi cards, but only with Additive White Gaussian Noise (AWGN) [7]. Comparing the measurement results to the error rate calculated by the simulator, the authors concluded that noise must have a more detrimental impact than OFDM interference. Network simulators would, therefore, produce overly pessimistic results in interference scenarios.

In fact, one could even argue in the opposite direction: Compared to noise, an OFDM signal could create more spotted interference exactly at the subcarrier frequencies and might, therefore, have a worse impact than a similar level of noise. These examples show that the relation between noise and interference is, at least, not trivial.

When consulting the available literature, we find that there are, on the one hand, papers that deal with physical layer performance in different channels, but without interference [8], [9] On the other hand, there are papers that consider interference, but do not target physical layer performance in general [6], [10]. Instead, they focus on characterizing the capture effect, i.e., the special case where a frame is interfered by a high power frame. In that case, the receiver might cancel reception of the initial frame and switch over to the high power frame to avoid losing both frames.

Here, we contribute to reconciling these different viewpoints on the impact of noise and OFDM interference on IEEE 802.11 a/g/p WLANs in the following aspects:

- we conduct an extensive set of physical layer simulations to investigate these questions (Section II);
- we systematically cross-validate these simulations with empirical results and over-the-air experiments using both lab equipment and off-the-shelf hardware and find that, on

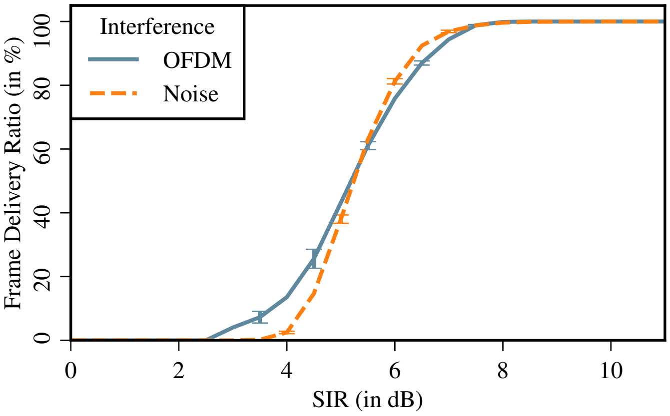

Figure 1. GNU Radio simulations of the frame delivery ratio of an OFDM frame that is interfered by OFDM frames or a similar level of noise.

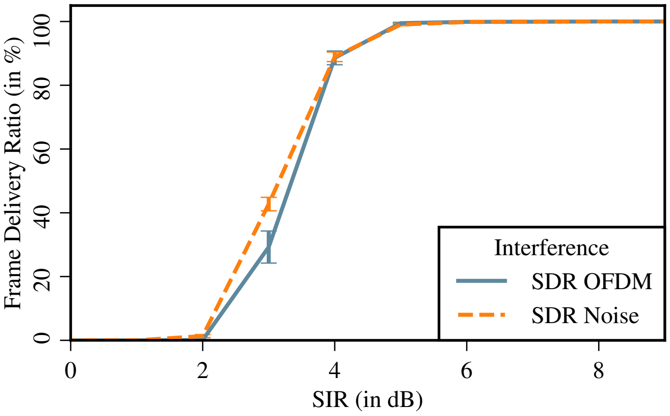

Figure 2. Frame delivery ratio of a Unex DCMA-86P2 when receiving a frame that is interfered by another OFDM frame or a similar level of noise.

a macroscopic scale, the impact of noise can be modeled by OFDM interference and vice versa (Section III);

- based on these findings, we illustrate a way of constructing a testbed for the controlled creation and evaluation of interference scenarios using only off-the-shelf WiFi cards (Section IV).

II. SIMULATIVE EVALUATION

To understand the impact of intra-technology interference and noise on IEEE 802.11p physical layer performance, we set up simulations with our well-validated GNU Radio based IEEE 802.11p transceiver implementation [11]. GNU Radio is a real-time signal processing framework for use in Software Defined Radio (SDR) systems based on General Purpose Processors (GPPs), where signal processing is implemented on a normal PC. Based on this, we can perform very detailed simulations of the IEEE 802.11p physical layer, as the SDR implementation works on the complex base band signal (a sampled version of the down-converted electro-magnetic waveform).

In our simulations, we send 546 Byte QPSK-½ frames that are interfered either by noise or by another IEEE 802.11p frame. Interference starts during the frame with a delay of 122 µs (corresponding to 31 % of the frame time). The parameters are chosen to allow crosschecking results and are similar to those employed by Fuxjaeger and Ruehrup [6]. To make the simulations as realistic as possible, we resampled one frame slightly to introduce sample, phase, and frequency offsets. Furthermore, we varied the alignment of OFDM symbols between the original and the interfering frame. In every simulation run, we sent 100 frames per run and made 80 runs per configuration.

Since we want to isolate the effect of interference, we set a high Signal to Noise Ratio (SNR) (over 40 dB) at the start of the frame, i.e., during the first 122 µs. Given the high SNR at frame start, the performance is determined by the power ratio of the interfered part of the frame. In the following, even though we alternate between adding noise and OFDM interference, we always denote this power ratio as the *Signal to Interference Ratio (SIR)* of the frame.

Figure 1 illustrates the frame delivery ratios for various SIRs; error bars indicate confidence intervals at a confidence level of 95 %. The results show that frame delivery ratios for noise and for OFDM are very similar. The sole difference is that OFDM shows a slightly more stretched curve. Looking at individual runs, we notice that this is because, with OFDM interference, we can clearly differentiate cases where the OFDM symbols were closely aligned (i.e., the FFT windows overlap) to when they were not. Still, with regard to packet level simulations, the results indicate that it is reasonable to treat noise and interference as similar.

To further validate our simulation results, we compared them to the NIST error model [7], a well-established error model for IEEE 802.11a/g/p that is used by popular network simulators like ns-3 and Veins. Compared to our simulations, the NIST model predicts a sharper transition from no reception to reception (data not shown); this is in line with observations when the model was validated against commercial cards [7].

III. EXPERIMENTAL EVALUATION

The simulations suggest that noise and interference may be treated as similar in network simulators. To back this important result by measurements, we set up radio transceivers in our lab. Since we already used an SDR implementation in our simulations, we can run the same software together with a radio frontend (we employed an Ettus Research B210 with a 9 dBi ECOM9-5500 dipole antenna) to transmit over the air. Like in the simulations, we generated the signal plus interference and transmitted the combined signal. The mixed signal is sent with high gain to minimize the impact of thermal noise. With this approach we know the SIR very precisely.

To assert that the effect of noise and interference is not specific to our SDR implementation, we used two off-the-shelf WiFi cards as receivers: a Unex DCMA-86P2, supported by the Linux *ath5k* driver; and a Netgear WNDA3200, supported by the Linux *ath9k_htc* driver. Especially the results from the Unex DCMA-86P2 are interesting, since the card is specifically designed for Vehicular Ad Hoc Networks (VANETs) and was already used in many field operational tests [12]–[15]. We

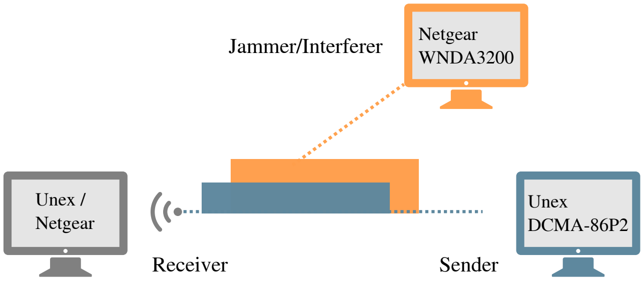

Figure 3. Based on commercial WiFi cards, our testbed provides a cheap and accessible solution.

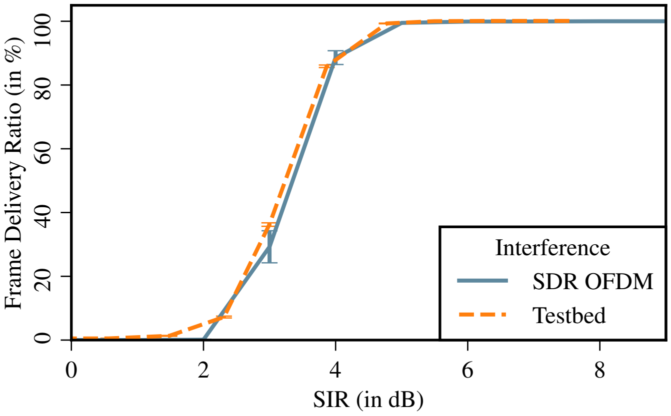

Figure 4. Frame delivery ratio of a Unex DCMA-86P2 under OFDM interference created with an SDR and our testbed.

modified both drivers to tune to the 5.9 GHz band allocated for VANETs. Given current limitations of the *ath9k_htc* driver, we performed the measurements at 20 MHz channel bandwidth.

Like in the simulations, we sent 546 Byte QPSK-½ frames and delayed the interfering signal by 122 µs. The experiments were conducted on channel 178 at 5.89 GHz. Using a spectrum analyzer, we made sure that this channel was vacant in our environment. Again, we repeated the experiment 80 times, sending 100 frames per run.

Figure 2 illustrates the frame delivery ratio for the Unex card. Also in this experiment, we can see that the performance of OFDM interference and noise is very similar, backing up our simulation results. For the *ath9k_htc* cards, we observed similar results (data not shown). Summing up, while different receivers showed different overall performance, they all coped equally well with noise and with interference.[1]

IV. OFF-THE-SHELF TESTBED

Prompted by these results, we now set out to check whether expensive SDR testbeds for investigating the robustness of Devices Under Test (DUTs) against noise might be substituted by a testbed made up of cheap off-the-shelf WiFi cards (which can only send OFDM frames). Such a testbed would be cheaper and easier to set up (not needing multiple closely synchronized SDRs), yielding a platform that is more accessible to research. Results gathered from such a testbed would also serve for further cross-validation of our simulations and measurements.

Figure 3 shows an overview of the resulting testbed. To create interference scenarios in a controlled and reproducible manner, we used a Netgear WNDA3200 USB WiFi dongle and flashed it with custom firmware, based on the work of Vanhoef and Piessens [16]. Developed in the security context, the firmware modifies the dongle to emit signals in response to transmissions, i.e., act as a reactive jammer. In a nutshell, the firmware constantly checks whether the transceiver is about to receive a frame and, if this is the case, interrupts reception and sends an interfering WLAN frame. Given the fact that we use a normal WiFi card as interferer, we cannot send arbitrary signals. However, we can send all kinds of IEEE 802.11a/g/p, i.e., OFDM, frames. As far as validating our measurements of previous sections goes, since in the testbed the data frame and the interfering frame are sent from different transmitters, we do not immediately know the SIR at the receiver. Therefore, we precede every interference experiment with a measurement run, where we sent frames in a similar configuration but one after the other, i.e., without interference. To ease these measurement runs, we extended the reactive jammer to wait for a configurable delay between sensing the frame transmission and transmission of the jam signal. Adjusting this delay with micro second resolution, we can, on the one hand, create different interference scenarios and, on the other hand, easily conduct the measurement runs. If we delay the jam signal long enough to be sent after the frame, the receiver can log frames from both the regular transmitter and the jammer, allowing us to calculate the SIR based on the received signal strength, which is annotated in the *radiotap* header of received frames. Like Pei and Henderson [7], we conducted experiments with different transmit gains and found that, at least, the relative power levels reported by the cards are very accurate. We made sure that the SNR is very high (above 40 dB) at parts that are not interfered. This way, we made sure that the performance in terms of packet delivery ratio is only determined by the SIR of the interfered part of the frame.

To validate our setup, we started by reproducing the capturing experiments of Fuxjaeger and Ruehrup [6] and Lee et al. [10]. Both papers show that a stronger interfering frame is captured, starting from a SIR of 8dB and is captured reliable with a SIR bigger than 12dB. (The data from these measurements is not shown, given the size constraints of the paper.)

After this preliminary validation, we configured the testbed to recreate the setup from previous measurements, i.e., we sent QPSK-½ frames on channel 178 and delayed the jammed frame by 122 µs. By adjusting the transmit power of the sender, we configured different SIRs and measured frame delivery ratios.

Figure 4 compares the results of these measurements against the SDR measurements reported in the previous section. The error bars, again, indicate the confidence intervals at a

---

[1] This might also explain the conclusions of Fuxjaeger and Ruehrup [6]. The authors compared the error model of the network simulator (i.e., noise in simulation) directly with interference measurements. Therefore, it could be possible that their particular receiver just did not match the error model; it does not necessarily imply that OFDM interference has a fundamentally different impact on physical layer performance.

confidence level of 95 %. The most important observation is that the curves are very similar, which backs up our previous results and validates the operation of the testbed.

To make sure that we do not merely observe the specific characteristics of one particular receiver (or receiver model), we conducted similar measurements with a Netgear WNDA3200. We confirmed that, also with this setup, the frame delivery ratio of the SDR experiments and our testbed matches very well (data not shown).

While this is very motivating as far as developing a cheap and effective testbed goes, the testbed setup is not without drawbacks. Since the jammer only *reacts* on transmissions, it has one limitation in that we cannot create arbitrary interference situations due to the delay between the jammer receiving a frame and it producing an interfering frame. Using a spectrum analyzer, we measured the delay at the receiver to be 122 µs (corresponding to about 15 OFDM symbols at 10 MHz bandwidth). Yet, considering that the testbed does not rely on SDRs (like [6]) and does not require complex synchronization and split interference domains (like [10]) – and considering that it uses real off-the-shelf WiFi hardware – it might provide an interesting avenue for future research. Moreover, apart from the fact that it shows how work from different areas can be combined, the fact that the testbed uses off-the-shelf WiFi cards also makes it accessible: It uses only regular hardware in a normal Linux WLAN setup, only using modified firmware. Using this firmware is straightforward, as it only requires replacing a file on the host computer as the operating system uploads the firmware to the WiFi dongle every time it is plugged. To allow experimenting with our testbed, we make the firmware with our modifications available for download.[2]

V. DISCUSSION AND CONCLUSION

To understand the impact of noise and interference on IEEE 802.11p, we conducted detailed physical layer simulations with a Software Defined Radio (SDR) implementation and performed over-the-air measurements in two very different testbeds. While we did not explore the full parameter space with all modulation and coding schemes, frame sizes, and interference situations, our experiments produced coherent results and a subset could even be crosschecked with the available literature. Overall, these results strongly suggest that intra-technology interference and noise have a similar impact on the packet delivery rate of IEEE 802.11p networks. The consequences for the network community are positive in the sense that the commonly adopted simplification of network simulators to treat noise just as interference is reasonable.

This result also has another interesting implication for constructing testbeds investigating the impact of noise on Devices Under Test (DUTs) designed for Vehicular Ad Hoc Network (VANET) operation. Considering that noise and interference have similar impact on the physical layer performance, we can build an accessible physical layer testbed based on normal WiFi cards. As future work, plan to extend the testbed to capture all interference scenarios.

ACKNOWLEDGMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 713567.

REFERENCES

[1] C. Sommer, J. Härri, F. Hrizi, B. Schünemann, and F. Dressler, "Simulation Tools and Techniques for Vehicular Communications and Applications," in *Vehicular ad hoc Networks - Standards, Solutions, and Research*, C. Campolo, A. Molinaro, and R. Scopigno, Eds., Springer, May 2015, pp. 365–392.

[2] M. Takai, J. Martin, and R. Bagrodia, "Effects of wireless physical layer modeling in mobile ad hoc networks," in *2nd ACM International Symposium on Mobile Ad hoc Networking and Computing (MobiHoc 2001)*, Long Beach, CA: ACM, Oct. 2001, pp. 87–94.

[3] R. Patidar, S. Roy, T. R. Henderson, and A. Chandramohan, "Link-to-System Mapping for Ns-3 Wi-Fi OFDM Error Models," in *Workshop on NS-3 (WNS3 2017)*, Porto, Portugal: ACM, 2017, pp. 31–38.

[4] C. Shahriar, M. L. Pan, M. Lichtman, T. C. Clancy, R. McGwier, R. Tandon, S. Sodagari, and J. H. Reed, "PHY-Layer Resiliency in OFDM Communications: A Tutorial," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 292–314, 2015.

[5] S. Papanastasiou, J. Mittag, E. G. Strom, and H. Hartenstein, "Bridging the Gap between Physical Layer Emulation and Network Simulation," in *IEEE Wireless Communications and Networking Conference (WCNC 2010)*, Sydney, Australia: IEEE, Apr. 2010, pp. 1–6.

[6] P. Fuxjaeger and S. Ruehrup, "Validation of the NS-3 Interference Model for IEEE802.11 Networks," in *8th IFIP Wireless and Mobile Networking Conference (WMNC 2015)*, Munich, Germany: IEEE, Oct. 2015, pp. 216–222.

[7] G. Pei and T. R. Henderson, "Validation of OFDM Error Rate Model in ns-3," Boeing Research & Technology, Tech. Rep., 2010.

[8] C. F. Mecklenbräuker, A. F. Molisch, J. Karedal, F. Tufvesson, A. Paier, L. Bernadó, T. Zemen, O. Klemp, and N. Czink, "Vehicular Channel Characterization and its Implications for Wireless System Design and Performance," *Proceedings of the IEEE*, vol. 99, no. 7, Jul. 2011.

[9] J. A. Fernandez, K. Borries, L. Cheng, B. V. K. Vijaya Kumar, D. D. Stancil, and F. Bai, "Performance of the 802.11p Physical Layer in Vehicle-to-Vehicle Environments," *IEEE Transactions on Vehicular Technology*, vol. 61, no. 1, pp. 3–14, Jan. 2012.

[10] J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. Choi, "An Experimental Study on the Capture Effect in 802.11a Networks," in *2nd ACM International Workshop on Wireless Network Testbeds, Experimental evaluation and Characterization (WiNTECH'07)*, Montreal, Canada: ACM, Sep. 2007, pp. 19–26.

[11] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, "Towards an Open Source IEEE 802.11p Stack: A Full SDR-based Transceiver in GNURadio," in *5th IEEE Vehicular Networking Conference (VNC 2013)*, Boston, MA: IEEE, Dec. 2013, pp. 143–149.

[12] F. A. Teixeira, V. F. e Silva, J. L. Leoni, D. F. Macedo, and J. M. Nogueira, "Vehicular networks using the IEEE 802.11p standard: An experimental analysis," *Elsevier Vehicular Communications*, vol. 1, no. 2, pp. 91–96, Apr. 2014.

[13] J. Santa, F. Pereñíguez, A. Moragón, and A. F. Skarmeta, "Experimental evaluation of CAM and DENM messaging services in vehicular communications," *Elsevier Transportation Research Part C: Emerging Technologies*, vol. 46, pp. 98–120, Sep. 2014.

[14] A. B. Reis, S. Sargento, F. Neves, and O. K. Tonguz, "Deploying Roadside Units in Sparse Vehicular Networks: What Really Works and What Does Not," *IEEE Transactions on Vehicular Technology*, vol. 63, no. 6, pp. 2794–2806, Jul. 2014.

[15] A. Geiger, M. Lauer, F. Moosmann, B. Ranft, H. Rapp, C. Stiller, and J. Ziegler, "Team AnnieWAY's entry to the 2011 Grand Cooperative Driving challenge," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 3, pp. 1008–1017, Apr. 2012.

[16] M. Vanhoef and F. Piessens, "Advanced Wi-Fi Attacks Using Commodity Hardware," in *30th Annual Computer Security Applications Conference (ACSAC 2014)*, New Orleans, LA: ACM, Dec. 2014, pp. 256–265.

---

[2] https://www.wime-project.net/projects/interference/