

Experimental Characterization and Modeling of RF Jamming Attacks on VANETs

Óscar Puñal, Carlos Pereira, Ana Aguiar, *Member, IEEE*, and James Gross, *Member, IEEE*

Abstract— In this work, we evaluate the performance of 802.11p-based vehicular communications in the presence of RF jamming attacks. Specifically, we characterize the transmission success rate of a car-to-car link subject to constant, periodic, and reactive RF jamming. First, we conduct extensive measurements in an anechoic chamber, where we study the benefits of built-in techniques for interference mitigation. In addition, we identify that the periodic transmission of preamble-like jamming signals can hinder successful communication despite being up to five orders of magnitude weaker than the signal of interest. We further provide the rationale behind this remarkably high jammer effectiveness. Additionally, we quantify the impact of reaction delay and interference signal length on the effectiveness of the reactive jammer. Next, by means of outdoor measurements, we evaluate the suitability of the indoor measurements for being used as a model to characterize the performance of car-to-car communications in the presence of RF jamming. Finally, we conduct outdoor measurements emulating a vehicular platoon and study the threats that RF jamming poses to this VANET application. We observe that constant, periodic, but also reactive jammer can hinder communication over large propagation areas, which would threaten road safety.

I. INTRODUCTION

Vehicular Ad-hoc Networks (VANETs) have recently attracted the interest of researchers and industry due to their potential to improve road safety [1] and traffic coordination [2]. The packets exchanged by these applications require timely and reliable delivery which poses a real challenge to VANETs due to the impairments of the vehicular wireless channel. To partially address these issues, standardization efforts have meanwhile lead to the approval of the IEEE 802.11p amendment [3]. For this amendment, the 802.11a Physical Layer (PHY) was modified by reducing the channel

bandwidth from 20 MHz to 10 MHz to better cope with multipath fading. Further features enhancing the reliability were the choice of the dedicated 5.9 GHz frequency band as well as the design of a prioritized channel access for critical messages. In the research domain, various works [4], [5] studied reliability enhancing measures, such as the use of short packets and robust modulation and coding schemes. Except for higher-layer security threats, which have been studied for example in [6] and [7], VANETs appear to be quite reliable from a standardization and research point of view with respect to safety-critical messaging. However, the impact of radio frequency (RF) jamming on VANETs has not been studied so far. With the proliferation of powerful software-define radio platforms that are capable of interfering 802.11p networks [8], RF jamming could compromise road safety. RF jamming has been extensively studied in the context of classical 802.11 networks without accounting for the particularities of car-to-car communications. Besides the differences in PHY design of 802.11p compared to other 802.11 amendments, the propagation conditions of VANET are fundamentally different due to the highly dispersive and rapidly changing vehicular environment. Hence, we expect differences in the impact of jamming and, therefore, experiments in representative vehicular scenarios are necessary to characterize the vulnerability of VANETs and its geographic extension. This paper addresses this short-coming and, in detail, we contribute the following:

- 1) We provide a thorough experimental evaluation of the performance of 802.11p devices under the impact of RF jamming in an anechoic chamber. We significantly extend our earlier work [8] by increasing the granularity of the measurements and by accounting for a larger variety of jamming signals. In particular, we consider a constant, a periodic, and a reactive jammer. Among other findings, we show that periodic jamming signals can impair communication up to an SINR of 56 dB.
- 2) We present a detailed description of the packet detection procedure of a reference 802.11p implementation in an Atheros chipset. We identify the elements that are most vulnerable to RF jamming, namely preamble-triggered false signal detections and dynamic range overflow at the analog-to-digital converter. Our measurements and observations extend earlier works [9].
- 3) We propose a methodology for using the performance characterization carried out indoor as a tool for modeling the behavior of 802.11p networks in the presence

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Manuscript received August 9, 2013; revised November 20, 2013 and February 17, 2014; accepted April 17, 2014. The review of this paper was coordinated by Prof. M. Fang.

This research was partially funded by the FCT projects PEst-OE/EEI/LA0008/2011 and PEst-OE/EEI/LA0008/2013. This research was partially funded by the DFG Cluster of Excellence on Ultra High-Speed Mobile Information and Communication (UMIC).

Ó. Puñal is with the UMIC Research Centre, RWTH Aachen University, Mies-van-der-Rohe-Str. 15, 52074 Aachen, Germany (e-mail: punal@umic.rwth-aachen.de).

Carlos Pereira and Ana Aguiar are with the Instituto de Telecomunicações, Rua Dr. Roberto Frias, s/n 4200-465 Porto, Portugal (e-mail: dee12014@fe.up.pt; ana.aguiar@fe.up.pt).

James Gross is with the KTH Royal Institute of Technology, Osquldas Väg 10, SE-100 44 Stockholm, Sweden (e-mail: james.gross@ee.kth.se).

of RF jamming. We apply the model to predict the performance of two communicating nodes in a vehicular environment and observe a high agreement between predicted and measured performance in the field.

- 4) We evaluate the impact of different RF jamming signals on a vehicular platoon by means of measurements. We confirm our earlier observations in [8] that any of the considered RF jamming attacks can effectively disrupt transmissions within a platoon over a large area. Particularly alarming are the dimensions of the communication *blackout area* caused by constant and periodic jammers, which can span more than 400 m in an open field environment.
- 5) We provide the outdoor data used in our earlier paper [8] and both indoor and outdoor data used in this paper. The data is available for download in *cawdad* [10] and [11] to foster the re-use and fast progress on the topic.

The rest of the paper is structured as follows. In Section II, we briefly describe the main PHY and Medium Access Control (MAC) layer characteristics of 802.11p. Furthermore, we introduce the selected 802.11p devices and review the main packet reception steps and vendor-specific interference mitigation techniques. In Section III, we describe the setup and methodology chosen for the measurements in the anechoic chamber and present the results obtained in the presence of constant, periodic, and reactive jammer. Section IV reproduces, in an open field environment, the measurements carried out in the anechoic chamber. We show that the indoor results can be used as a model to precisely predict the achievable performance of two vehicles in the presence of RF jamming. In Section V, we highlight the threats that RF jamming poses to a vehicle platoon and observe a complete communication disruption over large propagation areas, specially in the presence of non-reactive jamming signals. Finally, in Section VI we provide an overview of related work before we conclude the paper in Section VII.

II. PRELIMINARIES

In this section, we briefly summarize the IEEE 802.11p [3] standard which is part of the Wireless Access in Vehicular Environments (WAVE) standard [12] and defines PHY and MAC functionalities. We further provide a description of the devices used in our measurements. Background information and descriptions are provided for the later presentation and discussion of our results.

A. IEEE 802.11p

IEEE 802.11p is an amendment for vehicle-to-vehicle and vehicle-to-infrastructure communication. It is based on the 802.11a amendment with some modifications, as shown in Table I, which are intended to increase the robustness of the transmission in highly dispersive vehicular environments. Communication takes place in the 5.9 GHz band, which is also known as Dedicated Short-Range Communication

Parameter	802.11p	802.11a
Frequency band	5.9 GHz	5.2 GHz
Channel bandwidth	10 MHz	20 MHz
Subcarrier spacing	156.25 kHz	312.5 kHz
Data rates	3 to 27 Mbit/s	6 to 54 Mbit/s
Slot/SIFS/DIFS time	13/32/58 μ s	9/16/34 μ s
Preamble duration	32 μ s	16 μ s
PLCP header length	8 μ s	4 μ s
Symbol time	8 μ s	4 μ s

TABLE I
COMPARISON OF 802.11P AND 802.11A PHY AND MAC PARAMETERS.

(DSRC) band. The latter is divided into one control channel (CCH) and a variable country-specific number of service channels (SCH). For instance, six and four SCHs are defined in the US [13] and in Europe [14], respectively. Each channel is split into 64 OFDM subcarriers, of which 48 are used for transmitting data, while 4 are used for time/frequency synchronization and channel estimation (pilot subcarriers). The remaining 12 subcarriers are disabled to accommodate the frequency guard bands of the signal.

Beacon frames are broadcast over the CCH to advertise services offered on specific SCHs. In addition, the CCH is used for the transmission of safety messages. The access to CCH/SCH is done in a FDMA/TDMA fashion as specified by the *Multi-Channel Operations* of IEEE 1609.4 [12]. The latter mandates an equally distributed access time between CCH and SCH. Specifically, out of every 100 ms, 50 ms are allocated for transmitting and receiving on the CCH, while the remaining 50 ms are used for communication on the SCH. Hence, the overall network performance strongly depends on a correctly functioning CCH.

Every 802.11p packet consists of a preamble and PLCP, MAC, and WSMP (WAVE Short Message Protocol) headers, as illustrated in Figure 1. This control information is followed by a variable amount of payload data. 802.11p supports different modulation and coding combinations. Out of these different combinations, the PLCP header is transmitted using the most robust one, while MAC and WSMP headers are transmitted with the same modulation and coding as used for the payload. In general, the use of a robust modulation and coding combination is recommended to increase reliability at the cost of lower transmission rates [4]. The preamble consists of ten short training symbols, a guard interval, and two long training symbols and has a total duration of 32 μ s. The format of the preamble is shown in Figure 2. During the short training phase, a known bit sequence is transmitted with a periodicity of 1.6 μ s over 12 subcarriers evenly distributed over the bandwidth. This information is exploited to detect the signal, calibrate the automatic gain control (AGC), perform coarse frequency offset estimation, and synchronize the clocks of sender and receiver. The two long training symbols employ 52




Fig. 1. Default payload frame format in 802.11p.




Fig. 2. Format and timing of a default 802.11p preamble and PLCP header.

subcarriers to transmit a different bit sequence that is used to estimate the channel state and perform a fine frequency offset estimation and a fine time synchronization. The preamble is followed by the PLCP header which contains information about the modulation used to transmit MAC header and payload (*RATE* field). It further tells the receiver how long the remaining transmission is going to last (*LENGTH* field). A parity bit is also available to detect errors in the previous PLCP header fields. If the parity bit indicates that the PLCP header is free of errors, the receiver continues decoding user data for the time indicated in the *LENGTH* field.

The medium access in 802.11p is organized according to the standard CSMA/CA protocol. If a node wants to transmit a packet, it must first sense the medium idle for a certain time interval. The medium is considered idle if the detected energy is below the carrier sense threshold. Although the value of this threshold is not standardized, it is typically set equal to the receiver sensitivity. Some commercial 802.11 devices do not consider the medium busy if they cannot detect a legitimate signal [9], [15], which is the case with the 802.11p devices that were used in our experiments. This reported behavior does not conform with the 802.11 standard, as the latter indicates that the carrier sense threshold should be increased by 20 dB in the cases where the preamble portion has not been detected (see Clause 18.3.10.6 in [16]).

B. Measurement Equipment: IEEE 802.11p Device

We use NEC Linkbird 802.11p [17] devices as sender and receiver in all our experiments. These devices are the result of several vehicular communication projects funded by the European Commission and the German Federal Ministry of Education and Research [18]–[20]. The Linkbird devices are reference implementations of the WAVE standard and have been widely used in VANET experiments [21], [22]. The network interface cards of the devices feature an Atheros chipset (AR 512) with Hardware Abstraction Layer (HAL) version 0.9.17.1.

1) *Details of the Packet Reception Process:* Packet reception in current WLAN devices is a complex process that con-




Fig. 3. Typical analog reception chain in 802.11 devices [23].




Fig. 4. Automatic gain control logic blocks [23].

sists of various concatenated steps, namely automatic gain control, signal detection, time and frequency synchronization, and signal demodulation. In the following, we provide a detailed description of these steps in a commodity (Atheros) device. This information is relevant for understanding the performance impairments caused by interference signals.

a) *Automatic Gain Control (AGC):* Any signal that reaches the receive antenna (either noise, interference, or user information) enters the reception chain as shown in Figure 3. The radio signal is first amplified, then mixed down to intermediate frequency (IF) and amplified again. The signal is then split into complex components, which are mixed down to baseband (BB) where they are low-pass filtered and amplified. Both quadrature (Q) and in-phase (IP) components are digitized by the corresponding analog-to-digital (A/D) converters. Next, a power detector estimates the power after A/D conversion. This computation is completed rapidly (within one short training symbol [23]) and, if necessary, used to perform coarse gain adaptations in the analog domain, see Figure 4.

Although the dynamic range of the A/D converters is typically large (e.g., 70 dB [24]), the power of incoming 802.11 signals may span over an even larger range. Hence, additional adjustment steps are needed to bring the signal strength within the preferred range of the A/D converter, which is limited by the so-called *coarse-high* and *coarse-low* thresholds [23]. This range is significantly smaller than the full A/D converter range to avoid two potential problems: First, that the incoming power exceeds the upper A/D converter threshold leading to the distortion of the signal. Second, that the incoming power is not sufficiently above the lower A/D converter threshold, which reduces the reception quality due to a high quantization noise. If the detected power exceeds the coarse-high threshold, the analog gain is

significantly reduced to bring the signal back in range (e.g., by 17 dB [23]). Similarly, if the signal power falls below the coarse-low threshold additional amplification is triggered. In the particular case that the incoming power saturates the A/D converter often enough, a quick and aggressive analog gain reduction (e.g., 30 dB [23]) is performed to bring the signal back in range. The value for the coarse-high threshold is said to be between -70 and -60 dBm in [25] and between -65 and -61 dBm in [24]. Although we cannot confirm these values, the measurements presented in Section III suggest that our 802.11p devices use a threshold within the latter range.

b) Packet Detection: Once the incoming signal is within the preferred input range at the A/D converter, the receiver tries to detect the presence of a signal of interest. First, it has to be determined if the incoming signal is in-band or if it has been transmitted on a different channel and, hence, it is not intended for the receiver. This is done by comparing the digitized power at the output of the A/D converter with the power measured after low-pass and decimation filter [23], see Figure 4. If the signal is determined to be in-band, it can be detected by means of two independent methods, which are correspondingly triggered depending on the power of the signal.

Strong signal detection is triggered by the sudden increase in signal power that forces the reduction of analog gain as previously described. **Weak signal detection** is triggered to identify low-power signals that are in the range of the background noise power. The receiver looks for sequences with a periodicity of one short training symbol and compares the normalized self-correlation of a received sequence against a threshold. If the threshold is exceeded and a sudden increase in in-band power has been detected, the presence of a signal of interest is assumed. The packet detection phase is completed after a number of fine gain adaptations so that the signal is placed at the preferred level of the A/D converter. Finally, the selected gain is kept fixed for the remainder of the incoming signal.

c) Synchronization and Demodulation: Prior to extracting the payload, time and frequency synchronization are performed using the short and the long training sequences. Next, the receiver extracts information from the PLCP header about the modulation used for transmitting the payload and the total length of the remaining packet. The PLCP header further contains a parity bit used to detect errors in the header itself. If no errors are detected, the receiver demodulates payload bits for the specified time. As our jamming signal is an OFDM signal carrying modulated random bits (explained later in this section), the receiver extracts bits from the signal that do not carry meaningful information. Depending on the value of the bit decoded at the position that corresponds to the (expected) parity bit and on the value of the previously decoded bits, the receiver may declare the presence of non-valid or a valid OFDM transmission. In the latter case, the receiver can be kept busy for an unknown time.

2) Interference-related Adaptation Schemes for Packet Detection: Interference signals can block successful transmission and reception in WLAN. They refrain the transmitter

from accessing the medium, hamper an accurate amplifier gain configuration resulting in a reduction of signal quality or distortion of the signal, and impair signal detection. For instance, this latter issue is addressed by the Atheros proprietary *ambient noise immunity* (ANI) technique [26]. An interfering signal can trigger strong or weak signal detection as already described. In both cases, the receiver has means to determine the presence of interference (e.g., by detecting framing errors based on the parity bit of the PLCP header or by observing a low self-correlation value) and stop the reception process. However, by locking onto the interference signal, the receiver may miss the arrival of a legitimate packet. The ANI mechanism increases the robustness to this problem as follows: If the rate of *false packet detections* exceeds a certain threshold, the immunity is increased by reducing the sensitivity of the receiver. If, after the adaptation, the rate of false packet detections falls below a different threshold, the sensitivity can be progressively increased [26]. On the contrary, if the highest immunity level does not prevent the rate of false packet detections from exceeding the latter threshold, the OFDM weak signal detection block is disabled. Next, if the rate of false packet detections falls below the threshold, weak signal detection is enabled again, but a higher in-band power is mandated for triggering signal detection. Note that switching off the weak signal detection scheme may cause the receiver to ignore (weak) legitimate transmissions resulting in performance degradations [15].

3) Noise Floor Measurement and Adaptation: Atheros chipsets use the measured noise floor of the circuits as reference value to perform accurate measurements of the absolute signal strength. The noise floor present at the A/D converters consists of the thermal noise at the antenna plus the noise of the RF front end (or noise figure) [23]. The latter component is very stable against temperature changes and is measured during the AGC calibration phase, where the receiver circuits are isolated from the antenna. The periodicity of this calibration can be indicated in software and typically corresponds to 30 seconds [25]. Next, during operation, the receiver measures the environmental noise. Specifically, noise measurements are performed over short time spans during idle time, i.e., while the device is neither transmitting nor receiving any signal. To obtain the absolute power value at the antenna, the analog gain is subtracted from the power observed after A/D conversion and normalized based on the calibrated noise power. Then, when a signal of interest arrives at the receiver, the power observed at the A/D output is corrected by subtracting the analog gain and the measured noise floor. The signal power is reported via the received signal strength indicator (RSSI), which in Atheros chipsets is an expression of the signal-to-interference-and-noise ratio (SINR). The RSSI is expressed in dB (relative to the noise power) and obtained once per packet based on the power measurement performed after the AGC settling time.

C. Jamming Equipment and Profiles

There are different ways of implementing a jammer for 802.11 networks. If protocol-compliant jamming is studied, the typical approach is to use off-the-shelf network interface cards and tune protocol parameters accordingly, for example, by disabling CSMA/CA carrier sense or back-off deferrals [27], [28]. Alternatively, one can also study the impact of jamming devices that do not comply with the 802.11 standard, which is the approach taken in this work. We use Wireless Open-Access Research Platform (WARP) for implementing different *jamming profiles* [29]. The WARP boards are software-defined radios where the physical layer processing is realized on FPGA, while the higher layer processing is performed on an integrated PowerPC core. The WARP board provides an 802.11-like OFDM physical layer with a 10 MHz bandwidth. The transceiver of the boards [30] supports transmission at frequencies up to 5.875 GHz, hence, covering two 802.11p channels [3] (i.e., channels 172 and 174). The board is designed to provide a maximal output power of 18 dBm in the 2.4 GHz band. As the transmit power in the 5.9 GHz band is not specified, we measured it with a spectrum analyzer and found it to be 16.75 dBm. The OFDM signals transmitted in the default WARP implementation do not comply with the 802.11p standard as the PLCP header fields, among others, differ from the specifications. Furthermore, the transmitted WARP signals do not block medium access. It has been reported [9], [15] that certain commercial 802.11 devices (including our Atheros cards) consider the medium as busy only if they detect a standard compliant transmission. As already discussed, this design does not comply with the 802.11 standard.

Based on the WARP boards, we implement three different *jamming profiles*, namely **periodic**, **constant**, and **reactive**.

1) *Periodic Jammer*: The periodic jamming signal is characterized by a $64\mu\text{s}$ *ON* phase and a $10\mu\text{s}$ *OFF* phase, see Figure 5. The frame format of the periodic jammer consists of a basic zero-payload WARP frame as illustrated in Figure 6. The PLCP and MAC headers are both transmitted with a QPSK modulation.

2) *Constant Jammer*: The constant jammer is intended to be continuously transmitting a signal. Realizing this with the WARP boards is, however, not entirely possible. The amount of time that the device can be transmitting is upper-bounded. We determined this time with an oscilloscope to be 2.71 ms. Furthermore, there is an unavoidable $10\mu\text{s}$ idle gap between two consecutive transmissions required by the hardware to set up a new transmission. Still, this results in a jammer with very long active periods such that, in the following, we consider it as a constant jammer. To transmit the signal, random payload is generated to make up for the aforementioned 2.71 ms of transmission time. The random payload and the headers are transmitted with a BPSK modulation. Figure 5 illustrates the constant and periodic jamming profiles in the time domain compared to a regular 802.11p frame. Note that the time relations in the figures are not to scale.




Fig. 5. Constant and periodic jammer profiles in the time domain compared with a default 802.11p transmission.




Fig. 6. Frame format of the jamming signals generated by the WARP device. Short interference signal with a $64\mu\text{s}$ duration (Fig. 6(a)) for the periodic jammer, and long interference signal with a 2.71 ms duration (Fig. 6(b)) for the constant jammer.

3) *Reactive Jammer*: The reactive jammer is designed to start transmitting upon sensing energy above a certain threshold. The default OFDM design of the WARP platform features an energy detection block, which compares the instantaneous energy measured at the receive antenna with a threshold. We set the latter to -75 dBm as we empirically determined it to be a good trade-off between jammer sensitivity and false transmission detection rate¹. If the detected energy exceeds the threshold during a certain time span ($\mathcal{T}_{\text{detect}}$), an ongoing 802.11p transmission is assumed by the jammer. We set $\mathcal{T}_{\text{detect}}$ to $2\mu\text{s}$ to avoid reacting to sporadic noise power peaks. Then the WARP device has to switch from idle to transmit mode, which introduces a delay of $10\mu\text{s}$. Hence, the minimum reaction delay corresponds to $12\mu\text{s}$ and it can be increased in microsecond granularity. In comparison to previous reactive jammers [31], our design features the shortest reaction time. By appropriately tuning the reaction delay, the duration of the jamming signal ($\mathcal{T}_{\text{signal}}$), or by adding *sleep time phases* several reactive jamming patterns can be obtained, most of which are illustrated in Figure 7. In this work, we consider a reactive jammer that emits a single signal per (detected) 802.11p frame. The reactive jam-

¹Setting the detection threshold closer to the sensitivity of the device would have lead to the reactive jammer being constantly triggered by background noise, other sources of interference, and delayed copies of the jammer signal itself. We observed this behavior in outdoor measurements and increased the threshold to avoid the consideration of a reactive jammer that was practically acting as a constant one.




Fig. 7. Reactive jammer profiles in the time domain compared with a default 802.11p transmission.




Fig. 8. Point in time where the different reactive jammer configurations start and end their attacks.


ming signal is basically characterized by the tuple $(\mathcal{T}_{\text{reaction}}, \mathcal{T}_{\text{signal}})$. The different configurations result in different start and end points of the attacks as illustrated in Figure 8.

III. INDOOR EVALUATION


Before bringing our equipment to the field, we performed a set of measurements in an indoor environment to characterize the 802.11p devices. For this purpose we used a 30 m² big anechoic chamber at the Faculty of Engineering of the University of Porto (FEUP), which provides a multipath- and interference-free environment. The anechoic chamber is shown in Figure 9(a). In the following, we first give a detailed overview of the experimental design for these indoor measurements and later present our results.

A. Indoor Experimental Setup

We characterize the receiver response in terms of the average packet delivery rate (PDR) in the presence of various jamming profiles. Inside the chamber, we placed the antenna of the transmitting 802.11p device on a pole. The jammer antenna was placed on a second pole, while the antenna of the receiving device was located on a box at a similar height. The actual devices were placed outside the chamber and connected to the antennas via cables. Next, we started with the transmission of payload packets from the transmitter to the receiver. These packets had a size of 100 Byte and were transmitted with the QPSK modulation at a rate of 6 Mbit/s. Overall, one such packet occupies the channel for 232 μ s. These packets were generated at a rate of 100 packets per second. Hence, the wireless medium was idle most of the time with respect to 802.11p transmissions. All received packets were recorded by an application and the resulting traces were used to compute PDR and the reported RSSI from the receiver. Every measurement point was computed as the average PDR across a series of 10^4 packets. For more details on the setup of transmitter, receiver, and jammer,



(a) Experiment setup in the anechoic chamber.



(b) Path loss attenuation and distance between nodes.

Fig. 9. Indoor measurement environment.

we refer to Table II. We focused on the topology shown in Figure 9(b), and we additionally varied the attenuation of the jamming signal through the addition of RF attenuators. The signal and interference power at the receiver, as well as the resulting SINR values, are shown in Table III.

B. Calibration

Our 802.11p devices report an RSSI value for each successfully decoded packet. This RSSI value is designed to indicate the strength of the received signal in comparison to the anticipated noise floor, see Section II-B3. To simplify the analysis of the RSSI, especially when deducting the average received power strength from it, we initially performed a set of calibration measurements. We describe here the procedure used to obtain (1) the relationship between the transmit power set in software to the actual power of the transmitted signal and (2) the mapping between RSSI (in dB) and the actual received power (in dBm) above the noise floor.

1) *Calibration of the Transmit Power:* We carried out this calibration by connecting the transmitter to a spectrum analyzer through a coaxial cable with 2dB attenuation. Because we were measuring a pulsed modulated OFDM signal, the energy of each subcarrier varies from symbol to symbol. It is known that measuring the power of an OFDM signal is a challenging task [32]. Figure 11 shows a transmitted 802.11p frame in the time domain as captured by the spectrum analyzer. The figure illustrates the high variance of power levels within the OFDM transmission. We used the *maxhold* function of the spectrum analyzer to obtain the average power over the signal bandwidth across a sequence of 10^3 packets transmitted with QPSK modulation

802.11 Device	Data rate	6 Mbit/s (QPSK 1/2)
	P_{Tx}	17.48 dBm
	Payload length	100 Byte (232 μ s)
	Packet generation rate	100 [packets/s]
Constant Jammer	ON-phase / OFF-phase	2.71 ms / 10 μ s
Periodic Jammer	ON-phase / OFF-phase	64 μ s / 10 μ s
Reactive Jammer		(12 μ s, 47 μ s)
		(12 μ s, 64 μ s)
		(12 μ s, 500 μ s)
		(16 μ s, 500 μ s)
		(40 μ s, 500 μ s)
	Energy detection threshold	-75 dBm
General	Center frequency	5.860 GHz (Ch. 172)
	Jammer Power	16.75 dBm

TABLE II
SETUP PARAMETERS USED IN ALL OUR INDOOR EXPERIMENTS.


P_{Tx-Rx} [dBm]	P_{Tx-J} [dBm]	Att. [dB]	$P_{J-Rx_{const.}}$ [dBm]	$P_{J-Rx_{react.}}$ [dBm]	SINR _{const.} [dB]	SINR _{react.} [dB]
-38.9	-103.3	64	-94	—	55	57
-38.9	-95.3	56	-86	—	47	57
-38.9	-83.3	44	-74	—	35	57
-38.9	-73.3	34	-64	-64	25	25
-38.9	-63.3	24	-54	-54	15	15
-38.9	-58.3	19	-49	-49	10	10

TABLE III
POWER AND ATTENUATION SETTINGS IN ANECHOIC CHAMBER


that were carrying random payload. This method is considered in [32] as one of the most accurate approaches for measuring the power of an OFDM signal. In Figure 10(a) we show the resulting relationship obtained from the calibration measurement. Every point in the graph corresponds to a single value obtained using the previous method. We observe slight deviations from a linear relationship at the lowest and highest power values, while a linear behavior can be observed for the range from 9 to 16 dBm. When driven at full gain the measured power corresponds to 17.58 dBm, which is the configuration used in all our experiments. Note that the maximum power reported in the datasheet of the device is 21 dBm. Based on the measured values for the transmit power, we then proceeded to determine the mapping between RSSI and signal-to-interference-and-noise ratio (SINR).

2) *Calibration of the RSSI:* We connected transmitter and receiver via a coaxial cable (2 dB attenuation) and increased the attenuation between them step-wise by adding passive attenuator elements. This allowed us to precisely estimate the received signal power. Simultaneously, we recorded the average RSSI for a sequence of 10^4 packets. Figure 10(b)

shows the data samples obtained in the experiments and a linear and cubic fitting model for that data. Due to the lower complexity and high accuracy of the linear model, we use it for mapping RSSI samples σ to received power values γ (in dBm): $\gamma = 0.856 \cdot \sigma - 86.35$. By assuming the noise power equal to the receiver sensitivity (i.e., -86 dBm) and adding it to the estimated received power, we convert all RSSI values to SINR throughout the experiments. With respect to the receiver sensitivity, we never obtained RSSI values lower than -86 dBm after the above conversion on any received packet. Hence, we consider the receiver sensitivity to be -86 dBm. This is in accordance with discussions in the Madwifi mailing list [33] stating that the freeware version of the driver sets the card sensitivity to 10 dB higher than the commercial version. However, we suspect that the lack of RSSI reports below -86 dBm are in fact a limitation of the driver and not of the hardware itself. In Subsection III-D we conduct measurements that suggest that the hardware provides a sensitivity of -96 dBm, while the driver reports could be erroneous by 10 dB.



(a) Mapping the transmit power value set in software to the measured power.



(b) Mapping RSSI to received power.

Fig. 10. Fig. 10(a) shows measured transmit power values for various transmit power settings on the considered 802.11p devices. Fig. 10(b) shows RSSI values and the corresponding (measured) received power. We compare a linear and a third degree polynomial (i.e., cubic) fitting. The first yields an RMS error of 1.96, while the latter slightly reduces the error to 1.79. Due to the lower complexity of the linear model and the comparable accuracy, we select the least square linear model over more complex fitting models. The chosen model is given by: $P_{rx} = 0.856 \cdot \text{RSSI} - 86.35$ in dBm. Assuming a noise floor of about -86 dBm, the corresponding SINR can be obtained as: $\text{SINR} = 0.856 \cdot \text{RSSI}$.




Fig. 11. Transmitted 802.11p frame captured by the spectrum analyzer. The representation in the time domain shows that the preamble features stable power levels, while the OFDM signal (i.e., PLCP, MAC, and WSMP headers, as well as payload) exhibits irregular power patterns.

C. Measurement Results - Reactive Jammer

In the following, we present measurement results regarding the impact of reactive jamming on the performance of 802.11p communications. For that, we show (in Figure 12) the packet delivery ratio (PDR) for different SINR values, which are achieved by adding/removing passive attenuators to/from the RF output of the jammer device. Every depicted point corresponds to the average PDR value across a sequence of 10^4 packets. In addition, we compute the 95%

confidence intervals, but do not show them in the figures as they are below 1%. Small fluctuations of PDR and RSSI within a measurement are shown in Figure 13(a), where every depicted point corresponds to the average performance over 50 packet transmissions. As highlighted in Table III, a higher attenuation results in a lower interference power P_{J-Rx} and, thus, in a higher SINR and vice versa. On the other hand, the received signal strength of 802.11p data packets P_{T-X-Rx} was kept constant for all measured points. We considered several reactive jamming strategies, which are characterized by different reaction delays and signal durations as illustrated in Figure 7. Note that the reactive jammer is triggered by legitimate transmissions only if the sensed power is above -75 dBm, which happens when the applied attenuation is below 36 dB, see Table III. In that case, the jammer is not active and the legitimate communication is only disturbed by noise, which results in an SINR of 57 dB.

a) Impact of Reaction Delay: Figure 12(a) shows the PDR performance under the impact of the reactive jammer. In particular, we consider reaction delays of 12, 16, and 40 μs and the interference signal has a fixed length of 500 μs . Recall that this jamming configuration sends a new signal when energy is sensed, however, only once per detected 802.11p packet. For comparison, we also show the performance that is achieved when the jammer is disabled and observe that the communication is significantly impaired by all reactive jammer patterns. Specifically, 10 dB stronger signals are required to fully overcome the jammer. It can be further observed that a slight increase in reaction delay results in a significantly lower jammer effectiveness. For instance, the PDR improves by up to 2 dB when the jammer requires 16 μs to react instead of 12 μs . These additional 4 μs are particularly important as they correspond to the short training symbols (t_8-t_{10}) used for the coarse correction of frequency and time offsets, see Figure 8. Finally, a jammer that reacts with a delay of 40 μs misses the preamble and the PLCP header and has an up to 3.5 dB lower effectiveness.

b) Impact of Interference Duration: In Figure 12(b) we study the impact of the interference duration. For this, we consider the reactive jammer with 12 μs delay and three different durations, namely 47, 64, and 500 μs . As expected, the longer the interference, the more damage is inflicted on the communication success. However, reducing the interference signal from 500 to 64 μs has a negligible impact on the jammer effectiveness. Since the jammer already interferes the control information, only marginal gain is obtained from further interfering the payload part. However, an interference duration of 47 μs partially misses the MAC header and leads to a slightly lower effectiveness. In the selected configuration and for all the considered patterns, the impact of the reactive jammer degrades the 802.11p performance by at least 10 dB. When the signal of interest is 17 dB stronger than the reactive interference, the impact of the latter can be neglected. We conducted further measurements that consistently confirmed a 100% PDR under higher SINR conditions.




Fig. 12. Characterizing the impact of reaction delay and interference signal length on the effectiveness of the jammer. The reaction delay is a very important factor, as we observe that few microseconds larger delays result in a significantly lower jamming effectiveness. In contrast, the length of the interference signal has a limited impact on the performance. Important for the jammer is to interfere the preamble and relevant control information, e.g., PLCP and MAC headers. Prolonging the interference signal to also overlap with the payload does not increase the jammer effectiveness.

D. Measurement Results - Constant and Periodic Jammer

In contrast to the reactive jammer, constant and periodic interference signals do, most of the time, not overlap with the signals of interest and can be detected by the receiver as signals that do not comply with the standard. As introduced in Section II-B2, commodity WLAN devices are usually equipped with proprietary techniques to fight interference and can take advantage of identifying the presence of RF jamming signals. These interference mitigation algorithms require, however, a certain amount of time to converge. Therefore, the resulting performance in the presence of a jammer largely depends on whether the algorithms are still iterating or have already converged, which is highlighted by Figure 13(b). We differentiate between these two states and show the average PDR performance obtained during the initial transient and during the steady state, respectively. We define the initial transient as the time required by the receiver to bring the measured PDR to a stable value, which we observed to be typically in the range of 10 to 30 seconds. Correspondingly, we define the steady state as the time span where no appreciable PDR fluctuations are measured.

Initial Transient: Figure 14(a) shows the obtained PDR during the initial transient in the presence of the constant and periodic jammers. For comparison we also show the




Fig. 13. Time evolution of the PDR and RSSI values measured in the anechoic chamber.

performance obtained when the jammer is switched off and in the presence of the reactive jammer ($40\mu\text{s}$, $500\mu\text{s}$). The results show the average PDR obtained over 10^4 packet transmissions. Again, the 95% confidence intervals were below 1% for most of the considered points². The small fluctuations of the PDR in the time domain are illustrated in Figure 13(b). In the low SINR range, both the constant and periodic jammer are more effective than the reactive jammer by 3 and 7 dB, respectively. For higher SINR values, the constant and periodic jammer reduce the PDR to 0% after the nodes have reached a performance peak at 17 and 21 dB, respectively. We later discuss the reasons for this initially striking observation that the PDR decreases as the signal strength increases. The SINR range over which successful communication is completely blocked spans 18 dB in both cases. In the following we refer to this situation of consistently having a PDR of 0% under good signal conditions as *blackout phase*. Perfect communication (i.e., 100% PDR) is achieved at an SINR of 40 dB in the case of the constant jammer. The periodic jammer is significantly more effective and it blocks perfect communication up to 56 dB SINR.

Steady State: Once the algorithms for interference mitigation have converged, the resulting PDR changes significantly, as shown in Figure 14(b), which confirms the benefits of Atheros proprietary algorithms against interference. For instance, in the presence of the constant jammer the blackout phase completely disappears and the overall jammer effectiveness is comparable to the one of the reactive jammer.

²The only exception in the initial transient phase corresponds to the PDR measured under the periodic jammer at 46 dB SINR, where we obtained a confidence interval of 3.54%.




Fig. 14. PDR performance in the presence of constant and periodic jammers during initial transient and steady state.

These results indicate that the interference mitigation algorithms allow the receiver to synchronize to incoming signals of interest and avoid being triggered by constant jamming signals. In the presence of the periodic jammer, the PDR increases remarkably in the low and mid SINR range, but the blackout phase covers the same SINR range as previously observed. Satisfactory performance is achieved at an SINR of 45 dB and perfect communication only when the SINR reaches 56 dB. Again, the 95% confidence intervals were below 1% for most of the considered points³.

Blackout Phase: In the following, we elaborate on the rationale behind the observed blackout phase, focusing mainly on the steady state. In Figure 14(b), at an SINR of 21 dB the interference signal reaches the receiver with a power of -60 dBm, cf. Table III. This value is above the coarse high threshold and forces the receiver to perform a quick gain drop, as discussed in Section II-B1. This gain correction allows subsequent signals of interest (with a power of -38.9 dBm) to arrive within the range of the A/D converter without causing an overflow of the dynamic range [9]. If the jammer is further attenuated (i.e., for SINR values larger than 21 dB), then the interference signal reaches the receiver with a power that is below the coarse-high threshold, hence, no gain correction is performed. Next, if a signal of

interest arrives while the receiver is busy trying to decode the interference signal, the gain cannot be set appropriately and timely, which results in an overflow at the input of the A/D converter. This causes the observed decrease in PDR for an increasing SINR, which is highlighted by Figure 14(b), specifically in the SINR range from 21 dB to 44 dB.

The effectiveness of an RF jammer, especially in the presence of interference-mitigation techniques, highly depends on its ability to keep the receiver busy. In this respect, there are major differences between the constant and the periodic jammer considered in this work. This is mainly due to the weak signal detection procedure described in Section II-B1. Specifically, the constant jammer can trigger weak signal detection once per interference signal due to the initial preamble. Afterwards, weak signal detection is no longer triggered, since the random content carried by the interference packet is unlikely to self-correlate. Hence, the receiver is idle and can suitably perform AGC and synchronize with subsequent signals of interest, resulting in a satisfactory performance, see Figure 14(b). In contrast, the periodic interference transmits preambles about 50% of the time, thereby triggering weak signal detection at the receiver. Subsequent signals of interest induce a sudden increase in received power, so that the receiver tries to identify a valid signal. However, this happens without a prior adaptation of the gain. Therefore, the high incoming power of the legitimate signal falls outside the preferred A/D dynamic range and the receiver cannot extract the user information. This leads to the initially counterintuitive behaviour of observing a decreasing PDR for an increasing SINR as shown in Figure 14(b).

Finally, a PDR of 100% is reached at an SINR of 56 dB, that is, when the interference power is close to -96 dBm. At this point, the incoming power of the jamming signal does no longer trigger weak signal detection nor does it degrade the signal quality. Based on these results, we suspect that the sensitivity of our 802.11p devices corresponds to -96 dBm, as mentioned in Section III-B2.

Final Observations: There are three major observations that can be concluded from the results so far:

- The authors in [9] observed that an interference signal 20-30 dB weaker than the signal of interest can effectively disturb 802.11 transmissions. In our measurements we find that a periodic signal with a very high *on/off* rate and a preamble structure, can seriously hamper an 802.11p transmission despite being up to 56 dB weaker than the legitimate signal. Extrapolated to an outdoor environment, this means that a jammer located significantly far away from two communicating vehicles (that are close to each other) can still be very effective. We confirm this hypothesis in Section V.
- In the steady state, a periodic jammer can achieve a significantly better performance than a constant jammer, while the latter has a similar effectiveness as the reactive jammer. However, during the initial transient, constant and periodic jammers block communication in

³The only exceptions correspond to the PDR measured under the constant jammer at 11 dB SINR, which has confidence intervals of 1.38%. Furthermore, under the periodic jammer at 17 dB and 46 dB SINR we obtain confidence intervals of 3.29% and 3.54%, respectively.

a similar way. The performance of 802.11p devices in the presence of a jammer is particularly vulnerable to preamble-like structures that reach high levels of self-correlation and trigger (at least) weak signal detection events at the receiver.

- State-of-the-art 802.11 devices perform a set of adaptation techniques in the presence of non-reactive interference that have the potential of improving the communication performance. Reducing the convergence time of such techniques could increase their applicability for 802.11p vehicular networks by better adapting to highly dynamic vehicular environments.

IV. MODELLING THE IMPACT OF RF JAMMING IN A GENERIC VANET SCENARIO

In the previous section, we measured and analyzed the performance of 802.11p hardware in the presence of RF jamming in an anechoic chamber void of multipath propagation effects and influence of any external RF interference. We obtained a relationship between the PDR and the SINR in the presence of various jammer types. The results are VANET-specific in the sense that the hardware used is a reference implementation of the 802.11p amendment. In this section, we propose a method that uses those results as a model to study the impact of RF jamming on VANET communications in a generic setting. The method takes as input the positions of transmitter, receiver, and jammer, propagation models for the considered environment, and the PDR versus SINR results. The different steps are described as follows:

- 1) Obtain the distance from transmitter and jammer to the receiver.
- 2) Calculate the received power from the transmitter signal $P_{T,Rx}$ and from the jammer signal $P_{J,Rx}$,

$$P_{X,Rx} = \frac{P_{X,Tx} \cdot G_{X,Tx} \cdot G_{X,Rx}}{a_X}, \text{ for } X = \{T, J\}, \quad (1)$$

where $G_{X,Tx}$ and $G_{X,Rx}$ are the transmitter and receiver antenna gains, the attenuation a_X is obtained using propagation models suitable for the environment under consideration, and $P_{X,Tx}$ is the transmitted power. $P_{X,Rx}$ and $P_{X,Tx}$ are given in mW.

- 3) Calculate the SINR as $10 \log_{10} \left(\frac{P_{T,Rx}}{P_{J,Rx} + N} \right)$ [dB], where N is the noise floor. The variables are given in mW.
- 4) Finally, obtain the corresponding PDR at the calculated SINR using the indoor results.

In the remainder of this section we validate this method through outdoor measurements.

A. Outdoor Scenario Characterization

The scenario selected for the validation measurements is a rural area located in the periphery of Aachen in Germany, see Figure 15. This scenario provides two perpendicular roads. A *main road* that has a length of 600 m and a 120 m long *side*



Fig. 15. Satellite view of the scenario. In the basic topology, the jammer and the transmitter are static at the indicated positions, while the receiver is placed at different positions along the main road. The different positions of the receiver result in a varying received power from both transmitter and jammer. The terrain features a certain degree of inclination. The graph provides information about the ground height at three specific positions.

road. The line-of-sight along the main road (about 300 m) is shorter than the total length of the latter due to a slight descending slope of the road. Furthermore, the area between main road and side road exhibits a moderate elevation of the ground that can block line-of-sight along the diagonal path between both roads. The amount of traffic in the area is negligible with only few sporadic cars.

This open space scenario offers a low dispersive propagation environment, as there are no obstacles or buildings between or surrounding the vehicles. Therefore, shadowing, scattering, and reflection of the signals are expected to have a negligible impact on the performance. Hence, we select the Friis path loss model [34] to be an adequate abstraction of the propagation environment and use it to model the signal attenuations in Equation 1. The model assumes a logarithmic decay of the received power with increasing propagation distance. Due to the existing terrain asymmetries in the considered scenario, we characterize the path loss attenuation individually along the main road, side road, and along the diagonal connecting both roads. For estimating the path loss attenuation along the main and side road, the transmitter is kept static at the crossroad and the receiver moves along the corresponding road, respectively. Similarly, for estimating the attenuation along the diagonal path, the transmitter is static at the end of the side street (about 120 m away from the crossroad), while the receiver moves along the main road. To obtain samples of the received signal strength, we let transmitter and receiver communicate with the jammer switched off. In every measurement, the transmitter sends packets to the receiver at a rate of 100 packets per second for roughly 2 minutes. Transmission power, modulation, and packet format are the same as for the indoor measurements, cf. Table II. We compute the distances between the devices from their reported GPS positions, and use the signal strength values reported by the receiver and mapped according to the calibration results of Section III-B. Next, we obtain estimates of the model parameters by least squares fitting the measured distance and signal strength

Scenario	Path loss model	RMSE
Main road	$P_r = -2.026 \cdot 10 \log_{10}(d) - 26.10$	3.662
Side road	$P_r = -3.317 \cdot 10 \log_{10}(d) - 9.32$	0.906
Diagonal	$P_r = -5.263 \cdot 10 \log_{10}(d) + 47.47$	2.113

TABLE IV
PARAMETERIZATION OF PATH LOSS MODELS




Fig. 16. Measured and expected power as function of the distance between transmitter and receiver along the diagonal between main and side roads.

values to equation $P_r(d) = k - \alpha \cdot 10 \log_{10}(d)$ in dBm. Table IV shows the resulting path loss parameterization and the respective root mean square error (RMSE) obtained for each communication path. The parameterization for the diagonal path is out of the expectable range for such an environment, most likely due to the aforementioned terrain elevation. Nevertheless, Figure 16 shows that the model fits well the measured received power values. In the following we use the propagation models shown in Table IV to determine the expected received power of the transmitter and the jammer at the receiver.

B. Measurement Results

We conduct measurements in the previous scenario to validate the proposed method for modeling the impact of RF jamming on VANET communications. First, we describe the node topology and provide some details about the measurement methodology. Afterwards we present the results.

Setup and Methodology: We place the jammer at the end of the side street and the transmitter at the crossroad. The receiver is placed at different positions along the main road, as shown in Figure 15. The different node configurations result in varying SINR values. For every receiver position, the transmitter sends packets to the receiver at a rate of 100 packets per second over 2 minutes. This procedure is repeated for a representative subset of the available jamming patterns, namely constant, periodic, and two reactive jammers with a signal length of 500 μ s and reaction delays of 12 and 40 μ s, respectively. While processing the data, we differentiate between initial transient and steady state for the constant and periodic jammers. Notice that the results




Fig. 17. Real PDR as function of the SINR for the reactive jammers.

presented in this section refer exclusively to the performance obtained during the steady state. Each graph shows the measured PDR and the predicted PDR using the proposed model. For comparison we also show the PDR performance obtained when the jammer is switched off.

Reactive Jammer: Figures 17(a) and 17(b) show the PDR obtained in the presence of the reactive jammers. The figure shows the average PDR and the 95% confidence intervals. As opposed to the results obtained in the anechoic chamber, the PDR measured outdoors exhibited moderate fluctuations for a given SINR point due to the higher signal strength variability expected in an uncontrolled propagation environment. In both cases, the performance predicted by the model follows closely the PDR measured outdoors. The shape of both curves is similar, despite an SINR offset of 1-2 dB.

Constant and Periodic Jammer: Figure 18(a) shows the performance obtained in the presence of the constant jammer. The measured PDR follows the shape of the predicted graph, despite what seems to be an outlier at 11 dB SINR and a slight overestimation of the PDR for higher SINR values. Figure 18(b) confirms the presence of the blackout phase caused by the periodic jammer, since we consistently measured a PDR of 0% for SINR values below 57 dB.

C. Final Observations

In general, the presented results show that the performance of outdoor communication under jamming can be well approximated by applying the method proposed at the beginning of the section. This procedure empowers the community




Fig. 18. Real PDR as function of the SINR for the constant and periodic jammer.

to use the models to characterize VANET communications in the presence of a jammer and perform jamming-aware design of VANET applications and protocols.

While the validation and accuracy of our method has been exclusively evaluated in open space environments, we do not anticipate fundamental issues that limit its applicability in urban environments. We underline that the accuracy of the prediction is mostly dependent on the characterization of the propagation environment, that is, on an appropriate choice and calibration of the propagation model. However, in urban environments this is a challenging task, as shadowing and multi-path fading lead to signal strength fluctuations around the average value predicted by the path loss model. The stochastic nature of these processes can be hardly predicted and affects both legitimate and interference signals. Consequently, the instantaneous SINR fluctuates and so does the quality of the communication. Under highly variable conditions, our model can nevertheless be used to provide bounds of expectable performance given an educated guess regarding the magnitude of the signal strength fluctuations.

We have implemented and validated the model in ns3 [35] using the data from our previous work [8]. The simulation model for ns3 can be downloaded from [36] for experimentation and further development.

V. OUTDOOR PLATOONING EVALUATION

In this section, we evaluate the quality of the communication between two vehicles moving as a platoon in the presence of an RF jammer. The platooning movement is

characterized by a short and constant inter-vehicle distance and the same acceleration and speed across vehicles [37]. Platoons are intended to increase traffic efficiency, but they are vulnerable to jamming attacks. If the platoon is coordinated by a VANET connection, even short disruptions of the communication can have fatal consequences. For this study, we choose the same open space environment as in Section IV and an additional scenario, namely the large parking lot near the seafront in Porto used in our previous work [8]. For convenience, we name these scenarios **rural** and **seafront**. In the following, we provide a description of the environment and the setup for each scenario, and present the corresponding results.

A. Rural Environment

We place the jammer at the crossroads between the main and the side roads. The exact position is highlighted with a fire symbol in Figure 19. The communicating vehicles move at a constant speed of about 25 km/h along the main road, where they first approach the crossroad and later leave the jammed area. Note that, depending on the jamming signal, the vehicles start off at different positions. Specifically, in the presence of the constant and the periodic jammers the starting point is located far enough from the jammer to initially enable successful communication, while the required distance is shorter in the presence of the reactive jammer. Starting and ending points are highlighted in Figure 19 as well. In our measurements, the transmitter is closely followed by the receiver keeping an inter-vehicle distance of about 5 m. In this scenario, we evaluate the impact of constant, periodic, and two reactive jammers. Note that the relatively low moving speed is not expected to have a fundamental impact on the results, as mobility alone does not alter the characteristics of the propagation environment. However, higher speeds would affect the results, as both SINR and PDR would fluctuate over shorter time scales and so the transition phases in the performance would be steeper. From a geographical perspective though, the results are not expected to exhibit any substantial difference.

1) *Results*: Figure 19 illustrates the threat that an RF jammer (constant and reactive in this case) poses to a vehicular platoon. The areas over which communication is successful are highlighted in green⁴. Whenever the communication is *completely* disrupted (PDR equals 0%) there is no color being displayed. Figure 19(a) shows that the constant jammer effectively blocks the communication along the main road. Specifically, over a length of 465 m no single packet is successfully received. Normal operation is only possible when the vehicles are significantly far away from the jammer. Similarly, Figure 19(b) shows that the impact of the reactive jammer is significantly lower, since the blackout area spans *only* over a road segment of 70 m with the

⁴In fact, the color changes as function of the measured signal strength. As the transmission distance is almost constant, the measured signal strength fluctuates only slightly within certain bounds that are mapped to the green color.



(a) Constant jammer.




(b) Reactive ($12\mu\text{s}$, $500\mu\text{s}$) jammer.

Fig. 19. Satellite view of the rural scenario during the platooning measurements. The jammer is static at the indicated position, while transmitter and receiver move along the main road. Hence, the vehicles first approach the crossroad and then leave the jammed area. Successful communication is indicated by colored regions. The range of the blackout area is explicitly delimited with parentheses.


jammer at its center. The indoor measurements of Section III indicate that, once the adaptive mechanisms of the 802.11p devices have converged, both constant and reactive jammer have a comparable impact on the communication. Outdoors, however, we observe major differences in the effectiveness of these two jamming signals. Therefore, the constant jammer effectiveness obtained during the initial transient, see Figure 14(a), should be rather considered for predicting the achievable performance in the platoon configuration. This can be explained by the changes in topology and propagation conditions as the vehicles move, since in this situation of continuous change the 802.11p devices have difficulties setting their working point within the steady state.

In the following, we show the time evolution of the PDR in the presence of the considered jamming signals. In addition, we superpose the behavior of the *predicted* SINR, that is, the SINR computed based on the position of the vehicles and the jammer and the path loss model. Specifically, we consider the path loss model for the main road as defined in Section IV-A.

In Figure 20, it can be observed that the temporal behavior of PDR and SINR is similar for both reactive jammers. Only when the platoon is close to the jammer, a complete disruption of the communication occurs. The transition be-



(a) Platooning movement and reactive ($12\mu\text{s}$, $500\mu\text{s}$) jammer.




(b) Platooning movement and reactive ($40\mu\text{s}$, $500\mu\text{s}$) jammer.

Fig. 20. Platooning movement under the impact of the reactive jammer. The predicted PDR values are obtained by means of the model derived based on the indoor measurements.


tween a PDR of 100% and 0% (and vice versa) lasts for 5 s, which at the speed of travel corresponds to 35 m. The blackout area is in both cases of similar length, namely 10 s or 70 m. During the transition time a jamming detection strategy could extract and exploit important information from correctly received packets (RSSI, relative position of the nodes, PDR, among others) that could reveal the presence of a jammer as proposed, e.g., in [27]. This information may no longer be available when the PDR drops to 0% and jamming detection strategies may perform significantly worse or simply fail. The model used to map SINR to PDR shows a very good agreement with the actual measured performance. Despite the mobility of the nodes, the proposed approach can be applied to accurately model the impact of reactive jamming on VANET platooning in specific, and VANET communications in general.

Figure 21 shows that the blackout area is significantly larger around the crossroad in the presence of constant and periodic jammers compared to the reactive ones. Basically, both jammers exhibit a similar effectiveness by completely disrupting the communication over more than 450 m as shown in Figure 19(a)⁵. In both cases the transition time spans 15 s. The model introduced in Section IV to predict the resulting PDR is again used in Figure 21. In the presence of pro-active jammers, modeling the PDR in a scenario with mobility is considerably more complex than in the case of reactive jammers, which is reflected by a noticeable mismatch between predicted and measured PDR.

⁵Although the periodic jammer seems to disrupt the communication over a larger region than the constant jammer, this is mainly caused by a slightly lower speed of travel that prolongs the exposure time to the periodic jammer.



(a) Platooning movement and constant jammer.



(b) Platooning movement and periodic jammer.

Fig. 21. Platooning movement under the impact of the constant and periodic jammer. The predicted PDR values are obtained by means of the model derived based on the indoor measurements and differentiating between initial transient and steady state.


In Section III, we have shown that the 802.11p devices require a certain time (initial transient) before the adaptation mechanisms settle at a robust working point (steady state). In the indoor measurements we measured convergence times up to 30s. However, in outdoor scenarios with mobility it can be expected that the initial transient spans even larger periods. In the following, we show results for the predicted PDR based on both, initial transient and steady state values presented in Section III. In Figure 21(a) it can be observed that the initial phase provides a good match for the PDR performance obtained while the cars approach the jammer. When the cars leave the crossroads, that is, move away from the jammed area, the PDR prediction is improved by using the steady state values. Unfortunately, we are not able to determine the instant that separates initial phase from steady state. We believe that the mismatches between measured and predicted PDR are rather a consequence of the complex algorithms featured by the Atheros cards, than a fundamental problem with our modeling methodology. Figure 21(b) presents similar results obtained in the presence of the periodic jammer. Again, the performance during the initial phase shows an acceptable match for the measured PDR when the vehicles approach the jammer. When they move away, the steady state is a better option.

B. Seafront Environment

This scenario consists of a 500 m long parking lot near the seafront in Leça da Palmeira in the outskirts of Porto, see Figure 22(a). Due to the uniformity of the characteristics of the terrain and open space environment between the devices,



(a) Geographic and topology details.



(b) PDR and SINR temporal evolution.

Fig. 22. Platooning movement under the impact of the periodic jammer in the seafront scenario.

we obtain a single propagation model for this scenario with an RMSE of 3.932 [35]. The jammer, marked with a fire symbol on the figure, is located 180 m from the north end of the parking lot and slightly over 300 m from the south end. During the measurements, few vehicles sporadically drove by at speeds up to 60 km/h along a parallel street located 30 m away from the measurement area. In this scenario, we consider a periodic jammer and a reactive jammer⁶.

1) *Results:* A geographic visualization of the communication performance is illustrated in Figure 22(a), while Figure 22(b) shows the time evolution of PDR and SINR in the presence of the periodic jammer. It can be observed that the SINR decreases very fast within the first 15 s due to the increasing interference caused by the jammer. Shortly after leaving the north end until 80 m behind the jammer (towards the south end), the communication is completely blocked. The jammer creates a *blackout area* of about 250 m, which corresponds to 45 s at the speed of travel (about 20 km/h). While turning at the end of the parking lot, both vehicles can communicate again for nearly 20 s (range from 60 to 80 s). From there, they return to the north end experiencing again a large communication blackout around the jammer.

We conduct the same measurements with the reactive jammer. SINR and PDR are shown in Figure 23(b) and a geographic visualization is given in Figure 23(a). Recall that for the reactive jammer to be active it must first detect the transmitter. Then, to be able to impair packet delivery, it must also create sufficient interference power at the

⁶The reactive jammer has a slightly different pattern than the reactive jammers considered so far. It transmits a 64 μ s long signal after a delay of 16 μ s. From the measurements of Section III (Fig. 12) only slight differences in the jammer effectiveness are expected.




Fig. 23. Platooning movement under impact of the reactive jammer (16 μ s, 500 μ s) in the seafront scenario.

receiver. The communication is disrupted over 170 m in both directions around the reactive jammer, corresponding to 30 s, which is clearly visible in Figure 23(a). The interference range of the reactive jammer is shorter than for the constant jammer, namely 170 m compared to 250 m. Even though the reactive jammer is less effective, it still causes a significant damage to the communication in this scenario.

VI. RELATED WORK

To the best of our knowledge, the only previous works focusing on RF jamming attacks on VANETs are [8], [38], [39]. The work in [38] presents an approach for detecting reactive jamming in 802.11p networks, but it does not assess the impact of RF jamming, nor does it use realistic VANET scenarios for evaluation. The authors in [39] study, by means of network simulations, the impact of RF jamming on the dissemination of geocast messages. The attacker model consists of an 802.11p device tuned as a reactive jammer that sends a short signal upon sensing energy on the medium. The authors show that reacting upon sensing energy above the card sensitivity is significantly more effective (up to 20 %) than reacting only if the sensed SNR is sufficient to decode the incoming packet. Both attacks are able to block the dissemination of geocast messages in a simulated two lane highway. In a city scenario, however, the situation changes as there are alternative paths to disseminate the messages around the jammer. While these results are consistent with our observations in the field measurements and with our previous work [8], we cover a wider range of jamming signal profiles. Furthermore, we have shown in Section III that the behavior of real devices in the presence of RF jamming

is subject to complex performance issues. We believe that our measurements and proposed model can help researchers conduct realistic simulations of VANETs accounting for the impact of RF jamming.

Additionally, there is some related work on jamming at both MAC [40], [41] and PHY [42] for classical 802.11 WLANs. The authors in [9] study the vulnerabilities of 802.11b/g hardware to RF interference, which are associated to timing recovery and dynamic range selection issues. The authors consider a jammer that emits DSSS or OFDM modulated signals that do not comply with the standard and show that weak interference (30 dB less power than the legitimate signal) can significantly disrupt the communication by impeding time recovery. Our measurements show that, depending on the relative position of the nodes, a periodic jamming signal completely blocks communication up to an SINR of 56 dB. Particularly damaging is the case of weak signal detection triggered by the jammer followed by an overflow of the A/D dynamic range caused by the signal of interest. The observation and characterization of this event further extends the work in [9]. We also confirm the importance of correct timing recovery and show a degradation in (reactive) jammer effectiveness that correlates with the proportion of preamble that is missed. The authors in [43] present results for the error performance of 802.11b/g networks under the influence of RF jamming. Their results show that wide-band jamming damages OFDM-based 802.11g communications more severely (up to 7 dB) than they affect 802.11b (spread spectrum). The authors in [42] show that a constant wide-band noise signal is more effective (3-4 dB) than a constant wide-band digitally modulated signal at disrupting 802.11g communications.

VII. CONCLUSION

We have described the receiver structure of a reference 802.11p implementation and characterized how different RF jamming profiles impact the communication. Then, we have proposed and validated, by means of measurements, a procedure that uses the relationship between PDR and SINR measured in a controlled environment as a model to study the impact of RF jamming on VANET networks and applications. Finally, we have applied the proposed method to predict the behavior of a vehicular platoon under the influence of jamming, and have shown that the model provides realistic results. Furthermore, our results reveal that RF jamming poses a serious threat to VANET safety in general, and platooning applications in particular. In the latter case, reactive, constant, and periodic jammers can severely disrupt communication up to 465 m despite very short communication distances between legitimate devices. The significant impact of RF jamming reported in this work highlights the need for jamming-aware communications, protocols and applications, as well as effective jamming detection strategies. As a first step into that direction, we made our measurement data available for download in [10], [11] to be used by the community as input to network simulations and for further analysis.

As major lines of future work, we foresee the design of algorithms and protocols that detect and/or mitigate RF jamming attacks, as well as warning systems that alert the drivers about eventually malfunctioning safety applications. Additionally, more resilient PHY designs (e.g., with more frequent channel estimation) may increase the robustness of VANETs to jamming. Finally, protection measures should also be considered in system design by, for instance, implementing critical functionality directly into the firmware.

ACKNOWLEDGMENTS

We would like to thank Peter Steenkiste for his insightful comments. We thank Prof. Henrique Salgado and João Oliveira for the use of the spectrum analyzer and RF components. We also thank the anonymous reviewers for their careful reading of our manuscript and their insightful comments and suggestions.

REFERENCES

- [1] M. Segata and R. Lo Cigno, "Emergency Braking: A Study of Network and Application Performance," in *Proceedings of the 8th International ACM Workshop on Vehicular Inter-Networking, Systems, and Applications*, 2011.
- [2] Y. Zhang and G. Cao, "V-PADA: Vehicle-Platoon-Aware Data Access in VANETs," *IEEE Transactions on Vehicular Technology*, vol. 60, no. 5, pp. 2326–2339, 2011.
- [3] "IEEE 802.11p – Wireless Access in Vehicular Environments, Amendment 6 to 802.11," July 2010.
- [4] D. Jiang, Q. Chen, and L. Delgrossi, "Optimal Data Rate Selection for Vehicle Safety Communications," in *Proceedings of the 5th International ACM Workshop on Vehicular Inter-Networking, Systems, and Applications*, 2008.
- [5] A. Paier, D. Faetani, and C. Mecklenbräuker, "Performance Evaluation of IEEE 802.11p Physical Layer Infrastructure-to-Vehicle Real-World Measurements," in *Proceedings of International Symposium on Applied Sciences in Biomedical and Communication Technologies*, 2010.
- [6] A. Aijaz, B. Bochow, F. Dötzer, A. Festag, M. Gerlach, R. Kroh, and T. Leinmüller, "Attacks on Inter-Vehicle Communication Systems: An Analysis," in *Proceedings of the 3rd International Workshop on Intelligent Transportation*, 2006.
- [7] B. K. Chaurasia, S. Verma, and G. S. Tomar, "Attacks on Anonymity in VANET," in *Proceedings of the International IEEE Conference on Computational Intelligence and Communication Networks*, 2011.
- [8] O. Puñal, A. Aguiar, and J. Gross, "In VANETs We Trust?: Characterizing RF Jamming in Vehicular Networks," in *Proceedings of the 9th International ACM Workshop on Vehicular Inter-Networking, Systems, and Applications*, 2012.
- [9] R. Gummedi, D. Wetherall, B. Greenstein, and S. Seshan, "Understanding and Mitigating the Impact of RF Interference on 802.11 Networks," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4, pp. 386–396, 2007.
- [10] O. Puñal, C. Pereira, A. Aguiar, and J. Gross, "CRAWDAD data set uportorwthaachen/vanetjamming2012 (v. 2014-05-12)." Downloaded from <http://crawdada.org/uportorwthaachen/vanetjamming2012/>, May 2014.
- [11] O. Puñal, C. Pereira, A. Aguiar, and J. Gross, "CRAWDAD data set uportorwthaachen/vanetjamming2014 (v. 2014-05-12)." Downloaded from <http://crawdada.org/uportorwthaachen/vanetjamming2014/>, May 2014.
- [12] "IEEE Standard for Wireless Access in Vehicular Environments (WAVE) – Multi-Channel Operation," July 2011.
- [13] J. Kenney, "Dedicated Short Range Communication (DSRC) Applications Tutorial," in *IEEE 802.11 Wireless Next Generation Standing Committee*, no. 13/0541, 2013.
- [14] "Intelligent Transport Systems: Harmonized Channel Specifications for Intelligent Transport Systems Operating in the 5 GHz Frequency Band," October 2012. Number ETSI TS 102 724 V1.1.1.
- [15] I. Tinnirello, D. Giustiniano, L. Scalia, and G. Bianchi, "On the Side-Effects of Proprietary Solutions for Fading and Interference Mitigation in IEEE 802.11b/g Outdoor Links," *Computer Networks*, vol. 53, no. 2, pp. 141–152, 2009.
- [16] IEEE, *Standard for Information Technology 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*, 2012.
- [17] A. Festag, R. Baldessari, W. Zhang, and L. Le, "CAR-2-X Communication SDK - A Software Toolkit for Rapid Application Development and Experimentations," in *Proceedings of the International IEEE Vehicular Networking and Applications Workshop*, 2009.
- [18] G. F. M. of Education and Research, "Fleetnet." <http://www.neclab.eu/Projects/fleetnet.htm>. Last visit: November 29, 2013.
- [19] D. Reichardt, M. Miglietta, L. Moretti, P. Morsink, and W. Schulz, "CarTALK 2000: Safe and Comfortable Driving Based Upon Inter-Vehicle-Communication," in *Proceedings of the International IEEE Intelligent Vehicle Symposium*, vol. 2, pp. 545–550, 2002.
- [20] German Federal Ministry of Education and Research, "Network on Wheels." <http://www.network-on-wheels.de/>. Last visit: November 29, 2013.
- [21] M. Boban, T. T. V. Vinhoza, M. Ferreira, J. Barros, and O. Tonguz, "Impact of Vehicles as Obstacles in Vehicular Ad Hoc Networks," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 1, pp. 15–28, 2011.
- [22] T. Mangel, O. Klemp, and H. Hartenstein, "5.9 GHz Inter-Vehicle Communication at Intersections: A Validated Non-Line-of-Sight Path-Loss and Fading Model," *EURASIP Journal on Wireless Communications and Networking*, vol. 2011, no. 1, 2011.
- [23] "Atheros Communications Incorporated Patent: Method and System for Noise Floor Calibration and Receive Signal Strength Detection." <http://www.patentstorm.us/patents/7245893/description.html>. Last visit: November 29, 2013.
- [24] "QUALCOMM Incorporated Patent: Automatic Gain Control for a Wireless Receiver." <http://www.patentstorm.us/applications/20060222118/description.html>. Last visit: November 29, 2013.
- [25] S. Robitzsch, L. Murphy, and J. Fitzpatrick, "An Analysis of the Received Signal Strength Accuracy in 802.11a Networks Using Atheros Chipsets: A Solution Towards Self Configuration," in *Proceedings of the 7th International IEEE Broadband Wireless Access Workshop (GLOBECOM Workshop)*, 2011.
- [26] "Atheros Communications Incorporated Patent: Ambient Noise Immunity." <http://www.freepatentsonline.com/7349503.html>. Last visit: November 29, 2013.
- [27] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks," in *Proceedings of the 6th International ACM Symposium on Mobile Ad Hoc Networking and Computing*, 2005.
- [28] K. Pelechrinis, I. Broustis, S. Krishnamurthy, and C. Gkantsidis, "A Measurement-Driven Anti-Jamming System for 802.11 Networks," *IEEE/ACM Transactions on Networking*, vol. 19, no. 4, pp. 1208–1222, 2011.
- [29] "Rice University: The WARP Project." <http://warp.rice.edu>. Last visit: November 29, 2013.
- [30] "Maxim Integrated: MAX2829, Single-/Dual-Band 802.11a/b/g World-Band Transceiver ICs." <http://www.maximintegrated.com/datasheet/index.mvp/id/4532>. Last visit: November 29, 2013.
- [31] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "WiSec 2011 Demo: RFReact - A Real-Time Capable and Channel-Aware Jamming Platform," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 5, no. 3, pp. 41–42, 2011.
- [32] M. Briggs, J. Martinez, D. Bare, W. M. Ave, and A. R. Requirements, "Power Measurements of OFDM Signals," in *Proceedings of the 46th International IEEE Electronic Materials Conference*, 2004.
- [33] "The MadWifi Project: Multiband Atheros Driver for WiFi." <http://madwifi-project.org/>. Last visit: November 29, 2013.
- [34] H. Friis, "A Note on a Simple Transmission Formula," *Proceedings of the Institute of Radio Engineers*, vol. 34, no. 5, pp. 254–256, 1946.
- [35] C. Pereira and A. Aguiar, "A Realistic RF Jamming Model for Vehicular Networks: Design and Validation," in *Proceedings of the 24th International IEEE Symposium on Personal, Indoor, and Mobile Radio Communications*, 2013.
- [36] "A Realistic RF Jamming Model for Vehicular Networks." <http://>

- paginas.fe.up.pt/~dee12014/VANETjamming.html. Last visit: November 29, 2013.
- [37] L. Hobert, "A Study on Platoon Formations and Reliable Communication in Vehicle Platoons," Master's thesis, University of Twente, 2012.
 - [38] A. Hamieh, J. Ben-Othman, and L. Mokdad, "Detection of Radio Interference Attacks in VANET," in *Proceedings of the International IEEE Global Communication Conference*, 2009.
 - [39] E. Schoch, F. Kargl, and T. Leinmüller, "Vulnerabilities of Geocast Message Distribution," in *Proceedings of the 2nd International IEEE Workshop on Automotive Networking and Applications (GLOBECOM Workshop)*, 2007.
 - [40] J. Bellardo and S. Savage, "802.11 Denial-of-Service Attacks: Real Vulnerabilities and Practical Solutions," in *Proceedings of the 12th International USENIX Security Symposium*, 2003.
 - [41] K. Pelechrinis, M. Iliofotou, and S. Krishnamurthy, "Denial of Service Attacks in Wireless Networks: The Case of Jammers," *IEEE Communications Surveys Tutorials*, vol. 13, no. 2, pp. 245–257, 2011.
 - [42] I. Harjula, J. Pinola, and J. Prokkola, "Performance of IEEE 802.11-based WLAN Devices Under Various Jamming Signals," in *Proceedings of the 30th International IEEE Military Communications Conference*, 2011.
 - [43] T. Karhima, A. Silvennoinen, M. Hall, and S.-G. Haggman, "IEEE 802.11b/g WLAN Tolerance to Jamming," in *Proceedings of the 23rd International IEEE Military Communications Conference*, 2004.