ATCZ175 InterOP PROJECT

Interference Emulation Results

Institute of Electrodynamics, Microwave and Circuit Engineering
Technische Universität Wien

Advisor:
Assoc. Prof. Dipl.-Ing. Dr.techn. Holger Arthaber

by
Proj. Ass. Dipl.-Ing. Christian Spindelberger

October 9, 2019

Contents

Abbreviations

CCA clear channel assessment

ISM industrial, scientific, and medical

MTU maximum transmission unit

PER packet error rate

PLCP physical layer convergence procedure

QAM quadratue amplitude modulation

RF radio frequency

SINR signal to interference plus noise ratio

TCP transmission control protocol

UDP user datagram protocol

VSG vector signal generator

WLAN wireless local area network

1 Introduction

As final step, the introduced interference emulation methods, will be examined using the WLAN test setup. Performance tests will give a graphical output to evaluate the observed scenarios.

First of all, reference measurements utilizing ISM band records, replayed by a vector signal generator (VSG), are investigated in Section 2. These performance results are used to compare further techniques. Section 3 treats modulated noise as interfering signal. Two different models are applied and further improved by weighting power levels. It will be shown, that one decisive parameter influencing the comparison of interfering signals is caused by device dependent parameters. Furthermore, the capability of the low-cost interference source, is analyzed. Therefore, the main output of Section 4 covers the replacement of a VSG by this low-cost implementation.

2 Reference Measurements

Based on the measurement setup, reference measurements are made utilizing a VSG as an interference source. Such a device enables replaying customized baseband data. Hence, two recordings with a total length of 20 s are employed for interference injection[1]. Because of memory depth limitations of the VSG, the investigated records are divided into four 5 s long sequences. Before that, the data is processed by a 22 MHz broad filter. In addition to this, the output power is adjusted such that the received power by the server equals the characterized power levels. The implemented variable attenuator of the setup controls the signal to interference plus noise ratio (SINR) by adjusting the received average power from the client relative to the mean power of the recordings, replayed by the interference source. In order to adjust the packet length of TCP performance tests using *iperf*, the maximum segment size must be set to the desired value. This adaptation limits the TCP window size, decreases the throughput drastically, and consequently, almost no retransmissions occurred for small packet sizes ($<$ 200 Byte). As TCP performance tests did not lead to useful results, UDP is utilized in this section.

The packet error rate (PER) and throughput are parameters which have to be estimated precisely. Therefore, the sample size ensuring an appropriate confidence interval has to be determined. Based on the central-limit theorem such statistical problems assume a normally distributed expectation value. These assumptions have to be dropped regarding the UDP performance tests. The WLAN test setup is surrounded by an unknown amount of devices influencing the obtained results. Consequently, the measurement scenario is

[1] The recordings can be found on: `http://www.interreg-interop.eu/results/wlan_interference_analysis/measurement_campaign/`

affected by permanently changing ISM band traffic characteristics. Therefore, an empirical sample size of 100 performance tests, yielding a non-zero PER per SINR value, was set. This means that the PER computation for one operating point takes at least 16.7 min. Because of this time-consuming process, the two recordings with the highest traffic density have been investigated.

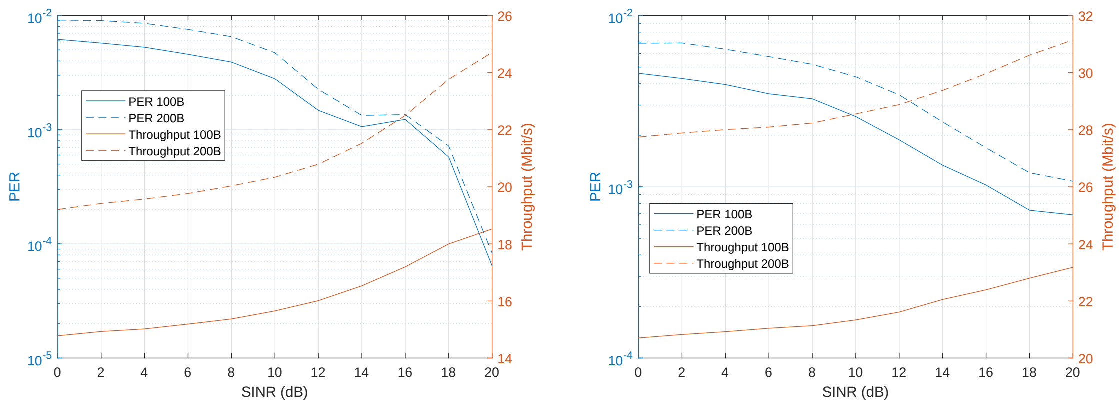

(a) HT-MCS-4

(b) HT-MCS-6

Figure 1: UDP performance test results for different modulations and packet lengths in WLAN channel 1

Figure 1 depicts the respective throughput and PER curves for a UDP performance test, varying the modulation index (HT-MCS-4/6: refer to WLAN instruction guide) and the packet length (100, 200 Byte). Due to longer time intervals, the probability of perturbations caused by interference becomes higher. Therefore, the PER is worse for longer packets, as expected. Although the PER of the short sequences is lower, the throughput is also decreased compared to the 200 Byte long sequences. The network layer is responsible for fragmenting data according to the MTU size. Hence, the amount of fragments increases for lower packet sizes, causing a higher data processing effort. Consequently, the throughput is decreased, while the PER is lower.

Another interesting outcome of these curves are the constant plateaus. The PER is constant for an SINR between 14 dB and 16 dB, depicted in Subfigure 1(a). Mainly, this behavior can be explained by clear channel assessment (CCA) decision threshold settings. The preamble detection threshold has a maximum value of −82 dBm. Considering the measurement setup, the hidden node model is established. If the attenuator is set to low values for achieving a higher SINR, the client is able to detect some of the interfering signals. Due to the collision avoidance scheme, the transmitter (client) is capable to keep the PER constant in this SINR range. For higher attenuations, the transparency gets worse and the PER is increased. As the threshold is specified by an upper bound only, manufacturers set decision constraints individually.

3 Modulated Noise

From former reports[2] it turned out that modulating noise is a promising technique to emulate ISM band interfering signals. According to the settings introduced in Section 2, two techniques creating random noise sequences are applied. Firstly, the time-quantized analysis will be investigated to model noise bursts.

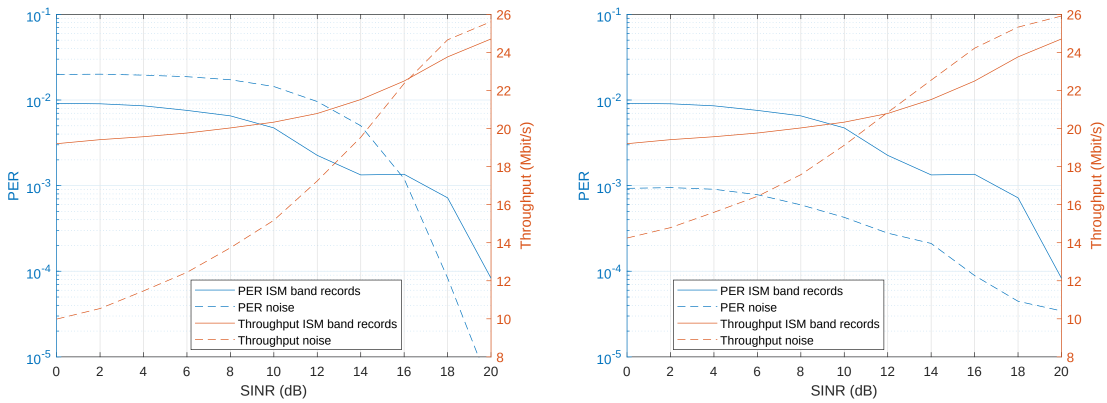

(a) Step size: 100 µs

(b) Step size: 1 ms

Figure 2: Time-quantized noise bursts: UDP performance test results, packet length 200 Byte, HT-MCS-4, WLAN channel 1

Figure 2 depicts the measurement results of the corresponding step sizes for an UDP performance test in WLAN channel 1. The packet size is set to 200 Byte, utilizing the HT-MCS-4 (16-QAM-3/4) modulation. The performance results are compared to the ISM band recordings in terms of PERs and throughputs. Obviously, the ISM band-induced PER behaves differently compared to noise. The first issue concerning this problem is the missing preamble in noise bursts. For instance, the WLAN related physical layer convergence procedure (PLCP) sublayer introduces a header for frame detection. The interfering ISM band records force the server station to demodulate all incoming preambles lying beyond the CCA preamble detection threshold. Thus, interfering WLAN packets with a small

[2] `http://www.interreg-interop.eu/results/wlan_interference_analysis/linux_based_wlan_test_setup/`

power level relative to the desired frame may also cause packet errors. Another difference caused by CCA constraints is the missing PER plateau for modulated noise. The plateau indicates the area, where the isolation maintained by the coupler is in a range where the client is able to react on ISM band interfering signals, holding the PER constant. Unlike the preamble detection threshold ($-82\,\mathrm{dBm}$), the energy detection threshold is defined to be maximally $-62\,\mathrm{dBm}$. This difference of $20\,\mathrm{dB}$ causes the missing plateau within the observed SINR range for modulated noise. Furthermore, the noise related throughput has a higher slope for both step sizes. Considering throughputs from Subfigure 2(b), one can notice that the data rate is still lower than for digital data, although the noise related PER is smaller. This behavior emphasizes that time-quantized modulated noise is not capable to approximate the reference throughputs in this configuration. Nevertheless, it is possible to fit a fractional part of the PER with noise utilizing an appropriate step size.

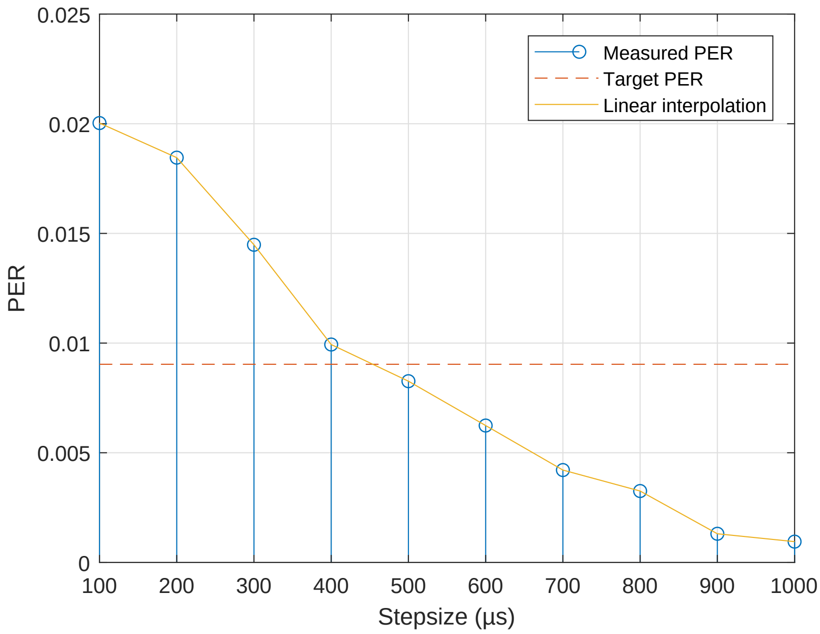

Figure 3: PER alignment function depending on the step size, SINR $= 2\,\mathrm{dB}$

In order to achieve a satisfying approximation, the main goal is to align the noise-induced PER curve with the one caused by ISM band traffic. Therefore, an SINR range has to be found where the PER curves have an almost constant shape, i.e., for SINR values below $6\,\mathrm{dB}$. Figure 3 depicts the measured PER at an SINR of $2\,\mathrm{dB}$ for several step sizes and the target PER. Obviously, the alignment function shows a monotonically decreasing behavior in the observed interval. Hence, it is possible to find an appropriate step size to fit the desired PER. In order to also fit the PER curve for higher SINR values, it seems natural to enhance the mean power of the interference signal. As a consequence of increasing the power, the PER curve is shifted to the right. Corresponding to Figure 3, an appropriate step size of 455 µs has been found. Examining Figure 4 shows the concerning performance test results. In the left Subfigure 4(a) the PER is well aligned for SINR values between 0 dB and 14 dB. In order to further improve the fit, the mean power was enhanced by 2.2 dB. These adjustments lead to an approximation of almost the whole desired PER. The only issue that remains is the constant plateau which could not be emulated.

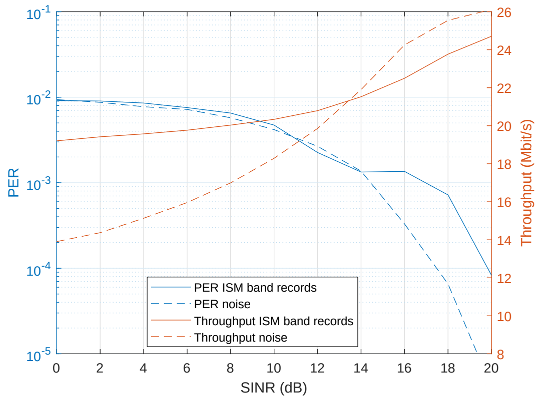

(a) Step size: 455 µs

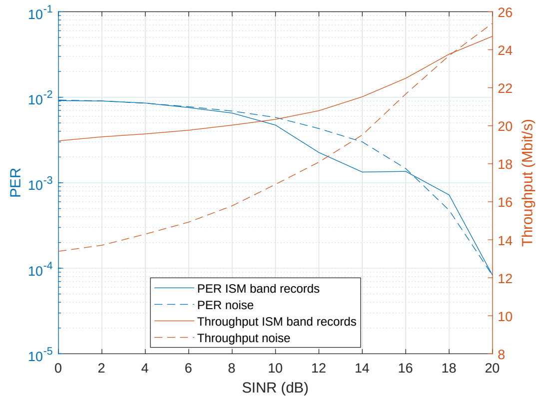

(b) Enhanced mean power by 2.2 dB

Figure 4: Time-quantized noise bursts: UDP performance test results, packet length 200 Byte, HT-MCS-4, WLAN channel 1

The time-quantized analysis can be further extended by varying on- and off-times corresponding to the random variables gained from the measurement campaign. The PER and throughput of this technique can be found in Subfigure 5(a). Varying on- and off-times yields a better fit concerning the throughput having a similar slope. This is not always the case, which will be discussed later in this section. Enhancing the mean power of the interfering signal by 3 dB results in a shift of the noise-induced PER to the right of the same range (Subfigure 5(b)). Hence, it is possible to fit the PER between 16 dB and 18 dB. One might notice the decay of PER and throughput for an SINR between 0 dB and 2 dB. Typically, an increased throughput is associated with a decay of the PER. As this is not the case in this SINR range, this behavior may be caused by crosstalk issues. WLAN modules based on MIMO have the opportunity to neglect receive paths when they are considered impractical. Presumably because of insufficient isolation, another received weak signal can be demodulated out of the receivers optimal operating point, resulting in a lower PER and throughput as well. Another technique to fit the PER curves is to omit power level values below a certain level with an exclusion rule. Subfigures 5(c) and 5(d) depict the respective performance results for an exclusion of power levels which are smaller than −70 dBm and −88 dBm.

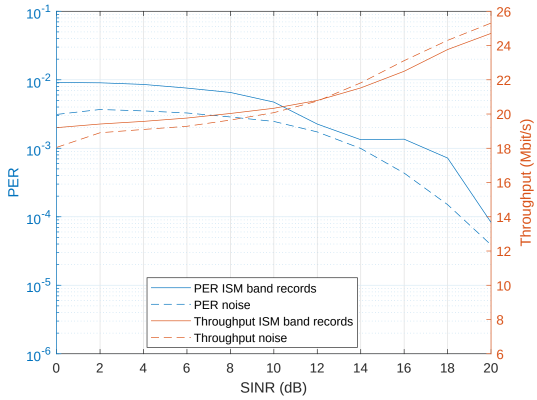

(a) Modulated noise

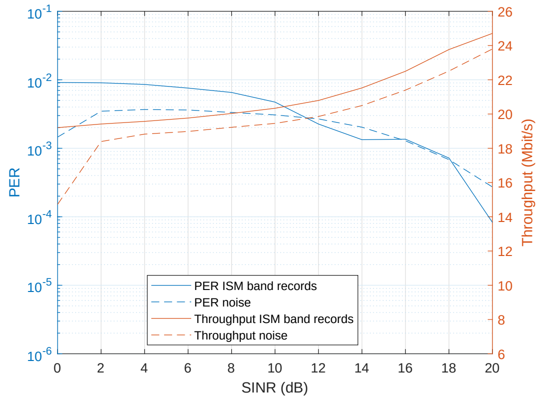

(b) Enhanced mean power by 3 dB

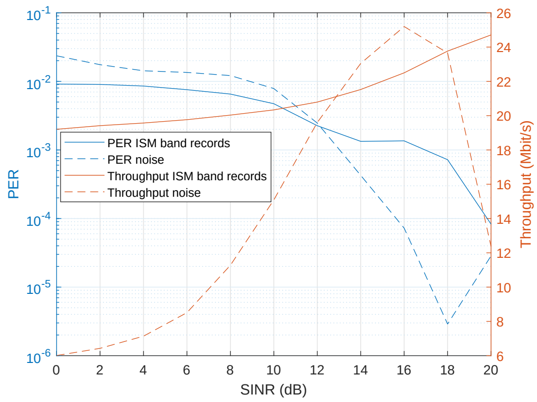

(c) Power threshold > −70 dBm

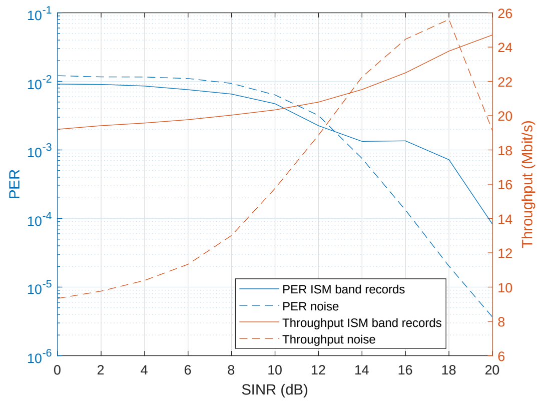

(d) Power threshold > −88 dBm

Figure 5: Modulated noise according to random variables from measurement campaign, packet length 200 Byte, HT-MCS-4, WLAN channel 1

Because of higher power levels, the PER becomes worse, enabling an approximation of the desired PER curve for SINR values between 0 dB and 12 dB. It is noteworthy to mention the strong throughput decay for an SINR beyond 16 dB. Adjusting a high SINR relates to a low attenuation between client and server. In this case, the power levels are high enough to be detected by the client yielding a lower throughput to keep the PER low. Thus, the working strategy of this implementation is oriented towards optimizing the PER by adjusting, for instance, the throughput. Due to blocking effects of the server, Subfigure 5(c) shows an increase of the PER between an SINR of 18 dB and 20 dB. At these power levels, the low-noise amplifiers are driven into their saturation region, causing an impractical performance.

As a consequence of saturation effects, it is not possible to align the noise-related PER curve by enhancing the mean power, like it was done in the time-quantized analysis. In addition to this, varying on- and off-times results in an even higher complexity, making this method impractical to approximate PER curves. Concerning all performance measurements in this section, the dilemma of throughput and PER approximation is present. Simply put, an approximation of the PER curve by noise does not lead to similar throughputs and vice versa. As already explained, this behavior is caused by CCA detection characteristics. Anyway, investigating Subfigure 5(a) might lead to the point that it is possible to fit throughput curves by modulating on- and off-times.

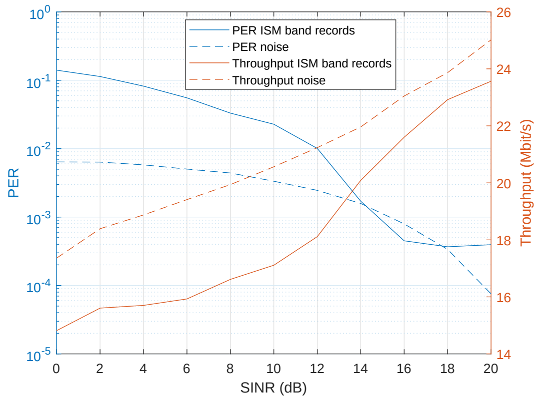

(a) HT-MCS-4, packet length 200 Byte

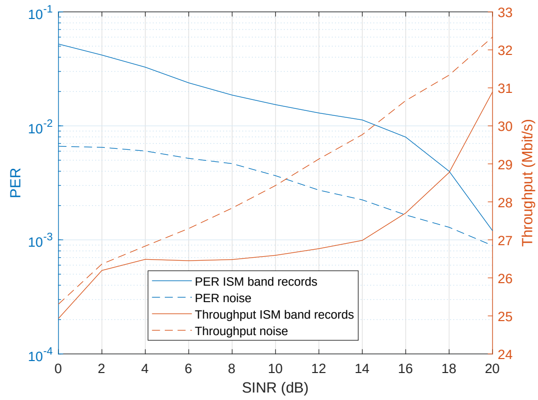

(b) HT-MCS-6, packet length 200 Byte

Figure 6: UDP performance test results: WLAN channel 36

Because of device-dependent characteristics, this is not always the case. Figure 6 depicts the performance measurement results compared to modulated noise in the 5 GHz band. Unlike the throughputs of Subfigure 5(a) in the 2.4 GHz band, the slopes appear in different shapes. The ISM band recordings show an area with a stagnation of traffic speed. This behavior gets even worse for higher modulations, while the noise related curve has a linear shape. Consequently, individual parameters, such as physical receiver design, or CCA threshold adaptation, influence the behavior of interfering signals in terms of PER and throughput stringently. For instance, one might think that the preamble detection threshold only works for WLAN related frames. Further investigations of the Schmidl and Cox algorithm showed that also BLE packets may cause a falsely detected WLAN packet. Figure 7 depicts the respective timing metric of a BLE packet from the measurement campaign. Although the signal is oscillating, the metric is beyond a level of 0.6.

Hence, BLE packets can also be detected if they are beyond the preamble threshold causing a decreased throughput.

In conclusion, two methods modulating noise have been presented. Time-quantized analysis yields an appropriate technique by adjusting the step size and mean power of the interference signal to approximate the PER curve. In contrast to the PER curve, it turned out that the throughput could not be fitted with this method. If the on- and off-times are also varied, it is possible to gain similar throughputs. Unfortunately, this characteristic depends on several individual device dependent properties probably causing different shapes of the throughput curves.

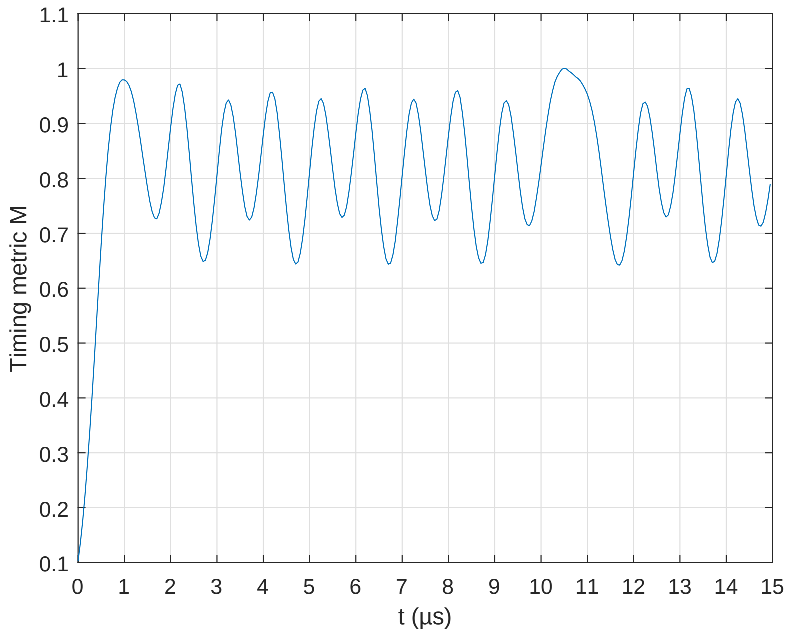

Figure 7: Coarse packet detection of a BLE packet utilizing Schmidl and Cox algorithm

4 Low-Cost Noise Source

A low-cost interference source was built to modulate noise by simply using off-the-shelve components[3]. In order to verify this concept, performance measurement results will be compared to those made by a VSG. In the following, the main aspects comparing the VSG with the low-cost noise source will be discussed.

[3] http://www.interreg-interop.eu/results/wlan_interference_analysis/low_cost_interference_emulator/

4.1 Memory Depth

The memory depth of the utilized VSG is limited to 256 MSa. At a sampling rate of, for instance, 51.2 MSa/s, the memory is capable to store a 5 s long recording. Considering the low-cost implementation, an *ATMEL 328P* microcontroller on an *Arduino Nano* board is used. The working principle of this device differs completely from the VSG. The noise source is not capable to repeat baseband data. It uses digitally-controlled attenuators to modulate amplified noise. Therefore, the memory depth of the microcontroller has to be examined. The on-/off-times, as well as the power level values, are generated by the corresponding random variables. The obtained values are then stored in the microcontroller. The total space for programming code is 30.72 kByte, and the dynamic storage used by variables offers 2,048 Byte. Obviously, the dynamic storage limits the maximum sequence length. The maximum amount of variables which can be stored without occurring instabilities is 500 values for the time-quantized method and 220 for modulated noise (on-/off-times). Due to these low values, several sequences have to be utilized to gain a fine resolution of the respective random variables. The measurement results will indicate if the reduced memory depth can be conquered by creating many smaller sequences and how high the resolution of the random variables must be to gain the same performance curves.

4.2 Reference Power Levels

As the low-cost noise source utilizes attenuators to modulate the output power over time, an absolute reference power level has to be defined. Therefore, the noise floor was investigated to calculate appropriate attenuation values relative to the randomly described power levels. The estimation of the noise floor density suffers from several effects. Hence, the calculation of the reference power level will include inaccuracies caused by interference signals. The survey of several sequences in the ISM band data, which were presumed to be noise, led to an average noise floor power level of −98 dBm (WLAN channel 1, BW = 22 MHz). Based on this reference level (P_{ref}), the attenuation values are calculated using the random power levels (P_{rand}) and the dynamic range of the source ($\mathrm{DR}_{\mathrm{ATT}} = 63.5$): $\mathrm{ATT} = \mathrm{DR}_{\mathrm{ATT}} - (P_{\mathrm{rand}} - P_{\mathrm{ref}})$.

4.3 Time Critical Behavior

Usually, standard radio frequency (RF) attenuators have an undefined output during state transitions, causing an overshoot of the envelope. Therefore, glitch-free attenuators have been utilized to conquer this effect. One negative aspect associated with such elements is the higher settling time compared to general purpose attenuators. The specified maximum settling time of the assembled attenuators is < 400 ns.

Another important aspect concerning the timing resolution of the low-cost noise source is the microcontroller. The digital output pins, which are controlling the attenuators, are separated into eight bits long ports. They can be easily set by simply manipulating the respective bits. In order to increase the switching speed, the attenuators are driven directly without a latch enable. Unfortunately, these ports are not well synchronized and may cause a temporal overshoot for a few hundred nanoseconds.

In contrast to the noise source, the VSG utilizes a digital-to-analog-converter with a resolution of 16 bits to replay baseband data directly from an internal storage. The sampling rate can be adjusted up to 120 MSa/s. In the following section, it will be investigated if such high timing resolutions are mandatory to achieve the same performance test results.

4.4 Comparison of Results

First of all, applying the time-quantized analysis to the low-cost noise source will be analyzed in contrast to the VSG. Therefore, several sequences with an amount of $N = 500$ samples each are utilized to jam the WLAN test setup.

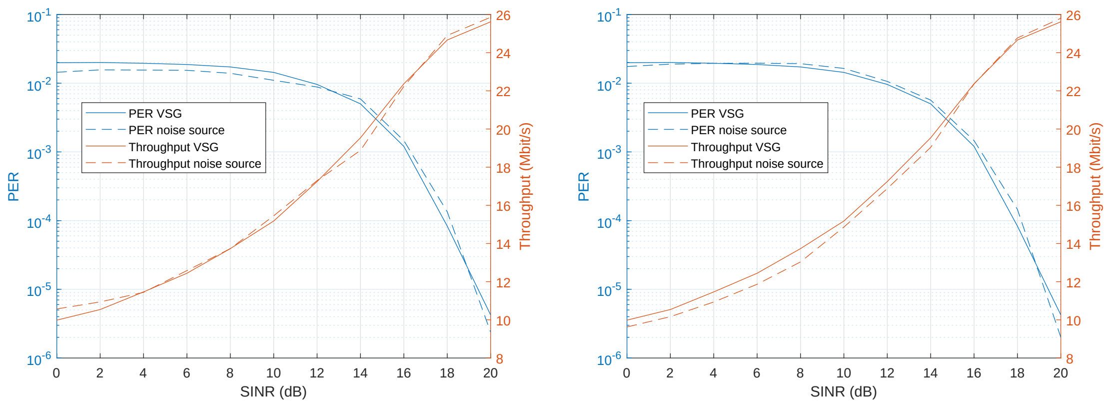

(a) $N = 500$ (b) $N = 2{,}000$

Figure 8: Comparing VSG with noise source: time-quantized noise, step size: 100 µs, HT-MCS-4, packet length 200 Byte, WLAN channel 1

In Figure 8, one can see UDP performance test results for different power level distribution resolutions (N). Obviously, an amount of $N = 500$ (one sequence) already yields similar performance curves, which can be found in Subfigure 8(a). Increasing the number of sequences fourfold (Subfigure 8(b)) fits the PER curve even better, but the throughput shows stronger deviations below an SINR of 10 dB.

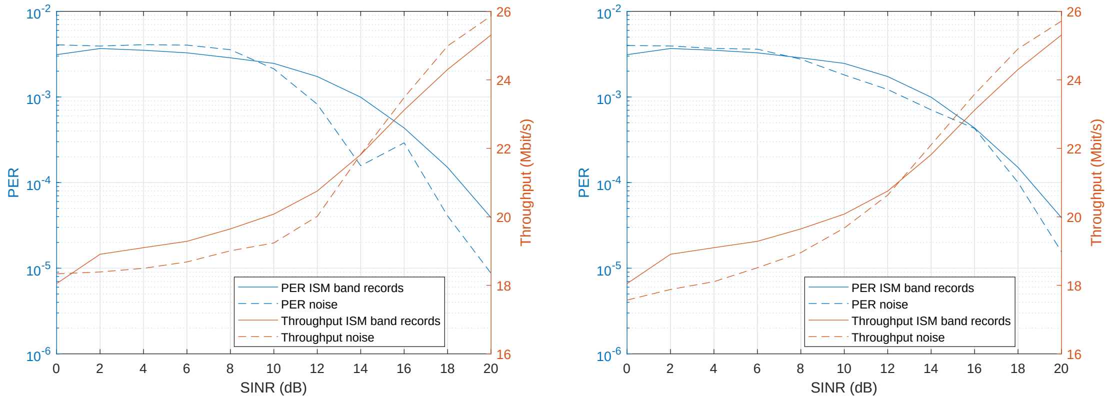

(a) $N = 220$ (b) $N = 2{,}200$

Figure 9: Comparing VSG with noise source: modulated noise, HT-MCS-4, packet length 200 Byte, WLAN channel 1

Utilizing on- and off-times allocates more memory of the microcontroller. Hence, one sequence consists of $N = 220$ on-/off-times and power level values, i.e., 660 values in total. Due to varying two more parameters, compared to the time quantized method, the amount of parameter combinations increases. Hence, the memory depth is a serious issue concerning this modulation technique. The performance curves from Figure 9 indicate that a high amount of parameter combinations are necessary to gain similar results. Even though the approximation gets better for an increasing amount of samples, the PER curve still shows a nonsmooth behavior for $N = 2{,}200$.

Recapitulating this section, regarding the realization issues of the noise source and performance test results, indicates that it is possible to replace a VSG by a low-cost implementation. The most serious impact turned out to be the memory depth of the microcontroller. As the time-quantized method led to a satisfactory approximation even for low resolutions, the utilization of on- and off-times demands an on-board logic with more memory space.