appeared in Proc. of European Wireless 2002, pp. 762-769, Florence, Italy, February 2002

Measurements of a Wireless Link in different RF-isolated Environments

James Gross, Andreas Willig
TU Berlin
Einsteinufer 25
10587 Berlin, Germany
{gross,awillig}@ee.tu-berlin.de
fon: +49 30 31423831, fax: +49 30 31423818

ABSTRACT

For the design and evaluation of wireless MAC and link layer protocols it is desirable to have some knowledge about the error patterns delivered by the wireless PHY. One way to obtain this knowledge is to perform measurements. In this paper we report on measurements taken in two almost ideal environments (no interferers, simple propagation environment) and using an IEEE 802.11 DSSS-Compliant physical layer (PHY). These measurements highlight the influence of selected parameters (Diversity, Distance) and can serve as "baseline" measurement for further studies.

1 INTRODUCTION

For the design and evaluation of wireless MAC and link layer protocols it is desirable to have some knowledge about the error patterns delivered by the wireless PHY, the latter being taken as the ensemble of baseband processing, high frequency circuitry, transmitters, receivers, and finally the channel. The statistics of these error patterns do not only help in the design of several MAC mechanisms (FEC coding, interleaving, retransmission schemes, choosing proper packet sizes), but can also serve to provide "realistic" parameters for wireless channel models, like, e.g., the popular Gilbert/Elliot model [6], [5].

One way to explore the error behavior of a wireless link is to do measurements. There are several measurement studies reported in the literature (see Section 5). This paper reports on several measurements taken in almost ideal, undistorted environments with very simple propagation characteristics. This serves two different goals: first, the simplicity of the chosen environments allows to highlight the influence of selected parameters more clearly, e.g., antenna diversity. In addition, since in the environments chosen the errors are dominated by multipath effects, we can specifically relate all the parameters to this specific source of errors. In contrast, in more complex environments it is often hard to associate observed error patterns to the chosen parameters. The second goal is to provide a "baseline" for other measurement studies in different environments. Having results for "clean room" environments can give better insights when explaining the results for more complex environments.

The measurements were taken with an IEEE 802.11 compliant radio modem employing a direct sequence spread spectrum (DSSS) PHY. Our setup uses a MAC-less version of the Harris/Intersil PRISM I chipset [1]. Because of its MAC-less operation (there is just a small engine for generating well-known packets) we have fine grained control about the timing and contents of the generated packets, and furthermore we avoid any interference with MAC mechanisms, e.g., discarding packets after checksum errors. In addition, the packet generation and reception process is not biased by any upper layer protocol or operating system behavior.

Clearly, the results reported here are specific for the scenario chosen, for the frequency band (2.4 GHz ISM band), for the particular wireless technology, and the properties and capabilities of the measurement setup. Nevertheless, we think that some qualitative results carry over to more complex environments, for example the phenomenon and burstiness characteristics of packet losses or the variability of a wireless link.

The remainder of the paper is structured as follows: in Section 2 we describe our measurement setup, including the most important characteristics of the PRISM I chip set, and a description of the evaluation methodology. In Section 3 we describe the chosen measurement environments and the set of fixed and variable parameters used throughout these measurements. In Section 4 we present the most important results for each of the three environments chosen. In Section 5 we give a brief overview on other packet- and bit-level measurement studies, and finally, in Section 6 we give our conclusions.

2 MEASUREMENT SETUP AND EVALUATION

In this section we give a brief overview of our measurement setup. More details can be found in [11].

2.1 IEEE 802.11 / PRISM I Radio Modem

In 1997, the IEEE 802.11 standard was finalized , describing a WLAN operating in the license-free 2.4 GHz ISM band (Industrial, Scientific and Medical band) and offering different bit rates: 1, 2, 5.5, and 11 MBit/s (see Table 1). We have used a MAC-less radio modem (based on Harris/Intersil PRISM I chipset [1]), which is compliant with IEEE 802.11 and uses the direct sequence spread spectrum (DSSS) PHY. Two antennas are attached to the modem to enable receiver diversity (i.e., the receiver selects the antenna with the maximum signal level). The transmitter power was fixed at 18 dBm, corresponding to 63 mWatt. The radio modem basically consists of high

Preamble	SFD	Signal	Service	Length	CRC16	Data = {chunk}

Figure 1: Format of a packet

frequency circuitry and a baseband processor. The latter accepts and delivers a serial bit stream from upper layers, optionally scrambles the data, and performs DSSS processing [8]. The characteristics of the serial bit stream is our focus of interest.

The baseband processor transmits and receives data in units of *packets*. A packet consists of a *header* and a *data part*, shown in Figure 1. The header fields control transmission and contain no MAC-related fields. A packet starts with a well-known preamble of fixed length, followed by a fixed value indicating the start of the header (start frame delimiter, SFD). The preamble and SFD allow the receiver to synchronize on the sender's clock (bit synchronization). The signal field indicates the modulation type used in the data portion of the packet, while the length field indicates the length of the data portion in microseconds (the service field has no significance). The CRC16 field contains a 16 bit cyclic redundancy check (CRC) checksum which is computed from the three previous values. If the checksum is wrong or the signal field carries an unknown value, the whole packet is discarded by the baseband processor. While the data part can use different modulation types, the header is always transmitted with BPSK modulation.

2.2 Measurement setup

We used two dedicated stations, a *transmitter station* and a *receiver station*, which do not change their roles during a measurement. The basic idea is that the transmitter station sends a well-known packet stream over the wireless link, which is captured and stored by the receiver station into a logfile. For generation and reception of the packets we use a microcontroller board carrying the radio modem and a separate processor. The coupling to the (Windows NT-based) host is achieved with a segment of 64 kByte shared memory, denoted as *host interface*. We call this board a *wireless NIC* (Network Interface Card). The wireless NIC contains a specific measurement application and neither MAC functionality nor any higher layer protocols. This way we have fine grained control over the packet generation and reception process and no bias is introduced by upper layer protocols.

Our setup enables variation of several parameters, which are related both to the properties of the radio modem and packet stream generation. The important adjustable parameters are shown in Table 1.

The transmitter station generates a *packet stream*. What the receiver captures is called a *trace*. If no errors occur, the trace is the same as the packet stream. The format of the packet stream was chosen such that: a) the number of 0's and 1's are equal; b) long runs of 0's or 1's are avoided; and c) it suffices to have a fraction of the packet (denoted as *chunk*) correctly received in order to determine which packet it originally was. Especially the last property enables bit-by-bit comparison of a received packet with the transmitted packet.

Parameter	**Description**
DiversityEnabled	determines whether receiver antenna diversity is used
ModulationCode	distinguishes modulation used for data portion: 1 MBit/s BPSK, 2 MBit/s QPSK, 5.5 MBit/s CCK, 5.5 MBit/s BMBOK, 11 MBit/s CCK, 11 MBit/s QMBOK
NumPackets	Number of Packets
GapTime	Time gap between two packets
NumChunks	Number of chunks per packet, Packet length = *NumChunks* times 288 bits

Table 1: Adjustable parameters

The generated packet stream consists of *NumPackets packets*, which are transmitted at equidistant start times, and all packets having the same parameters and packet size. The data part of a packet consists of an integral number of *chunks*. For generating a chunk, every bit of a 32 bit sequence number is mapped to eight bits (with $0 \mapsto 11000011$ and $1 \mapsto 00111100$), giving 256 bits. Additionally, a header (`0xffff`) and trailer (`0x0000`) are generated, giving an overall chunk size of 288 bits.

2.3 Measurement Evaluation

For this paper we focus merely on *packet losses* and *bit errors*. Packet losses occur due to the receiver failing to acquire bit synchronisation. To acquire bit synchronisation it is necessary to detect the preamble and the SFD field. If either fails, the packet is lost in total. Packet losses are detected using the timestamps generated by the receiver and the property that the transmitter sends these packets at equidistant start times. Clearly, bit errors can occur only in received packets. When comparing an actually received packet with its corresponding transmitted packet, it can be said precisely, on which position bit errors occur.

For a given trace, both the packet losses and bit errors are displayed and evaluated in the form of *binary indicator sequences*. In general, this is a finite sequence of zeros and ones. As a convention, in binary indicator sequences we associate with a 1 an error event (e.g., an erroneous bit or a lost packet) and with a 0 the correct event.

We subdivide binary indicator sequences into *error bursts* and *error-free bursts* according to a *burst order* k_0. We define an error-free burst of order k_0 to be a contiguous all-zero subsequence with a length of at least $k_0 + 1$. In contrast, an error burst of order k_0 is a subsequence of at least one bit length, and with ones at its fringes, furthermore, within an error burst at most $k_0 - 1$ consecutive zeros are allowed.

For every trace two important binary indicator sequences were computed. The *packet loss indicator sequence* (PLIS) of a single trace is constructed by marking lost packets with a 1 and received packets (no matter whether they show bit errors or not) with a 0.

The *bit error indicator sequence* (BEIS) of a single trace is constructed by `XOR`ing every received (but possibly erroneous) packet with its corresponding expected (error-free) packet, and simply concatenating the results in the order of increasing packet numbers. Please note that in the BEIS any information about packet boundaries, lost packets, or packet gap times is completely ignored.

The BEIS can be seen as part of the available input of a MAC protocol or a coding scheme.

3 Measurement Environments and Parameters

In all measurements we wanted to observe wireless errors in the absence of other RF-sources and for simple multipath scenarios. We picked two environments and performed several measurements for each environment. Each measurement is characterized by specific parameter settings and a clear intention. However, not all performed measurements are presented here. For a complete overview refer to [7].

3.1 Environments

As first RF-isolated environment an open field outside of Berlin was chosen. It was located far away from any village. Within a distance of one kilometer no buildings were present. The measurement setup was placed on a gravel street. On both sides of the street were large open fields. No traffic passed by during the measurements. Furthermore the participants did not move during a trace. This environment could be considered as a best-case environment, since other RF-sources, specifically in the ISM-band, were not present. Furthermore, the multipath scenario was extremely simple, since no buildings or other structures were present within a large range. The only possibility for wave reflections was given through the surface.

The second environment was chosen to have simple wave reflecting structures, while maintaining the property of having no other RF-interferers closeby. The location meeting these criterias is a sports gym in the suburbs of Berlin. The gym is 30 meters wide and 60 meters long. The floor consisted of usual sports rubber whereas the walls were made of stone covered by wooden segments. The gym had a height of 20 meters and the ceiling consisted of stone and glas. During our measurements no occurrences like people walking through the gym disturbed the equipment. The participants were able to activate the equipment and get out of the gym before the first traces were transmitted. This environment may serve as a best-case indoor scenario, since no other RF-sources were active and in addition wave reflecting structures are kept as simple as possible for an indoor place.

3.2 Parameters

At both places a couple of measurements were performed from which we present the most interesting ones.

One focus of the project was to investigate the impact of receiver antenna diversity on the error behavior of the wireless link. Within each environment measurements were designed to capture this.

At the open field location two measurements were of particularly interest. The first measurement intended to observe error behavior at increasing distances between the Rx and Tx units. For this purpose traces were recorded at certain distance points between 30 meters and 60 meters. At each distance point traces for 3 different modulation types (1MBit BPSK, 2 MBit QPSK and 11MBit CCK) were activated. We will refer to this measurement as **distance** measurement. The environment may be viewed as best-case scenario, and we expected to see increasing error rates with increasing distance and change of modulation schemes. Table 2 shows the variable and fixed parameters of this measurement.

Parameter	Setting
Rx-Tx-Distance	30, 40, 45, 50, 53, 55, and 57 meters
DiversityEnabled	False
Packet Size	1008 Bytes
Modulation Type	1 MBit/s BPSK, 2 MBit/s QPSK, 11 MBit/s CCK
Line-of-Sight ?	True

Table 2: Fixed and *variable* parameters of the **distance** measurement

Parameter	Setting
Rx-Tx-Distance	20 meters
DiversityEnabled	True vs. False
Packet Size	2016 Byte vs. 216 Byte
Modulation Type	1 MBit/s BPSK, 2 MBit/s QPSK, 11 MBit/s CCK
Line-of-Sight ?	True

Table 3: Fixed and *variable* parameters of the **packetsize** measurement

The second measurement performed at the open field kept the distance fixed at 20 meters. Instead the packet size was toggled between a large size (2016 Byte = 56 Chunks) and a small size (216 Byte = 6 Chunks). This was done for the three already mentioned modulation types while for each setting diversity was activated for one trace and deactivated for a second. We will refer to this measurement as **packetsize** measurement. Refer to Table 3 for the fixed and variable parameter settings of measurement **packetsize**.

In the third presented measurement we studied the influence of a non line-of-sight (NLOS) setting versus a line-of-sight (LOS) setting. The Tx-unit was placed inside of the sports gym. The Rx-unit was placed in a hallway leading to the sports gym. The LOS path between Tx- and Rx-unit could be interrupted by a wooden door. The modulation scheme was kept fixed to 2 MBit QPSK, and the distance between the units was fixed at 20 meters. The hallway consisted of the same material as the sports gym. Diversity was activated for five recorded traces and then deactivated for the same amount of traces. We will refer to this measurement as **hallway** measurement. Here the main focus was on how the attenuation in the NLOS setting influences the error behavior compared to the LOS setting. Refer to Table 4 for the fixed and variable parameter settings.

4 Measurement Results

We introduce the following definitions: the *packet loss rate* (PLR) of a given trace denotes the fraction of lost packets (technically speaking, the number of ones in the packet loss indicator sequence PLIS compared to the to-

Parameter	Setting
Rx-Tx-Distance	20 meters
DiversityEnabled	True vs. False
Packet Size	1008 Byte
Modulation Type	2 MBit/s QPSK
Line-of-Sight ?	True vs. False

Table 4: Fixed and *variable* parameters of the **hallway** measurement

tal length of the sequence). The *bit error rate* (BER) of a single trace denotes the fraction of the number of erroneous bits as compared to the overall number of bits in all received packets (clearly not taking packet losses into account). The *packet error rate* (PER) of a single trace gives the fraction of all received packets, which show at least one bit error (and hence would be considered erroneous by a MAC protocol employing a checksum scheme). Finally, the *packet bit error rate* (PBR) is the fraction of erroneous bits with respect to a single packet.

4.1 Distance Measurement

As expected we encounter a situation where the lower modulation types are much more reliable than the 11 MBit CCK modulation type. Up to 50 meters, the 1 MBit BPSK and the 2 MBit QPSK modulation types behave stable in terms of bit error rates and packet error rates, see Figures 2 and 3, where the BER, PER and PLR is shown for varying distance. The PER and PLR seem to be correlated, and both are varying over several orders of magnitude. The BER is always below 10^{-3} for BPSK, while it reaches up to 10^{-2} for QPSK. Having comparably low BERs and high PERs together, we conclude that the bit errors do not typically occur in large clusters, but only in small clusters or as single-/two bit errors (compare Figure 5, where the PBR is shown for a representative trace). For the 11 MBit CCK modulation the picture changes (Figure 4). The BERs and PERs have already high values even for the shortest distances. Furthermore, there is no clear coupling between PER and PLR.

At a distance greater than 50 meters, error rates are quite variable for all modulation types. We encounter distance points such as the 57 meters mark, where the two lower modulation types behave again quite stable, but there are also distance points such as the 53 meter mark, where error rates are very high for all modulation types. A likely explanation is that we placed the setup on points of constructive and destructive interference, respectively.

The error characteristics change for different modulation types and distances. While for the lower modulation types within the first 50 meters errors occur only in single packets (in Figure 6 the PBR vs. the received packet number for a representative trace is shown), for the 11 MBit CCK modulation errors occur throughout almost every packet (Figure 5). This behavior changes after we cross the 50 meters mark. Now errors occur in many packets also for the lower modulation types (Figure 7). Nevertheless the lower the modulation type is the more seldom do errors in packets occur.

As explained in Section 2.3, for a single trace we form its *bit error indicator sequence* (BEIS) and, for a given

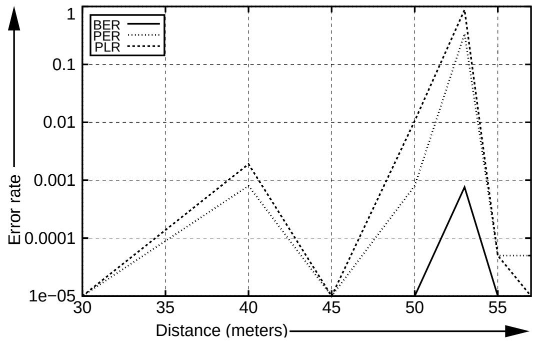

Figure 2: BER, PER and PLR ($\log_{10}$ scale) vs. distance in meters for BPSK modulation (**distance** measurement)

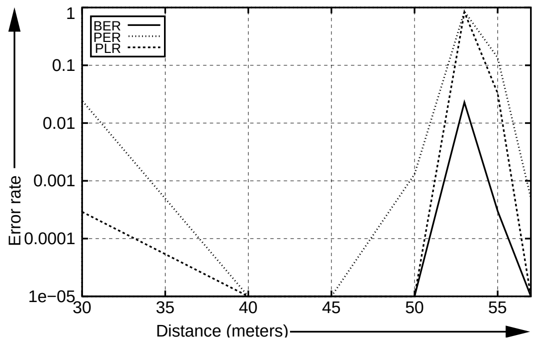

Figure 3: BER, PER and PLR ($\log_{10}$ scale) vs. distance in meters for QPSK modulation (**distance** measurement)

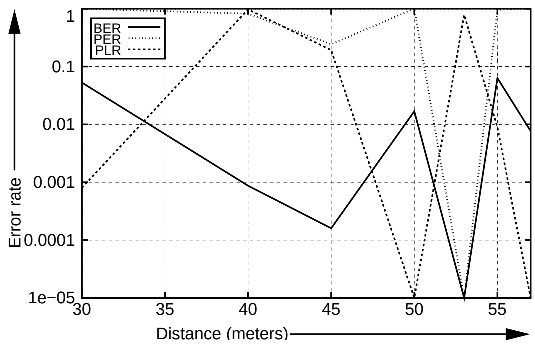

Figure 4: BER, PER and PLR ($\log_{10}$ scale) vs. distance in meters for 11 MBit CCK modulation (**distance** measurement)

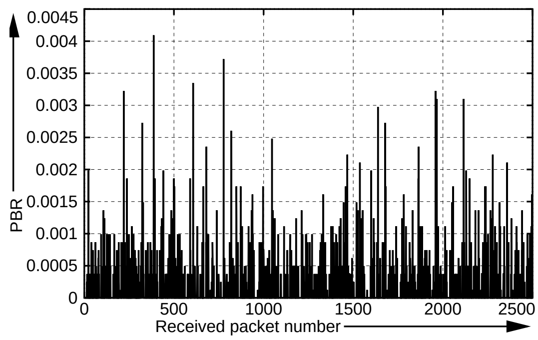

Figure 5: PBR vs. received packet number of **distance** measurement (11 MBit CCK modulation, below 50 meters)

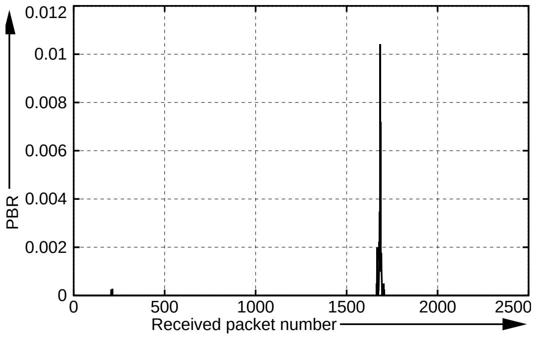

Figure 6: PBR vs. received packet number of **distance** measurement (QPSK modulation, below 50 meters)

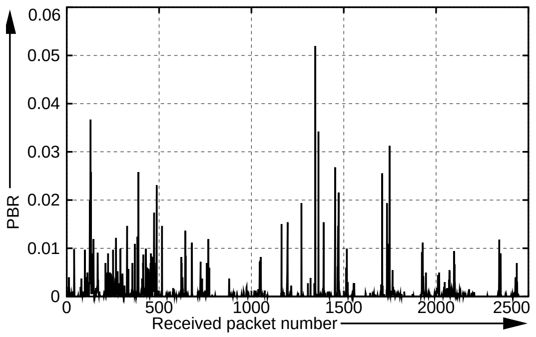

Figure 7: PBR vs. received packet number of **distance** measurement (BPSK modulation, above 50 meters)

burst order k_0, its according *error bursts* and *error-free bursts*, which happen to alternate within a BEIS. One area of interest are the *burst lengths* of these respective bursts, and the question, whether the burst length sequences show any correlation.

For the lower modulation types error burst lengths and therefore error burst behavior shows often a correlated structure, as can be seen in Figure 9, where the autocovariance function of the error burst length sequence for different burst orders k_0 is shown. In contrast, for the 11 MBit CCK modulation type the error burst lengths tend to be uncorrelated (Figure 8). This is independent of the chosen distance point.

4.2 Packetsize Measurement

We did not encounter any error for the 1 MBit BPSK modulation. This is probably due to the small distance

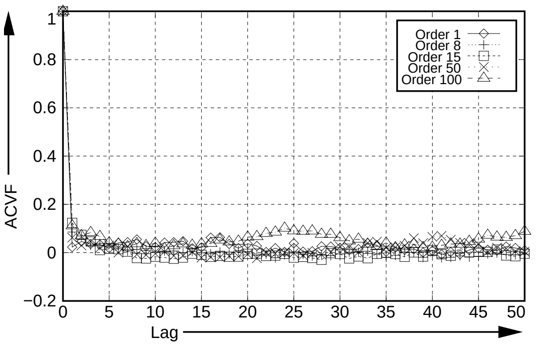

Figure 8: Autocovariance function (ACVF) vs. lag of the error burst lengths for trace 12 of **distance** measurement (11 MBit CCK modulation)

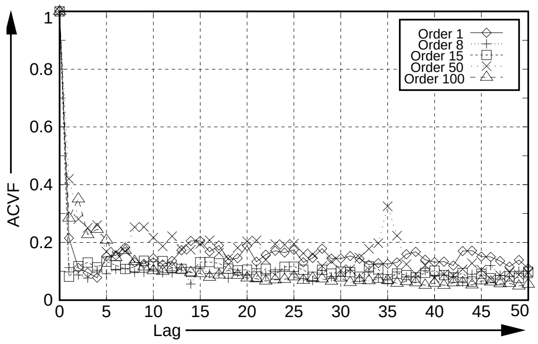

Figure 9: Autocovariance function (ACVF) vs. lag of the error burst lengths for trace 14 of **distance** measurement (QPSK modulation)

Trace	BER	PER	PLR
11 MBit, Large P.	0.01939	0.97687	0.00334
11 MBit, Small P.	0.15491	0.98996	0.1328
2 MBit, Large P.	0	0	0.00029
2MBit, Small P.	0.00024	0.05836	0.05765

Table 5: Results of BER, PER and PLR for traces without diversity of the **packetsize** measurement

between transmitter and receiver unit. For the 2 MBit QPSK modulation scheme, there were only few errors. However, for this case errors tend to be more severe for smaller packets (see Tables 5 and 6). This is also true for the 11 MBit CCK modulation type, where error rates decrease while switching from small to large packetsizes. A possible explanation for the lower BER's for large packets is offered in Figures 14 and 15. Both figures show an *error position histogram* where for a single trace and a fixed bit position within a packet it is counted how many packets show actually a bit error on this position. The bit positions are on the x-axis of these figures. It can be seen that bit errors do not occur at all positions with equal probability. Instead, errors tend to occur at the beginning of a packet, while having a lower density at higher bit positions. Hence, for longer packets the peak at the beginning can be better compensated. Another view on the differences between small and large packets is shown in Figures 10 and 11, where it can be seen that the PBR show a different behavior.

Although receiver antenna diversity leads to a decrease of this effect, error rates do not vanish for the 11 MBit CCK modulation. For the 2 MBit QPSK modulation however, receiver antenna diversity improves the error behavior such that no errors occur any more for both packet types. Furthermore enabling diversity suppresses all packet losses.

The impact of receiver antenna diversity can be stud-

Trace	BER	PER	PLR
11 MBit, Large P.	0.00033	0.4748	0
11 MBit, Small P.	0.00177	0.4077	0
2 MBit, Large P.	0	0	0
2MBit, Small P.	0	0	0

Table 6: Results of BER, PER and PLR for traces with diversity of the **packetsize** measurement

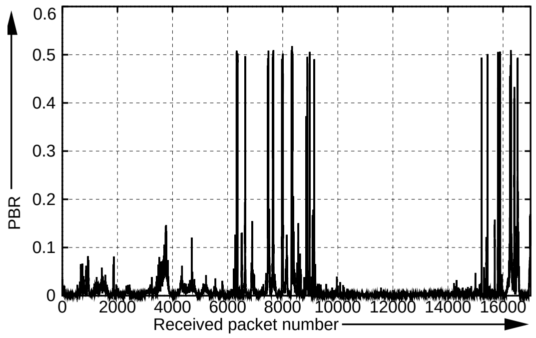

Figure 10: PBR vs. received packet number for a trace of **packetsize** measurement (Large packets, 11 MBit CCK modulation, without diversity)

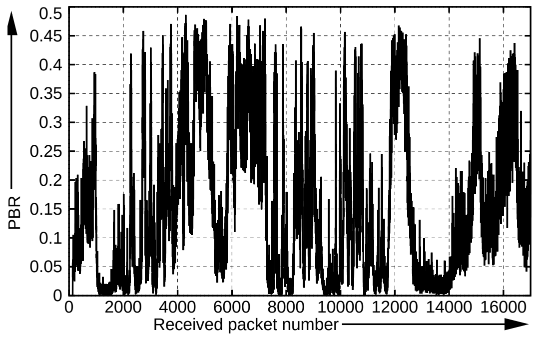

Figure 11: PBR vs. received packet number for a trace of **packetsize** measurement (Small packets, 11 MBit CCK modulation, without diversity)

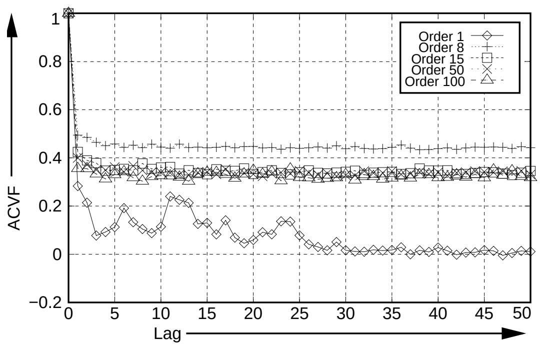

Figure 12: Autocovariance function (ACVF) of the error burst lengths for a trace of the **packetsize** measurement (Large packets, 11 Mbit CCK modulation, without diversity)

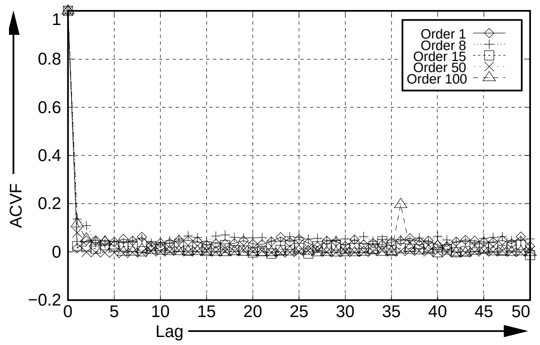

Figure 13: Autocovariance function (ACVF) of the error burst lengths for a trace of the **packetsize** measurement (Large packets, 11 MBit CCK modulation, with diversity)

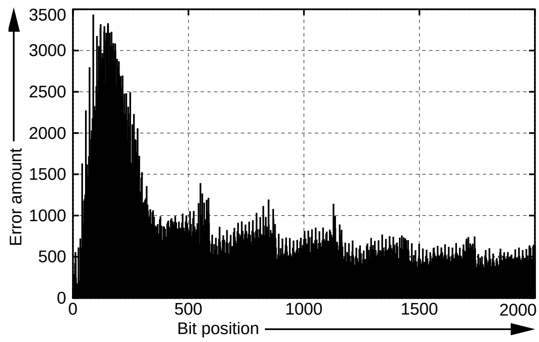

Figure 14: Error position histogram for a trace of the **packetsize** measurement (Large packets, 11 MBit CCK modulation, without diversity)

ied nicely in this measurement. The obvious impact is of course the significant reduction of error rates such as the BER, PER, and PLR. But also receiver antenna diversity decorrelates the error burst lengths. This can be observed in the autocovariance functions of the error burst lengths. In Figures 12 and 13 we show that the autocovariance function of error burst lengths displays more than weak correlation for the case without diversity, while for the case with diversity the burst lengths are nearly uncorrelated. Also, if considering the error position histogram for traces with and without diversity (Figures 14 and 15), errors only occur in the beginning of a packet in the case of active receiver antenna diversity. In contrast, substantial amounts of errors occur also in higher bit positions in the case without diversity.

4.3 Hallway Measurement

Unfortunately, receiver antenna diversity is not always as efficient as presented in Section 4.2, as is indicated by the results of the **hallway** measurement. Typical error rates are shown in Table 7. From this table it can be observed that the BER has comparable orders of magnitude for both the LOS and the NLOS scenario. However, the PERs are quite different, which points to rather different error characteristics. In fact, for the NLOS scenario the PER is much higher than for the LOS scenario. Receiver antenna diversity has no strong impact on this behavior. In Figures 16 and 17 we show for both the NLOS and LOS scenario the PBR vs. received packet

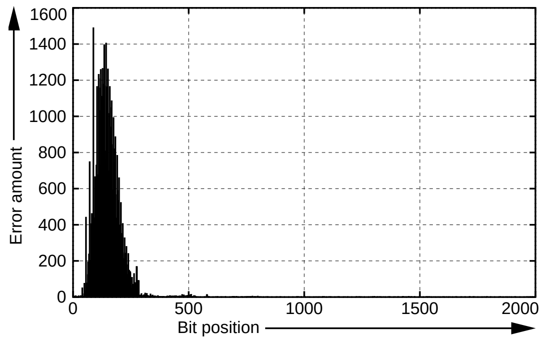

Figure 15: Error position histogram for a trace of the **packetsize** measurement (Large packets, 11 MBit CCK modulation, with diversity)

Trace	BER	PER	PLR
LOS, no diversity	0.00032	0.00125	0
LOS, diversity	0.00036	0.00145	0
NLOS, no diversity	0.00045	0.73345	0
NLOS, diversity	0.00045	0.7342	0

Table 7: BER, PER, and PLR for LOS and NLOS traces of **hallway** measurement with and without Diversity

number. In the LOS-case we have a few erroneous packets, but each one exhibiting a high fraction of erroneous bits, while in the NLOS-case the PBRs are quite low, but almost every packet is erroneous. The high fraction of erroneous bits per packet is due to a phenomenon called "bit-shifted" packets, where the receiver randomly inserts or deletes a bit from the received packet (likely due to problems with achieving and maintaining bit synchronisation). Since this happens typically at the beginning of a packet, the measurement software sees a left- or right-shifted packet, and counts lots of errors. The result of this can be explained by an example (compare the description of the format of the generated packet stream in Section 2.2): consider 11000011 00111100 is transmitted, and the receiver inserts a random bit at the beginning while taking away one bit at the end, finally giving 011000011 0011110. If we XOR both sequences, we obtain 6 bit errors during BEIS computation. This phenomenon is discussed more deeply in [11].

These observed error characteristics are even more obvious if we consider the correlation structure and the error

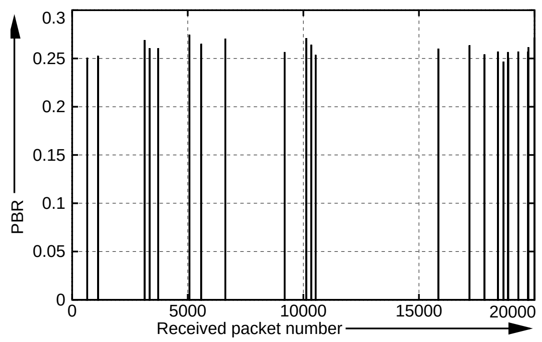

Figure 16: PBR vs. received packet number for a trace of **hallway** measurement (LOS, 2 MBit QPSK modulation, with diversity)

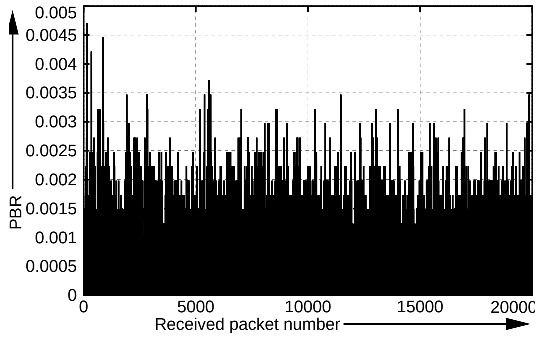

Figure 17: PBR vs. received packet number for a trace of **hallway** measurement (NLOS, 2 MBit modulation, with diversity)

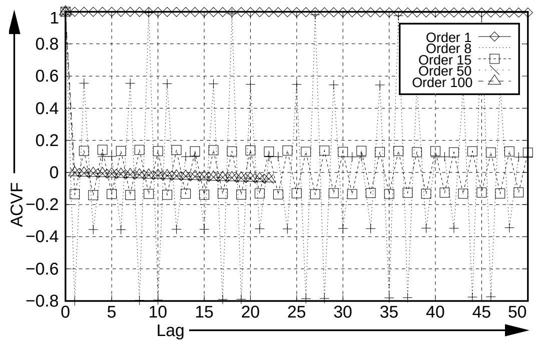

Figure 18: Autocovariance function (ACVF) of the error burst lengths for a trace of **hallway** measurement (LOS, 2 MBit QPSK modulation, with diversity)

position histogram. As can be seen from Figures 18 and 19, interrupting the LOS path leads to a much weaker correlation of error burst lengths. In the LOS setting, error burst lengths are periodically strongly correlated. This can be explained by taking the error position histograms (not shown here) into account: while for the LOS case the bit errors are evenly distributed over the whole packet length (due to many bit-shifted packets), the NLOS case shows pattern of clustering: peaks occur throughout the packet, but on preferred positions, namely multiples of 128.

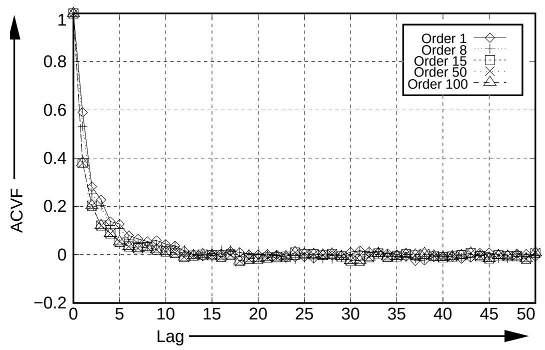

Figure 19: Autocovariance function (ACVF) of the error burst lengths for a trace of **hallway** measurement (NLOS, 2 MBit QPSK modulation, with diversity)

5 RELATED WORK

In a recent paper of Eckhardt and Steenkiste [4] adaptive error correction techniques are applied to WLAN traces, recorded in measurements using WaveLAN (902-928 MHz frequency band, 2 MBit/s QPSK modulation, receiver antenna diversity, see also [3]). They generate a specific UDP/IP packet stream, the underlying WaveLAN uses a CSMA/CA variant without retransmission on the MAC level. An important finding is that at short distances with no interferers the packet loss rate is zero and the packet error rate (PER, rate of packets with at least one bit error) is negligible, while with co-channel interferers the packet loss rates go up to 31%, a lot of truncated packets occur, and the PER is strongly varying. Almost all packets with corrupted bits have fewer than 5% of their bits corrupted. Errors tend to occur in bursts, which are most often restricted to one or two bytes length. The packet loss rate and bit error rate are insensitive to the packet size.

The work described in reference [9] concentrates on tracing and modeling of wireless channel errors on a packet level, incorporating a full UDP/IP protocol stack over WaveLAN (902-928 MHz frequency band, DSSS, QPSK, 2 MBit/s). All interference sources are suppressed. When only the load is varied (in terms of interarrival times for packets of fixed size 1400 bytes), the PER does not change. When varying the packet size, the PER doubles with every 300 byte increase of packet size, reaching $\approx 10^{-3}$ for 1400 bytes. When only varying the distance, the PER doubles every 17 feet, up to ≈ 0.08 at 130 feet.

One of the earliest WLAN packet-level studies is [2]. Again, a 902-928 MHz WaveLAN with 2 MBit/s QPSK, DSSS, and receiver antenna diversity was used. The authors have focused on varying the distance. For increasing distance the PER increases, however, there is a sharp cutoff, since it increases dramatically within a few meters, while before the increase rate was low. They found that bit errors tend to be non-consecutive. They defined two erroneous bits to belong to the same error burst, if they are located in neighbored bytes. Typically only the minimum number of bits for constituting an error burst is erroneous (only one erroneous bit per byte). Furthermore, some error burst lengths are strongly preferred at all distances and packet sizes, e.g. 13 or 14 bits long. The mean bit error rates are found to be "roughly constant" over all packet sizes and distances.

A similar measurement study performed by our group is presented in [10].

6 CONCLUSIONS

In this paper we have presented results from measurements obtained with an IEEE 802.11 DSSS-Compliant radio modem in two different "best-case" environments. These results show clearly, that: a) the bit error behavior depends on the modulation scheme, b) in the simplest multipath environment (**distance** measurement) receiver diversity is of significant help, while in slightly more complex environments (**hallway** measurement) the influence of diversity gets less visible, and c) packet losses are often a serious problem. Specifically the issue of packet losses (see Section 2.3) is an interesting challenge for the design of MAC protocols. In our further research we will explore other environments.

REFERENCES

[1] Richard L. Abrahams. *2.4GHz 11Mbps MACless DSSS Radio HWB1151 Users Guide - AN9835.1*. Intersil, 1999.

[2] D. Duchamp and N.F.Reynolds. Measured performance of wireless lan. In *Proc. of 17th Conf. on Local Computer Networks, Minneapolis*, 1992.

[3] David Eckhard and Peter Steenkiste. Measurement and analysis of the error characteristics of an in-building wireless network. In *Proc. of ACM SIGCOMM'96 Conference*,, pages 243–254, Stanford University, California, August 1996.

[4] David A. Eckhardt and Peter Steenkiste. A trace-based evaluation of adaptive error correction for a wireless local area network. *MONET - Mobile Networks and Applications*, 4:273–287, 1999.

[5] E. O. Elliot. Estimates of error rates for codes on burst-noise channels. *Bell Systems Technical Journal*, 42:1977–1997, September 1963.

[6] E. N. Gilbert. Capacity of a burst-noise channel. *Bell Systems Technical Journal*, 39:1253–1265, September 1960.

[7] James Gross, Michael Jaeger, and Andreas Willig. Measurements of a Wireless Link in different RF-isolated Environments. TKN Technical Report Series TKN-01-005, Telecommunication Networks Group, Technical University Berlin, June 2001. http://www-tkn.ee.tu-berlin.de/publications/tknrreports.html.

[8] Intersil. *HFA3860B Data Sheet, File Number 4594.1*, 1999.

[9] Giao T. Nguyen, , Randy H. Katz, Brian Noble, , and Mahadev Satyanarayanan. A trace-based approach for modeling wireless channel behavior. In *Proceedings of the Winter Simulation Conference*, Coronado, CA, December 1996.

[10] Andreas Willig, Martin Kubisch, Christian Hoene, and Adam Wolisz. Measurements of a Wireless Link in an Industrial Environment using an IEEE 802.11-Compliant Physical Layer. *IEEE Transactions on Industrial Electronics*, 2001. accepted for publication.

[11] Andreas Willig, Martin Kubisch, and Adam Wolisz. Measurements and Stochastic Modeling of a Wireless Link in an Industrial Environment. TKN Technical Report Series TKN-01-001, Telecommunication Networks Group, Technical University Berlin, March 2001. http://www-tkn.ee.tu-berlin.de/publications/tknrreports.html.