Performance Evaluation of the IEEE 802.11 Infrared Physical Layer

Adriano J. C. Moreira [(*)]
Rui T. Valadas and A. M. de Oliveira Duarte [(**)]
(*) Dept. de Informática, Campus de Azurém, Universidade do Minho, 4800 Guimarães, Portugal, Email:adriano@uminho.pt, tel: +351 53 510148, fax: +351 53 510250
(**) Dept. de Electrónica e Telecomunicações, Universidade de Aveiro, Portugal

Abstract

Since 1990, the IEEE 802.11 working group is preparing a standard for wireless local area networks. The current version of the standard includes two radio spread spectrum physical layers and one infrared physical layer, all sharing a common medium access control layer (MAC).

This paper describes some of the work performed during the specification of the infrared (IR) physical layer (PHY) that resulted in many contributions to the definition of the standard.

Among the many technical solutions that have impact in the performance of the physical layer, the choice of a modulation method and the frame format are the most important. For the modulation method, Pulse Position Modulation (PPM) was adopted because of its high power efficiency.

This paper describes the work performed regarding the definition of the frame format that better matches the particular characteristics of the indoor wireless infrared channel. The proposed frame format is presented and its performance is evaluated through analytical calculation of the frame error rate. For each field of the frame there is an explanation of the proposed format based on the probabilities of imitation, non-detection and/or error.

The frame format here described was proposed to the IEEE 802.11 working group and was adopted and included in the current version of the standard.

1. Introduction

In June, the IEEE 802.11 working group started the definition of a standard for wireless local area networks. In June 1997, the IEEE Draft 6.1 of the standard was approved for publication [1]. This standard specifies a Medium Access Control layer (MAC) that can be used with one of three Physical Layers (PHY): two of them using radio with spread spectrum modulation and one using infrared signals.

The Infrared Physical Layer (IR PHY) relies on optical signals in the 800-900 nm band and direct detection of the optical signals to transmit data at 1 or 2 Mbps using the diffuse mode of propagation [2]. The physical layer specification defines two radiation patterns for the emitter and one for the receiver. The two specifications are intended for two different scenarios: one for desktop units and the other for handheld portable devices.

The modulation method adopted for this PHY is Pulse Position Modulation (PPM). PPM was adopted because it is one of the most power efficient modulation methods, which is appropriate for a channel where the propagation losses are very high [3]. For each data rate a different order of PPM is used: 16-PPM for 1 Mbps and 4-PPM for 2 Mbps [4]. This scheme allows that the same basic pulse is used for the two data rates, which facilitates the transceiver implementation, especially its synchronisation functions. The same optical peak power is used for both data rates, allowing the same maximum distance between any pair of transceivers to be attained for the two specified data rates. This feature can be used to trade-off communication speed for power consumption, which is of particular importance for battery operated terminals.

In this paper the IR PHY frame format is described (section 2) and its impact on the physical layer performance is evaluated through the analytical calculation of the frame error rate (section 3). In section 4 some results are presented, which have been used to find the receiver sensitivity specification to include in the IEEE 802.11 IR PHY standard.

2. The infrared physical layer frame format

Three frame formats have been proposed during the definition of the IEEE 802.11 IR PHY, and are shown in Figure 1.

First International Symposium on Communication Systems and Digital Signal Processing, Sheffield Hallam University, UK, 6-8 April 1998

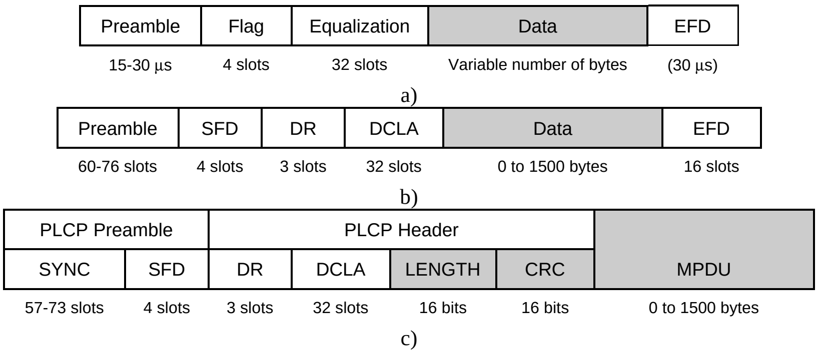

Figure 1. The three proposed frame formats for the IR PHY.

The frame format shown in Figure 1a) was proposed as part of a complete specification for the IR PHY [5]. This specification was for a physical layer supporting transmission at 1 Mbps only and using 16-PPM. The second frame format was proposed as part of a second complete specification for the IR PHY [4]. This latter specification introduced some new features such as: (i) support for multiple data rates and specification for 1 and 2 Mbps; (ii) the use of different orders of PPM for each data rate; (iii) a shorter preamble; (iv) a frame field to carry the PHY type or data rate (two types defined for 1 and 2 Mbps); (v) explicit End-of-Frame Delimiter in contrast with the implicit one specified in the first proposal. The third frame format [6] (Figure 1c) was proposed to make the PHY specification conformant with the basic rules of the IEEE 802.11 working group [7]. The major differences compared with the second frame format are: (i) the inclusion of a field to carry the frame length; (ii) the inclusion of a field to implement error detection over the length field; (iii) the absence of an End-of-Frame Delimiter field.

One characteristic common to the three frame formats is that not all fields are transmitted using PPM. Only the shadowed fields in Figure 1 are transmitted using PPM. All the other fields are transmitted using OOK-NRZ and the same basic pulse as PPM.

2.1. The first frame format

The first frame format is composed by the following fields:

a) `Preamble`: the purpose of this field is to carry synchronisation information at the slot level, to allow the receiver to detect the presence of a frame being transmitted (a function similar to carrier sense) and to allow automatic gain control adjustment. This field consists of a sequence of alternated presence and absence of a pulse in consecutive time slots. The duration of this field is variable between 15 and 30 µs.

b) `Flag`: the purpose of this field is to carry synchronisation information at the PPM symbol, bit and frame levels and, simultaneously, to carry the type of physical layer that generated the frame. For this field four words are defined (`0000`, `0001`, `0010`, `0011`), with the word `0010` being used for the proposed PHY.

c) `Equalization`: this field is used to allow the receiver circuits to adjust to the high difference between the average signal level of the Preamble and of the Data fields (the first has an average signal level equal to half the peak level, while the last has an average signal level which is 1/16 of the peak level). Its format is a sequence of two 16-PPM symbols with the pulse in slot 4.

d) `Data`: this field carries the MAC Protocol Data Unit, and no limits are imposed on its length. The bits are transmitted using 16-PPM.

e) `End-of-Frame Delimiter` (EFD): for this PHY, the frame end is found by waiting for a period of silence longer than 30 µs.

2.2. The second frame format

Two major reasons induced the proposal of this second frame format: the support for multiple data rates and performance. While based on the first frame format, it presents the following differences:

a) The length of the SYNC field (former Preamble) is shorter to reduce the overhead, to reduce the optical emitter stress and to improve the frame error rate by reducing the probability of imitation of the SFD field during the Preamble;

b) The former Flag field was replaced by two other fields: Start-of-Frame Delimiter (SFD) and Data Rate (DR), which separates the two functions previously assigned to the Flag field;

c) The PPM symbol, bit and frame synchronisation functions are implemented through the dedicated SFD field. The format of this field is the word `1001`, and should be used for all present and future PHYs.
d) A new DR field is included to carry the data rate at which the frame information fields are transmitted, or the PHY type. For this PHY two words are defined for 1 and 2 Mbps: `000` and `001`, respectively.
e) Since different orders of PPM are used at 1 and 2 Mbps, the DCLA (d.c. Level Adjustment, former Equalization field) field should have a different format for each order of PPM. For 1 Mbps its format is two 16-PPM symbols with the pulse at the 9th slot whereas for 2 Mbps its format is four 4-PPM symbols with the pulse at the 3rd slot.
f) An explicit End-of-Frame Delimiter is used. Its format is a 16 slot word, which reduces the time needed to find the frame end and also reduces the probability of error in its detection. Its format is the word `0000011011011011`.

2.3. The third frame format

The third frame format was proposed to make the IR PHY conformant with the IEEE 802.11 basic rules [7]. One of these rules says that, while looking for the beginning of a frame, the minimum Hamming distance between the SFD pattern and the bit sequence being checked is 4. As a result, a false frame start should not be declared for any pattern of 3 or less errors. While the SFD proposed with the second frame format was chosen in order to minimise the frame error rate (FER), it does not respect the previously described rule since the minimum Hamming distance was only 2. For a Hamming distance of 4, the SFD should be at least 9 slots long, which in turn would increase significantly the FER.

The adopted solution is shown in Figure 1c). The SFD field was left intact and two new fields were included before the Data field. The first one is a 16 bit Length field carrying the Data field length and the second one a 16 bit CRC field to implement error detection over the Length field. This solution solves the previously described problem as follows: if a false frame start is declared due to imitation of the SFD during the Preamble or due to noise, both the Length and CRC fields are incorrect and the error is detected by the CRC.

Since a Length field is included, no EFD is necessary, and it was removed from this third frame format.

3. Performance calculation

Most of the work around the frame format design was accomplished to improve the physical layer performance. In this section the performance of the IR PHY is estimated through the calculation of the FER of the adopted frame format (Figure 1c). Where relevant, the methods used to define the format of the fields are described.

For the frame format shown in Figure 1c), the FER is given by:

$$FER = 1 - P_{SYNC} \cdot P_{SFD} \cdot P_{DR} \cdot P_{LENGTH} \cdot P_{CRC} \cdot P_{MPDU} \tag{1}$$

where P_{SYNC}, P_{SFD}, P_{DR}, P_{LENGTH}, P_{CRC} and P_{MPDU} are the probabilities that the fields *SYNC*, *SFD*, *DR*, *LENGTH*, *CRC* and *MPDU* are correctly detected.

In the receiver, the frame detection process is triggered by the Carrier Sense (CS) mechanism. After a signal is detected in the medium and slot synchronisation is acquired, the receiver starts to search for a valid *SYNC* field. A valid *SYNC* is declared when a sequence of *L1* bits with a pattern similar to *SYNC* is found. This function can be implemented by digital correlation of the received signal with the *SYNC* pattern. The longer the bit sequence (*L1*) the lower will be the probability of false *SYNC* detection (the SYNC field can be imitated by noise), and the higher will be the probability that the *SYNC* is not detected due to errors. The probability that a valid *SYNC* is not detected (and the frame is missed) is a function of the probability of error for each position checked, and these probabilities are not statistically independent. However, a simple lower bound for the probability that the *SYNC* field is correctly detect can be computed and is given by:

$$P_{SYNC} > 1 - \left[1 - (1 - BER)^{L1}\right]^{Nt1} \tag{2}$$

where *BER* is the bit error rate, *L1* is the correlation window length and *Nt1* is the number of statistically independent trials available to search for the Preamble. *Nt1* is a function of the *SYNC* field length (*L0*), correlation window length (*L1*) and of the part of the *SYNC* field wasted in the CS and synchronisation processes (*L2*), and is given by:

$$Nt1 = MaxInteger[(L0 - L2 - L1 + 1) / L1] \tag{3}$$

After a valid *SYNC* is detected, the receiver starts to search for the *SFD* field. This process starts during the *SYNC* field. For the *SFD* field to be correctly detected, it should not be imitated during the *SYNC* due to errors and should be detected while looking at the correct position. The probability of correct detection is then given by:

$$P_{SFD} = (1 - P_{imit\_SFD}) \cdot (1 - P_{miss\_SFD}) \tag{4}$$

where P_{imit_SFD} is the probability that the *SFD* is imitated during the *SYNC* and P_{miss_SFD} is the probability that the *SFD* is missed while looking at the correct position due to errors.

The probability that the *SFD* is missed while looking at the correct position is given by:

$$P_{miss\_SFD} = 1 - (1 - BER)^{L3} \tag{5}$$

where *L3* is the *SFD* field length. The probability of imitation is given approximately by:

$$P_{imit\_SFD} \approx 1 - \prod_{i=-Nt2}^{-1} [1 - P_{imit}(dif(i))] \tag{6}$$

where the imitation probabilities for each trial are considered statistically independent, and with:

$$P_{imit}(dif) = BER^{dif} \cdot (1 - BER)^{L3 - dif} \tag{7}$$

where *Nt2* is the number of trials and *dif* is the Hamming distance between the *SFD* word and the bit sequence being inspected.

The probability that the *SFD* field is correctly detected depends on the field length (*L3*) and on its pattern. A long field reduces the probability of imitation but increases the probability of missing its detection. A study was carried out to find the field format that maximises the probability of correct detection and the results are shown in Table 1.

SFD pattern	Hamming distance	P_{imit_SFD}	P_{miss_SFD}	$1-P_{SFD}$
`001, 100`	1	5.73×10^{-6}	8.60×10^{-7}	6.59×10^{-6}
`110`	1	5.45×10^{-6}	8.60×10^{-7}	6.31×10^{-6}
`1001, 1100`	2	3.20×10^{-12}	1.15×10^{-6}	**$1.15´10^{-6}$**
`00110`	2	1.64×10^{-12}	1.43×10^{-6}	1.43×10^{-6}
`01100`	2	1.73×10^{-13}	1.43×10^{-6}	1.43×10^{-6}
`100111`	3	$<1\times10^{-15}$	1.72×10^{-6}	1.72×10^{-6}

Table 1. Probabilities of incorrect detection of the SFD field, for several field formats (*Nt2*=40; SNR=3.19).

Based on these results, the pattern `1001` was adopted for the *SFD* field.

As soon the *SFD* field is detected, no more searching processes are required. The next three slots are for the *DR* field, and the probability of correct detection for this field is given by:

$$P_{DR} = (1 - BER)^{L4} \tag{8}$$

where *L4* is the field length. After the detection of both the *SFD* and *DR* fields, bit, PPM symbol and frame synchronisation are achieved. The validation that this is not a false frame is achieved by detecting, demodulating and checking the *LENGTH* and *CRC* fields. The probabilities of correct detection of these fields are:

$$P_{LENGTH} = (1 - BER2)^{L5} \tag{9}$$

$$P_{CRC} = (1 - BER2)^{L6} \tag{10}$$

where *BER2* is the bit error rate for the L-PPM modulated fields and *L5* and *L6* are the length of the two fields. The probability that the CRC field does not detect a false SFD is very low (in practice it is neglected), since the probability that errors in a 16 bit word not being detected by a 16 bit CRC is negligible [9].

The probability of correct detection of the MPDU field is given by:

$$P_{MPDU} = (1 - BER2)^{L7} \tag{11}$$

where *L7* is the *MPDU* field length.

As stated earlier, the *SFD*, *DR* and *DCLA* fields are transmitted using OOK-NRZ while the other fields are transmitted using L-PPM modulation. Since different modulation methods are used, the bit error rate is different for these two field groups. Considering an AWGN channel without optical interference, as specified in the IEEE 802.11 IR PHY "Receiver Sensitivity" section, for OOK-NRZ, the BER is given by:

$$BER_{OOK} = \frac{1}{2} Erfc\left[\frac{n_T}{\sqrt{2} \cdot s_T}\right] \tag{12}$$

with $s_T = \sqrt{q \cdot I_B / T_s}$ and $n_T = P_{avr} \cdot R$, where P_{avr} is the average received optical power, R is the photodetector responsivity, T is the slot period, q is the electron charge and I_B is the average background photocurrent [2].

For L-PPM two types of receivers can be used: a more complex maximum likelihood (ML) receiver [10] or a simpler threshold (TH) receiver. The gain of the ML receiver over the TH receiver is about 1.5 dB. For the ML receiver, the BER is given by [8]:

$$BER_{PPM\_ML} = \frac{2^{k-1}}{2^k - 1} \cdot \left[1 - \frac{1}{\sqrt{p}} \cdot \left(\frac{1}{2}\right)^{L-1} \cdot \int_{-\infty}^{+\infty} e^{-x^2} \cdot \left(1 + Erf\left[\frac{\sqrt{2} \cdot s_T \cdot x + L \cdot n_T}{\sqrt{2} \cdot s_T}\right]\right)^{L-1} \cdot dx\right] \tag{13}$$

where k is the PPM order and L is the number of slots per PPM symbol ($L = 2^k$). For the TH receiver, the BER is given by [8]:

$$BER_{PPM\_TH} = \frac{2^{k-1}}{2^k - 1} \cdot \left[1 - \left(\frac{1}{L} \cdot P_{01} \cdot (1 - P_{10})^{L-1} + \sum_{n=1}^{L} \frac{1}{n} \cdot \binom{L-1}{n-1} \cdot (1 - P_{01}) \cdot P_{10}^{n-1} \cdot (1 - P_{10})^{L-n}\right)\right] \tag{14}$$

where:

$$P_{01} = P_{10} = \frac{1}{2} \cdot Erfc\left(\frac{L \cdot n_T}{2 \cdot \sqrt{2} \cdot s_T}\right) \tag{15}$$

are the probabilities of non-detection of a transmitted pulse and of detection of a pulse not transmitted, respectively, which are equal if the threshold level is half the peak level at the sampling instant.

4. Results

The receiver sensitivity is specified as the minimum average irradiance (power per unit area) required for a frame error rate of 4×10^{-5}, for a frame with 512 bytes of data. For the IR PHY, the required signal-to-noise ratio (SNR) for a FER of 4×10^{-5} was calculated and the results are shown in Table 2. These results show a gain of the ML receiver over the TH receiver of about 1.23 dB (and not 1.5 dB, since the first 3 frame fields use OOK-NRZ and not L-PPM). They also show that, for the TH receiver, the FER is dominated by the probability of error in the MPDU field, while for the ML receiver the FER is dominated by the probabilities of error in the SFD and DR fields.

Similar calculations were performed for the first and second frame formats described in section 2. When using a TH receiver, the three formats have the same performance since the MPDU error probability dominates the FER. When using a ML receiver, the third format shows an improvement of 0.37 dB over the first format and an improvement of 0.13 dB over the second one.

Based on the above analysis, the receiver sensitivity was calculated for the conditions described in the IEEE 802.11 IR PHY specification, which consider a channel with an average background light level of 0.1 mW/cm^2 and without optical interference [11] (due to artificial light). The results are shown in Figure 2 as a function of the photodetector active area, for a frame with 512 bytes of data transmitted at 1 Mbps. Two curves are shown in Figure 2: one for a receiver with a negligible front-end noise and another

Probabilities	**Conditions**	**TH receiver** (SNR=3.89 dB)	**ML receiver** (SNR=2.66 dB)
$1 – P_{SYNC}$	*L0=57; L1*=8; *L2*=20; *Nt1*=3	≈0.0	5.76×10^{-14}
$1 - P_{SFD}$	*L3*=4; *Nt2*=20	8.61×10^{-9}	1.93×10^{-5} •
P_{miss_SFD}	"	8.61×10^{-9}	1.93×10^{-5}
P_{imit_SFD}	"	≈0.0	4.43×10^{-10}
$1 - P_{DR}$	*L4*=3	6.46×10^{-9}	1.45×10^{-5} •
$1 - P_{LENGTH}$	*L5*=16 bits	1.55×10^{-7}	2.43×10^{-8}
$1 - P_{CRC}$	*L6*=16 bits	1.55×10^{-7}	2.43×10^{-8}
$1 – P_{MPDU}$	*L7*=4096 bits (512 bytes)	3.97×10^{-5} •	6.22×10^{-6}
FER	-	4.00×10^{-5}	4.00×10^{-5}

Table 2. Probabilities of error in the detection of the frame fields.

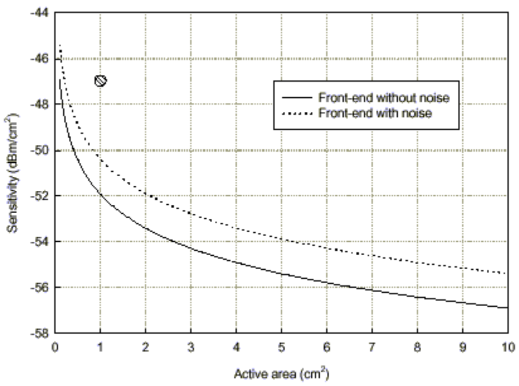

Figure 2. Receiver sensitivity versus detector active area.

for a receiver producing a noise level equal to the noise induced by the background light. These two curves limit the zone for the receiver sensitivity of a practical receiver.

These results were used to estimate a reasonable value for the receiver sensitivity to include in the IR PHY specification. Assuming an active area of 1 cm^2 (to limit the receiver cost), a value of –47 dBm/cm^2 was adopted for the receiver, which includes a 3 dB margin for implementation imperfections and to allow smaller photodetectors to be used. For the 2 Mbps reception mode, the receiver sensitivity specification is 6 dB higher.

5. Conclusions

The frame format of the IEEE 802.11 infrared physical layer was described. This format was designed through analytical calculation of the frame error rate. The format was adopted by the IEEE 802.11 working group and is part of the current standard for wireless LANs.

6. References

1 IEEE 802.11 Draft 6.1 of the "IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specification", 26 June 1997

2 Fritz R. Gfeller and Urs Bapst, "Wireless In-House Data Communication via Diffuse Infrared Radiation", Proceedings of the IEEE, Vol. 67, No. 11, November 1979

3 John Robert Barry, "Wireless Communication Using Non-Directed Infrared Radiation", PhD Dissertation, University of California at Berkeley, USA, 1992

4 Adriano J. C. Moreira, Cipriano Lomba, Rui Aguiar, Rui T. Valadas, A. M. de Oliveira Duarte, "IR PHY Proposal", Submission to the IEEE 802.11 Standardisation Project, Doc: IEEE P802.11 - 94/96, May 1994

5 Roger Samdahl - Photonics Corporation, "Baseband IR PHY Proposal", Submission to the IEEE 802.11 Standardisation Project, Doc: IEEE P802.11 - 94/56, 1994

6 Adriano J. C. Moreira, Rui T. Valadas, A. M. de Oliveira Duarte, "Proposed Revision of the Infrared Baseband Frame Format", Submission to the IEEE 802.11 Standardisation Project, Doc: IEEE P802.11 - 94/153, July 1994

7 "Project Authorization Request", Doc: IEEE P802.11-91/58, May 1991

8 Adriano J. C. Moreira, "Sistemas de Transmissão Ópticos em Espaço Livre para Ambientes Interiores", PhD Thesis, University of Aveiro, Portugal, February 1997

9 Mischa Schwartz, "Information Transmission, Modulation, and Noise", Third Edition, McGraw-Hill, 1985

10 John G. Proakis, Digital Communications, McGraw-Hill International Editions, Second Edition, 1989

11 Adriano J. C. Moreira, Rui T. Valadas, A. M. de Oliveira Duarte, "Optical Interference Produced by Artificial Light", ACM Wireless Networks, Nº 3, 1997, pp. 131-140