The Infrared Physical Layer of the IEEE 802.11 Standard for Wireless Local Area Networks

Rui T. Valadas, António R. Tavares, and A. M. de Oliveira Duarte
University of Aveiro, Portugal
Adriano C. Moreira, University of Minho, Portugal
Cipriano T. Lomba, ENT Empresa Nacional de Telecomunicações, SA

ABSTRACT The new IEEE 802.11 standard for wireless local area networks defines a specification for an infrared physical layer. This article gives an overview of infrared technology and describes the IEEE 802.11 specification in detail, presenting a historical perspective of its development. The infrared physical layer was designed for diffuse systems supporting two data rates (1 and 2 Mb/s) and includes provisions for a smooth migration to higher data rates. The specification is suitable for low-cost transceivers but allows interoperability with higher-performance systems. The main application envisaged for IEEE 802.11 infrared wireless local area networks is ad hoc networks.

There is a growing interest in wireless local area networks [1] as a consequence of the large-scale utilization of personal computers and mobile communications. The portable personal computer integrates computational power and mobility in a single platform and introduces the need for accessing communication networks without the restrictions imposed by cables.

The interest in wireless local area networks (WLANs) has recently found an echo in international standardization bodies. In the United States, the executive committee of the IEEE 802 project created, in July 1990, the IEEE 802.11 group to work on the specification of a WLAN for different technologies, including radio and infrared. The standard was approved in June 1997 [2]. An essential characteristic of the IEEE 802.11 specification is that there is a single medium access control (MAC) sublayer common to all physical (PHY) layers. This feature will allow easier interoperability among the many physical layers that are expected to be defined in the future, driven by the fast technological progress in this field. There are presently three PHY layers in the standard: infrared, frequency hopping spread-spectrum, and direct sequence spread-spectrum. Infrared and radio can be considered complementary technologies for the support of WLANs. Infrared technology [3] is well suited for low-cost low-range

applications, such as ad hoc networks (small area networks set up for a short period only).

IEEE 802.11 NETWORKS

Figure 1 illustrates a scenario, in an educational environment, for the application of an IEEE 802.11 network. In this scenario, still a bit futuristic, each student owns a portable computer that he uses as a notebook and, in general, as a learning tool, exploring the applications resident on disk. During classes the portables of the students and the teacher are interconnected through an IEEE 802.11 network. This will allow for cooperative work within the classroom, and accessing applications and files resident on the school server. Although many of these features are already provided by the school backbone (eventually a cabled network such as Ethernet or token ring), the fact that portables are no longer dependent on a cable to access the network enables these facilities to be provided irrespective of where the classes take place, allowing for a more flexible and efficient working environment. Moreover, it promotes better integration of the work performed inside and outside the classroom.

IEEE 802.11 networks are part of the overall IEEE 802 architecture. An important consequence is that all functions unique to WLANs must be assigned to the PHY or MAC layers. The entities that compose an IEEE 802.11 network are represented in Fig. 2. A group of stations associated to establish direct communication form a basic service set (BSS). The area occupied by stations from a BSS is the basic service area (BSA) or cell. Cells can overlap partially or totally, or be physically disjoint. The coverage of a cell depends on the propagation environment and the transceiver characteristics. Large areas can be covered using multiple cells, and thus multiple BSSs. The system used for the interconnection of a group of BSSs is the distribution system (DS). The DS can be, for example, an Ethernet or asynchronous transfer mode (ATM) network. The entity that allows stations to access the DS is the access point (AP). The AP includes all functions of a station. A group of BSSs interconnected by a DS forms an extended service set (ESS).

The simplest IEEE 802.11 network is an independent BSS. This type of network is sometimes called an ad hoc network. Examples of ad hoc networks are classrooms (as in Fig. 1), small meetings, and conference rooms. IEEE
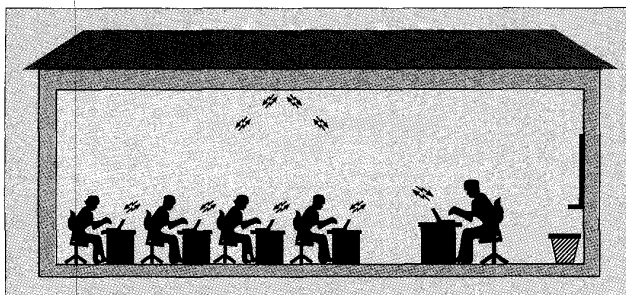

■ **Figure 1.** *A scenario for an IEEE802.11 network.*

802.11 networks based on infrared technology are well adapted to the requirements of ad hoc networks.

An Overview of the IEEE 802.11 Medium Access Protocol

Wireless networks require several additional functions to those in cabled networks, to adapt to the particular characteristics of the transmission channel and the mobility of the stations. The IEEE 802.11 MAC sublayer includes, besides the basic MAC function, additional functions for fragmentation and reassembly of frames when the quality of the transmission channel is poor, association and reassociation of stations with APs, temporal synchronization for the support of delay-sensitive applications, power management for battery-operated stations, and data rate adaptation.

Medium access is based on a carrier sense multiple access with collision avoidance (CSMA/CA) protocol. According to CSMA/CA, a station that has a packet ready for transmission starts by sensing the channel (*carrier sense*). If the channel is free for a period longer than an interframe space, denoted distributed interframe spacing (DIFS), the station transmits immediately. If the channel is busy, the station keeps sensing the channel until the channel is free for a period longer than DIFS. When the channel finally becomes free, the station further delays its transmission using a binary backoff algorithm, similar to that of Ethernet, before making another transmission attempt. This is the *collision avoidance* aspect of the protocol. The CSMA/CA protocol is also enhanced with *immediate positive acknowledgment* and *reservation* mechanisms. According to the acknowledgment mechanism, when a station receives a valid data packet (DATA) addressed to itself, it must confirm the reception of this packet by sending an acknowledgment (ACK) mini-packet to the source station. To ensure that the destination station does not have to contend for medium access when sending the ACK, the destination station uses an interframe space shorter than DIFS, called short interframe spacing (SIFS). The use of immediate positive acknowledgment as part of medium access allows for error detection by the source station and is required due to the high levels of noise and interference impairing the transmission channel. Positive acknowledgment can only be used with directed packets. Therefore, the transmission of multicast and broadcast packets is less robust. The reservation mechanism is based on the exchange of Request-To-Send (RTS) and Clear-To-Send (CTS) mini-packets between the source and destination stations prior to sending DATA packets. RTS and
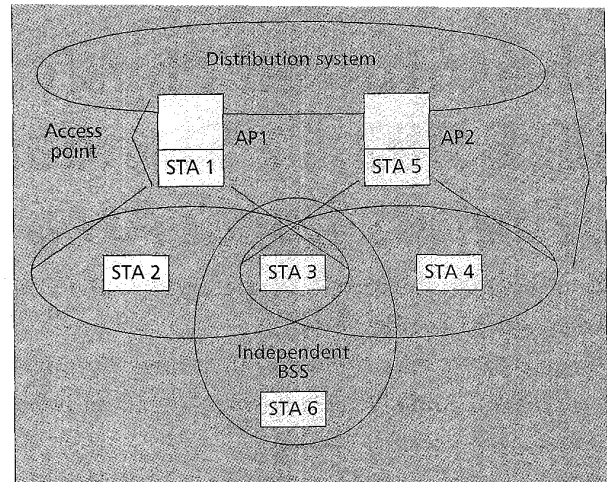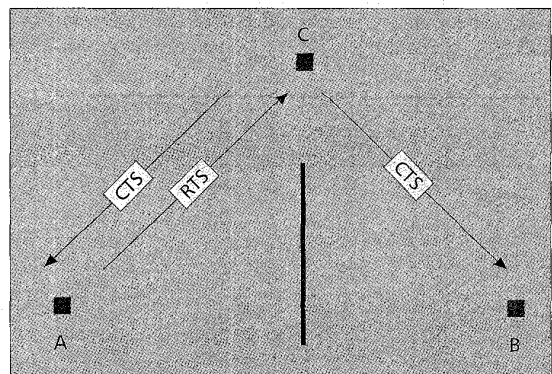

■ **Figure 2.** *IEEE 802.11 network components.*

CTS broadcast information on the time interval the channel will be occupied from RTS (or CTS) ending to ACK ending. Each station maintains a timer called a network allocation vector (NAV) that indicates in each instant the remaining channel occupation time. The NAV at all BSS stations must be updated whenever an RTS or CTS packet is received. Stations can only transmit after the NAV countdown has expired. In this way, RTS and CTS mini-packets reserve the channel for data transmission. The medium access mechanism is illustrated in Fig. 3.

The reservation mechanism can be used to combat the hidden station problem. If station B is not able to sense the activity on the channel produced by station A, it is said to be *hidden* from A. In Fig. 4 station B is hidden from station A (but not from C). Station A transmits to C, but station B will not sense this activity. In this case, station B can transmit freely, thus interfering with the transmission from A. However, if A and C exchange mini-packets RTS and CTS with an indication of the time the channel will be occupied, station B, although not directly sensing station A, will be informed through the CTS mini-packet of the time the channel will be occupied by A and C, and will not interfere with the transmission of the DATA packet sent by A. Note that this mechanism will not avoid collisions since RTS mini-packets can still be sent simultaneously from A and B. However, RTS collisions are less harmful than DATA collisions in terms of throughput performance since RTS mini-packets are relatively short.
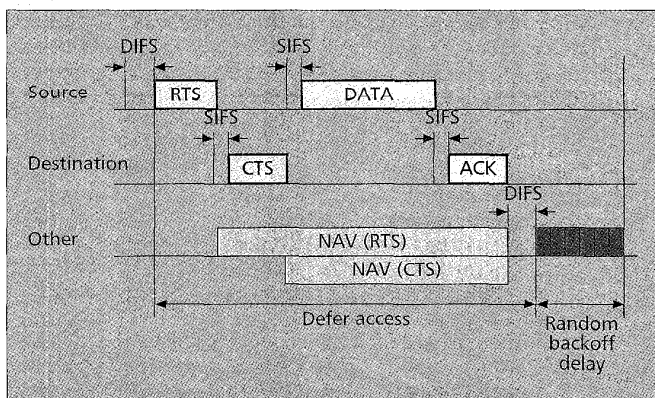
■ **Figure 3.** *Medium access protocol.*

■ **Figure 4.** *The use of RTS and CTS mini-packets to combat the hidden station problem.*

4 bits @ 1 Mb/s
1 symbol 16-PPM
8 bits @ 2 Mb/s
4 symbols 4-PPM
250 ns
4-PPM symbol
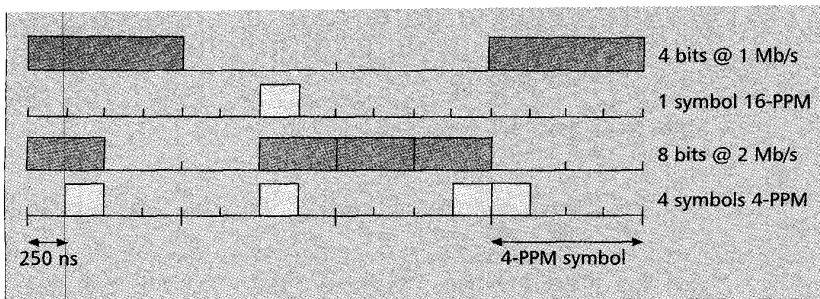
■ **Figure 5.** *PPM signals at 1 and 2 Mb/s.*

INFRARED TECHNOLOGY

Wireless infrared links are based on *intensity modulation* and *direct detection* of the optical carrier. This is similar to many optical fiber links. Intensity modulation is performed by varying the current of a laser diode or an LED. Direct detection is performed by PIN photodiodes or APDs that produce an electrical current proportional to the incident optical power. The opto-electronic components used more frequently are LEDs and PIN photodiodes.

Infrared radiation has properties very similar to visible radiation. Typical surfaces of indoor environments are, in general, good reflectors of infrared radiation (the main exceptions are dark and transparent surfaces). This property has two important consequences. First, infrared radiation can propagate through multiple reflections on the propagation environment. Propagation through multiple paths can provoke time dispersion of the received pulses, which is called *multipath dispersion*. The effects of multipath dispersion are observed as intersymbol interference. Second, having obstacles between the emitter and the receiver can provoke a significant attenuation of the collected optical power, which is called *shadowing*.

The bandwidth of an infrared link is mainly determined by multipath dispersion; intersymbol interference is not significant for data rates up to approximately 10 Mb/s.

The illumination sources of indoor environments (sunlight and artificial light generated by incandescent and fluorescent lamps) radiate in the same wavelengths as the infrared data signal. Also, typical intensity levels of the ambient light collected at the photodetector are usually much higher than data signal intensity levels. Ambient light provokes shot noise due to the random nature of the photodetection process. Moreover, artificial light provokes interference due to periodic variations of light intensity [4]. These variations can occur at the double of the main frequency or at the switching frequency of the electronic ballasts of fluorescent lamps. In general, for low and moderate data rates the ambient noise is the main factor degrading the performance of wireless infrared links [5].

Besides ambient noise, wireless infrared links can also be affected by front-end noise and by nonoptical electromagnetic interference of various types, such as that provoked by power supplies. In general, nonoptical electromagnetic interference can be significantly reduced by properly shielding the front-end and photodetectors, filtering the power supply, and using differential configurations at the front-end pre-amplifier. The front-end noise has two components:
• Thermal noise associated with the photodetector bias resistor
• Thermal noise associated with the channel impedance, in case of an FET-based front-end, or shot noise associated with the base current, in case of a BJT-based front-end
In general, these two degrading factors can be made insignificant relative to ambient noise.

When operating an infrared wireless system, users become directly exposed to IR radiation. Excessive exposure to IR radiation may originate ocular lesions. The extension of these injuries depends on several factors: the intensity and duration of exposure, the wavelength of the radiation, and the dimensions of the emitting area. There are international safety regulations that define the maximum levels to which a user can be exposed without suffering any lesion. In practice, these standards limit the average optical power emitted by the source to a few hundred milliwatts. Other factors pushing for the use of low levels of emitted optical power are the power consumption of battery-operated stations and the poor conversion efficiency of LEDs.

Shadowing and ambient noise drive the need to emit high levels of optical power. However, as discussed before, the optical power level is restricted a priori by international safety regulations and by the power consumption of the stations. Therefore, the transmitted signal must be processed to allow its detection with the lowest possible signal-to-noise ratio. Pulse position modulation (PPM) is generally accepted as the technique that offers the best characteristics for transmission in this type of transmission channel [3]. PPM maps binary words into pulse positions. In L-PPM, a word with k bits is coded into one of the $L = 2^k$ positions of an L-PPM symbol.

Infrared systems can be classified according to the way stations establish an optical path between themselves. There are three types of systems: point-to-point, quasi-diffuse, and diffuse [3]. The IEEE 802.11 specification was developed for diffuse systems. In a diffuse system the optical path between emitter and receiver is established through multiple reflections in the surfaces of the propagation environment. Both emitter and receiver are nondirective and have arbitrary orientation. Ideally, the received optical level must be independent of the position and orientation of the receiver.

THE INFRARED PHYSICAL LAYER SPECIFICATION

The infrared PHY layer supports two data rates: 1 and 2 Mb/s. The specification of two data rates is aimed at allowing:
• A smooth migration to higher data rates
• Asymmetric operation of the BSS
There is a different PPM scheme for each data rate: 16-PPM for 1 Mb/s and 4-PPM for 2 Mb/s. The purpose of this feature is to ensure that the basic pulse is the same at both data rates, which minimizes the additional complexity introduced by the 2 Mb/s data rate. The emitter and receiver circuits can be almost identical (in particular, the same front-end can be used at both data rates), and the most significant enhancements are required on the synchronization circuits [6]. The PPM signals at 1 and 2 Mb/s are represented in Fig. 5. The duration of each pulse is 250 ns, and the peak optical power is 2 W. Therefore, the average optical power is 250 mW at 2 Mb/s and 125 mW at 1 Mb/s.

THE FRAME FORMAT

The main function of the physical layer frame is to carry the MAC protocol data unit (MPDU) while supporting the various features of the physical layer. When developing the

infrared PHY-layer frame format, several aspects were taken into account. The frame format should:

■ **Figure 6.** *Infrared physical-layer frame format.*

- Allow interoperability with future infrared PHY layers (most probably at higher data rates)
- Support multiple data rates
- Be optimized for performance through minimization of the frame error rate (FER)

Three infrared PHY-layer frame formats were submitted to the IEEE 802.11 Working Group [6–8]. The format finally adopted is shown in Fig. 6.

Only the shadowed fields in Fig. 6 are transmitted using L-PPM. The first three fields are transmitted as on-off keying (OOK) signals, using the same 250 ns pulse as the shadowed fields. Although this solution penalizes the FER performance, it minimizes the overhead introduced by the PHY header. In addition, this feature enables interoperability with future PHY layers if they choose to adhere to the format of the first three fields. The Direct Current Level Adjustment (DCLA) field is used to allow the receiver circuits to adjust to the difference between the average signal level of the OOK and L-PPM-based fields. A different DCLA format is specified for each data rate.

The SYNC field is used for clock recovery (i.e., slot synchronization), for carrier sense and to allow stabilization of the receiver circuits (e.g., the AGC). Its format corresponds to a 2 MHz clock signal. The Start of Frame Delimiter (SFD) field performs frame alignment (i.e., it delimits the beginning of subsequent fields). The Data Rate (DR) field is used to identify the type of PHY layer that transmitted the frame. In the current standard there are two types, corresponding to the 1 and 2 Mb/s data rates. Since different orders of PPM are used for each data rate, this field implicitly carries synchronization information for decoding the PPM symbols. The DR field is three slots long, allowing for specification of six additional PHY layers.

At the receiver, carrier sense declaration requires the detection of a signal with a predefined power level and the detection of a valid SYNC pattern. The SYNC pattern is a fragment of the overall SYNC field sequence. After a signal with a predefined power level is detected in the medium and slot synchronization is acquired, the receiver starts searching for a valid SYNC pattern. This function can be implemented by digital correlation of the received signal with the SYNC pattern. The pattern can be imitated by noise, which may result in frame misalignment. The longer the pattern, the lower the probability of false SYNC detection but the higher the probability of no SYNC detection due to errors. The probability that a valid SYNC is not detected is negligible if a short pattern is used (eight slots) [9].

The SFD pattern requires careful design due to its strong impact on the FER. The search for a valid SFD field follows carrier sense declaration and should start during the SYNC field. During the search process, the SFD can be imitated due to errors in the SYNC field. Also, the SFD will not be detected if it suffers from errors. Thus, the probability that the SFD is correctly detected depends on the probability of imitation and the probability of error of the SFD. The format that maximizes the probability that this field is correctly detected can be
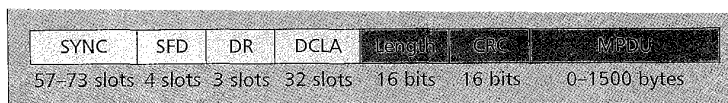
either word "1001" or "1100" [9]. The 1001 word was adopted in the standard.

The previous solution for the SFD format, although optimized in terms of FER, still does not comply with one of the basic rules of IEEE 802 networks, which imposes that the Hamming distance of the word used for frame alignment purposes should be at least 4. The adopted SFD pattern has a Hamming distance of only 2. For a Hamming distance of 4, the SFD should be at least nine slots long, which would significantly increase the FER. In order to comply with this rule without degrading the FER, it was decided to include a frame length field (LENGTH) protected by an error detection field (CRC). If a false SFD is declared, the LENGTH and CRC fields will have incorrect values with very high probability, which will be detected by the CRC. If the declaration of a valid frame is delayed until these two fields are correctly detected, the misalignment probability is very low, while keeping the FER at a minimum. The LENGTH field is also used to delimit the end of the frame.

RECEIVER SENSITIVITY AND PERFORMANCE ISSUES

The performance of the infrared PHY layer can be estimated through the calculation of the FER of the adopted frame format. Simple calculation of the BER can lead to erroneous conclusions. For the frame format shown in Fig. 6, the FER is given by [9]

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

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.

The basic requirements of IEEE 802.11 networks mandate that the FER should be lower than 4×10^{-5} for frames with 512 bytes of data. In the infrared PHY layer, receiver sensitivity is defined as the minimum irradiance (optical power per unit area) required to achieve this FER specification, under a stationary ambient light irradiance level of 0.1 mW/cm². This FER value is achieved at a signal-to-noise ratio of 2.66 dB, assuming an active area of 1 cm² and the use of a maximum likelihood PPM receiver. The resulting error probabilities of the frame fields are presented in Table 1.

The results show that the FER is dominated by the probabilities of error of the SFD and DR fields and not of the MPDU field, even after the optimization of the SFD field. The receiver sensitivity was specified at 2×10^{-5} mW/cm² for 1 Mb/s and 8×10^{-5} mW/cm² for 2 Mb/s. These values include a margin for implementation imperfections and factors not included in the calculations.

The receiver sensitivity specification considers only stationary ambient noise, not the interference produced by artificial light. This results from the need to produce a specification where the ambient light conditions could easily be reproduced for conformance testing purposes. Clearly, the definition and reproduction of interference light conditions would be a difficult task. However, an optical receiver developed with little or no attention to the optical interference prob-

Probabilities	ML receiver (SNR = 2.66 dB)
$1 - P_{SYNC}$	5.76×10^{-14}
$1 - P_{SFD}$	1.93×10^{-5}
$1 - P_{DR}$	1.45×10^{-5}
$1 - P_{LENGTH}$	2.43×10^{-8}
$1 - P_{CRC}$	2.43×10^{-9}
$1 - P_{MPDU}$	6.22×10^{-6}
FER	4.00×10^{-5}

■ **Table 1.** *Probabilities of error in the detection of the frame fields.*
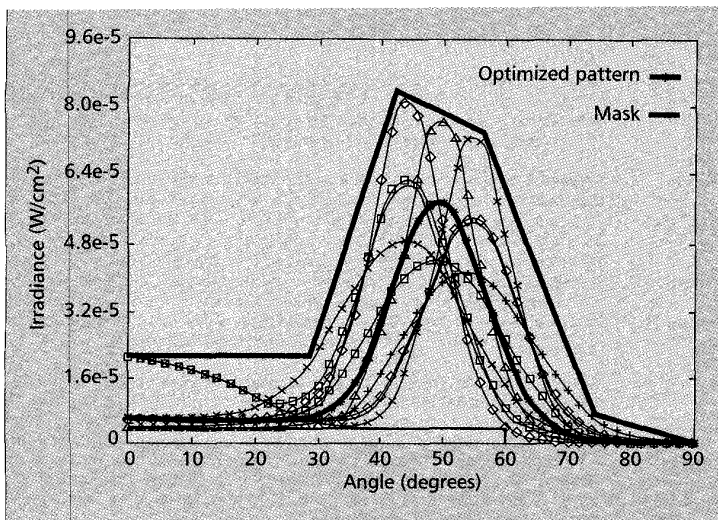
■ **Figure 7.** *The mask of the emitter radiation pattern.*

lem may be a conformant receiver, but with very degraded performance operating under artificial light.

There is no maximum receiver sensitivity value specified in the standard, opening the possibility for competition among different manufacturers. There are several options for increasing the link performance while maintaining conformance with the IEEE 802.11 specification:
• Increase the receiver active area
• Use of angular diversity at both the emitter and receiver [3, 10]
• Use of interference cancellation techniques [11]

EMITTER RADIATION PATTERN

The specification of an emitter radiation pattern [12] had in view:
• Minimizing propagation losses
• Allowing system operation in a large set of dissimilar propagation environments
• Ensuring its conformance with the safety standards for infrared radiation [13]

The low cost requirement restricted attention to the use of LEDs on the emitter array. An array with all LEDs oriented vertically would result in an excess of irradiance around the source position. Clearly, the number, orientation, radiation pattern, and emitted power of each LED in the array are parameters that can be optimized to make power distribution as uniform as possible, thus minimizing channel propagation losses. Ideally, the irradiance should have a constant value slightly higher than the receiver sensitivity over the whole cell area. The emitter radiation pattern was specified assuming that stations would be moving on a plane parallel to the ceiling.

The emitter radiation pattern is specified in terms of a mask that defines the irradiance bounds as measured by a test receiver. The algorithm used to define the mask [11] searched for an optimized radiation pattern, while accounting for manufacturing tolerances on the orientation and optical characteristics of the LEDs and, for the diversity of propagation environments, which was considered by restricting to two extreme cases in terms of propagation losses: large open plants and walled rooms.

The emitter radiation pattern mask is shown in Fig. 7. The mask is defined by lower and upper irradiance limits, for each angle between the emitter axis and the axis of a test receiver, positioned 1 m away from the emitter. The irradiance is normalized to the optical peak power.

The optimized emitter radiation pattern corresponds to an array of 11 commercially available LEDs: one LED with half-power angle (hpa) = 41° and 15 mW of total optical power, vertically oriented, and 10 LEDs, with hpa = 9° and 11 mW of total optical power, oriented at 50° with the vertical. Figure 7 also shows several radiation patterns corresponding to different tolerances: 10 percent for the angle of the LEDs with the vertical, 25 percent for the hpa of the narrower LEDs, and 50 percent for the hpa of the larger LEDs.

SUPPORT OF THE OPTIONAL DATA RATE

For the purpose of minimizing the hidden station problem, the same cell coverage is specified for 1 and 2 Mb/s. This requires approximately the same energy per symbol at 1 and 2 Mb/s. Since the pulse density of a 4-PPM signal is four times that of a 16-PPM signal, the average optical power emitted at 2 Mb/s is approximately 6 dB higher than that emitted at 1 Mb/s. However, since frames at 2 Mb/s require half the time to be transmitted, the penalty in terms of optical energy per frame is approximately 3 dB. Nonetheless, the additional power consumption associated with transmission at 2 Mb/s led to specifying this data rate as optional. In summary, transmission at 2 Mb/s is optional, while transmission at 1 Mb/s and reception at both rates is mandatory.

CONCLUSIONS

The new IEEE 802.11 standard for wireless local area networks defines a specification for an infrared physical layer. We have described this specification in detail, giving a historical perspective of its development and providing some background on infrared-technology-specific issues. The infrared physical layer was designed for diffuse systems supporting two data rates (1 and 2 Mb/s) and includes provisions for smooth migration to higher data rates. The specification is suitable for low-cost transceivers but allows interoperability with higher-performance systems. The main application envisaged for IEEE 802.11 infrared wireless local area networks is ad hoc networks.

ACKNOWLEDGEMENTS

REFERENCES

[1] K. Pahlavan, T. Probert, and M. Chase, "Trends in Local Wireless Networks," *IEEE Commun. Mag.*, Mar. 1995, pp. 88–95.
[2] Project IEEE 802.11, "IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," Draft 6.1, June 1997.
[3] J. Kahn and J. Barry, "Wireless Infrared Communications," *Proc. IEEE*, vol. 85, no. 2, Feb. 1997, pp. 265–98.
[4] A. Moreira, R. Valadas and A. Duarte, "Optical Interference Produced by Artificial Light," *ACM Wireless Networks J.*, vol. 3, May 1997, pp. 131–40.
[5] A. Moreira, R. Valadas, and A. Duarte, "Performance of Infrared Transmission Systems Under Ambient Light Interference," *IEE Proc. J. Optoelect.*, vol. 143, no. 6, Dec. 1996, pp. 339–46.
[6] A. Moreira et al., "IR PHY Proposal," Doc. IEEE P802.11 — 94/96, 1994.
[7] A. Moreira, R. Valadas, and A. Duarte, "Proposed Revision of the Infrared Baseband Frame Format," Doc. IEEE P802.11 — 94/153, 1994.
[8] R. Samdahl, "Baseband IR PHY Proposal," Doc. IEEE P802.11 94/56, 1994.

[9] A. Moreira, R. Valadas, and A. Duarte, "Performance Evaluation of the IEEE 802.11 Infrared Physical Layer," *Proc. 1st Int'l. Symp. Commun. Sys. and Digital Sig. Processing*, Sheffield, U.K., Apr. 6–8, 1998, pp. 10–15.

[10] R. Valadas, A. Tavares, and A. Duarte, "Angle Diversity to Combat the Ambient Noise in Indoor Optical Wireless Communication Systems," *Int'l. J. Wireless Info. Networks*, vol. 4, no. 4, Oct. 1997, pp. 275–88.

[11] A. Moreira, R. Valadas, and A. Duarte, "Reducing the Effects of Artificial Light Interference in Wireless Infrared Transmission Systems," *Proc. IEE Colloq. Opt. Free Space Commun. Links*, Savoy Place, London, U.K., Feb. 19, 1996, pp. 5/1–5/9.

[12] C. Lomba, R. Valadas, and A. Duarte, "Radiation Pattern Specification for the Baseband IR PHY," Doc. IEEE P802.11 — 94 / 173, 1994.

[13] C. Lomba, R. Valadas and A. Duarte, "Safety Issues of the Baseband IR PHY," Doc. IEEE P802.11 — 94/174, 1994.

BIOGRAPHIES

RUI T. VALADAS (rv@ua.pt) received a Licenciatura degree from Instituto Superior Tecnico, Lisbon, Portugal, in 1986, and a Ph.D. degree from the University of Aveiro, Portugal, in 1996, all in electrical engineering. He joined the University of Aveiro in 1986, where he is now an assistant professor in the Department of Electronics and Telecommunications. He participated in the specification of the IEEE 802.11 "Wireless Access Method and Physical Layer Specification" standard. He led the development of a network interface card for a wireless infrared local area network, part of the ESPRIT.6892 POWER project. He is scientific leader of a project to develop advanced techniques for wireless infrared systems. His main interests are in the research of angle diversity and adaptive rate techniques to overcome the effects of ambient light and also in the effect of multi-user interference on the performance of multiple access protocols for wireless infrared systems.

ADRIANO C. MOREIRA (adriano@uminho.pt) received a Licenciatura degree in electronics and telecommunications engineering in 1989 and a Ph.D. degree in electrical engineering in 1997, both from the University of Aveiro, Portugal. He joined the University of Minho, Portugal, in 1996, where he is now an assistant professor in the Informatics Department. He was a voting member of the IEEE 802.11 working group (Wireless Access Method and Physical Layer Specification) where he participated in the specification of the infrared physical layer. His main areas of interest are in indoor optical channel characterization, modulation methods for optical wireless transmission systems, wireless local area networks, and telemetry for the industrial environment.

CIPRIANO T. LOMBA (cipriano.lomba@ent.efacec.pt) received a Ph.D. degree in electrical engineering from the University of Aveiro, Portugal, in 1997, an M.Sc. degree in electronics engineering from the University College of North Wales, U.K., in 1991, and a Licenciatura in electronics and telecommunications, from the University of Aveiro in 1989. He contributed actively in the IEEE 802.11 standardization activities, and has published several papers in international journals and conferences. Currently, he is project leader at the R&D Department of ENT Empresa Nacional de Telecomunicações, SA., a manufacturer of telecom equipment in Portugal.

ANTONIO R. TAVARES (tavares@ua.pt) received a Licenciatura degree in electronics and telecommunications engineering from the University of Aveiro in 1991. From 1991 to 1993 he worked as a post-graduate student in the Department of Electronics and Telecommunications at the University of Aveiro. Since November 1993 he has been a Ph.D. student in the same department. His research interests include wireless infrared communication systems, angle diversity receivers, and the characterization of optical noise and interference due to ambient light.

A. MANUEL DE OLIVEIRA DUARTE (duarte@ua.pt) received a Licenciatura degree in electrical engineering from the University of Coimbra in 1976, M.Sc. and Ph.D. degrees in telecommunications systems and electrical engineering, sciences in 1981 and 1984, respectively, from the University of Essex, U.K. He joined the University of Aveiro in 1978 where he is now an associate professor with the Department of Electronics and Telecommunications. In addition to his lecturing activities, he conducts research in several telecommunications-related domains (broadband networks and services, techno-economics of telecommunications, social impact of telecommunications, etc.). He coordinates the Broadband Systems Group of the University of Aveiro and is co-responsible for Integrated Communications and Multimedia Systems at the Institute of Telecommunications at Aveiro, a joint venture between Portugal Telecom and the University of Aveiro.