IEEE 802.11
Wireless Access Methods and Physical Layer Specifications

PPM INFRARED PHY STRAWMAN PROPOSAL

Roger Samdahl

Photonics Corporation

Introduction

In recent submissions, Moreira et al [1] and Lomba et al [2] have established that 16 PPM can be used to advantage in an infrared network provided that data rates are kept low enough to avoid unacceptable levels of pulse spreading due to multipath dispersion. Our development work in this field supports this conclusion. We believe that 16 PPM can be successfully employed at data rates up to approximately 5 MHz in systems designed to cover a standard conference room or classroom environment.

On the theoretical side, we have developed a radiosity model (depending only on Lambertian, diffuse reflections and without any time of flight capability) which allows calculation of steady state optical intensity levels within a bounded area. Analyses using this model have produced results similar to those presented in the referenced paper [2].

We too have been interested in analyzing the effects of radiator and detector orientation and interactions of room IR power, room size and reflectivities. Figures 1 and 2 show the differences between operation with a vertically oriented transmitter array and one arranged so that the LED's are tilted at 70 degrees to the normal (and spaced uniformly in the horizontal plane). I have little to add to the excellent work that has been presented other than the following observations:

- The effects of increasing the field of view of the detector can be significant, even in the limit as the FOV approaches 90 degrees.

- The reflectivity of wall, ceiling and floor surfaces is of great importance. As shown in Figure 3, with the proper conditions of transmitted power, room size and reflectivity, the IR attenuation with distance has a fairly gentle slope, apparently due in part to diffuse reflections. Under these conditions, a small change in reflectivity can have the effect of dropping signal levels below the limit of detectability over a large part of the room. The solid base line represents the effective detector sensitivity; the two curves show IR optical signal levels for wall and floor reflectivities that vary by 2 to 1. Such variations in reflectivity have been measured during test runs that have been performed at a variety of operating sites.

Typical IR Distributions

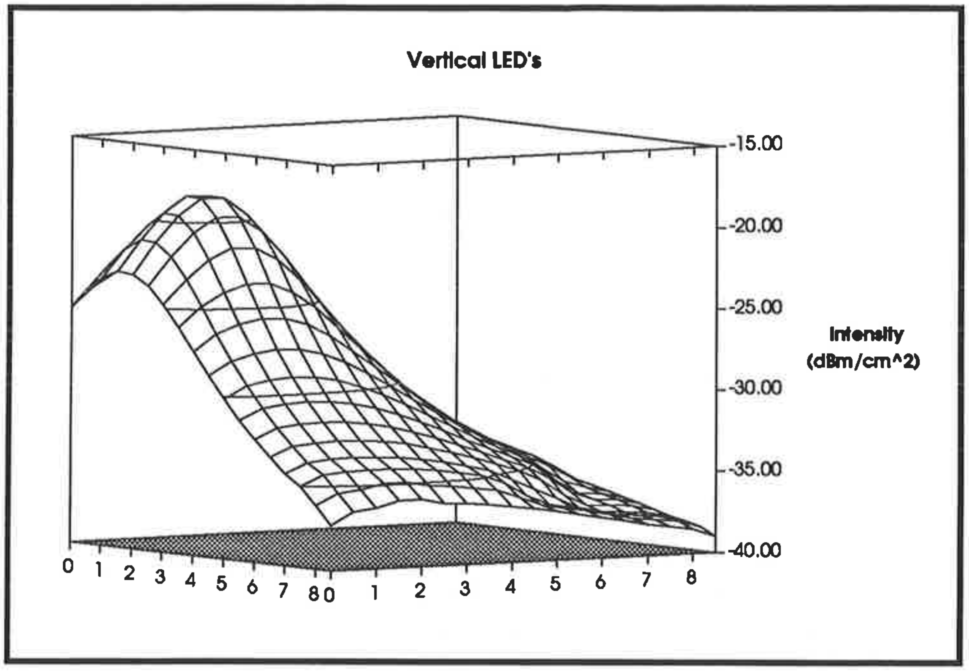

Figure 1.

The graph shown in Figure 2 represents the IR intensity developed by the PHY device within a 30 by 30 foot room with optimized LED positions. It corresponds to the intensity per unit area, normal to the horizontal, measured one meter above the floor.

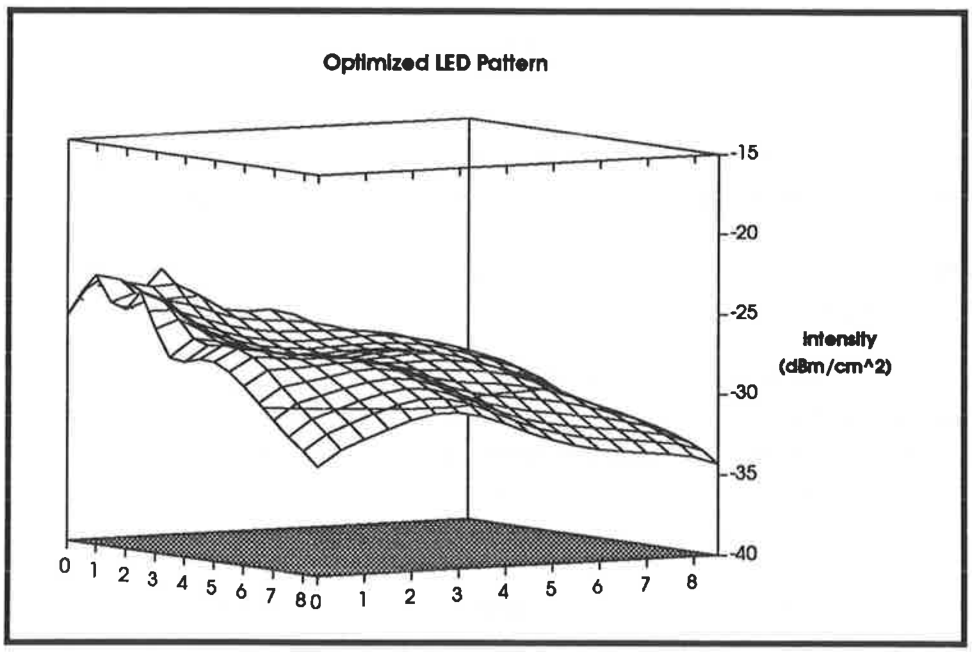

Figure 2.

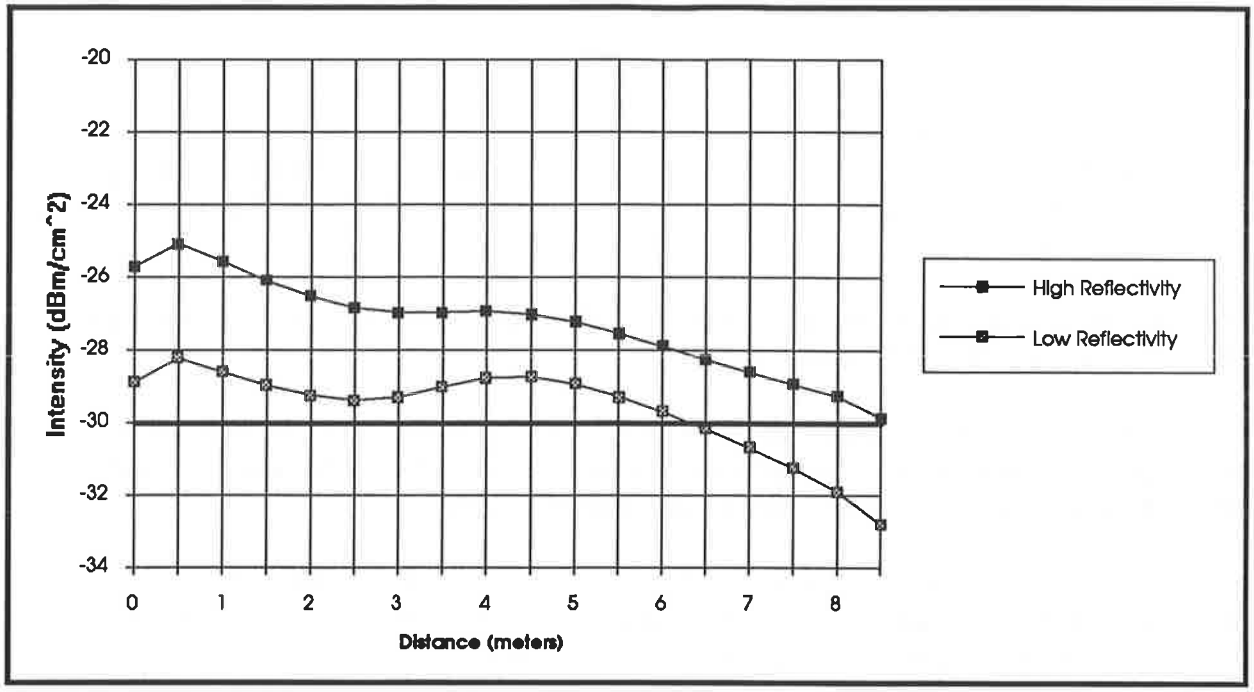

Figure 3.

- It is typically true that as the maximum operating range limit is approached, a large part of the link losses are due to the cosine dependence of the receiving PIN diode detector(s). Under these conditions, an array of receiving PIN diode detectors, tilted away from the horizontal, can provide improved performance.

- Care must be taken to model the expected distribution of transmitter and receiver positions when units are used in a working environment. This is especially true of pen based notepads which are frequently operated with the writing surface held as a clipboard. Optimization that is tuned to one particular orientation will probably fail under these conditions.

We have confirmed the IR intensity predictions of our model with practical tests both in open areas where wall reflections are minimized, and in smaller, closed areas where the wall effects are more important. We have also been able to demonstrate bit error rates that closely match the predictions made in the referenced paper.

Based on the excellent analysis provided by Moreira et al, and supported by our own modeling and practical testing, we are proposing a strawman specification for a PPM PHY layer for IR systems.

A Strawman PPM PHY Specification

Overview

The PHY layer that is described here is of relatively simple design, but it provides a minimum level of performance and several features that we feel will be important as IR networks begin to appear. This is particularly true of features that allow overlapping, dissimilar IR systems to operate in a friendly fashion.

This PHY layer was developed for use with a CSMA/CA MAC protocol, but it is not limited in application to any specific MAC architecture.

We have chosen to implement a 16 state PPM (or PPM(4)) system as a compromise between minimum practical pulse widths, allowable current limits for available LED's, and improved signal to noise characteristics afforded by PPM systems with more states.

The systems that have been built around this PHY architecture have operated at 1 MBPS. Based on experimental evidence, we believe that the same techniques can be extended to operation approaching 5 MBPS in agreement with reported pulse spreading characteristics for modeled systems operating within similar sized rooms. Therefore, this PHY is being presented as a specific solution for IR PPM systems operating at data rates 1, 2 and 4 MBPS. Beyond that point, inter-symbol interference will make simple PPM systems unworkable, in spite of their signal to noise advantages, and other encoding schemes will provide better performance and cost characteristics.

Several assumptions have been made in the development of this specification:

- The proposed PHY is assumed to be implemented in an electrical device that is connected to a network interface controller (NIC) that can be 'flow controlled' by the PHY. The choice of the NIC and its configuration are more related to the MAC specification than to the PHY, but there are a couple of minimum requirements that it must meet. The NIC must include Request to Send and Clear to Send lines or their equivalents. The Clear to Send line must be controlled by the PHY to halt traffic from the NIC to the PHY. Typically, the delays that will be introduced by the PHY will be short (few tens of microseconds), but this capability is required to minimize the PHY collision window. Secondly, the NIC must be clocked by the PHY for transfers in either direction. This allows the PHY to synchronize the delivery of bytes or bits (as required) to and from the PHY.

- The interface is assumed to be bi-directional but not full duplex. The non-channelized nature of the IR system makes duplex operation effectively impossible and results in the inability of IR systems to provide any sort of collision detection.

- It is assumed that incoming IR traffic will always be given priority over outbound IR traffic and that every effort should be made to avoid damaging incoming IR frames. The system assumes that transactions are by frame or packet transfers.

- The proposed PHY provides a carrier sense output that can export its state to the attached MAC. Generally, the MAC should provide flow control to the outgoing packet stream using this carrier sense capability to detect incoming traffic. The ability of the PHY to provide its own flow control is intended to allow minimizing the collision window inherent in CSMA/CA systems. The PHY can operate with a collision window of less than 1 usec while a software only based MAC control would be limited to several tens of usecs.

PHY Layer Specification Summary

Data Rate	1, 2 or 4 MBPS	
Data Encoding	Modified Gray Code	See Table 1
Transmitter IR Freqency	850 to 900 nm	
Transmitted Pulse Width	230 ± 10 nsec	
Transmitted Pulse Rise Time	75 ± 5 nsec	
Transmitted Pulse Fall Time	75 ± 5 nsec	
Transmitter Position Jitter	75 ± 5 nsec	
Receiver Pulse Filtering	125 ns < Tpw < 320 ns	
Error Detection	Illegal Symbol	
Receiver Sensitivity	-30 dBm/cm^2	With -10 dBm/cm^2 Ambient
Max. Received Pulse Width	320 nsec	
Max. Receive Pulse Rise Time	100 nsec	
Max. Receive Pulse Fall Time	100 nsec	
Max. Receive Pulse Jitter	100 nsec	
Preamble	4 Symbols of Clock w/Mod	See Table 2
Equalization	2 Symbols of 02h	
Data	Any length	
End of Frame Delimiter	None	
Interframe Gap	Minimum 96 usec	
Transmission Control	RTS from MAC, CTS to MAC	
Collision Avoidance	Transmission holdoff for 1 usec	
Carrier Sense	After 7 Pulses in 2 Symbols	
Error Recovery	None	

Data Encoding

By definition, each symbol in the data part of the frame encodes four bits of data: The transmitted symbol contains one pulse in one of sixteen possible slot positions. The first slot of a transmitted symbol represents the hexadecimal digit 3. The last slot in a symbol represent hexadecimal 0. Each symbol encodes one hexadecimal character, with a 1 in the slot corresponding to the hexadecimal representation of the character. For example, a symbol encoding a hexadecimal 7 would have a 1 in slot 3. (See Table 1 below.)

hex input value	Symbol pulse position
0	0000000000000001
8	0000000000000010
c	0000000000000100
4	0000000000001000
6	0000000000010000
f	0000000000100000
a	0000000001000000
2	0000000010000000
1	0000000100000000
9	0000001000000000
d	0000010000000000
5	0000100000000000
7	0001000000000000
f	0010000000000000
b	0100000000000000
3	1000000000000000

Table 1.

Transmitted IR

The IR pulse shape generated by the PHY is shown in figure 4. The specified tolerances on pulse width and rise and fall time must be met by the PHY in normal operation. Variations outside these limits will cause unreliable operation at the receiving unit and can result in valid data being rejected. Note that the rise and fall times are measured from the 10% and 90% points and the pulse width is the full width at half maximum value.

The minimum optical output power is 100 mW, the maximum is 200 mW (averaged over one symbol) including all the energy radiated by the PHY device into the full half-plane oriented nominally in the positive vertical direction above the device. The maximum IR peak radiance along the normal to the top surface of the device will be less than 15 Watts per steradian, measured with a spot metering device. The IR wavelength will be between 850 and 900 nm.

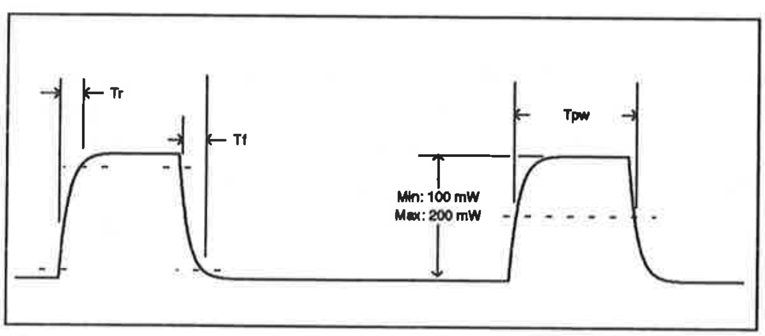

	1 MBPS ns	2 MBPS ns	4 MBPS ns
Tr	75 ± 5	35 ± 5	18 ± 5
Tf	75 ± 5	35 ± 5	18 ± 5
Tpw	230 ± 10	115 ± 7	57 ± 5

Figure 4. Transmitted IR Pulse Shape

In addition to the conditions on the width, rise and fall times of individual pulses, the pulse to pulse position variations of the transmitted signal, or timing jitter, must meet certain conditions. (Refer to figure 5.) Within a single frame, there is an ideal pulse position for every received symbol which is defined by the preamble clock that is transmitted with each frame. Errors in pulse position can be caused by timing errors in the transmitting system. The resulting jitter must meet the maximum jitter specification for proper data recovery.

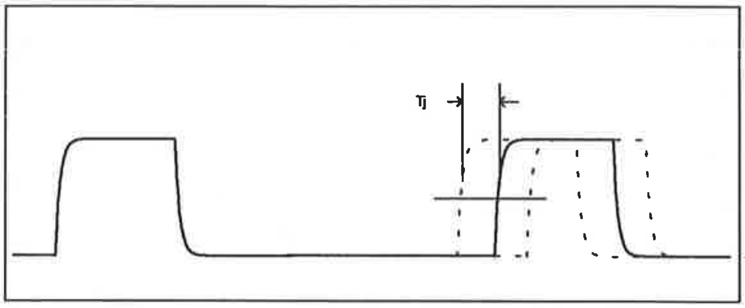

	1 MBPS ns	2 MBPS ns	4 MBPS ns
Tj	75 ± 5	35 ± 5	18 ± 5

Figure 5. Worst Case Transmitter Position Jitter

Receiver Pulse Filtering

Two mechanisms in the PHY operate to remove pulses that have a width that falls outside of expected limits. First, incoming pulses are synchronized and then checked against a minimum width requirement. For guaranteed detection, pulses must be wider than 125, 63, or 36 nsec (FWHH) respectively for 1, 2, or 4 MBPS systems . Pulses that fail this test are rejected as noise. Second, the symbol identification process will reject symbols that are ill formed because of pulses that are too wide. Guaranteed acceptance as a valid symbol requires pulses to be less than 320 (160, 80) nsec wide. Again, failure to meet this requirement causes a symbol to be rejected and the frame is flagged as bad.

Error Detection and Processing

Error detection on the symbol level is accomplished by verifying that symbols are properly formed. After applying the detectable pulse tests described above, each symbol is examined to verify that one and only one slot time is filled. Any symbol failing this test is flagged as bad.

Any bad symbol within the data portion of the frame causes the frame to be aborted. No further symbols are decoded and no further data is sent from the PHY to the MAC. Typically this will cause the MAC (via the operation of the NIC most likely) to detect a faulty packet, probably as a CRC error.

Receiver sensitivity

Minimum conforming receiver sensitivity is -30 dBm/cm^2 for 10^{-9} BER (normal incidence). (-27 dBm/cm^2 at 2 MBPS and -24 dBm/cm^2 at 4 MBPS).

Actual performance will be affected by uncontrollable factors such as ceiling height, wall and ceiling material and texture, and transceiver orientation. Hence for a given installation it will be difficult to predict maximum operating range. With this in mind, the PHY device is specified to be capable of 10^{-9} bit error rate (BER) at a signal irradiance of -30 (-27, -24) dBm/cm^2 incident on the infrared PIN diode filter with an ambient IR light level of -10 dBm/cm^2. The PHY device electronics may incorporate a gain compression circuit to allow the unit to continue to operate in the presence of higher ambient light levels from sunlight and very bright indoor lighting, but with a correspondingly downgraded sensitivity and BER. The receiver's required minimum input signal sensitivity threshold versus ambient IR levels is shown in figure 6.

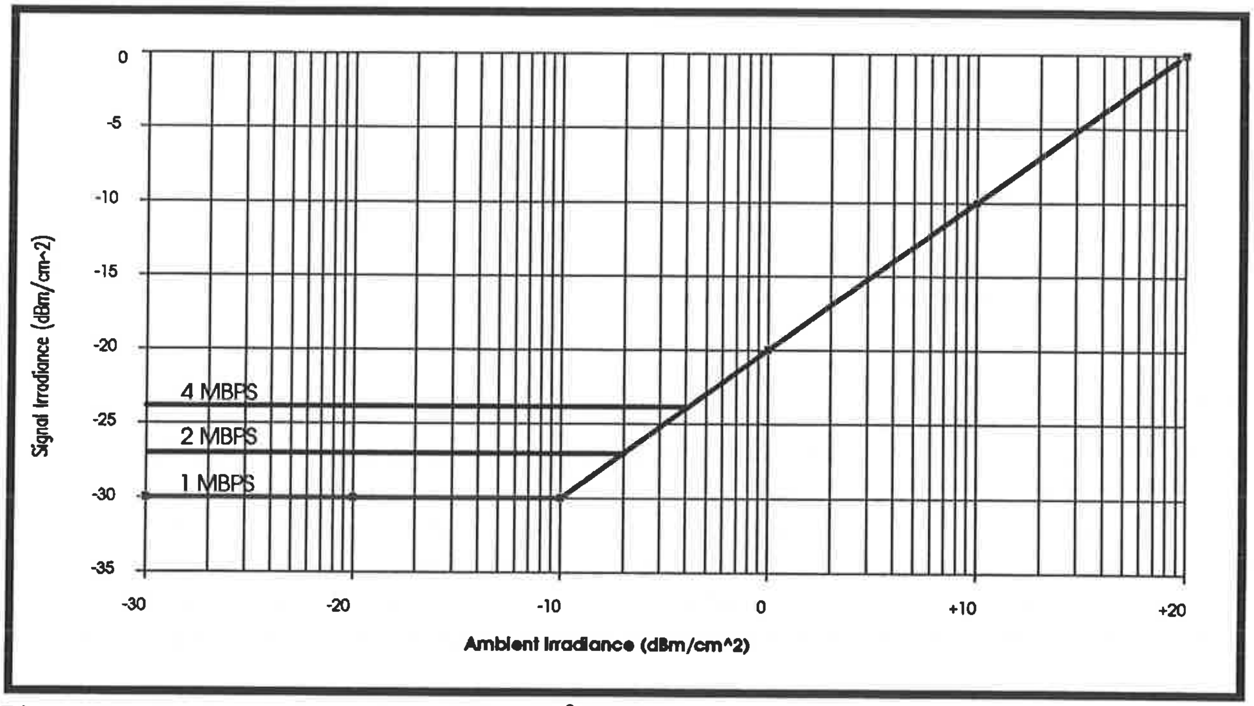

Figure 6. Minimum Signal Level for 10^{-9} BER

Receiver Pulse Recovery Limitations

The received pulse will be a distorted reflection of the transmitted pulse due to noise and pulse spreading effects. Conforming devices must be able to properly recover a pulse string that has been subjected to the following levels of pulse position jitter and pulse spreading.
See Figures 7 and 8.

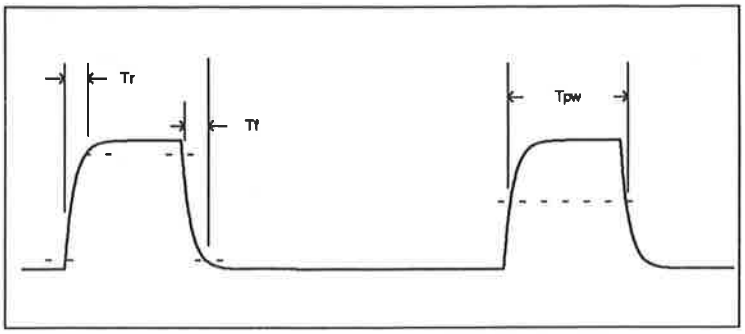

	1 MBPS ns	2 MBPS ns	4 MBPS ns
Tr	100	50	25
Tf	100	50	25
Tpw	320	160	80

Figure 7. Recoverable Received IR Pulse Shape

The clock that is recovered at the receiver establishes both slot and symbol boundaries for the data that follow. Received pulse jitter is measure relative to that idealized pulse position. The conforming PHY receiver must properly decode single pulses that are displaced up to +/- 75 nsec from the normal position.

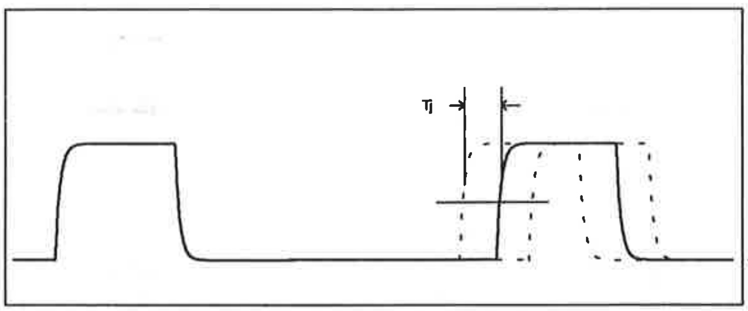

	1 MBPS ns	2 MBPS ns	4 MBPS ns
TJ	100	50	25

Figure 8. Received Pulse Position Jitter

Message Model

Transmitted information will be organized as information frames of indefinite length. Frame error rate considerations will probably limit frame lengths to a few hundred bytes, but this characteristic is not critical to, or part of, the PHY specification.

Transmitted frames will consist of the following:

Preamble:

> The preamble lasts 4 symbols times. Each symbol will be filled with an alternating bit pattern at the basic system clock frequency. This clock frequency will be equal to twice the bit transmission rate, e.g. 2 MHz for a 1 MBPS design. The last 4 slots in the last preamble symbol will be coded as shown below in Table 2. This pattern serves as the delimiter for the end of the sync period provided by the preamble and provides necessary word synchronization. During the preamble, the PHY must perform any AGC acquisition, receiver clock recovery, and carrier sense detection required for proper PPM detection.

During this synchronization interval the receiver gain in the PHY device must compensate for an input signal intensity range of at least 50 dB in order to cover the required dynamic range.

Obviously, the AGC adjustment will have to precede the process of clock recovery. Typically, AGC acquisition needs to be accomplished within the first 8 usec of the preamble for proper operation.

Equalization pattern:

The equalization period lasts two symbol times. The equalization field consists of two symbols containing hex 02 encoded in the standard pulse position code described above. The equalization symbols allow time for the receiver system to come into voltage equilibrium with the 6.25% duty cycle represented by the 16 state PPM. This can simplify the design of AC coupled systems. These two symbols are ignored by the rest of the system.

Data:

The data field consisting of an indeterminate number of symbols. Data characters are formatted as shown above in Table 1.

Preamble	0101010101010101
.	0101010101010101
.	0101010101010101
.	0101010101010010
1st equalization symbol	0000000010000000
2nd equalization symbol	0000000010000000
data	xxxxxxxxxxxxxxxx
.	xxxxxxxxxxxxxxxx
.	xxxxxxxxxxxxxxxx
.	xxxxxxxxxxxxxxxx

Table 2. Frame Structure

Note that there is no end-of-packet delimiter. Also note that the preamble pattern consists of illegal symbols for 16 state PPM, i.e. more than one slot in each symbol is occupied.

Interframe gap

The minimum interframe gap for successive frames will be 96 usec. This should be controlled by the MAC protocol in use, but this spacing determines that maximum allowable receiver recovery time, following either the reception or transmission of a frame from a PHY device. Obviously, the

MAC may cause the actual interframe gaps to be longer, or variable, but the PHY must be able to support operation at the minimum gap time of 96 usec.

Transmission Control

Frame transmissions are initiated by the NIC asserting a Request to Send line (or equivalent). In response, the PHY must begin sending the preamble and equalization symbols. During that time, the PHY must assert the Clear to Send line (or equivalent) and begin clocking in data. The MAC and associated NIC must be able to provide a data stream at the necessary rate to maintain symbol generation and transmission without underruns. If underruns occur, the frame transmission should be immediately terminated and the rest of the frame ignored as it is clocked in from the NIC.

Completion of frames during transmission is signaled to the PHY by the release of the Request to Send line (or equivalent) from the NIC associated with the attached MAC. It is essential that the PHY be able to terminate transmission without generating any kind of error flag after the last symbol has been sent. Typically, the Request to Send line will be released by the NIC just before, during, or just after the final symbol is transferred to the PHY. Buffering and controls within the PHY must be able to terminate the frame after the last symbol without introducing any spurious IR pulses that will cause receiving units to generate error conditions. The PHY must also accept a release of the Request to Send line at other, abnormal times, with a guarantee that no more than one completely formed symbol will be transmitted. This is required in order to support abort operations during which the MAC instructs the PHY to terminate a transmission at some abnormal time. Most commonly, this will happen when the MAC has instructed the PHY to transmit a frame and then determines that a receiving operation has begun with an incoming frame.

Collision Avoidance

Collisions are avoided in the PHY layer through the following two mechanisms. First, the PHY must lock-out any transmissions for a period of 1 usec following the receipt of any incoming IR edge. The received edge may be the beginning of a frame from another unit, a noise spike generated in the PHY itself or the result of some other IR signal generated in the local environment. Regardless, the PHY will block, or hold off the Request to Send for 1 usec before acknowledgment. Second, the PHY must assert Carrier Sense true if it has received a valid preamble signature. The valid preamble signature is defined as a sequence of 7 detectable edges within any 8 usec period. The Carrier Sense output is available to the MAC layer protocol and will cause the MAC to avoid initiating frame transmissions until the IR is clear and any other MAC dependent conditions are satisfied. Normally these MAC conditions with include a random delay before transmission is attempted. Keep in mind that the 7 edges that trigger this condition must be 'detectable', i.e. they must meet the minimum width criterion defined earlier. This has the effect of setting a minimum width threshold for pulses that will be allowed to interfere with normal transactions. This is a compromise, if narrower pulses were allowed to trigger the lock-out, more delays would be generated due to noise effects. Conversely, if a tighter filter was used, more collisions would occur due to marginal reception conditions, were a real signal is incoming, but fails to meet the filter requirements due to weak signal strength.

The PHY hardware lock-out provides for a minimum collision window, regardless of MAC protocols. Typically, a collision window of less than 0.5 usec can be achieved using this mechanism, much shorter than the time required for the MAC to recognize an incoming frame in order to abort a transmission initiation. The downside of such a mechanism is that it can be triggered by isolated noise spikes; for which reason the hold-off interval is limited to only 1 usec.

The Carrier Sense line is deasserted 32 usecs after all IR traffic ceases to be visible to a PHY device. Typically this will be 32 usecs after the completion of a normal frame, but this delay interval will also be enforced after a data error is detected, due either to a damaged packet caused by noise or collision, or simply a unit drifting out of range. The fact that there is no end-of-frame delimiter implies that noise will not be able to trigger the signal detector at normal operating gain levels.

Error recovery

No error recovery capability is incorporated in the PHY.

Compatibility and Interoperability

Fully conforming units operating at the same nominal data rate should be able to interoperate, assuming that a common MAC is in use. Conforming units operating at different data rates must be able to tolerate each other, without being able to exchange information. 'Tolerate' in this case means detection and PHY dependent hold-off of interfering transmissions in the presence of a foreign packet, and adherence to the assertion of Carrier Sense and response by the associated MAC.

References

[1] Moreira, A. J., et al, "Modulation / Encoding Techniques for Wireless Infrared Transmission", Contribution to the IEEE 802.11-93/79

[2] Lomba, C. R., et al, "Update of Propagation Losses and Impulse Response of the Indoor Optical Channel", Contribution to the IEEE 802.11-93/142, September 1993

PPM INFRARED PHY STRAWMAN PROPOSAL

Roger Samdahl
Photonics Corporation

- Recent contributions have established PPM performance at moderate data rates.

- Model studies show that FOV improvements are significant out to 90°.

- Reflectivity and size of environment is critical for low power, low cost systems.

- Model studies indicate that transmitter LED's should be directed away from the vertical to minimize distance dependent dynamic range.

- Care must be taken to accommodate the actual orientation of notebooks, PDA's etc.

- Based on the work described in the noted papers and our own studies, we would like to propose the following PHY layer for PPM at 1,2 and 4 MBPS.

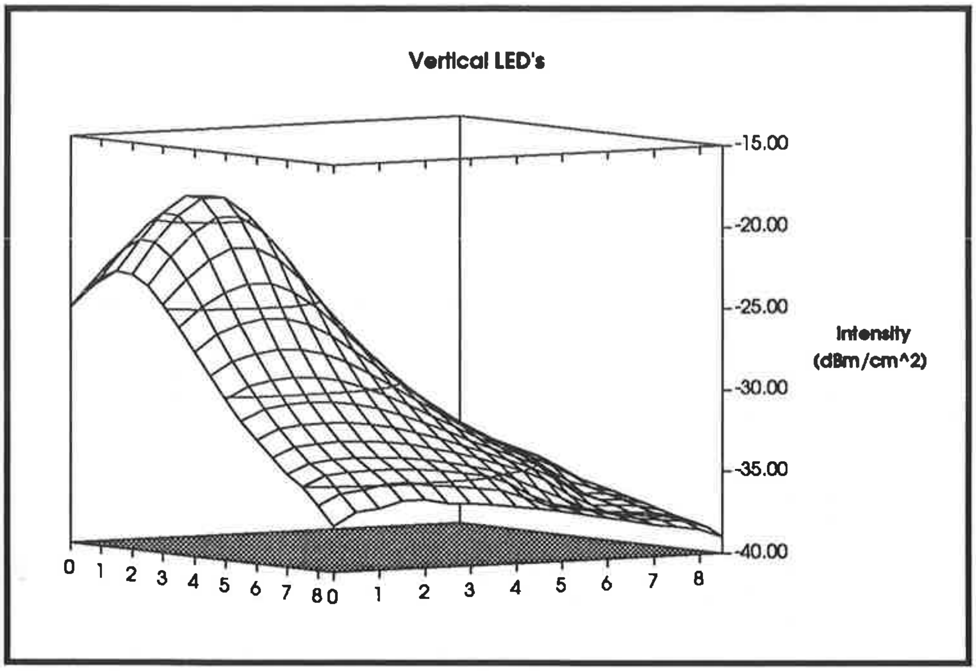

Figure 1.

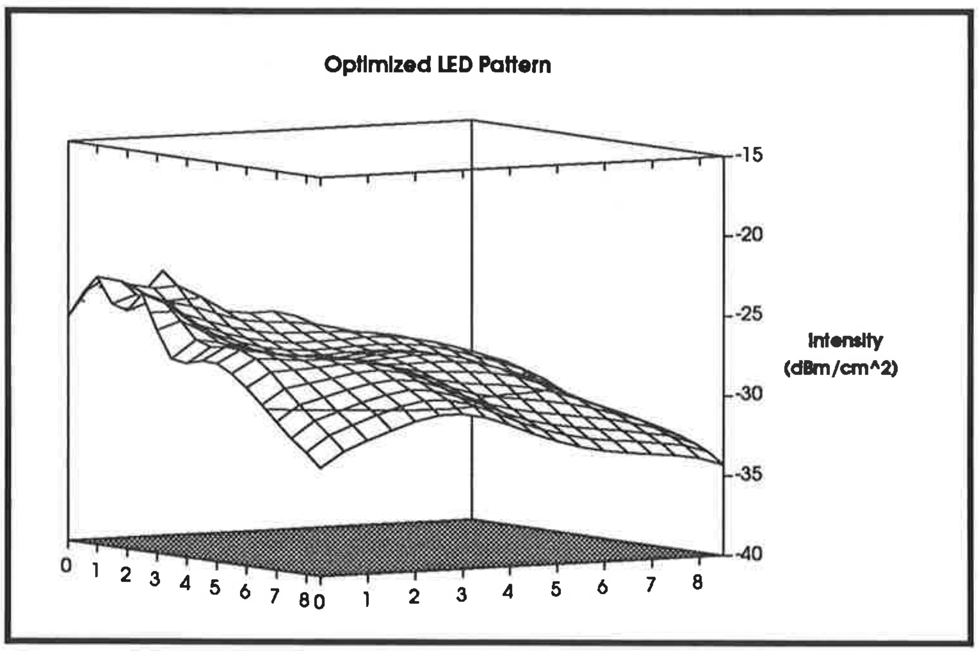

Figure 2.

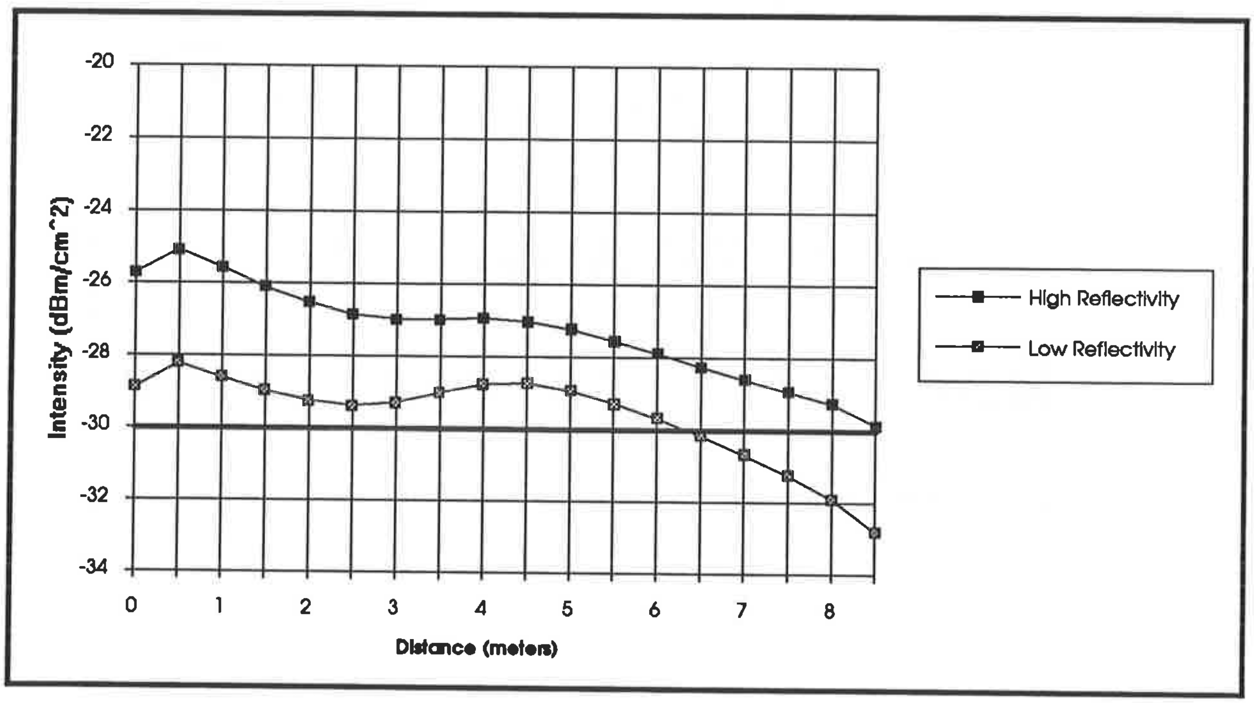
-20
-22
-24
-26
-28
-30
-32
-34
Intensity (dBm/cm^2)
0
1
2
3
4
5
6
7
8
Distance (meters)
High Reflectivity
Low Reflectivity

Figure 3.

PHY Layer Specification Summary

Data Rate	1, 2 or 4 MBPS	
Data Encoding	Modified Gray Code	See Table 1
Transmitter IR Frequency	850 to 900 nm	
Transmitted Pulse Width	230 ± 10 nsec	
Transmitted Pulse Rise Time	75 ± 5 nsec	
Transmitted Pulse Fall Time	75 ± 5 nsec	
Transmitter Position Jitter	75 ± 5 nsec	
Receiver Pulse Filtering	125 ns < Tpw < 320 ns	
Error Detection	Illegal Symbol	
Receiver Sensitivity	-30 dBm/cm^2	With -10 dBm/cm^2 Ambient
Max. Received Pulse Width	320 nsec	
Max. Receive Pulse Rise Time	100 nsec	
Max. Receive Pulse Fall Time	100 nsec	
Max. Receive Pulse Jitter	100 nsec	
Preamble	4 Symbols of Clock w/Mod	See Table 2
Equalization	2 Symbols of 02h	
Data	Any length	
End of Frame Delimiter	None	
Interframe Gap	Minimum 96 usec	
Transmission Control	RTS from MAC, CTS to MAC	
Collision Avoidance	Transmission holdoff for 1 usec	
Carrier Sense	After 7 Pulses in 2 Symbols	
Error Recovery	None	

Data Encoding

hex input value	Symbol pulse position
0	0000000000000001
8	0000000000000010
c	0000000000000100
4	0000000000001000
6	0000000000010000
f	0000000000100000
a	0000000001000000
2	0000000010000000
1	0000000100000000
9	0000001000000000
d	0000010000000000
5	0000100000000000
7	0001000000000000
f	0010000000000000
b	0100000000000000
3	1000000000000000

Table 1.

Transmitted IR Patterns

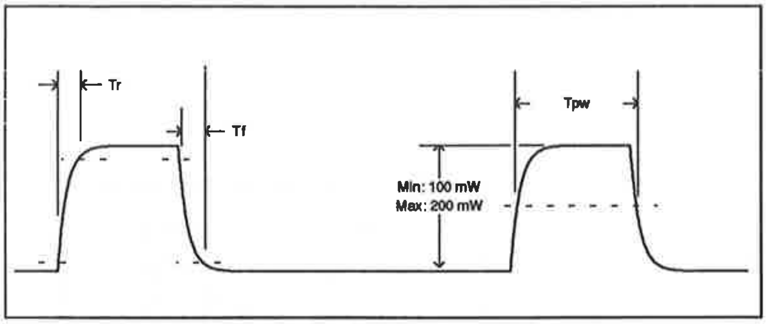

Figure 4. Transmitted IR Pulse Shape

	1 MBPS ns	2 MBPS ns	4 MBPS ns
Tr	75 ± 5	35 ± 5	18 ± 5
Tf	75 ± 5	35 ± 5	18 ± 5
Tpw	230 ± 10	115 ± 7	57 ± 5

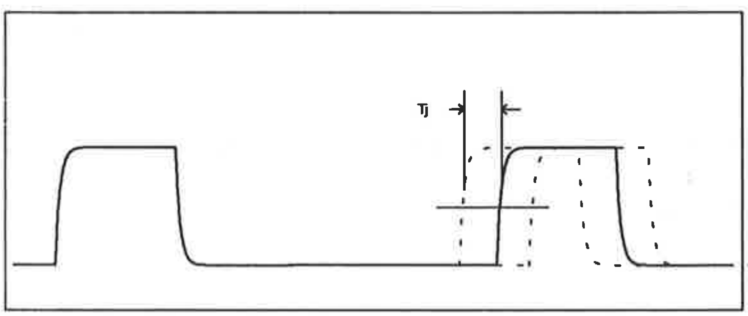

Figure 5. Worst Case Transmitter Position Jitter

	1 MBPS ns	2 MBPS ns	4 MBPS ns
Tj	75 ± 5	35 ± 5	18 ± 5

Receiver Operation

- Receiver sensitivity is -30 dBm/cm^2 for 10^{-9} BER.
- Receiver must operate in the presence of high ambient IR (100 uw/cm^2)

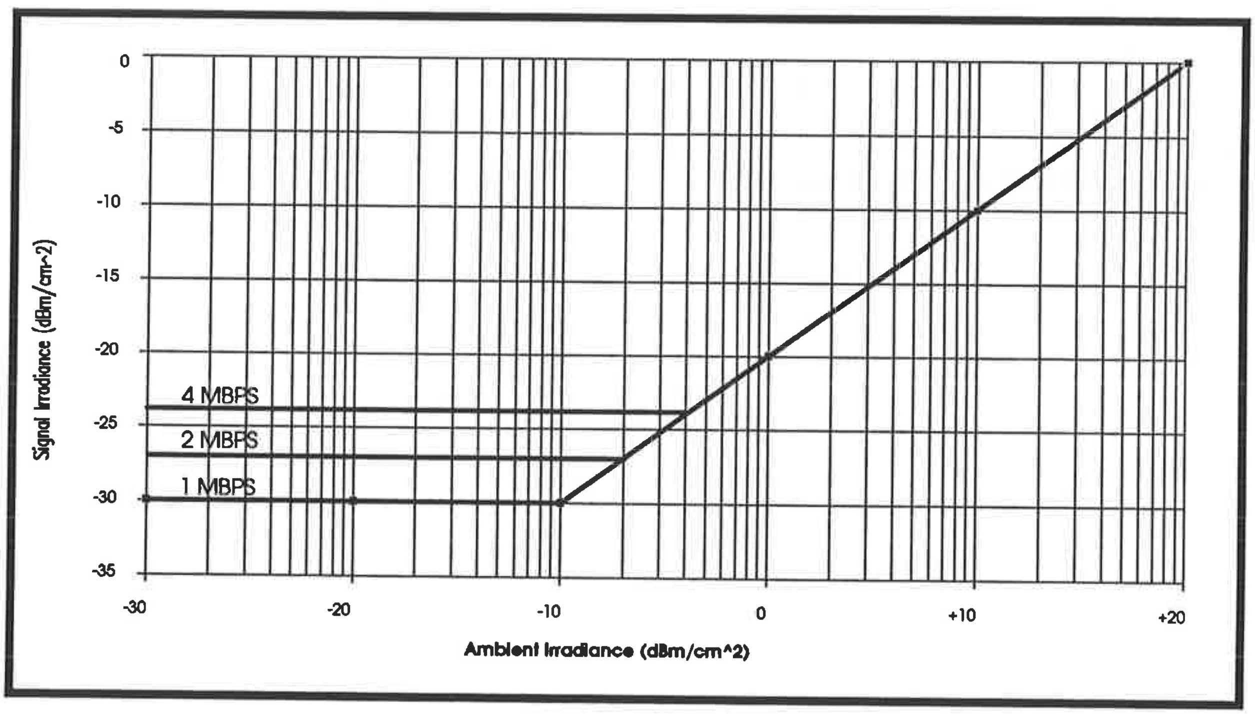

Figure 6. Minimum Signal Level for 10^{-9} BER

Receiver Operations

- The receiver rejects pulses that are too narrow (125, 63, 36 nsec).
- Pulses that are too wide will be rejected as doublets (320, 160, 80 nsec).

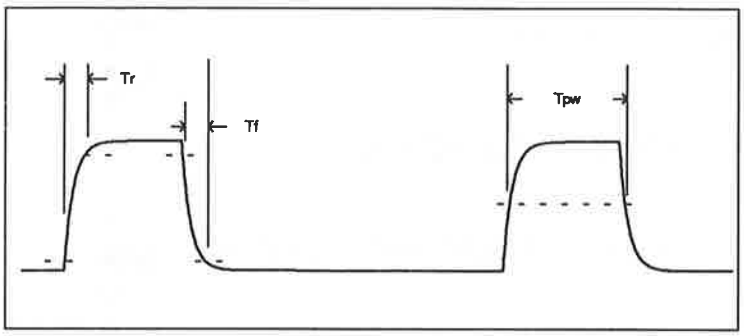

	1 MBPS ns	2 MBPS ns	4 MBPS ns
Tr	100	50	25
Tf	100	50	25
Tpw	320	160	80

Figure 7. Recoverable Received IR Pulse Shape

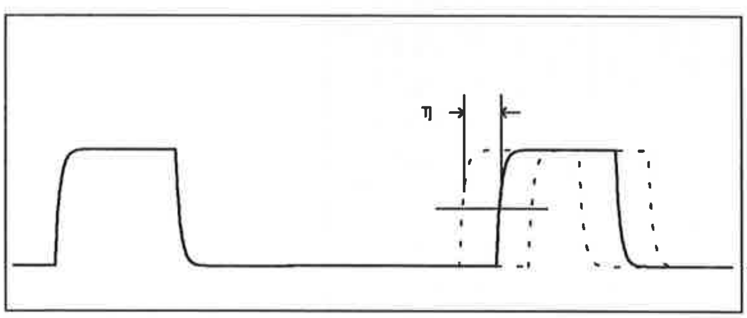

	1 MBPS ns	2 MBPS ns	4 MBPS ns
Tj	100	50	25

Figure 8. Received Pulse Position Jitter

Message Model

- Preamble with end flag.
- Equalization period for receiver recovery.
- Data.
- No end-of-frame delimiter.
- Interframe gap of 96 usec.
- Flow control by both PHY and MAC
- Frame termination by absence of IR for 32 usec.
- Collision avoidance by hardware lock-out and MAC flow control
- Carrier sense requires 7 legitimate pulses in 2 symbol times.
- No error recovery.

Preamble	0101010101010101
.	0101010101010101
.	0101010101010101
.	0101010101010010
1st equalization symbol	0000000010000000
2nd equalization symbol	0000000010000000
data	xxxxxxxxxxxxxxxx
.	xxxxxxxxxxxxxxxx
.	xxxxxxxxxxxxxxxx
.	xxxxxxxxxxxxxxxx

Table 2. Frame Structure

Compatibility and Interoperability

- Conforming units operating at the same data rate will interoperate.
- Conforming units operating at different data rates will recognize each other and properly execute collision avoidance through the carrier sense mechanism.