IEEE 802.11
Wireless Access Methods and Physical Layer Specifications

Baseband IR PHY Proposal

Roger Samdahl
Photonics Corporation
2940 N. First Street
San Jose, CA 95134
Phone: (408) 955-7930
FAX: (408) 955-7950
EMAIL: 70110.360@compuserve.com

Introduction

This proposal is based on an earlier submission titled: "PPM INFRARED PHY STRAWMAN PROPOSAL", R. Samdahl, presented to the IEEE 802.11 Working Group as document 802.11-93/212. The proposed PHY has been simplified in several areas, and various omissions and errors have been corrected.

I propose that Pulse Position Modulation (PPM), and specifically 16 state PPM, be adopted as the standard encoding method for a low bit rate, baseband IR PHY standard. It has been established [1,2] that 16 state PPM can provide very acceptable operation when pulse widths are not unacceptably broadened due to multipath dispersion. Within a target environment that includes conference room and classroom sized areas, pulse spreading will be within acceptable limits for 16 state PPM at data rates in excess of 4 MBPS.

This proposal deals both with the definition of a PHY/IR Interface which will allow interoperability between units developed by different manufactures, and a MAC/PHY interface which will encourage the development of integrated circuits that implement the PHY signaling protocol. Although both are covered here, these two interfaces are separate and can be modified independently.

This proposal incorporates the following features:
- *A common preamble, independent of the data rate of the PHY.*
- *A variable preamble length to support advanced receiver technologies.*
- *A preamble termination character that defines the type of the PHY in use.*
- *A low level hardware lockout to minimize the CSMA/CA collision window.*
- *Fixed and limited transmitter output power and receiver sensitivity.*

A 16 State PPM IR PHY Specification

This specification describes a practical IR PHY device that addresses the low end of the IR performance range. Compromises have been consciously made to allow implementations that are simple and inexpensive. The choice of 16 state PPM is particularly suitable for operation in the 1 to 4 MBPS range and has been shown to be an effective encoding scheme when performance is measured in bits of throughput per watt of input power. The IR pulse patterns required for 16 state PPM are easily generated and detected using commercially available optical components that are designed for the home consumer (remote control) market.

The specified PHY lacks several sophisticated features that might provide superior performance. The most notable compromise is in the ratio of transmitter output power to receiver sensitivity. Today's reality is that the lowest cost per unit area of coverage can be achieved with larger numbers of more powerful LED's rather than with more sensitive IR detector systems. Future designs with more sensitive detectors will improve battery lifetime and will reduce the intensity of IR to which the human eye is exposed. These are important goals, but I believe that cost pressures must push such a design standard off into the future. (See Appendix A. for a note on the need for power and sensitivity limits).

Working transceiver designs meeting most of the specifications included in this proposal have demonstrated satisfactory operation at 1 MBPS in rooms measuring up to 30 x 30 feet, depending on the reflectivity of wall and ceiling surfaces. Analysis of pulse spreading characteristics in rooms of this size suggests that satisfactory performance can also be obtained at 2 and 4 MBPS using similar techniques with modified timing parameters. Beyond 4 MBPS, 16 state PPM becomes progressively less satisfactory due to intersymbol interference created by pulse spreading, and it is thought that other encoding schemes will provide better performance and cost characteristics

This proposal breaks the PHY layer into three parts as shown in the simple block diagram (Figure 1); a body, which contains the actual process, an interface between the PHY and the MAC, and an interface between the PHY and the IR medium. The specification addresses the two interfaces separately, the PHY/IR interface is described in terms of the PHY's emitted and received IR characteristics, and the MAC/PHY interface is described in terms of the handshaking that needs to be provided to support the 802.11 Foundation MAC.

Interoperability depends on compatibility at the PHY/IR interface. This proposal does not address the possibility that a physical PHY device might be an isolated component that could move from one implementation of the MAC/PHY interface to another, and no attempt is made to define an exposed interface point between the MAC and PHY.

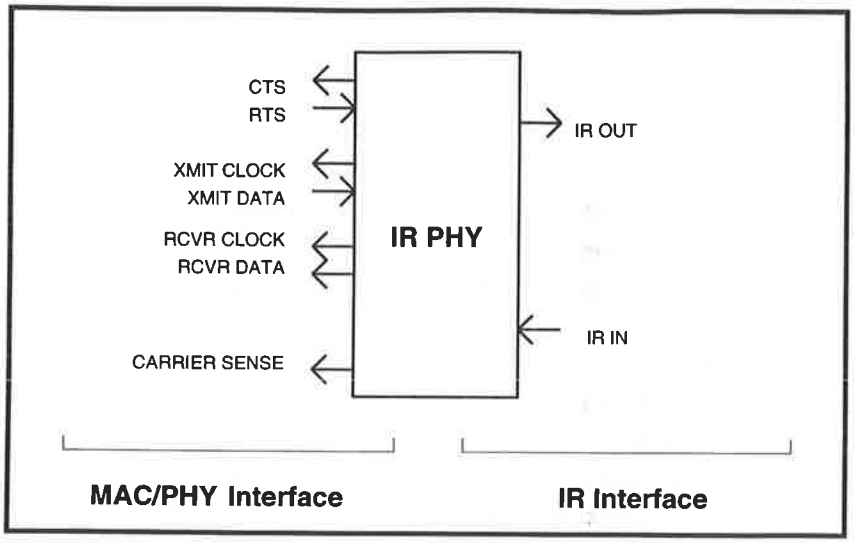

Figure 1. Simplified PHY Block Diagram

General Specification

Data Rate	1 MBPS	
Number of channels	1	

IR Transmitter Specification

Transmitted Pulse Power	2 Watt peak ± 20%	
Transmitter IR Wavelength	850 to 900 nm	
Transmitted Pulse Width	230 nsec ± 5%	FWHM
Transmitted Pulse Rise Time	75 nsec ± 5%	10% - 90%
Transmitted Pulse Fall Time	75 nsec ± 5%	10% - 90%
Transmitter Position Jitter	75 nsec ± 5%	
Min. IR Output Cone - FWHM	90°	Opening vertically
Maximum IR Irradiance	15 Watts/sr	

IR Receiver Specification

Tolerated Pulse Spreading	± 80 nsec	
Sensitivity for 10^{-9} BER	-33 ± 3 dBm/cm^2	with -10 dBm/cm^2 Ambient
Dynamic Range	50 dB	
Max. Recovery Time	30 usec	To 90% of nominal sens.
Error Detection	Illegal Symbol	
End of Frame Delimiter	Illegal Symbol	
Receiver Field of View - FWHS	120°	Full Width Half Sensitivity
Maximum Field of View	170°	At physical limit

IR Interface Protocol Specification

PHY Preamble	2 MHz IR clock pattern	See Table 2
Min. Preamble length	15 usec	See Table 2
Max. Preamble length	30 usec	See Table 2
Preamble Termination	1 usec Flag pattern	See Table 3
Equalization Pattern	2 Symbols of 02h	
Data Field Length	0 to 4000 octets	
Data Encoding	Modified Gray Code	Table 4

MAC/PHY Interface Specification

Transmission Control	Falling edge of RTS from MAC	
PHY Trans/Rec Data	Serial Data Stream at the nominal data rate	Clocked by the PHY
Carrier Sense output to MAC	After 12 usec of valid preamble	Used by MAC for clear channel assessment (CCA)
Error Recovery	None	
PHY Collision Avoidance	2 usec hardware lockout	On any IR pulse
PHY Level Collision Window	1 usec or less	

Table 1.

PHY/IR Interface Specification

Transmitted IR Pulse Shape and Jitter

The IR pulse shape generated by the PHY is shown in Figure 2. 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.

The minimum optical output power is 100 mW, the maximum is 150 mW (averaged over one symbol) including all the energy radiated by the PHY device into the full half-plane oriented 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.

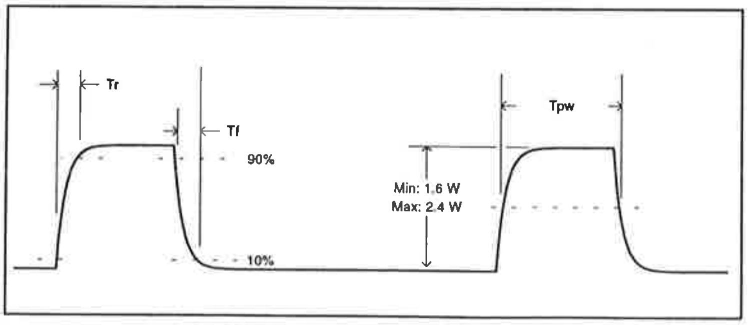

	1 MBPS ns ± 5%
Tr	75
Tf	75
Tpw	230

Figure 2. 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 3.) 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.

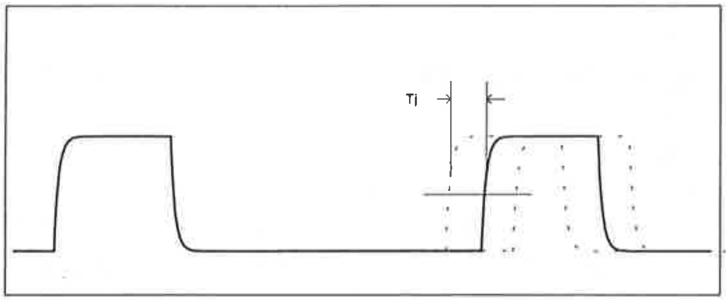

	1 MBPS ns ± 5%
Tj	75

}Figure 3. Worst Case Transmitter Position Jitter

Tolerated Received Pulse Spreading

The PHY receiver must tolerate pulses that have be spread or compressed by reflections in the room. The required spread of pulse widths is 150 to 310 nsec. Conforming PHY's must guarantee that pulses with amplitude above the minimum sensitivity limits are properly detected in the receiver. Refer to Figure 4.

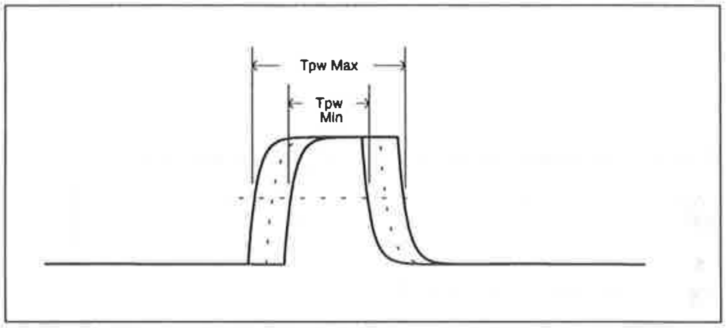

	1 MBPS ns
Tpw-min	150
Tpw-max	310

Figure 4. Tolerated Pulse Spreading

Error Detection and Processing

Any bad symbol within the data portion of a 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) to detect a faulty packet, probably as a CRC error.

Receiver sensitivity

Minimum conforming receiver sensitivity is -33 ± 3 dBm/cm^2 for 10^{-9} BER (normal incidence), Conforming devices will achieve 10^{-9} BER at some point within this range and will fall below 10^{-9} BER before reaching the lower extreme. This effectively limits both the upper and lower values of receiver sensitivity. For an explanation of this requirement, see Appendix A. The minimum receiver sensitivity specification includes the presence of a background IR radiator generating -10 dBm/cm^2 of IR at 850 nm. This interfering source is assumed to be free of modulation.

Dynamic Range and Gain Compression

The dynamic range of the receiver is required to be greater than 50 dB, allowing input signal levels up to a maximum of +20 dBm/cm^2. Conforming PHY's will recover data at a BER better than 10^{-9} at this signal level.

Although not contained within the body of the formal specifications, the PHY device electronics may incorporate a gain compression circuit that will allow the unit to continue to operate with downgraded sensitivity and BER in the presence of higher ambient light levels from sunlight and very bright indoor lighting. The receiver's required minimum input signal sensitivity threshold versus ambient IR levels is shown in Figure 5.

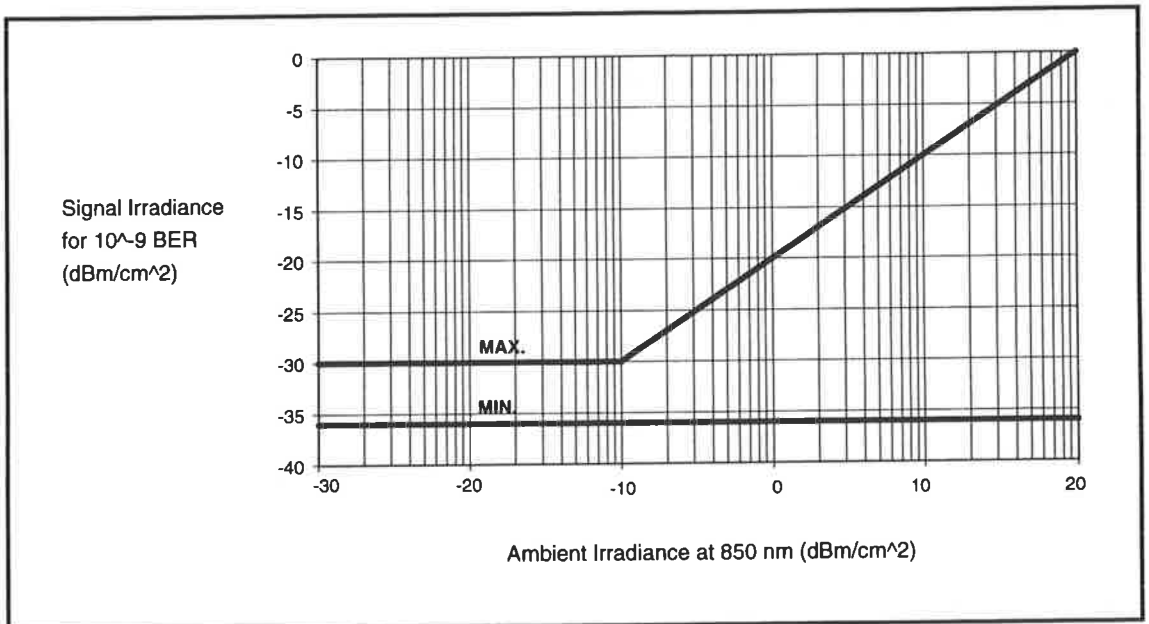

Figure 5. Allowed Signal Level Range for 10^{-9} BER (typical)

Receiver Recovery Time

This specification defines the minimum interval between the completion of a transmitted (or received) frame and the time at which the receiver's gain has recovered to 90% of its peak value. It is assumed that the receiver is operating with AGC; if not this specification is not applicable.

The minimum recovery time for successive frames will be 30 usec. This specification puts a lower bound on the MAC protocol's selection of an interframe gap to something in excess of 30 usec.

Frame 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 provides for recovery and synchronization of the receivers slot clock, symbol clock and byte (word) clock.
>
> It is intended that the preamble defined here be used for all IR PHY's regardless of encoding method or data rate. For an explanation and a description of the preamble's use in other PHY's, see Appendix B.
>
> The preamble lasts a minimum of 15 usec and a maximum of 30 usec. It consists of an alternating bit pattern of filled and empty slot times, essentially a 2 MHz clock pattern with pulses that conform to the basic transmitter pulse shape specified above. In addition to being specified for this PHY, it is proposed that all IR PHY's use this identical pattern. Refer to Appendix B for a discussion of the need for a common preamble.
>
> The 4 slots immediately following the preamble act as a delimiter which defines the frame symbol boundaries. This delimiter is only partially defined in Table 2, allowing for several different patterns to be used. It is proposed that these variations be used as flags to signal the type of PHY that generated the preamble. Again, refer to Appendix B.
>
> During the preamble, the PHY must perform any AGC acquisition and clock recovery and, during the first 12 usec of the pattern, it must assert Carrier Sense. The receiver AGC is required to operate over a 50 dB range of input signal intensity.

Equalization pattern:

> The equalization period lasts two symbol and consists of two symbols containing hex 02 encoded in the standard pulse position code described below. 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
.	010101010101nnnn
1st equalization symbol	0000000010000000
2nd equalization symbol	0000000010000000
data	xxxxxxxxxxxxxxxx
.	xxxxxxxxxxxxxxxx
.	xxxxxxxxxxxxxxxx
.	xxxxxxxxxxxxxxxx

Table 2. Frame Structure

nnnn	PHY Type
0000	TBD
0001	TBD
0010	1 MBPS 16 PPM
0011	TBD

Table 3. Flag Bits

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.

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 4 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 4.

MAC/PHY Interface Specification

This PHY layer is intended for use with a CSMA/CA MAC protocol such as the 802.11 Foundation MAC, but it is also applicable to other MAC architectures.

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 associated with the MAC rather than the PHY, but there are a couple of minimum requirements that it must meet. The NIC must include a four wire handshake consisting of a Request to Send line (RTS), a Clear to Send line (CTS), a Transmit Data line (TXD) and a Receive Data line (RXD). 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 (a 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 bits 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 a single frequency IR systems to provide any sort of collision detection.

- The proposed PHY provides a carrier sense output that can export its state to the attached MAC. Generally, the MAC will do a Clear Channel Assessment (CCA) based on the state of this output and will control the outgoing data stream using this CCA to detect incoming traffic. This capability is augmented by a hardware collision avoidance mechanism described below.

Transmission Control

Frame transmissions are initiated by the NIC asserting a RTS line. In response, the PHY must begin sending the preamble and equalization symbols. During that time, the PHY must assert the CTS line 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. This will result in the reception of a runt frame by the receiving node which will be discarded.

Completion of frames during transmission is signaled to the PHY by the release of the RTS line 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 RTS 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 RTS 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.

PHY Level Collision Avoidance

In addition to the CCA function that is executed in the MAC in response to the state of the Carrier Sense output from the PHY, the PHY also provides a low level collision avoidance mechanism that operates with a much shorter collision window to minimize the collisions inherent in a CSMA/CA system. 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.

A conforming PHY will lock out any transmissions for a period of 2 usec following the receipt of any incoming IR pulse of width greater than 125 nsec. The received pulse 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 any RTS for 1 usec before giving a CTS acknowledgment. As a result, incoming IR traffic will always be given priority over outbound IR traffic under the assumption that every effort should be made to avoid damaging incoming IR frames. This is a compromise; if narrower pulses are allowed to trigger the lock-out, more delays will be generated due to random IR noise. Conversely, if a tighter filter is used, more collisions will occur due to marginal reception conditions, where a real signal is arriving but fails to meet the filter requirements due to weak signal strength.

Carrier Sense for MAC Level Collision Avoidance

The Carrier Sense output to the MAC is initiated by a valid preamble signature defined as a sequence of 7 detectable edges within any 8 usec period. The Carrier Sense output remains true until the PHY detects no discernible IR traffic for a period of 30 usec. Typically this will be 30 usec after the completion of a normal frame, but this delay interval will also be enforced after a data error is detected. Note that while the PHY level avoidance mechanism it triggered by any detected pulse in the IR, the Carrier Sense output may require that detected incoming pulses meet a more rigorous test. Typically, the incoming pulses must be of a quality that allows clock synchronization in the receiver.

Note that because the flag pattern at the end of the preamble is expected to vary in future PHY implementations, it cannot be used to trigger the Carrier Sense mechanism.

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 1994

Appendix A. -- The Requirement for Power and Sensitivity Limitations

It has been suggested that bounds on the maximum radiated IR power and receiver sensitivity should not be included in a PHY specification. The argument is that any restriction on these parameters will stifle advancement in the state of the art. While this may be a valid argument, I believe that the harm caused by not bounding the loop gain and transmitter output level far outweigh the advantages.

The performance of wireless networks, whether IR or radio based, depends on uniform bilateral loop gain between all participating nodes in a BSS. Nodes with dissimilar output power or receiver sensitivity create hidden nodes in the BSS. Graphically the problem can be viewed as shown below: First, assume that nodes A, B, and C all share common receiving sensitivity characteristics, i.e. they can just recover incoming signals at a specific signal level of S uW/cm^2. Second, assume that when transmitting, nodes B and C are just able to generate S uW/cm^2 at the periphery of the heavy circles in the diagram below. Node A, however, is more powerful, and is able to generate S uW/cm2 at the periphery of the dotted circle.

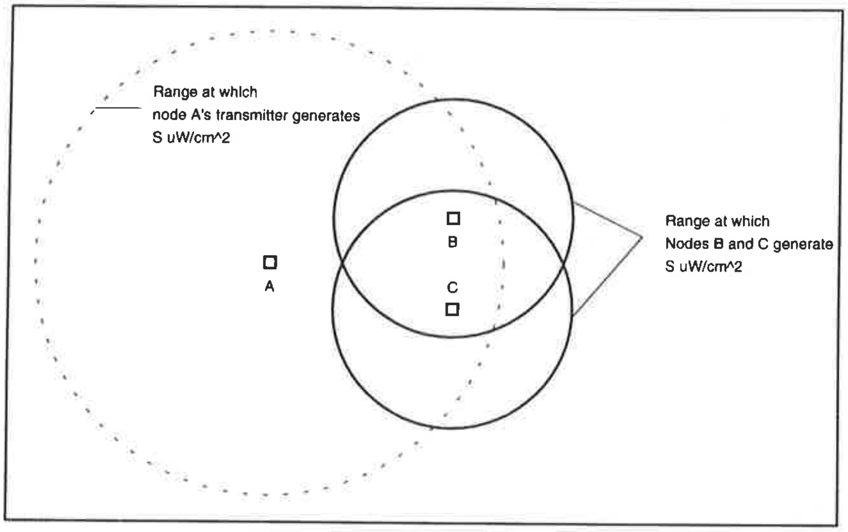

This results in an asymmetry in the operating ranges for the three units. Node A will be unable to detect a conversation going on between nodes B and C, and hence will be unable to properly determine whether the medium is available for transmission. Unaware of traffic between B and C, A's transmissions will frequently collide with any traffic between B and C. Node A becomes a hidden node by virtue of its improved transmitting range.

Consider next the opposite condition in which node A has a less powerful transmitter the B and C. If node A is positioned as shown below, nodes B and C will not be able to detect its transmissions and will collide with them without the benefit of the collision avoidance mechanisms that are otherwise available.

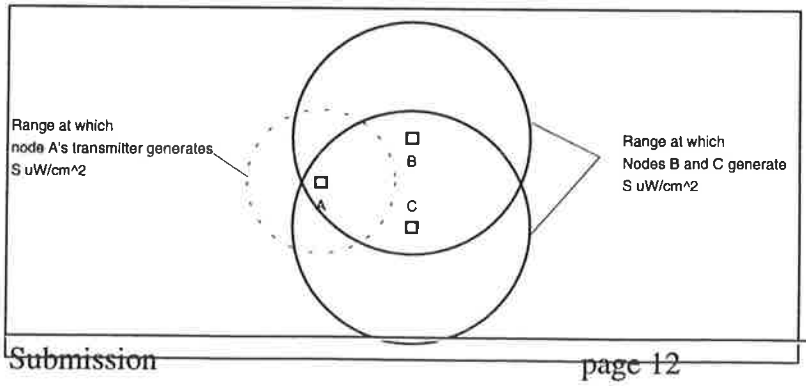

A little inspection will show that when receiver sensitivities are mismatched a similar condition occurs, although in this case the more sensitive receiver is blocked from transmitting through the collision avoidance mechanism under conditions where it would actually be able to transmit without interfering with other, distant, conversations.

Note that this problem is not necessarily related to that actual operating range of any pair of similar devices. Two pairs of units, each with the same maximum range, can exhibit this problem when they try to interoperate. There must be not only similarity in range, there must also be similarity in radiated output power and receiver sensitivity to avoid the generation of unintentional hidden nodes.

To avoid this difficulty, the PHY specification must mandate a minimum and maximum transmitter power output and a minimum and maximum receiver sensitivity. Singular units with higher than normal output power will appear to the network as hidden nodes and will cause collisions; singular units with higher than normal receiver sensitivity will not damage the network, but will defer their transmissions in an unfair fashion.

Appendix B. -- The Need for a Constant Preamble Design

An important feature of this proposal is that a common preamble will be used for all IR PHY's. The common preamble is one that does not change character from one PHY to another. Because the specified mechanism for Carrier Sense discrimination depends only on the detection of a properly formed preamble, the existence of a common preamble will allow a conforming PHY to properly interpret the presence of an alien IR network.

Note that this specification requires that Carrier Sense be asserted after 12 usec of a properly formed preamble is detected. Carrier Sense assertion does not depend on identifying the flag bits that terminate the preamble.

In additions to allowing an accurate generation of the required Carrier Sense output, the use of a standard preamble will allow BSS's of differing types of IR networks to physically overlap. Such overlaps are going to be common with wireless networks, either because of co-located infrastructures or because of informal ad hoc networks that form in common areas, e.g. school libraries. In the instance of overlapped dissimilar PHY BSS's, the quality of the MAC's clear channel assessment, as well as the fast acting PHY level lockout will determine whether either network can continue to operate.

It is also possible that future PHY's will make use of the flag bits that are embedded in the delimiter that marks the end of the preamble. One serious use for this capability is the design of a multi-lingual Access Point that is capable of handling traffic from several generations of IR PHY's.