DRAFT STRAWMAN INFRARED PHY INTERFACE SPECIFICATION

I. Overview

1. This document provides information on building infrared wireless networks having data rates from 0.1 to 10 Mbit/second within a room. Medium access is intended to be mediated by a carrier sense protocol which permits, but does not require, repeaters for range extension. A separate proposal on MAC layer protocols for infrared networks will be submitted in the future. An objective of the MAC layer is to permit cooperative coexistence of multi-speed, multi-higher-level protocol networks.
2. Infrared wireless networks employing diffuse (unaimed) infrared energy have a number of advantages over their RF counterparts. They can provide very robust communication over short ranges (9 - 10 meters) without requiring user aiming. Infrared networks can be packaged in small physical sizes (about that of a credit card) permitting incorporation in portable computers. Power requirements can be quite modest (50 mW listening power), imposing minimum impact on portable operation battery life (1% to 2%). Infrared networks may be constructed from readily-available, low cost consumer optoelectronic components and can be incorporated in portable computers for around $20. Infrared networks are immune to radio-frequency interference. Finally, because infrared networks operate using the same infrared energy employed by TV remote controls, user concerns over safety are absent.
3. Diffuse infrared networks operate by flooding an area with infrared light (Fig. 1). To minimize blockage, the infrared signals are directed over a relatively wide angle and are diffusely (not specularly) reflected from ceilings, walls and other parts of a room (Fig. 2). Thus, no single optical ray path is critical to establishing a link. This permits reliable operation over a wide range of transceiver orientations. Accidental blockage is quite improbable.
4. Because light cannot penetrate walls, infrared networks are inherently very secure. This confinement of optical energy to a single room permits interference-free operation of many such networks in different areas of a building.
5. Infrared signals are rapidly attenuated over relatively short distances, permitting a micro-cellular approach to building larger networks, maximizing "frequency" reuse. This provides a highly parallel approach to providing network bandwidth. This rapid signal attenuation with distance, combined with the incoherent nature of the optical energy employed, also minimizes the intersymbol interference effects of multiple reflection paths, as the strongest signal path tends to dominate received signal waveform and phase cancellation does not occur.

II. Goals

1. The proposed physical layer infrared interface is intended to define the common air interface for all infrared networks having data rates from 0.1 to 10 Mbit/second. Goals include minimization of power consumption, cooperative interference avoidance, interoperability of units employing different speeds and protocols, low cost of baseline unit compliance and enhanced operation compatible with baseline units.
2. This interface should support multiple speeds of operation to permit the greatest flexibility for manufacturers to match performance and cost to the specific application. In some cases a manufacturer may choose data rates compatible with older, wired networks to minimize buffering requirements when wireless networks are linked to the wired infrastructure.
3. A great many wireless applications are, by their nature, portable, dictating a requirement for low power consumption to maximize battery life. This specification defines a means of implementing a "deep sleep" mode, compatible with very low power, and a means to "wake-up" sleeping nodes when required for communications. This same means is employed to implement a robust "channel busy" or carrier sense function.
4. By providing for the ability of transceivers operating at any speed to detect the transmissions of transceivers operating at any speed, destructive interference between network nodes may be avoided.
5. Implementation cost is a critical factor. This specification permits design of compliant units at minimum cost while providing the capability for enhanced performance, higher cost operating modes. Regardless of the capabilities of enhanced units, they will be capable of interoperating with units which have been designed to meet the baseline specification.

6. While repeaters may be used to extend range and support different uplink and downlink speeds for matching transceivers operating at different speeds, no repeaters are required for any reasonable number of transceivers to establish an ad hoc network.

III. Data Rate

1. Data transfer rates anywhere in the range of 0.1 to 10 Mbits/second are supported by this specification.

IV. Optical Wavelength

1. Choice of optical wavelength to be employed in optical networks is dictated by physics and economics. Silicon photodetectors exhibit maximum efficiency at wavelengths around 900 nm. Fortunately, light emitting diodes fabricated from Gallium Aluminum Arsenide emit optical energy at this wavelength. Both transmitting and receiving devices are widely employed in consumer electronics equipment for remote control and are available at low cost.
2. Inexpensive optical filtering can eliminate shorter wavelength optical energy generated by room lighting, minimizing receiver shot noise.
3. Optical wavelengths transmitted shall be between 800 and 900 nm.

V. Optical Power and Spatial Distribution

1. To provide maximum design flexibility, optical power transmitted must be limited only by minimums required for cooperative medium-sharing and maximums dictated by safety. Fortunately, this provides a range of power levels which permits great flexibility in choice of design points.
2. In order for transceivers to cooperatively share a common physical space of maximum dimensions of 10 M, they must be capable of detecting transmissions from each other. Measurements in typical office/conference room environments indicate a need for a minimum peak transmitted optical power level of 1.75 watts with energy distributed spatially over as wide an angle as possible. Because of the use of pulse position modulation, average transmitted optical power will be approximately 100 mw. Such power levels may be achieved using arrays of eight to ten inexpensive light emitting diodes for data rates in the Mbit/second range.
3. Peak emitted optical power shall be a minimum of 1.75 watts measured into a 2 inch aperture centered on and in contact with the transmitter's emitting surface.
4. Average emitted optical irradiance measured at a distance of 1 cm from the transmitter surface shall be less than 10 mw/cm^2.
5. Even allowing for multiple diffuse reflections from room surfaces, to eliminate any need for user aiming, the optical power must be distributed as widely as possible.
6. Optical intensity measured at any angle less than 45 degrees from a line normal to the surface of the transmitter and centered on the transmitter's emitting surface shall be at least 70% of the peak intensity measured anywhere in that conical region (Fig. 4).
7. Optical intensity measured at any angle less than 60 degrees from a line normal to the surface of the transmitter and centered on the transmitter's emitting surface shall be at least 50% of the peak intensity measured anywhere in that conical region (Fig. 4).

VI. Optical Rise Time and Fall Time

1. To permit the use of inexpensive light emitting diodes designed for use in consumer electronics equipment, while permitting operation over a wide range of speeds, optical rise and fall time specifications should be as loose as possible.
2. Use of conventional rise and fall time measurements for these kinds of light emitting diodes at higher data rates is inappropriate. At higher data rates, light emitting diodes designed for use in consumer electronics equipment may produce pulses which resemble triangular shapes more than rectangles. They begin to turn off before they reach what would be 90% of an equilibrium level if the rising waveform were to continue beyond the time of a single pulse slot.
3. Instead of a rise time or fall time specification, the optical data or synchronizing pulses are specified to have a defined pulse width measured at their 50% of peak intensity point.

DRAFT STRAWMAN INFRARED PHY INTERFACE SPECIFICATION

4. Data, synchronizing, busy and not-busy pulses are defined in sections VIII, IX and X.
5. Optical data or synchronizing pulses (but not busy or not-busy pulses) shall have their amplitudes above 50% of their peak amplitude level for a time greater than 90% and less than 110% of one pulse width. A pulse width is defined as the reciprocal of four times the data rate.
6. Optical data or synchronizing pulses (but not busy or not-busy pulses) shall not have amplitudes above 10% of their peak amplitude level for a time greater than two pulse widths (Fig. 5).
7. Busy or not-busy pulses shall reach 90% of their peak intensity within 350 nsec and fall to less than 10% of their peak intensity within 350 nsec (Fig. 5).

VII. Receiver Sensitivity

1. Receivers must be capable of reliably detecting the busy and not-busy pulses to avoid missing transmissions intended for their network node as well as for avoiding initiating transmissions while another transceiver is transmitting (carrier sense).
2. Receivers shall be capable of detecting busy or not busy pulses with a probability of missing or false detection of less than 10^-3 when the busy or not-busy optical pulse peak irradiance measured at the receiving surface is greater than 7 uW/cm^2.

VIII. Modulation and Data Encoding

1. Quantized (one out of sixteen) pulse position modulation is used to maximize signal to noise ratio for direct detection optical systems while maintaining simple encoding and decoding. The characteristics of infrared optical sources and receivers directs this choice. This encoding technique also permits identification of symbols received with errors as these will contain more than or less than one pulse per sixteen slot symbol (Fig. 3).
2. Light emitting diodes are average power limited. Optical output power is sublinearly related to drive current.
3. Photodiodes produce an electrical current proportional to the optical power intercepted. Received electrical signal power is therefore proportional to the square of the optical power. Noise power (from most sources) is proportional to the receiver bandwidth. To maximize received electrical signal to noise ratio, it is advantageous to operate at high peak power.
4. By encoding four bits into one optical pulse located in one of sixteen possible positions the peak optical power received is approximately sixteen times that which could be obtained with non-return to zero (NRZ) transmission and so the electrical signal power is increased by sixty-four times. The bandwidth is increased by four times that of NRZ. Thus the signal to noise ratio is improved relative to NRZ by a factor of sixteen, or 12 dB.
5. The optical signal encodes four bits in a straightforward manner. An optical pulse in the first time slot of a sixteen slot symbol represents a bit pattern of '0000'. A pulse in the second slot represents '0001', and so on through all sixteen four-bit combinations.
6. Interpretation of the meaning of four bit symbols in different sections of data packets is defined at higher protocol levels. MAC layer interpretation will maximize design flexibility, leaving designers free to incorporate a wide variety of error correction and out-of-band signaling capabilities.
7. The standard will permit code violations consisting of two or more optical pulses in a single symbol time followed by an appropriate number of symbols times containing no pulses (to maintain constant average power). Interpretation of the meaning of these code violations is left to higher protocol levels.
8. End of the data field is signified by receipt of the unique code violation caused by the ending not-busy pulse.

IX. Carrier Sense, Busy and Not-Busy

1. Receivers must be capable of detecting transmissions from transmitters operating at any permitted speed. Receivers designed for minimum "listen" power consumption may need to switch to a higher-power, more sensitive mode for accurate reception of data packets.
2. Each packet transmitted at any speed starts with a pulse lasting 5 usec and ends with a pulse lasting 2.5 usec.. The starting busy pulse begins the packet and is followed immediately by synchronizing

pulses. The ending not-busy pulse is sent immediately following the last data symbol. These fixed-width pulses permit receivers designed for any speed to detect the beginning and end of the transmissions of transmitters operating at any data rate as well as the end of a packet sent by a data rate compatible transmitter. Having two different pulse widths for busy and not-busy pulses permits distinguishing busy from not-busy and avoids extended "out-of-phase" operation. The busy pulse is wider than the not-busy pulse since the consequences of failure to recognize busy are more severe.

3. Use of a 5 usec wide busy pulse permits receivers to operate in a low power, narrow bandwidth mode while waiting for transmissions. At the maximum data rate of 10 Mbit/second, this pulse adds the equivalent of a 50 bit latency. The latency introduced is correspondingly less at lower speeds. Choice of a 5 usec pulse is a compromise between minimizing latency and maximizing reception reliability for low power receivers.
4. A receiver detecting a busy pulse at the beginning of a transmission asserts its carrier sense indication to the connected node. A receiver detecting a not-busy pulse at the end of a transmission resets its carrier sense indication. To prevent "lock-ups" caused by a receiver missing an end-of-packet pulse, connected nodes must implement a timeout function causing carrier sense to be deasserted after passage of time sufficient for a maximum length packet plus a small safety margin. This timeout will remove carrier sense after the timeout has elapsed without receipt of a not-busy pulse.
5. Receivers are free to re-enter their low-power, narrow bandwidth mode once higher level protocols determine that the transmission is not intended for the receiver's connected network node.
6. Since transmitters may produce a busy pulse of higher amplitude than their normal data stream (due to light emitting diode rise time limitations), receivers must reset their automatic gain control to a maximum gain level upon detecting the end of the busy pulse. If they were to allow their gain to be set by the wider, stronger pulse, the receive level could be improperly set for the subsequent data stream.

X. Synchronization

1. Receivers require a "training pattern" before transmission of data begins in order to set automatic gain control levels, to permit the receiver to obtain bit synchronization and detect the beginning of data transmission.
2. The required synchronizing pattern is 60 time slots of a 0101... pattern terminated by a 0010 pattern. All pulses following this are data pulses until the end of packet.

XI. Packet Format

Packets consist of the following fields (Fig. 6):

1. Starting busy pulse
2. Synchronizing pattern
3. One symbol for use by the MAC layer
4. Zero to 4096 data symbols (including code violations if employed)
5. Ending not-busy pulse

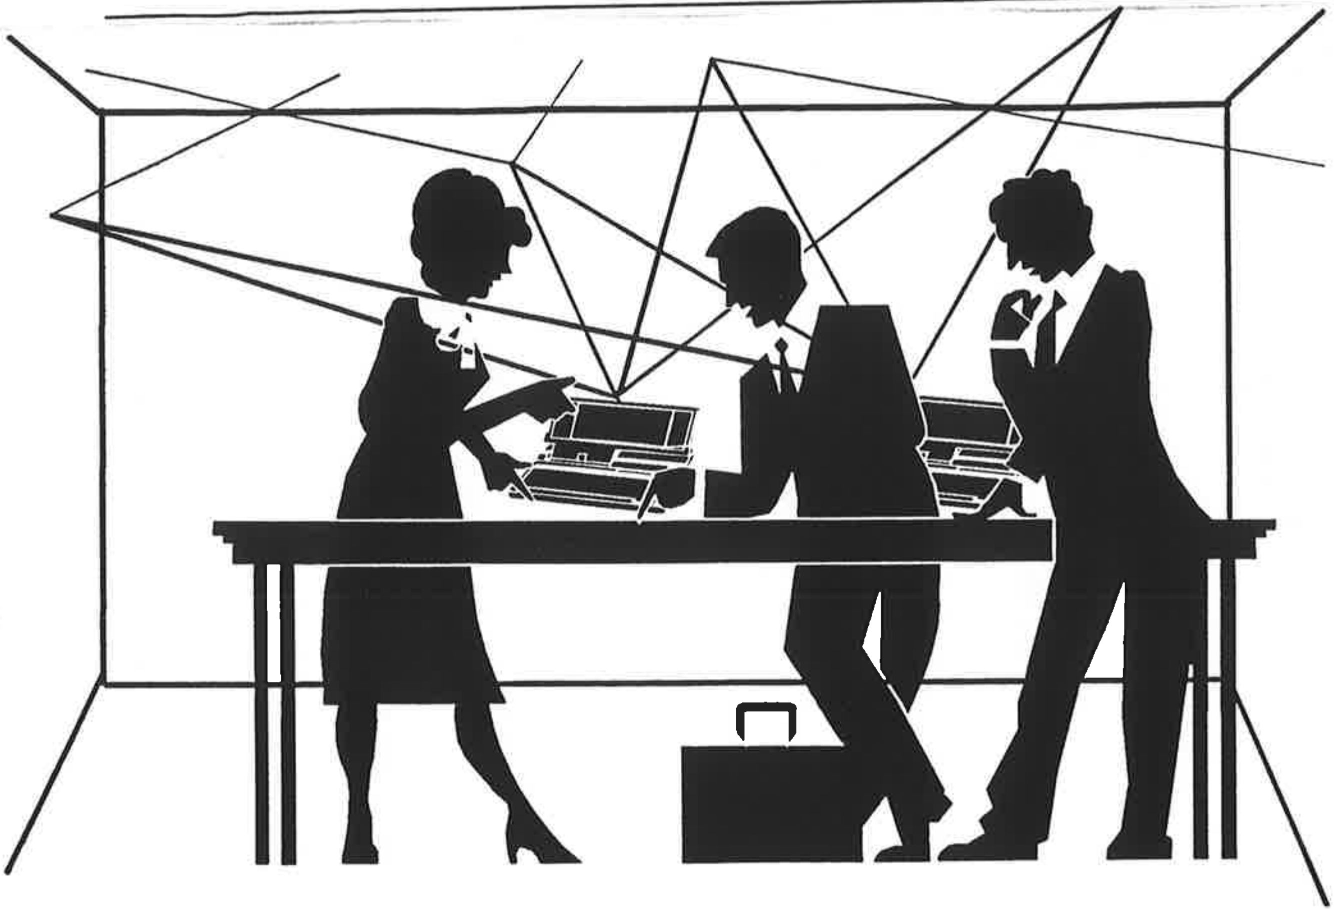

Fig. 1 Diffuse Infrared Network

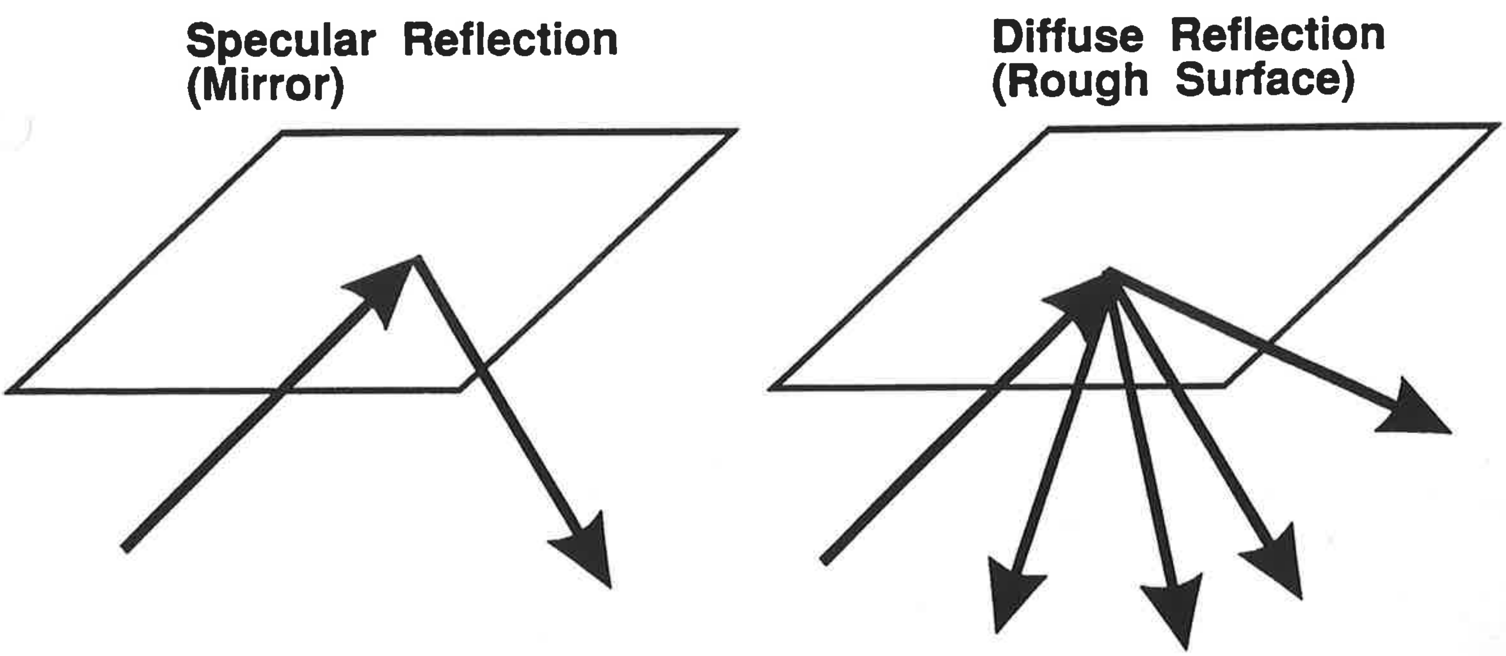

Fig. 2 Specular vs Diffuse Reflection

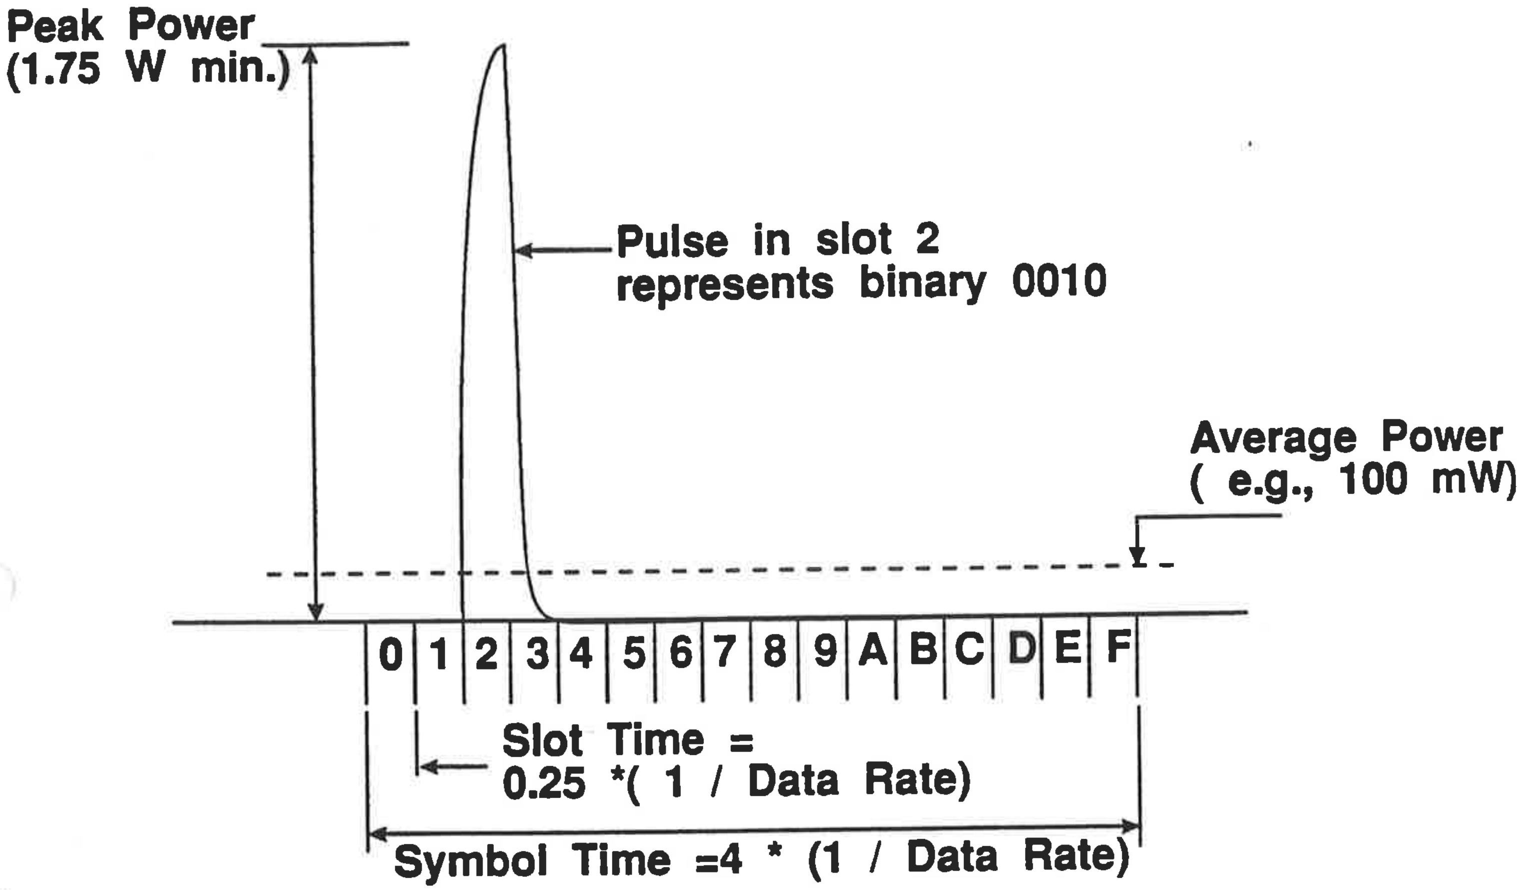

Fig. 3 Transmitted Optical Pulse

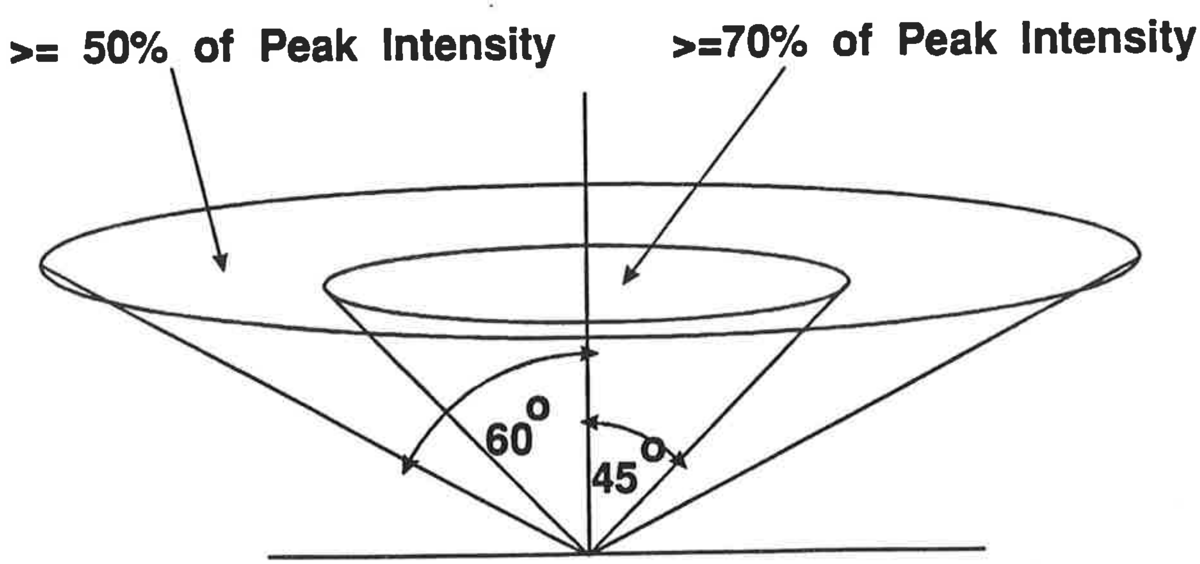

Fig. 4 Spatial Distribution of Optical Power

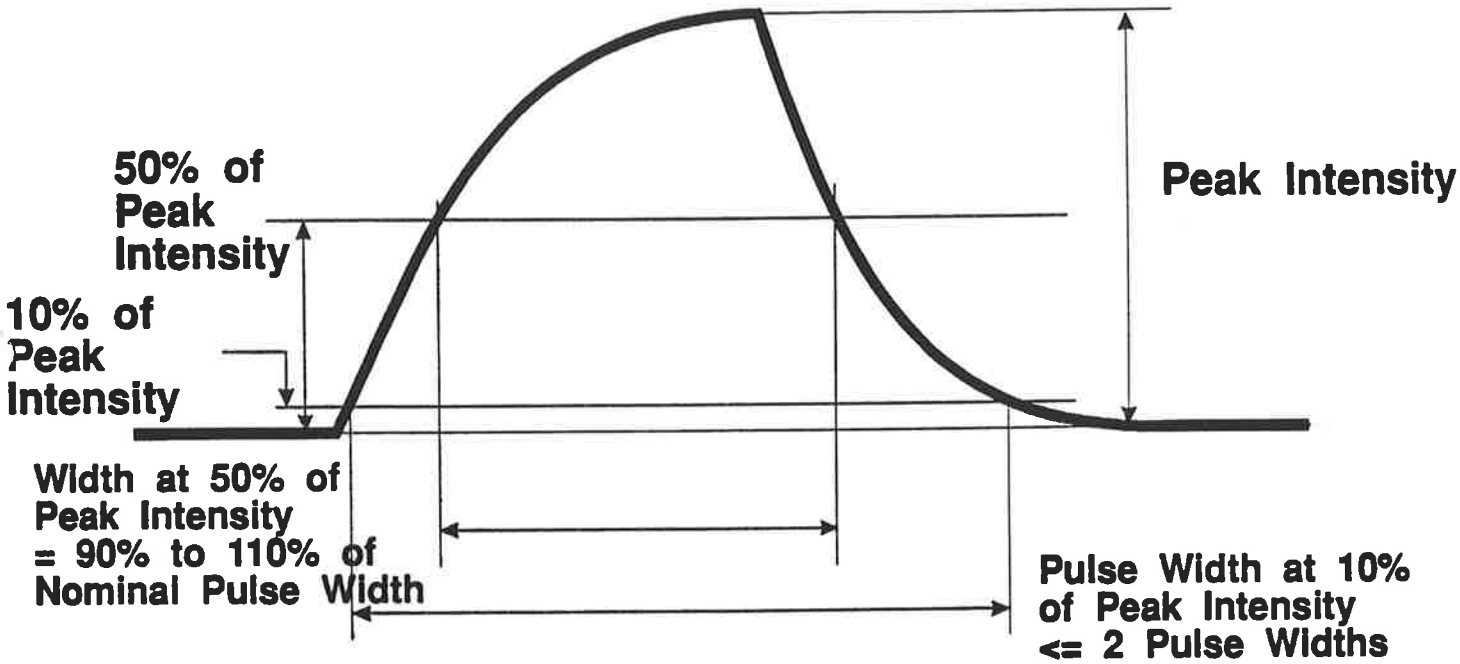

Fig. 5 Optical Pulse Shape

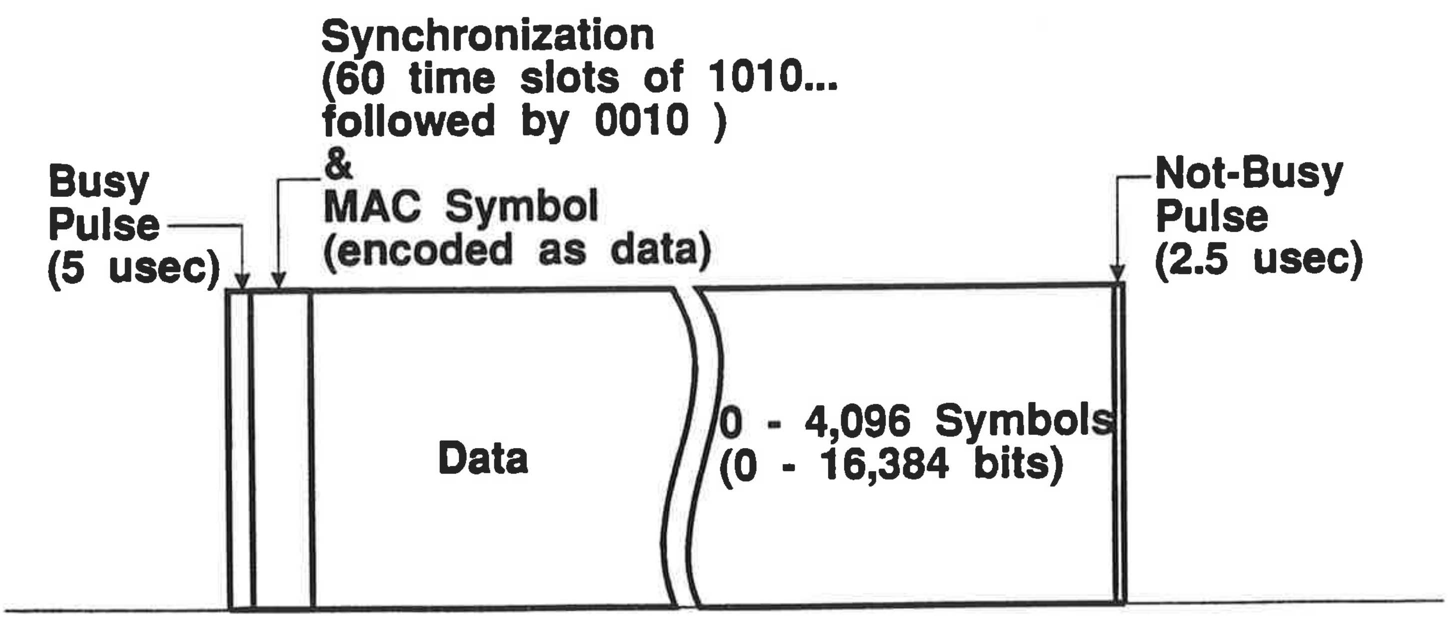

Fig. 6 Packet Structure