CIRCUIT CELLAR ONLINE

Technically Speaking

James Antonakos

Ethernet Technology

Part 2: The 100- and 1000-Mbps Ethernet

In this portion of his series, James walks us through the operation of Fast Ethernet and Gigabit Ethernet. Join him as he discusses each 100- and 1000-Mbps technology, looks at their individual encoding methods, and takes us through the basics of wireless Ethernet.

In the March issue of *Circuit Cellar Online* (www.chipcenter.com), I discussed the basic operation of the 10-Mbps Ethernet, including the format of an Ethernet frame, collision detection, and Ethernet LANs. This month, I'll walk you through the operation of Fast Ethernet (100 Mbps) and Gigabit Ethernet (1000 Mbps).

FAST ETHERNET

Demand for bandwidth quickly exceeded the capacity of the 10-Mbps Ethernet. Even moving from hub-based 10BaseT networks to switch-based networks provided only temporary relief. Fast Ethernet (100-Mbps Ethernet) is implemented in several different ways, all collectively referred to as 100BaseT technology.

One disadvantage of Fast Ethernet is its smaller network diameter, which is typically one-tenth that of 10-Mbps Ethernet, or around 200 m in a non-switched network. This reduction in network diameter is necessary in order to maintain the parameters of CSMA/CD at the faster data rate. The reason being, the signals still move at the same speed in the cable, but the frame times are shorter by a factor of 10.

Figure 1 shows the 100-Mbps Ethernet interface definition. Several new sublayers have been added as a result of the requirements of 100-Mbps transmission. For example, the Manchester encoding used in the 10-Mbps Ethernet is not well suited to high-frequency operation. Other data

100BaseT4	100BaseTX	100BaseFX	100BaseT2
Cable type: three pairs of Cat 3 (or higher) UTP	Cable type: two pairs of Cat 5 UTP	Single-mode fiber limit: 10,000 m	Cable type: two pairs of Cat 3 (or higher) UTP
Data rate: 33.3 Mbps (25 Mbps, 12.5 MHz) on each pair (total of 100 s)	Data rate: 125Mbps on each pair (at 31.25 MHz)	Multi-mode fiber limit: 2000 m (412 m for half duplex)	PAM5x5 encoding
8B6T encoding	4B5B encoding	4B5B encoding	Connector: RJ-45
Fourth pair is used for collision detection	Half- and full-duplex operation	Connector: duplex SC, ST and FDDI MIC allowed	
Half duplex only	Maximum segment length: 100 m		
Maximum segment length: 100 m	Maximum network diameter: 205 m		
Maximum network diameter: 205 m	Maximum number of repeaters: two		
Maximum number of repeaters: two	Connector: RJ-45		
Connector: RJ-45			

Table 1—*Here are the media parameters for each 100-Mbps Ethernet technology.*

encoding and signaling techniques are used instead, using special bit patterns and multilevel signaling to transfer data in 4-bit chunks instead of one bit at a time (as in the 10-Mbps Ethernet). Let's take a look at each 100-Mbps technology (see Table 1).

100BASET4

In the 100BaseT4 technology, 8B6T coding replaces 8-bit data values with six ternary codes, which may have the values –, +, or 0. Table 2 shows a small sample of the 256 code patterns used in 8B6T encoding. The patterns are chosen to provide good DC characteristics, error detection, and reduced high-frequency effects. Special patterns can also be used as markers or control codes.

A multilevel signaling scheme is used, which allows more than one bit of data to be encoded into a signal transition. This is why a 12.5-MHz frequency carries a 33.3-Mbps stream. Think about it this way: each cycle of the 12.5-MHz carrier contains two levels. This gives 25 million level changes per second on a single UTP pair. The signals on each of the three UTP pairs change a total of 75 million times each second. Dividing 75 million levels per second by six levels per 8B6T symbol gives 12.5 million symbols per second. Each symbol is equivalent to a unique 8-bit pattern, so multiplying 12.5 million symbols per second by 8 bits per symbol gives 100-million bps, the required data rate. Figure 2 shows a sample 8B6T encoded waveform. Note that the 12.5-MHz signaling frequency is within the 16-MHz limit for the Cat 3 cable.

Hex value	6T Code group					
00	+	–	0	0	+	–
01	0	+	–	+	–	0
02	+	–	0	+	–	0
03	–	0	+	+	–	0
04	–	0	+	0	+	–
05	0	+	–	–	0	+
06	+	–	0	–	0	+
07	–	0	+	–	0	+
08	–	+	0	0	+	–
10	+	0	+	–	–	0
3F	+	0	–	+	0	–
5E	–	–	+	+	+	0
7E	0	0	+	–	–	+
80	+	–	+	0	0	–
C0	+	–	+	0	+	–
FF	+	0	–	+	0	0

Table 2—*Eight bits of data are encoded as six ternary levels.*

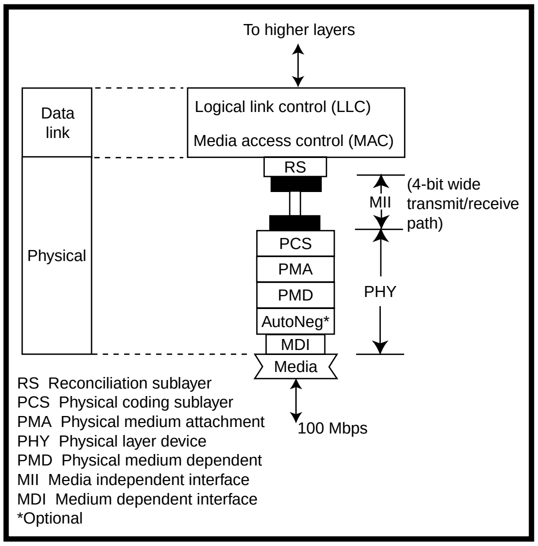

Figure 1—*Here you can see the organization of the 100BaseT architecture. Note that the data is transferred in 4-bit chunks instead of one bit at a time.*

100BASETX

Table 3 shows the 4B5B encoding for all 16 4-bit data patterns. Notice that there is always a mixture of ones and zeros in each 5-bit pattern. This is done to prevent long strings of ones or zeros from being encoded, which contributes to loss of synchronization on the signal.

Figure 3 shows how a three-level signal called multiple level transition (MLT-3) is used to represent the 4B5B bitstream. Each 4-bit data value is replaced by its 5-bit 4B5B counterpart. Thus, the 100-Mbps datastream becomes a 125-Mbps 4B5B encoded datastream. Using MLT-3 allows the 125-Mbps 4B5B datastream to be carried using a signal rate of 31.25 MHz (31.25 MHz × 4 bits per cycle = 125 Mbps). Because the signaling frequency of 31.25 MHz is greater than the 16-MHz limit of the Cat 3 cable, a better cable, Cat 5, is required. Cat 5 cable has a frequency limit of 100 MHz.

100BASEFX

In this technology, the 4B5B encoded data is transmitted using non-return-to-zero, invert-on-one (NRZI) signaling. A 4B5B data rate of 125 Mbps is obtained using a 62.5-MHz carrier. Figure 4 illustrates a sample encoding and waveform. NRZI is well suited for fiber because of its bilevel nature.

100BASET2

Sending 100 Mbps over only two pairs of UTP requires yet another encoding and signaling scheme. In the 100BaseT2 technology, two five-level pulse amplitude modulation (PAM) signals are sent over the UTP pairs, with a signaling rate of 12.5 MHz. Each cycle of the signal provides two PAM5x5 level changes, so there are 25 million level changes per UTP pair. Each pair of PAM signals (called A and B) encode a different 4-bit pattern (along with other, special patterns for Idle mode) using combinations of these levels: +2, +1, 0, –1, –2. So, 25-million PAM5x5 pairs × 4 bits per pair = 100 Mbps.

Figures 5a–c show the symbol constellations found in PAM5x5 encoding, as well as a sample pair of waveforms. Note that when it is not transmitting data, the 100BaseT2 link transmits an idle signal to maintain synchronization. During Idle mode, the signals on A and B alternate between +1 and –1, and +2, 0, and –2.

FAST LINK PULSES

With 100BaseT technology came the ability to perform auto-negotiation between each end of a 100BaseT connection. When the connection is established (plugging both ends of the UTP cable into their respective ports), a series of fast link pulses (FLP) are exchanged between the ports. The 33 pulses contain 17 clock pulses and 16 data pulses. The 16 data pulses form a 16-bit code indicating the capabilities of the port, such as Communication mode (half duplex or full duplex) and

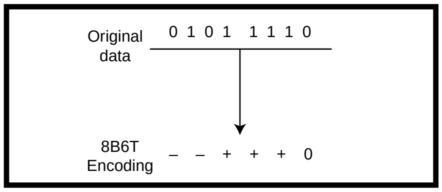

Figure 2—*An example of how 8 bits of data are encoded into a sequence of six ternary codes can be seen here.*

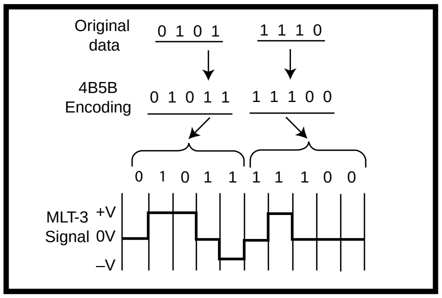

Figure 3—*The 4B5B encoding shows how digital data is represented using a multilevel signal.*

speed (10, 100, 10/100).

Originally, 10BaseT NICs used a single normal link pulse (NLP) to perform a link integrity test. An indicator LED on the NIC showed the status of the link. NLP pulses are typically generated every 16 ms when the transmitter is idle, as indicated in Figure 6a. NICs that support FLPs send a burst containing 2 ms of pulses, as illustrated in Figure 6b.

Note that the even pulses are the data pulses (one when a pulse is present 62.5 µs after a clock pulse, zero if there is no data pulse). The 125-µs spacing between the clock pulses allows the entire burst to complete in 2 ms. After the FLP bursts have been exchanged, the ports decide on the best capabilities for the link according to the priority shown in Table 4. It's interesting to note that all this goes on as soon as you plug a live UTP cable into a port.

TWO REPEATER TYPES

The original 10-Mbps Ethernet specified a single repeater type to propagate frames between segments. Fast Ethernet contains specifications for two types of repeaters, Class I and Class II. Class I repeaters are slower (140 bit times for its round-trip delay) than Class II repeaters (92 bits times or less), but provide functions such as translation between the many different 100BaseT technologies. Class II repeaters, although faster, support only a single technology.

Standard topologies for 100BaseT networks are one Class I repeater, which provides a network diameter of 200 m using copper cable and stations that may be 100 m from the repeater, and two Class II repeaters. The latter are connected via a 5-m cable that provides a diameter of 205 m and stations that may be 100 m from each repeater.

100VG-ANYLAN

Developed by Hewlett-Packard, the IEEE 802.12 Standard 100VG-AnyLAN technology is a 100-Mbps LAN technology capable of handling both Ethernet and token-ring frames. 100VG-AnyLAN uses domain-based priority access, a method whereby stations are polled in a round-robin fashion. Each polled station may make a normal-priority or high-priority request, which is then processed by a higher-level 100VG-AnyLAN hub (not the same as an Ethernet hub). This access method eliminates collisions, allowing 100VGAnyLAN networks to have a larger diameter than a 100BaseT network. The high-priority requests are intended to support real-time multimedia applications such as voice and video.

Figure 7 shows a sample 100VG-AnyLAN network. The IEEE 802.12 specification provides for three levels of 100VG-AnyLAN hubs connected in a hierarchical manner. The entire network must support either Ethernet frames or token-ring frames. For more information on 100VG-AnyLAN, visit www.100vg.com. This site contains detailed documentation and product information.

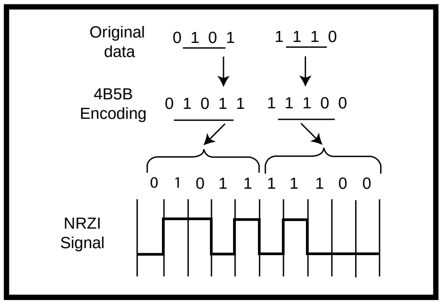

Figure 4—*Eight bits of data are 4B5B encoded and transmitted using NRZI signaling.*

GIGABIT ETHERNET

When the demand for bandwidth exceeded the 100-Mbps Ethernet, it was natural to think about extending Ethernet to 1000 Mbps, rather than use some other, incompatible technology like ATM or FDDI. Extending the data rate, however, led to a decrease in the network diameter (so that CSMA/CD is maintained). You might agree that a network diameter of 25 m or less is not practical. Two techniques used to increase the data rate from 100 Mbps to 1000 Mbps, but still maintain a reasonable network diameter, are carrier extension and frame bursting.

Figure 8 shows an Ethernet frame with a 0- to 448-byte carrier extension field. The carrier extension is used to maintain a minimum 512-byte Ethernet frame (not including the preamble and SFD). So, a 10/100-Mbps Ethernet frame containing only 100 bytes would require a carrier extension field of 412 bytes to be used over Gigabit Ethernet.

Frame bursting involves sending

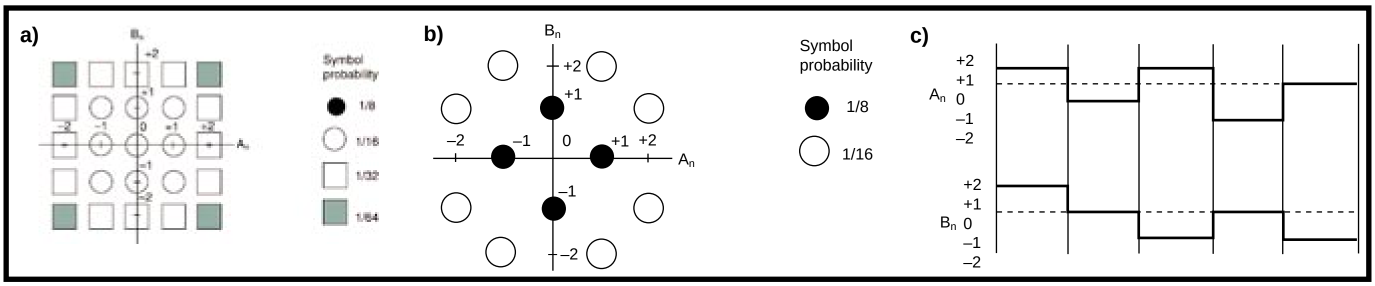

Figure 5a—*The data mode constellation and probabilities for PAM5x5 encoding can be seen here.* **b—***Here you can see the Idle mode constellation and probabilities for PAM5x5 enoding.* **c—***In two PAM5x5 encoded waveforms, each pair of levels represents four bits of data.*

4-bit	Data 5-bit Data				
0000	1	1	1	1	0
0001	0	1	0	0	1
0010	1	0	1	0	0
0011	1	0	1	0	1
0100	0	1	0	1	0
0101	0	1	0	1	1
0110	0	1	1	1	0
0111	0	1	1	1	1
1000	1	0	0	1	0
1001	1	0	0	1	1
1010	1	0	1	1	0
1011	1	0	1	1	1
1100	1	1	0	1	0
1101	1	1	0	1	1
1110	1	1	1	0	0
1111	1	1	1	0	1

Table 3—*Each 4-bit data group is replaced with a unique 5-bit group.*

multiple frames in a burst of transmission. The first frame in the burst must be carrier extended if its length is less than 512 bytes. Additional frames in the burst do not require carrier extensions, but an interframe gap of 0.096 µs (still 96 bit times) is needed between frames. The transmitter continues transmitting during the interframe gap to maintain its hold on the network media. A burst timer starts when the first frame is transmitted limiting the length of the burst to a maximum of 65,536 bits.

Unlike Fast Ethernet, Gigabit Ethernet goes back to a single repeater type. A Gigabit Ethernet repeater must support all 1000-Mbps technologies and operate at a fixed speed of 1000 Mbps (no 10/100/1000 or 100/1000 capabilities are defined, although there are multi-speed Gigabit switches available on the market).

GIGABIT ETHERNET ARCHITECTURE

Figure 9 shows the multilayer architectural model for Gigabit Ethernet. You may notice the similarities to the 10- and 100-Mbps Ethernet technologies. One notable difference is the new 8-bit-wide transmit and receive paths. In addition, notice that full-duplex operation is available in every Gigabit technology, which was not the case for the 10-Mbps Ethernet and Fast Ethernet technologies. The MAC layer, described in IEE 802.3z, deals with issues such as half- and full-duplex operation, carrier extension, and frame bursting. Let's examine the different Gigabit Ethernet technologies
(see Table 5).

1000BASET/CX/LX

The two-dimensional PAM5x5 encoding you were introduced to in 100BaseT2 is extended here to a four-dimensional PAM5x5 system.

Transferring 8 bits of data reliably at Gigabit speeds requires another change in the method used to encode data. The 8B10B coding method (originally developed for FibreChannel) replaces 8-bit data values with 10-bit code words. The code words are chosen from groups in such a way that the number of ones and zeros transmitted is kept in balance. A signaling rate of 1.25 Gbps is required to encode the 1-Gbps datastream.

Table 6 summarizes the 10-, 100-, and 1000-Mbps Ethernet copper and fiber technologies. Even with this large number of technologies, there is still one more to examine—Wireless Ethernet.

WIRELESS ETHERNET

Wireless Ethernet, the use of Ethernet over radio frequency (RF) or infrared (IR), is covered by the IEEE 802.11 Wireless LAN Standard. A

Priority	Link Choice
1 (Highest)	100BaseT2 (Full duplex)
2	100BaseTX (Full duplex)
3	100BaseT2(Half duplex)
4	100BaseT4
5	100BaseTX (Half duplex)
6	10BaseT (Full duplex)
7	(Lowest) 10BaseT (Half duplex)

Table 4—*Here you can see the auto negotiation priorities for the 100BaseT Ethernet.*

wireless Ethernet network consists of one or more fixed stations (base stations), which service multiple mobile stations. Some implementation details are that the frame formats for IEEE 802.3 (Ethernet) and IEEE 802.5 (token-ring) remain the same, CSMA/CA is used, and faster speeds (1- and 2-Mbps operation) are becoming available.

Carrier sense multiple access with collision avoidance (CSMA/CA) differs from CSMA/CD in that the wireless transceiver cannot listen to the network for other transmissions while it's transmitting. Its transmitter simply drowns out any other signal that may be present. Instead, stations attempt to avoid collisions by using random back-off delays to delay transmission when the network is busy (some other station is transmitting). A handshaking sequence is used between communicating stations (ready and acknowledge packets) to help maintain reliable delivery of messages over the air.

TYPES OF WIRELESS LANS

There are two primary types of Wireless Ethernet LANs: RF-based and IR-based.

In RF-based wireless Ethernet, signals can propagate through objects, such as walls, reducing security. The industrial, scientific, medical (ISM) band is used for transmission at the following frequencies:

- 902 to 928 MHz
- 2.40 to 2.4835 GHz
- 5.725 to 5.850 GHz

Data is transmitted using the spread spectrum technologies called frequency

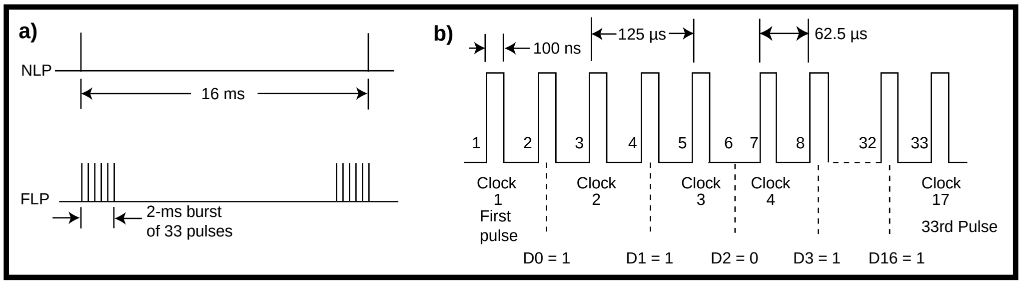

Figure 6a—*General timing of the fast link pulses used in the 100BaseT Ethernet can be seen here.* **b—***The FLP burst organization contains 33 pulses, 16 of which are used to specify capabilities.*

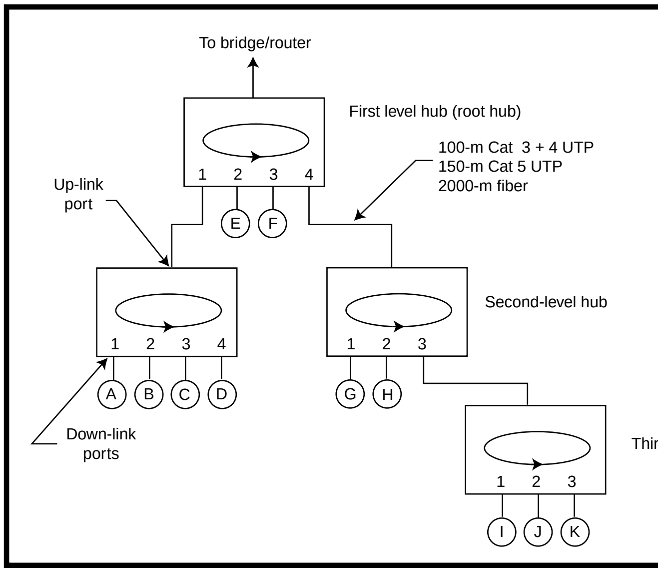

Figure 7—*In the 100VG-AnyLAN architecture, the root hub is at the top of the network tree.*

hopping and direct sequence. In frequency hopping, the transmitter hops from frequency to frequency, seemingly at random, transmitting a portion of each frame at each frequency. The receiver hops to the same frequencies, using the same pseudo-random sequence as the transmitter. A measure of security is added to the data because it's difficult to eavesdrop on all of the associated frequencies and reassemble the frame fragments.

The direct sequence method involves exclusive-ORing a pseudo-random bitstream with the data before transmission. The same pseudo-random bit sequence is used in the receiver to get the original data back.

Two types of IR-based wireless Ethernet are used: diffused IR and point-to-point IR. Diffused IR bounces signals off walls, ceiling, and floors. The data rate is limited by the multipath effect, whereby multiple signals radiate from a single transmission, each taking a different path to the receiving stations (see Figure 10). Point-to-point IR uses line-of-sight IR lasers and provides a faster data rate than diffused IR. It also works over longer distances (up to 1 mile).

1000BaseT	1000BaseCX	1000BaseSX	1000BaseLX
Cable type: four pairs of Cat 5 UTP	Cable type: shielded balanced copper (twinax)	Short wavelength laser	Long wavelength laser
Data rate: 1000 Mbps (2000 Mbps for for full duplex)	Data rate: 1000 Mbps (200 Mbps for full duplex)	Wavelength: 770 to 860 nm	Wavelength: 1270 to 1355 nm
PAM5x5 encoding	8B10B encoding	Data rate: 1000 Mbps 2000 Mbps for full duplex)	Data rate: 1000 Mbps (2000 Mbps for full duplex)
Connector: RJ-45	Connector: 9-pin D or 8-pin FibreChannel Type 2 (HSSC)	8B10B encoding	8B10B encoding
Maximum cable length: 100 m	Maximum cable length: 25 m	62.5/125 Multi-mode fiber length limit: 275 m (half and full duplex)	Single-mode fiber length limit: 5000 m (full duplex), 316 m (half duplex)
IEEE Standard 802.3ab	Short haul copper	50/125 Multi-mode fiber length limit: 550 m (full duplex),	Multi-mode fiber length limit: 550 m (full duplex) duplex), 316 m (half duplex)
		Connector: SC	Connector: SC

Table 5—*Here are the media parameters for each 1000BaseT Ethernet technology.*

NEXT MONTH

Next month, I conclude this series by examining the operation and differences of the components required to construct an Ethernet LAN. These components include network interface cards, transceivers, hubs, switches, and routers. I'll also look inside an Internet service provider to see how its network provides Internet access to multiple users. ◼

James Antonakos is a professor in the Department of Electrical Engineering Technology at Broome Community College, with over 25 years of experience designing digital and analog circuitry and developing software. He is also the author of numerous textbooks on microprocessors, programming, and microcomputer systems. You may reach him at antonakos_j@sunybroome.edu or visit his web site at www.sunybroome.edu/~antonakos_j.

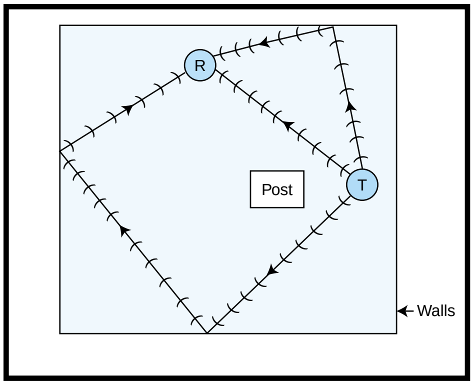

Figure 10—*Here you can see the multipath effect in diffused IR wireless Ethernet.*

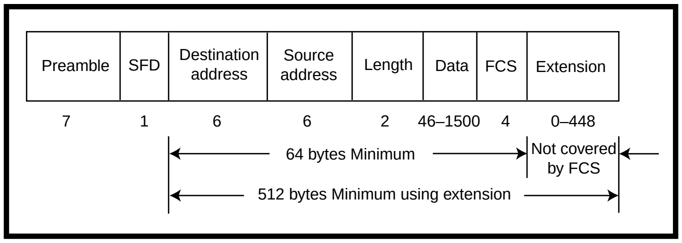

Figure 8—*Carrier extension modification to the Ethernet frame can be seen here.*

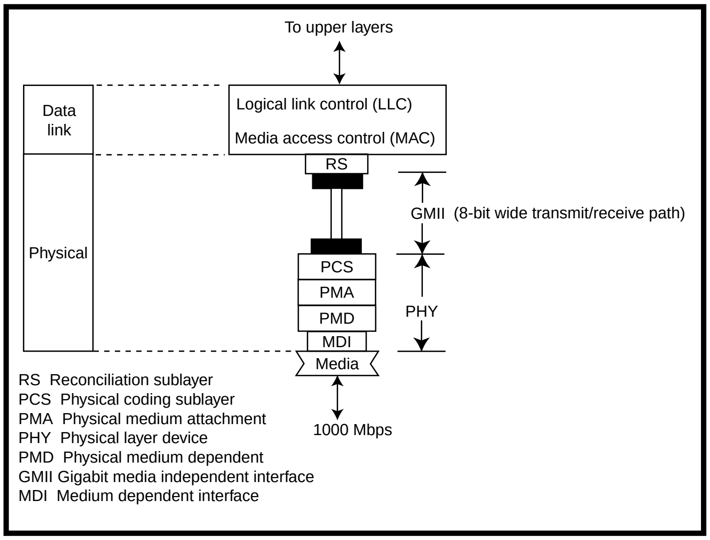

Figure 9—*In the 1000BaseT architecture shown here, the data is transferred in 8-bit chunks.*

Technology	Maximum segment length	Encoding Method	Topology media		Bit rate (bps)
10Base5	500 m	Manchester	bus	50-ohm coax	10 M
10Base2	185 m	Manchester	bus	50-ohm coax	10 M
10BaseT	100 m	Manchester	star	2 pair UTP cat. 3,4,5	10 M
100BaseFL	2000 m	Manchester	star	Multi-mode fiber*	10 M
100BaseT2	100 m	PAM 5x5	star	2 pairs UTP cat. 3,4,5	100 M
100BaseT4	100 m	8B/6T	star	4 pairs UTP cat. 3,4,5	100 M
100BaseTX	100 m	4B/5B with MLT-3	star	2 pairs UTP cat. 5	100 M
100BaseFX	412/2000 m	4B/5B with NRZI	star	Multi-mode fiber*	100 M
1000BaseT	100 m	PAM 5x5	star	4 pairs UTP Cat 5	1000 M
1000BaseSX	275 m	8B/10B	star	Multi-mode fiber†	1000 M
1000BaseLX	316/550 m	8B/10B	star	Multi-modeFiber‡	1000 M
1000BaseCX	25 m	8B/10B	star	Twinax	1000 M

* Fiber is duplex 62.5/125 µm multi-mode fiber
† Max segment length is 316/550 m with 50/125 µm multi-mode fiber
‡ Max segment length is 316/550 m with 50/125 µm multi-mode fiber or 316/5000 m with 10/125-µm single-mode fiber

Table 6—*A summary of the many different Ethernet copper and fiber technologies is shown here.*