TABLE 9.12 Step-by-Step CRC-32 Calculation

Operation	Data
Initialize CRC-32 state bits	0xFFFFFFFF
Word to be calculated	0x4D41524B
Reorder bytes to end high-byte first after bit-flipping	0x4B52414D
Flip bits for LSB-first transmission of high-byte then low-byte	0xB2824AD2
Clock word through XOR logic	0x5C0778F5
Flip bits of CRC	0xAF1EE03A
Optionally swap bytes of CRC for final result	0x3AE01EAF
Invert CRC state bits when input stream is completed	0xC51FE150

3 Mbps), including 10-, 100-, and 1,000-Mbps varieties. Ten-gigabit Ethernet is just now beginning to emerge. Ethernet originally ran over single shared segments of coaxial cabling, but most modern installations use twisted pair wiring in a physical star configuration. The familiar standards for Ethernet over twisted pair are 10BASE-T, 100BASE-T, and 1000BASE-T.

There is a whole family of Ethernet and related standards defined by the IEEE under the 802 LAN/MAN (local area network/metropolitan area network) Standards Committee. More specifically, the 802.3 CSMA/CD (carrier sense, multiple access, collision detect) Working Group defines Ethernet in its many forms. The 802.3 Ethernet frame format is shown in Table 9.13. A seven-byte preamble and a start of frame delimiter (collectively, a preamble) precede the main portion of the frame, which includes the header, payload, and trailer. The purpose of the preamble is to assist receivers in recognizing that a new frame is being sent so that it is ready to capture the main portion of the frame when it propagates through the wire. Not including the preamble, a traditional Ethernet frame ranges from 64 to 1,518 bytes. Two 48-bit Ethernet, or MAC, addresses are located at the start of the header: a destination address followed by a source address. The MSB of the address, bit 47, defines whether the address is unicast (0) or multicast (1). A unicast address defines a single source or destination node. A multicast address defines a group of destination nodes. The remaining address bits are broken into a 23-bit vendor block code (bits 46 through 24) and a 24-bit vendor-specific unique identifier (23 through 0). Manufacturers of Ethernet equipment license a unique vendor block code from the IEEE and then are responsible for assigning unique MAC addresses for all of their products. Each vendor block code covers 16 million (2^{24}) unique addresses.

Following the addresses is a length/type field that has two possible uses, for historical reasons. Prior to IEEE standardization, Xerox got together with Intel and Digital Equipment Corporation to agree on a standard Ethernet frame called *DIX*. DIX defines a type field that uniquely identifies the type of payload (e.g., IP) to enable easier parsing of the frame. When the IEEE first standardized Ethernet, it decided to implement a length field in place of a type field to more easily handle situations wherein payloads were less than the minimum 46 bytes allowed by the standard. This bifurcation of Ethernet caused interoperability problems. Years later, in 1997, the IEEE changed the field to be a combined length/type field. Values up to 1500 are considered lengths, and 1501 and above are considered types. Most Ethernet implementations use the original DIX-type field scheme. The IEEE has standardized a variety of type values to identify IP and certain other protocol extensions. Payloads with fewer than 46 bytes must be *padded* with extra data to meet the minimum frame size. The

TABLE 9.13 IEEE 802.3 Ethernet Frame Format

Field	Bytes	Fixed Value
Preamble	7	0x55
Start of frame delimiter	1	0xD5
Destination address	6	No
Source address	6	No
Length/type	2	No
Payload data	46–1500	No
Frame check sequence (CRC)	4	No

resolution of how many real data bytes are actually in an Ethernet frame is typically handled by higher-level protocols, such as IP, that contain their own length fields.

Modern Ethernet frames can be longer than 1,518 bytes for a couple of reasons. First, the IEEE has defined various data fields that can be thought of as extensions to the traditional Ethernet header. These include VLAN (virtual LAN) and MPLS (multiprotocol label switching) tags. Each of these extensions provides additional addressing and routing information for more advanced networking devices and adds length to the frame. Second, the industry began supporting *jumbo frames* in the late 1990s to extend an Ethernet frame to 9 kB. The advantage of a jumbo frame is that the same amount of data can be transferred with fewer individual frames, reducing overhead. Jumbo frame support is not universal, however. Older Ethernet equipment most likely will not handle such frames.

The frame check sequence is a 32-bit CRC that is computed across the entire main portion of the frame from the destination address through the last payload byte.

Because of its original topology as a shared bus, Ethernet employs a fairly simple yet effective arbitration mechanism to share access to the physical medium. This scheme is collision detection with random back-off, which was discussed earlier. Ethernet is referred to as CSMA/CD because of its access sharing mechanism. Frame size plays a role in the operation of CSMA/CD. A minimum frame size is necessary to ensure that, for a given physical network size, all nodes are capable of properly detecting a collision in time to take the correct action. Electrical signals propagate through copper wire at a finite velocity. Therefore, if two nodes at opposite ends of a bus begin transmitting at the same time, it will take a finite time for each to recognize a collision. Once a frame is successfully in progress, all other nodes must wait for that frame to end before they can transmit. A maximum frame size limits the time that a single node can occupy the shared network. Additionally, a maximum frame size limits the buffer size within Ethernet MAC logic. In the 1970s and early 1980s, the cost of memory was so high as to justify relatively small maximum frame sizes. Today, this is not a significant concern in most products, hence the emergence of jumbo frames.

Even when Ethernet networks are deployed in physical star configurations, they are often connected to hubs that electrically merge the star segments into a single, logically shared medium. A traditional Ethernet network is half-duplex, because only one frame can be in transit at any instant in time. Hubs are the least expensive way to connect several computers via Ethernet, because they do little more than merge star segments into a bus. Bus topologies present a fixed pool of bandwidth that must be shared by all nodes on that bus. As the number of nodes on a network increases, the traffic load is likely to increase as well. Therefore, there is a practical limit on the size of a bussed Ethernet network. *Bridges* were developed to mitigate Ethernet congestion problems by connecting