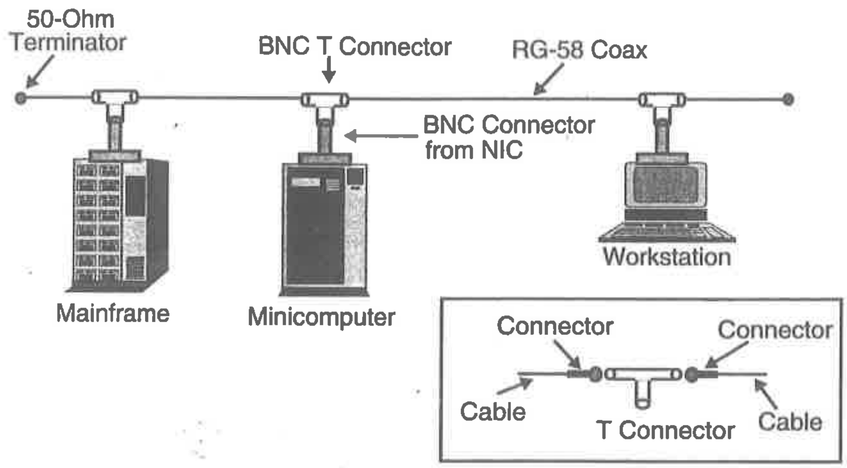

FIGURE 9.3 Example of a 10BASE2 (ThinWire Ethernet) LAN. The topology is a bus configuration, the maximum length of a single segment is 185 m, and a single segment supports up to 30 nodes. Nodes are connected to the cable via BNC T connectors, which must be spaced at least 0.5 m apart to prevent signal interference. Each end of the cable is terminated with a 50-ohm resistor, and one end of the cable must be grounded. A segment is composed of several pieces of cable, with each piece being connected via a T connector (see inset).

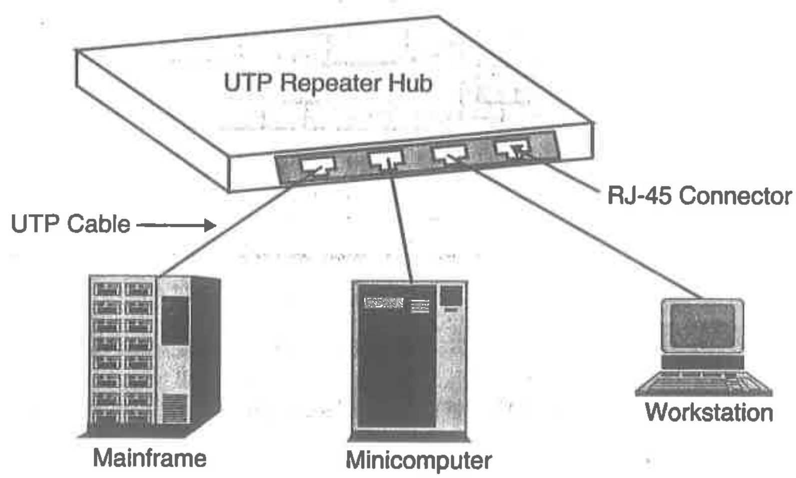

FIGURE 9.4 Example of a typical 10BASE-T (Twisted-Pair Ethernet) LAN. The topology is a star configuration, and the maximum length of a segment is 100 m. Note that only two nodes can be connected per segment—a workstation and a repeater. Both the hub and a device's network interface card (NIC) contain eight-pin modular (RJ-45) connectors. Cable can be Category 3, 4, or 5; higher grade cables provide better performance.

repeaters, delay is how long it takes a repeater to transmit a regenerated signal after receiving it. Of the various devices available for extending an Ethernet/802.3 LAN, only repeaters need careful consideration because they are layer-1 devices that regenerate all incoming signals, including collisions, and propagate these regenerated signals to all the segments connected to its ports. Repeaters extend the diameter of a network but are considered to be part of the same collision domain of networks designed using only repeaters. Bridges, switches, and routers, however, are layer-2 or layer-3 devices. They perform filtering and frame translations (e.g., from an 802.3 format to an 802.5 format). They do not propagate collision signals from one segment to another. Hence, these devices effectively partition a network into multiple collision domains. For example, in Figure 9.8, if the third repeater from the left were a bridge, then there would be two separate collision domains. The first contains the first two repeaters, and the second contains the last repeater. (The network diameter remains the same.) Thus, if we intend to extend our network using repeaters, we have to limit how many repeaters are connected to the network so we can maintain that 512 bit-time upper limit.

The 5-4-3 Repeater Placement Rule

Since the amount of delay introduced by repeaters varies from one vendor to another, a general rule of thumb called the *5-4-3 repeater placement rule* can be followed if we do not want to calculate the total delay ourselves. The 5-4-3 rule requires no more than five segments of up to 500 m each, no more than four repeaters, and no more than three segments that have end nodes connected to them. This rule is also known as the 4-repeater rule or the 5-4-3-2-1 rule. In the latter, the 2 implies that two of the five segments are used as interrepeater links, and the 1 implies that a configuration using the maximum parameters permitted results into one collision domain. If a LAN consists of a single segment (Figures 9.1 and 9.2) or a single UTP repeater hub (Figures 9.4, 9.5, and 9.6), then the 5-4-3 repeater placement rule is of no practical concern. However, if a LAN's diameter is going to be extended using repeaters, then the 5-4-3 repeater placement rule becomes extremely critical.

In lieu of this general rule, we can always calculate the actual delay of our network. If we want to do this, then we must: (a) get the vendor latency specifications for all the repeaters, NICs, and cable used on our network; (b) add the delays; (c) multiply the sum by 2 (we want round-trip values); and (d) see if the product is less than or equal to 51.2 μs. If it isn't, then our network is not Ethernet/802.3 compliant. It is usually a lot easier just to follow the 5-4-3 repeater placement rule.

If we choose neither to follow the 4-repeater rule nor calculate the actual delay, then our network might be noncompliant and could possibly develop problems. Some problems manifest themselves as bad or invalid frames. These include oversized frames, runt frames, jabbers, alignment or frame errors, and CRC errors; all were discussed in Chapter 5. Timing errors can also lead to retransmissions. For example, a sending node that does not receive an acknowledgment from the receiving node within a prescribed time period assumes the frame was lost and retransmits it. Continual retransmissions can ultimately lead to degradation of network performance. Probably the biggest problem with Ethernet networks is collisions. Box 9.1 contains a description of some common causes of collisions, and Box 9.2 addresses Ethernet performance issues.

BOX 9.1 Common Causes of Collisions

Common causes of collisions in Ethernet/802.3 include the following:

Propagation Delay. Different types of media have different propagation delays, as do repeaters and bridges. If nodes are far apart on long segments, there is a greater chance of collisions due to propagation delay, especially in high-traffic environments. Add repeaters or bridges to the configuration and the probability increases. Thus, the length of certain types of Ethernet/802.3 media can have a dramatic effect on whether collisions occur.

Nodes Not Following the Rules. The IEEE specifications establish specific rules for collision detection and retransmission. Unfortunately, not all vendors follow these rules. When a vendor violates the rules on collision detection, it can affect the performance of the network. The 802.3 standard specifies that after a collision, retransmission of a packet should occur after generation of a random amount of timer delay not to exceed 1024 slot-times, which occurs on the tenth to fifteenth consecutive collision. (One slot-time is defined as 51.2 μs for 10-Mbps cable plants.) This means that if a collision occurs, a jamming signal is sent, a random number is generated, and the controller waits that long before retransmission. The controller then seizes the cable (after sensing it) and sends the packet again (if the cable is idle). Although the wait interval is supposed to be between 1 and 1024, some vendors violate this rule and set a ceiling lower than 1024; doing so does not allow random numbers higher than a predetermined value, thus allowing a system to acquire the network quicker than those that generate a higher random number in accordance with the IEEE standard.

Noise. Noise is a pretty obvious source of collisions. Recall that noise is any type of undesirable signal. Noise can come from a variety of locations, including external sources or harmonic distortion. Various everyday office equipment such as copiers, laser printers, ballast transformers on fluorescent tube lighting, and HVAC motors also causes noise problems. Noise is more problematic with UTP cable.

Improper Segmentation of Cable. This reason is restricted to coaxial cables. ThickWire cable should be cut in accordance with the standard and at specific lengths (e.g., 23.4 m, 70.2 m, etc.); ThinWire networks should have at least 0.5 m distance between nodes. Improper segmentation can cause noise and harmonic distortion problems, thus increasing the likelihood of collisions.

Babbling Transceivers. When a transceiver fails, it begins to spew all kinds of trash on a cable; collisions inevitably occur.

Workarounds to the 5-4-3 Repeater Placement Rule

One way to increase the number of Ethernet ports and still comply with the 5-4-3 repeater placement rule is to use chassis-based repeater hubs (see Figure 6.8). These hubs have individual slots in which multiport Ethernet interface boards (also called *blades, cards,* or *modules*) are installed. The primary feature of a chassis hub is that each board is connected to the same backplane to which the repeater unit is connected. Since all boards share the same backplane with the repeater, all of the devices connected to the hub use only one repeater.

Another workaround is stackable repeater hubs (see Figure 6.9). These devices consist of individual hubs "stacked" one on top of another. Instead of a common chassis backplane, stackable hubs use a "pseudobackplane" based on a common connector interface. An external cable interconnects the individual hubs in a daisy-chain. Once interconnected, the entire