TABLE 3.12 INTEROPERABILITY OF PRESTANDARD AUTO-SPEED AND AUTO-NEGOTIATION SWITCHES AND NICs

NIC-Only Hub	10BASE-T–Only Hub	100BASE-TX-10/100TX Hub
10BASE-T–only NIC	No choices can be made.	N/A (better buy a new NIC!)
Pre-standard auto-sensing 10/100TX NIC[2]	NIC automatically selects 10 mode.	NIC automatically selects 100 mode.
New auto-negotiation 10/100TX NIC[3]	NIC auto-negotiates to 10 mode.	NIC auto-negotiates to 100 mode.

1 *For example, the SynOptics 28115 hub will operate at either 10 or 100Mbps, but does not automatically adjust its speed. The speed adjustment has to happen manually via the switch management software.*
2 *Examples are the Intel EtherExpress Pro/100, or the 3Com Fast EtherLink.*
3 *The Intel EtherExpress Pro/100 Model B and the 3Com EtherLink XL include auto-negotiation.*

Ethernet Repeaters

Repeaters extend the size of a network by joining multiple segments into a larger segment. A repeater works at the physical layer (Layer 1) of the OSI model. That means the repeater does not process the data at all. A repeater has no MAC, only PHYs, and merely receives the incoming signals and reconditions them for immediate retransmission on all ports. Repeaters are synonymous with shared media. A repeater is invisible to all nodes on a repeated LAN. It appears as though all nodes are connected via one cable.

Repeaters are almost as old as Ethernet itself. Fast and Gigabit Ethernet make provisions for repeaters, although the trend toward dedicated media and switching has been accelerating for the past five years. With Gigabit Ethernet in particular, repeaters might soon be outdated. Figure 3.13 shows the internal workings of a repeater.

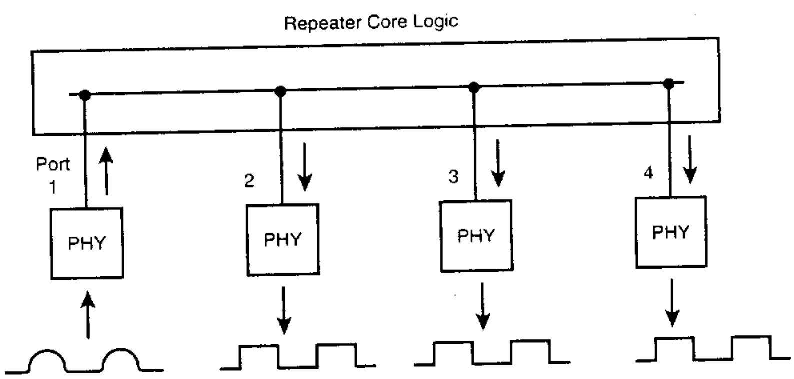

FIGURE 3.13 *A repeater operates at Layer 1 of the OSI model: it only contains PHYs and some core logic, but no MACs. This figure shows a repeater receiving a degraded signal on port 1 and retransmitting a clean signal on ports 2–4.*

1st -Generation 10/100 Hub (No Auto-Sensing)[1]	New Auto-Negotiation
Hub manually set to 10 mode (with management software).	Hub auto-negotiates to 10 mode.
Manually set hub to 100 mode, NIC will automatically adjust to 100 mode also.	Manually set hub and NIC to 100 mode.
Manually set hub and NIC to 100 mode.	Both hub and NIC auto-negotiate to 100 mode.

How Repeaters Work

All Ethernet repeaters work as follows:

1. An encoded data signal is received on a particular port, either from a node, a switch, or another repeater. This signal has been electrically degraded because it has traveled some distance from the source to the repeater over less-than-perfect cable. (This applies to both UTP and fiber; the degradation just takes longer over fiber.)
2. The incoming port, PHY, processes the data and re-creates the perfect digital data internally.
3. Finally, the digital signal is forwarded to all ports, where the PHY on each port converts the data into an encoded signal, to be transmitted over the wire again.

Note that no data is stored, no frames are re-created, and every received bit is sent on, irrespective. Figure 3.14 illustrates the flow diagram of a repeater.

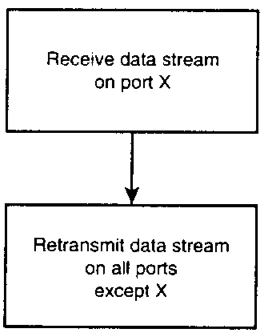

FIGURE 3.14 *Flow diagram of the repeater forwarding logic: There isn't much to it. Later on you will see the same diagram for a switch, which is much more complex.*

The Repeater Collision Domain

A *segment* is defined as a group of nodes connected to the same repeater. Multiple segments can be connected via repeaters. A repeater propagates all network traffic present on one segment to all other segments to which it is connected. All interconnected segments are in one *electrical collision domain*.

Electrical signals take a certain time to travel across a cable. In addition, all repeater hops introduce a small delay or latency. This is the delay between the time an incoming signal is received and the time that signal is transmitted again to all ports. These two forwarding delays and their impact on Ethernet collision detection is a key factor for determining 10, 100, and 1000Mbps Ethernet network design and diameter rules.

A collision occurs when two nodes simultaneously sense that the media is available and then attempt to transmit at the same time. Somewhere in the shared collision domain the two data patterns literally collide, causing a voltage surge on the cable. The voltage surge travels back along the cable to the two sources, where it is detected by the transmitters. A collision must be detected by the nodes causing the collision before they stop transmitting, including the one at the farthest end of the cable; otherwise, the nodes would never know that their transmission had been corrupted and would proceed as though nothing had happened. After the transmitting station recognizes that a collision has occurred, a JAM signal is generated, which is equivalent to canceling the last transmission. The worst case for collision detection is very short frames, because they leave very little time for the collision signal to travel back to the transmitter. Figure 3.15 shows the collision domain and minimum frame size.

Let's look at how the maximum network diameter depends on the minimum frame size. First, we need to define a variable called the transmission time. This is the amount of time that the transmission takes, from beginning to end. It is also equal to the time that any collision has to travel back to the transmitting station.

The minimum transmission time can be calculated from the minimum frame size. For 10Mbps Ethernet, the minimum transmission time is

```
Minimum transmission time = Minimum frame size (512 bits) x Bit-time (100ns) = 51.2µs
```

The signal needs to travel from the transmitter to the point where the collision occurs. Then the collision signal needs to travel back to the transmitter again. The minimum transmission time number in the preceding equation needs to be divided by 2 to account for round-trip delays. Therefore, the minimum transmission time is equal to 25.6µs (round-trip).

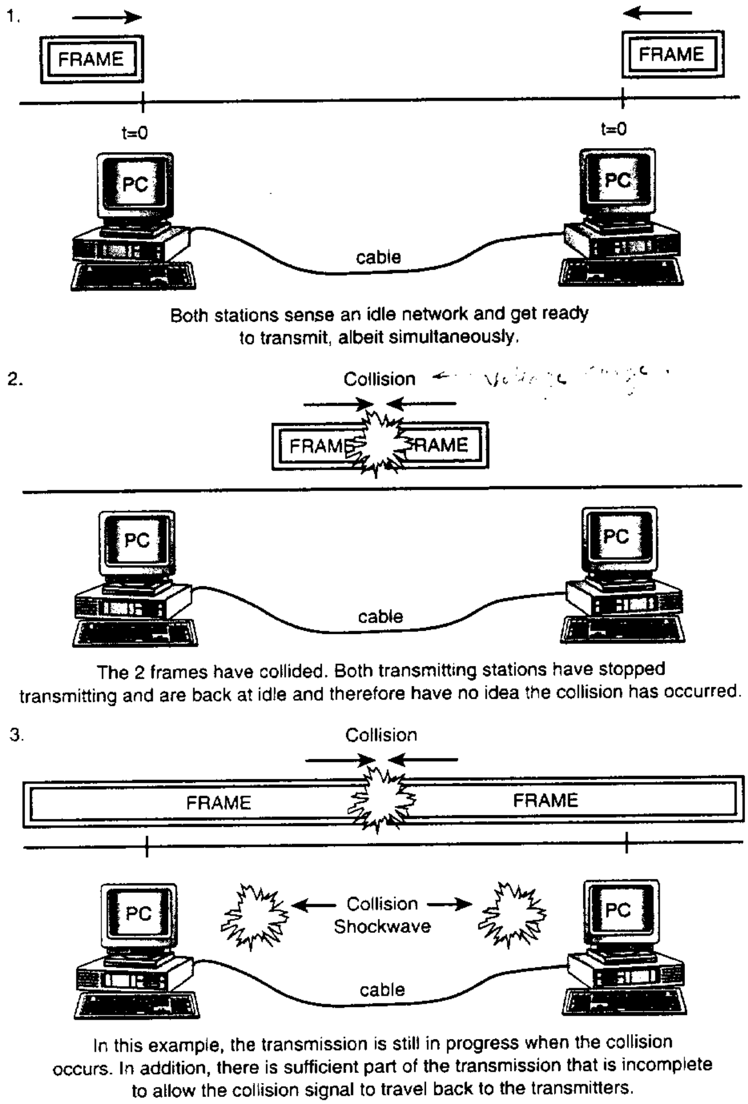

FIGURE 3.15 *Collisions occur when two nodes attempt to transmit simultaneously. The maximum network diameter is directly related to the minimum frame size.*

Copper and fiber cable have a delay of between 0.5 and 0.6µs/100m. To calculate the maximum network diameter, we merely need to divide the minimum transmission time of 25.6 µs by the cable delay.

$$\text{Network diameter} = \frac{25.6\mu s}{0.6\mu s/100m} > 4{,}000m$$

This calculation still excludes any delays introduced by repeaters, nodes, and so on. Many manufacturers specify the latency of their repeaters, so you could calculate the exact network diameter depending on the manufacturer's numbers. This is very cumbersome. To make our lives a little easier, the IEEE came up with a quick and easy-to-remember rule that limits the 10Mbps network size to 2500 meters. The rule is described in more detail later.

When 10Mbps Ethernet was originally designed, the minimum frame size of 512 bits was chosen to come up with a realistic network diameter. The problem with Fast Ethernet and Gigabit Ethernet is the bit time. The network diameter is directly related to the minimum frame size, as well as the bit time.

$$\text{Network diameter} = K \times \frac{\text{Frame size}}{\text{Bit-time}}$$

As we move from 10 to 100 to 1000Mbps, the bit time decreases by an order of 10 every time. If we have a network diameter of 2500 meters for 10BASE-T, this shrinks to one-tenth, or approximately 250 meters for 100BASE-T, which is already marginal. For Gigabit Ethernet, the network diameter would shrink to 25 meters, which is clearly not a viable option. Therefore, the IEEE decided to increase the minimum frame size for Gigabit Ethernet to 512 bytes (or 4096bits), as opposed to 64 bytes for 10 and 100Mbps Ethernet. This, as you might notice, is only an eight-fold improvement. This yields a network diameter of around 200 meters for Gigabit Ethernet. Table 3.13 summarizes the different collision diameter variables for 10, 100, and 1000Mbps Ethernets.

You can calculate the exact maximum network diameter for any network using the delays for cable and manufacturers specifications for latencies, but nobody would really want to do that. (For you die-hards who insist, we have put down some guidelines. See the section "Calculating Your Own Collision Diameter," later in this chapter.)

The IEEE made things easy and developed a set of guidelines. We discuss these cheat sheet guidelines for 10, 100, and 1000Mbps repeater segments in the following section.

10Mbps Ethernet Repeater Rules

You can easily memorize the 10Mbps Ethernet golden rule as the *5-4-3-2-1* rule:

- *Five* segments are allowed (of 500-meter diameter each).
- This implies *four* repeater hops in the data path.
- *Three* of these segments may be populated with nodes.

- *Two* segments cannot be populated but are only inter-repeater links.
- All of this makes *one* large collision domain with a maximum of 1024 stations. Total network diameter can be up to 2500 meters.

TABLE 3.13 THE COLLISION DIAMETER FOR 10, 100, AND 1000MBPS ETHERNETS ARE A FUNCTION OF MINIMUM FRAME SIZE AND BIT TIME

Parameter	10Mbps	100Mbps	1000Mbps
Minimum frame size	512 bits (64 bytes)	512 bits (64 bytes)	4096 bits (512 bytes)[1]
Collision diameter (bit times)[2]	512 bit times	512 bit times	4096 bit times
Bit time (μs)[3]	0.1μs	0.01μs	0.001μs
Maximum round-trip delay (μs)[4]	51.2μs	5.12μs	4.096μs
Maximum network diameter, without repeaters [5]	Approximately 45710m	457m (1/10 of Ethernet)	3661 m (8/10 of Fast Ethernet)
Maximum IEEE-specified network diameter, with 100-meter UTP connections [6]	2500m	205m	200m
Maximum number of repeaters with specified network diameter[7]	~5	2 or 1	1

1 *This number is eight times the regular or Fast Ethernet minimum frame size. This doesn't entirely make up for the one-tenth reduction in diameter due to the speed increase, but it comes close. To be totally correct, the actual Ethernet frame size has not increased. Instead, the minimum Gigabit Ethernet frame size is still 64 bytes, just like 10 and 100Mbps Ethernet frames. The 802.3z specification requires a carrier extension to be added to small frames to keep the wire busy for a minimum of 512 bytes. In this way, it looks like the minimum frame size is 512 bytes.*

2 *Collision diameter = minimum frame size. The IEEE specifies the minimum frame size.*

3 *Bit time = 1/data rate. Depends on the speed of Ethernet.*

4 *Round-trip delay = bit time * minimum frame size.*

5 *Absolute best case, calculated using a copper delay of 0.56μs/100m segment. Ignores 100m UTP limitation. Also assumes no repeater delays, so this is a theoretical number. Because fiber cabling has similar delays to copper, this number comes close to the half-duplex fiber network diameter specifications.*

6 *This is the real-world number that includes a budget for repeaters, too.*

7 *We discuss exact repeater rules later in the chapter.*

Figure 3.16 illustrates this 5-4-3-2-1 design rule, which applies to 10Mbps Ethernet only. This 5-4-3-2-1 rule is only a rough guideline, but it works very well in most cases. Actual numbers vary by manufacturer. We have seen from four to seven repeater hops being specified. (In theory, as discussed previously, a 10Mbps Ethernet LAN could have a collision diameter of at least 4000 meters, if no repeaters were present. This is a number that can be attained only by using a point-to-point single fiber connection because all copper-based cabling will need to be repeated multiple times to go this far.)

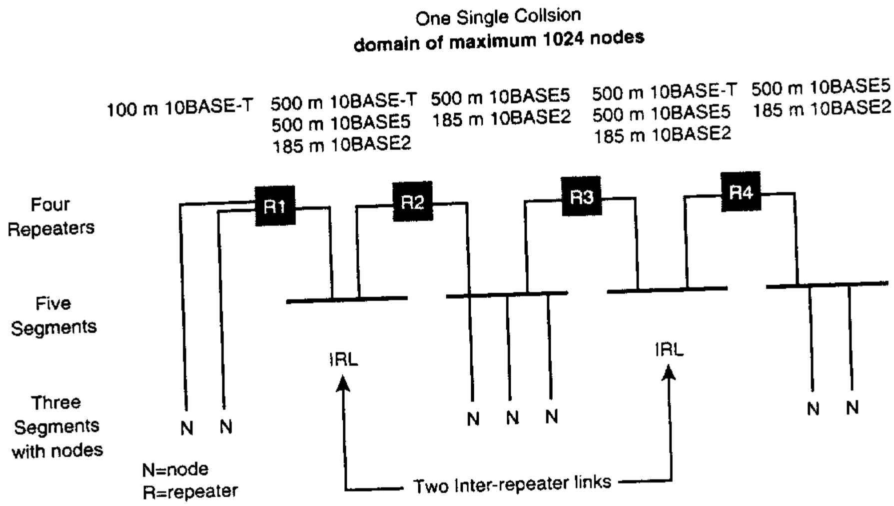

FIGURE 3.16 *The Ethernet 5-4-3-2-1 repeater rule stipulates a maximum of five repeaters, four segments, three populated, and two repeater links, making up one collision domain.*

100BASE-T Repeater Rules

The IEEE repeater rules have hence been significantly changed for 100BASE-T. The IEEE 802.3u Fast Ethernet standard contains some specifications for Fast Ethernet repeaters. Due to the faster speed, the Fast Ethernet network diameter has been reduced from 2500 meters to approximately one-tenth that size.

The actual diameter depends on two things:

- *The type of cable used*—Fiber cable has a slightly smaller delay: about 0.5µs/100, as opposed to UTP, which has a delay of 0.56µs/100m. (Note that all UTP segments still need to be 100 meters, as this is an EIA/TIA spec.)
- *The type of repeater used*—With 10BASE-T, latency was not an issue, but with 100BASE-T, it becomes a major differentiating feature for a repeater. As the network diameter is rather limited for Fast Ethernet, the IEEE has defined two different classes of repeaters, called Class I and Class II:
    - A Class I repeater can have a relatively large port-port timing delay of 0.7µs or less. It operates by translating line signals on an incoming port to digital form and then retranslating them to line signals when sending them out on the other ports. This makes it possible to repeat signals between different media segments, such as FX and TX, within the same repeater hub. The larger delay means the 2.56µs collision diameter only allows for one Class I repeater to be used with two 100-meter cable lengths. For two fiber links, the maximum cable diameter is 272 meters.

- Class II repeaters have a lower latency of 0.46μs or less. Class II repeaters do not regenerate the digital signal like a Class I repeater does. A Class II repeater immediately repeats the incoming signal to all other ports without a complete translation process. Class II repeaters are preferred from a network design perspective because the lower latency allows for two repeater hops. The drawback of Class II repeaters is a single repeater cannot mix different media PHYs.

Class I and Class II differ with respect to their internal design and latency characteristics. This allows a more accurate network diameter calculation to get closer to the limit.

If two 100-meter UTP links are used, this configuration allows for an inter-repeater link of 5-meter UTP to yield a total network diameter of 205 meters. Alternatively, two fiber connections and an inter-repeater fiber link with a total length of 228 meters are permitted. Most Fast Ethernet repeaters shipping today are Class II.

The two cases just described both use only one media type for the entire network diameter. The IEEE has also laid down some guidelines for mixed copper-fiber installations, but these situations are rare. Table 3.14 illustrates the network diameter guidelines for these situations. Figure 3.17 graphically depicts Table 3.14.

TABLE 3.14 FAST ETHERNET COLLISION DIAMETER RULES

Connection	All UTP	Mixed-Media 100BASE-TX[1] and 100BASE-FX	All Fiber
Repeated Network Segment Diameter			
One Class II repeater	200m	309m	320m
One Class I repeater	200m	261m	272m
Two Class II repeaters	205m	216m	228m
Point-to-Point Connections[2]			
Node-node or switch-switch half-duplex	100m	N/A	412m
Node-node or switch-switch full-duplex	100m	N/A	2000m

1 *T4 allows for a few meters more.*

2 *We have shown the point-to-point distance limitations for reference only. For single-mode fiber connections operating in full-duplex, even longer segments are possible, but the IEEE didn't specify SMF as part of the Fast Ethernet standard.*

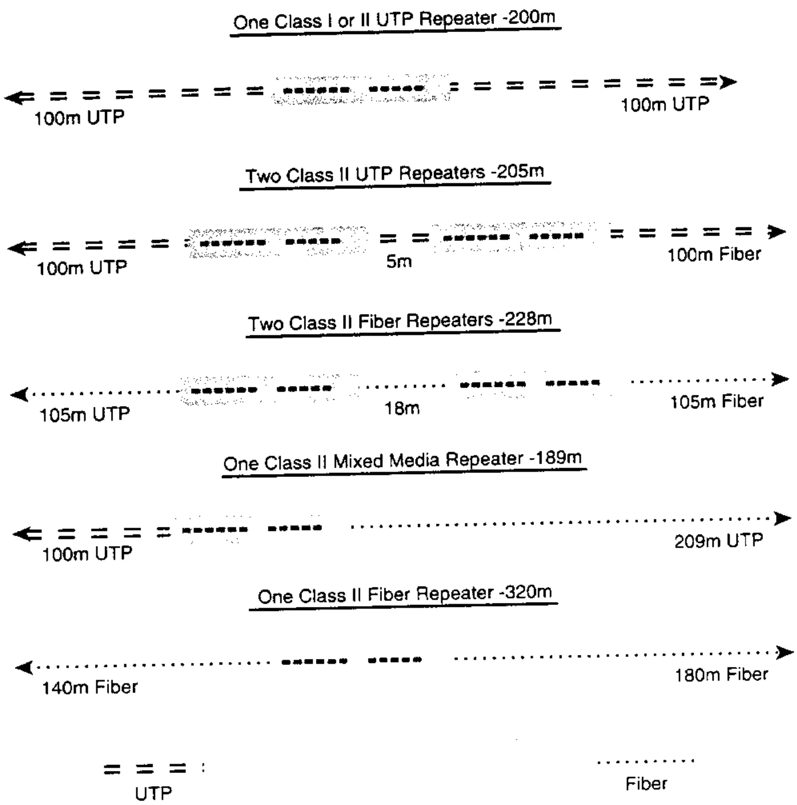

FIGURE 3.17 *The network diameter for a 100BASE-T network can range from 200 meters for UTP to 320 meters for fiber, assuming one repeater.*

Some people originally viewed the Fast Ethernet repeater specification (two repeater hops and a 205-meter UTP network diameter) as a big issue. Most people soon realized, however, that one could effectively work around both issues with the use of stackable repeaters and switching hubs to get to any network diameter and an infinite number of hops. Refer to the deployment chapters for more details.

Five years ago, a network diameter of 205 meters combined with one repeater hop would have made 100BASE-T impractical if not altogether useless for most LAN managers. Two technological developments, however, have made it possible for 100BASE-T repeaters to work well even within these tight restrictions:

- Most 100BASE-T networks are additions to an existing 10BASE-T network. Every new 100BASE-T segment that is added to a 10Mbps network will require a switch to get from 10Mbps to 100Mbps. The 205-meter calculation is started at the switch, and 205 meters is sufficient to get to most nodes with one or two intermediate repeater hops. If the distance from switch to node is longer than 205 meters, another switch must be added to further extend the network diameter. Alternatively, 100BASE-FX could be run to repeaters or nodes that are longer than 205 meters away from the switch.

- Standalone, unmanaged 100BASE-T repeaters will be rare. Most 100BASE-T repeaters will be stackable, meaning that many repeaters can be physically placed on top of each other and connected via a fast backplane bus. The fast backplane bus does not count as a repeater hop and makes the entire stack look like one larger repeater. In fact, most 100BASE-T stackables can be stacked four or more high. Electrically, the repeater stack appears as one larger repeater.

Calculating Your Own Collision or Network Diameter

You can calculate your own network diameter if you like. Use the following equation to make sure your repeated Ethernet network segment will function satisfactorily:

```
(repeater delays + cable delays + NIC delays) x 2 < Maximum round-trip delay
```

Note the following:

- You can obtain the maximum round-trip delay from Table 3.15. It is 51.6μs for 10Mbps Ethernet, 5.12μs for Fast Ethernet, and 4.096μs for Gigabit Ethernet.
- The factor 2 accounts for round-trip delays as the frame needs to travel from transmitter, to the collision point and back to the transmitter again.
- We can obtain repeater and NIC delays from the manufacturer. Ethernet repeaters typically have latencies of 2μs or less. Fast Ethernet Class I repeater delays are less than 0.7μs; Class II delays are less than 0.46μs. A 10Mbps NIC delay is about 1μs; Fast Ethernet NIC delay is 0.25μs. All networks are terminated on both ends: two NICs or a NIC and a switch port.
- We can obtain the cable delays from Table 3.15. Cable delays do not at all depend on the Ethernet speed. Cable delays typically are measured as a fraction of the speed of light and mostly depend on the insulation material used. A 100-meter UTP section still takes about 0.55μs one way.

- For 10Mbps networks, 51.2µs allows for quite a few repeaters, cable segments, and DTEs to exist without exceeding the collision domain restrictions. In most real-world 10BASE-T networks, you will find it very difficult to exceed this number.

TABLE 3.15 INDIVIDUAL NETWORK COMPONENT DELAYS FOR CALCULATION OF NETWORK DIAMETER

Component	Maximum Delay Times in µs
Two Fast Ethernet NICs or switch ports	0.5
Fast Ethernet MII	0.2
100m Category 5 cable segment	0.556
1m Category 5 cable	0.00556
100m Fiber optic cable segment	0.5
1m fiber cable	0.005
Fast Ethernet Class I repeater	0.7 max
Fast Ethernet Class II repeater	0.46 max
Two Gigabit Ethernet NICs	0.864 max
Gigabit Ethernet repeater	0.488 max

Figure 3.18 illustrates the different delay components in a shared-media network.

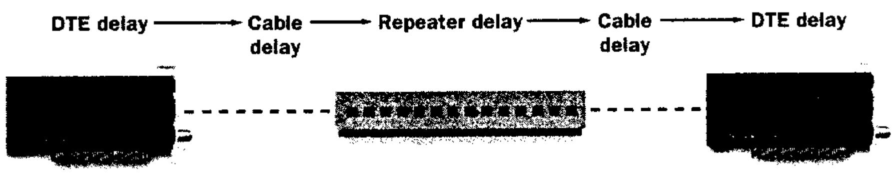

FIGURE 3.18 *If you want to calculate your own maximum network diameter, you need to add up the delays of the individual components and make sure they don't exceed half the minimum frame size.*

Some additional points to note:

- If you are using an MII cable, you need to add it to the cable segment length.
- Decide on an appropriate safety margin. (We recommend five bit times.)
- You will see that it is actually possible to build 100BASE-T networks with three or even four repeater hops if the individual cable segments are significantly less than 100 meters.