

What You Will Learn

After reading this chapter, you should be able to

- ✓ Explain how computers can communicate binary 1s and 0s using electricity
- ✓ List the components of popular LAN cables and connectors
- ✓ Compare and contrast straight-through and cross-over Ethernet cables
- ✓ Explain the operation of an Ethernet hub
- ✓ Summarize the benefits of using a structured cabling system



CHAPTER 4

How to Build a Local (Network) Roadway

So far in this book, you have read about how networks allow computers to communicate. Networks include software, some of which sits in the computers, and some of which sits in routers. The network also includes hardware, such as the network interface cards (NICs) in the PCs, mentioned in Chapter 3, “Building a Network: It All Starts with a Plan.” Finally, networks include cabling, which provide a physical means to transmit bits across a network.

This chapter is the first one in this book’s second major part: “Running the Local Department of (Network) Transportation.” In the United States, most cities, all states, and the U.S. federal government have a department of transportation (DOT). Each DOT plans, builds, and fixes problems with roadways. You can learn a lot about networking by comparing networks to roads.

Driving Bits Across the Network Roadway

I just got back from lunch at one of my favorite lunch places: La Frontera Mex-Mex Grill. When it was time for lunch, I got in my car, drove to the restaurant, ordered my usual—*numero dos, con pollo, por favor*—and drove back home.

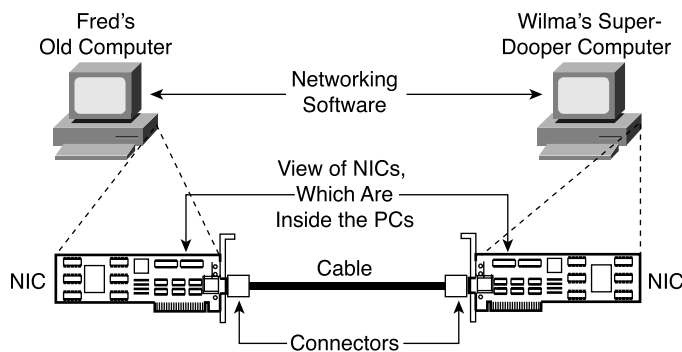
When I got home, my wife said, “So you drove over that really cool street outside our house on your way to lunch, huh?” Yeah, right. Who cares what roads I drove on? It’s where I drive *to* that’s important.

Well, to know networking well, you need to know some of the basics about the networking equivalent of roads. If you find yourself really getting interested in what's in the next few pages, you might just be one of those people who really would like a career working with the technical side of networking. If not, you should at least know the concepts so that you can communicate with the networking geeks of the world and have a firm understanding of networking.

What's a Local-Area Network?

Chapter 3 defined a local-area network (LAN) as a network in which the devices are relatively close together. Of course, a network, once again according to this book, includes computers, software, hardware, and cabling that allow the computers to communicate. Although this definition is accurate, you really do not get a detailed picture of a LAN this way. So, Figure 4-1 shows a LAN, which is one you've seen before in Chapter 3.

Figure 4-1 Components of a Simple LAN



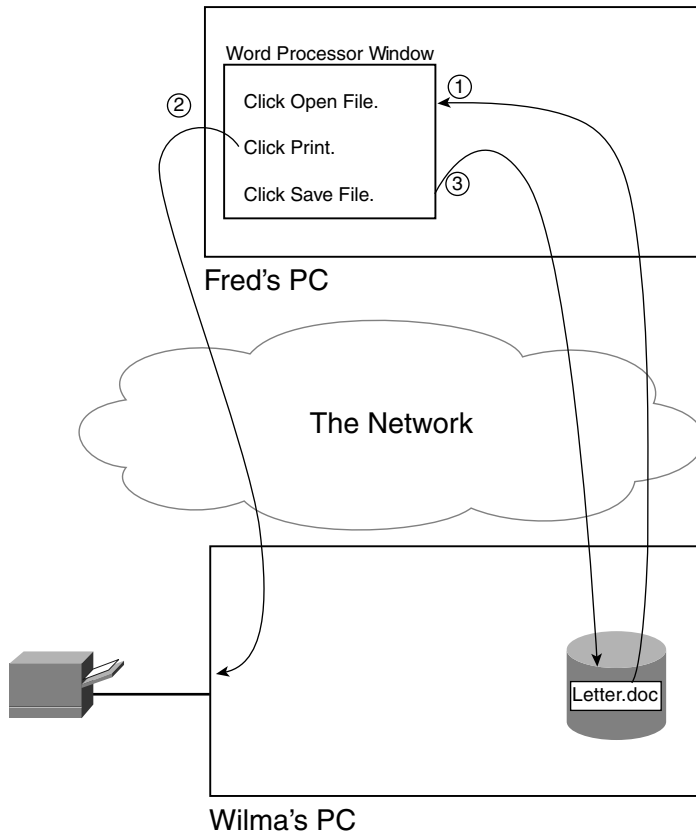
In this case, the LAN consists of some obvious elements. First, you need at least two computers. The computers need to have networking software; otherwise, they will never attempt to communicate. They also need the physical ability to transmit bits from one computer to the other—hence the need for the *cable* and the NICs in each computer. (Figure 4-1 shows the NICs outside the PCs so that you can see them, but they are typically inside the PCs.)

LANs do not get any simpler than this one. Larger LANs can get much more complex, with lots more components like networking hubs and switches, which you will learn more about in the next few chapters.

In this chapter, you learn about how computers can transfer bits across the network roadway. The topics covered here might be the equivalent of what a DOT engineer might talk about over lunch with a stranger: “Hey, you know we’re going to be paving Parker-Puckett Parkway pretty soon. Pretty cool, huh?” The reply, “Hey, isn’t that near that Mexican place, La Frontera?” And the response from a true DOT engineer: “I don’t know—don’t care. That’s *beside* the road. I just care about the road.” Likewise, some network engineers think of the LAN as the cable, possibly the NICs, but they typically don’t care a lot about the computers that happen to be connected to the cabling. Likewise, this chapter focuses more on how computers send bits to each other, rather than the applications that run on those computers.

Transmitting Bits Across the Local Network Roadway

Back in Chapter 2, “A Network’s Reason for Existence,” you saw an example where Fred opened a file that sat on Wilma’s disk drive. Later, Fred printed the file on the printer connected to Wilma’s computer, and finally, Fred saved the file back on Wilma’s disk drive. Figure 4-2 shows the process, in sequence.

Figure 4-2 Basic Flow with Fred Using a File/Print Server


For this process to work, Fred and Wilma must be able to cause a *bunch of bits* that sit in memory in one computer to be sent to the other computer. A file is just a bunch of bits. So, imagine that by using a word processing program on a single computer, you can open a file, read the contents, and edit and change the file. Similarly, if you can send the same bits in that file to another computer across the network, you or someone else can edit the file by using the same word processing program on that computer. So, for a network to work, the network needs to be able to get a bunch of bits from one computer to another.

Driving Bits Across a Wire

To send one binary code from one device to another, the sending device puts some electricity on the wire. Electrical signals have many characteristics that a NIC can control and vary. By varying one of these features to two different values, with one value meaning binary 1 and one meaning binary 0, you can transfer data over the wire.

For example, imagine that both PC1 and PC2 have a NIC, and there is a single wire connecting the two cards. The wire is just a skinny, long piece of copper, and copper conducts electricity very well. Now, imagine that the *encoding standard* used by the company that made the NIC defines that a binary 0 is represented by a voltage of ± 5 volts, and a binary 1 is represented with ± 10 volts. Encoding is the term that refers to a set of rules that defines what a sender should make the electrical signal look like to imply a binary 0 or a binary 1. Figure 4-3 depicts the general idea.

Figure 4-3 Basics of Data Transmission Across a Wire





In the figure, PC2 generates some electricity on the wire. In this case, PC2 wants to send the binary value 0101. So, it sends a 5-volt signal, then 10 volts, then 5 volts, and then 10, because the imaginary encoding scheme in this example states that 5 volts means 0, and 10 volts means 1. PC1, on the other end of the wire, senses the incoming electrical signal and interprets the electricity, using that same set of encoding rules to mean 0101, exactly as PC2 intended.

Note that the graph shown in Figure 4-3 shows a discrete, or constant, voltage. Because the X-axis (horizontal axis) represents time, when the voltage changes, it changes immediately to the next value. The use of discrete, constant values, which are then instantly changed to other possible discrete values (as in Figure 4-3), is called *digital transmission*. To transmit binary numbers, or *binary digits*, it is useful to transmit the data using digital transmission.

For the digital transmission of data to work correctly, not only must the sender and receiver agree to what electrical characteristics mean a binary 0 or 1, but they also must agree to the rate at which the bits are transmitted over the wire. In Figure 4-3, the receiver (PC1) must think about the electrical signal at different points in time, on a regular interval. Likewise, the sender (PC2) must use this same regular time interval to decide when it should change the digital electrical signal. For instance, if PC2 varied the voltage to mean either 0 or 1 every .1 seconds, and PC1 sampled the incoming electrical signal every .1 seconds, they could transfer 10 bits in a second. The speed of this network connection would be 10 bits per second.

If the two PCs did not agree on the transmission speed, the devices couldn't transfer the binary information. For instance, imagine that PC2 thought the speed was 10 bits per second, meaning it should encode a new bit every 1/10 of a second. If PC1 thought that it should be receiving a bit 20 times per second, it would sample the incoming electrical signal every 1/20 of a second. PC1 would think it was sending 10 bits each second, and PC2 would think it received 20 bits.

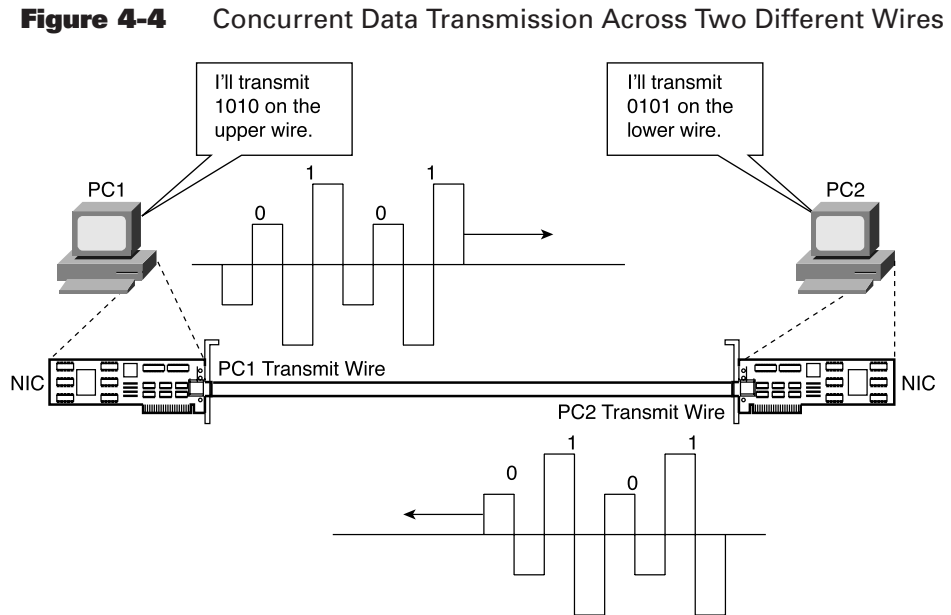
The term *bps* (short for bits per second) often refers to the speed of networking connections. Note that the unit is bits, not bytes. In real life, LANs typically run at much higher speeds, with a slow LAN transmitting at 10 million bits per second (Mbps, also called megabits per second).

Notice that Figure 4-3 represents electricity as a square *waveform*, with positive and negative voltages. You don't really need to worry about the electrical details, but as you progress through learning about networking, you will see other drawings like this one. The networking cards use an alternating current, or AC. The positive voltage means the current is in one direction, and the negative current means the current runs in the opposite direction.

The Need for a Two-Lane (Network) Road

In Figure 4-3, PC2 sends an electrical signal to PC1. As it turns out, if PC1 tried to send some electricity to PC2 at the same time, over the same wire, the electrical signals would overlap, and neither PC1 nor PC2 could understand what was sent.

To solve the problem, PC1 and PC2 need to use two wires—one for PC1 to send bits to PC2, and one for PC2 to send to PC1. Figure 4-4 shows two wires, with PC1 and PC2 sending and receiving at the same time.




The Equivalent of Asphalt: Cables

You can get in your car and drive around your yard, the sidewalk, in your neighbor's yard, or through the park. Of course, it's better if you drive on the road! The earlier examples in this chapter showed NICs using a single copper wire for data transmission in each direction. But rather than just have a couple of wires somehow stuck into the side of the cards, in real life, we use cabling and connectors to manage the wires, making the job of the electrician much more comfortable and convenient.

The copper wires that networking cards use are encased inside a cable. The cable is made from plastic, with the copper wires inside the cable. Figure 4-5 shows a drawing of the most popular type of cabling used for LANs today.


Figure 4-5 Typical LAN Cable



If you look closely at this figure, you can see each copper wire, as well as the plastic coating on the wire. The copper wire is thin, making it brittle. In fact, the wire could easily break in your hand. To help prevent the wire from breaking, a thin plastic coating is painted onto each wire. Conveniently, each wire uses a different color of plastic coating, so you can look at each end of the cable and figure out which wire is which. As you might guess, and as you will learn more about in the next few pages, it is important that you can identify a particular wire on each end of the cable.

Also note that the wires in Figure 4-5 are twisted together in pairs. Each pair of wires is cleverly called a *twisted pair*. The term refers to a pair of wires twisted around each other to reduce the amount of electrical interference on the wires. In layman's terms, *electromagnetic interference (EMI)* occurs when electrical signals that exist in the air—caused by other wires or other nearby electrically powered devices—change the electrical currents on the wire. If outside EMI changes the signal on the wire, the receiving computer might misinterpret a 0 as a 1 or a 1 as a 0, or it might not have a clue what the sender really sent. Sending the electrical signals over a twisted pair rather than a single wire eliminates a lot of EMI effects. (Besides, the wires are pretty skinny anyway, and copper is cheap, so why not use two?)

Another thing that can be added to the cable to reduce EMI is shielding. Shielding, as the name implies, shields the wires inside the cable from the effects of EMI. However, shielded cabling has more stuff in it, making the cable less bendable and more expensive to produce. Shielded cables are called *shielded twisted pair (STP)*, and you could probably guess that unshielded cables are called *unshielded twisted pair (UTP)*. LAN technology has evolved to the point where less expensive UTP cabling can be used in most environments, with STP cabling being used in environments where significant EMI issues exist. Figure 4-5 showed UTP cabling; Figure 4-6 shows an example of an STP cable.


Figure 4-6 Shielded Twisted Pair (STP) Cabling

Painting the Lines on the Road: Connectors

Imagine that you see a new road that has just been paved. The road is so new that the DOT hasn't even painted the lines yet. So, you turn onto the road and enjoy the ride. After a couple of miles, you look up and see that someone else is on the road—a huge, speeding truck—and the driver wants to take his half of the road in the middle of the road. After swerving off the road, you might think to yourself, “Boy, that reminds me of one of the reasons that you use connectors on the end of networking cables!”

Okay, that's far fetched, but it does lead into a key point about connectors. If roads had no lanes, and there were no traffic laws, the roads would be pretty dangerous so the DOT paints lines on the road to create traffic lanes. Similarly, connectors line up the wires on the end of a cable into well-defined physical locations inside the connector. Essentially, each wire in the cable is identified by color, and each colored wire has a specific reserved place inside the connector that the electrician attaches to the end of the cable. The connectors put the wires into the right place, just like the lines on a road guide cars into the right place.


An electrician can take a cable and attach a connector to the end of it. When he attaches the cable to a connector, each of the wires protrudes into the connector so that the electricity can flow when connected to a device. The tip of the exposed wire in the connector is called a *pin*. A *pin* is nothing more than a physical position in the end of the connector in which the copper part of the wire sits. You can think of a pin like you think of a lane on a road. Figure 4-7 shows a photo of a typical connector, called an *RJ-45 connector*, along with a drawing of the same connector.

Figure 4-7 Typical Networking Connector (RJ-45)

The Telecommunications Industry Association (TIA) and Electrical Industries Alliance (EIA) define standards for cables and connectors. For instance, they define how to use an RJ-45 connector for LANs and options for several types of UTP cabling. (You can learn more about TIA and EIA by going to their websites, at <http://www.tiaonline.org> and <http://www.eia.org>.)

If you have never seen an RJ-45 connector, and you use Ethernet LANs at work or school, you could remove the cable from your PC's network card and look at it. In most cases, the cable will be using an RJ-45 connector. (Ethernet is by far the most popular type of LAN today. You'll read more about it in the next several chapters.)

The EIA/TIA defines standards for which wires fit into which pins when you make a cable for use with Ethernet. Two of those standards are shown in Figure 4-8.

Figure 4-8 Pinout Options for RJ-45 Connectors



Each of the drawings in the figure represents an RJ-45 connector. The RJ-45 connector allows eight wires to be inserted into it. EIA/TIA standards suggest the numbering schemes for the eight pin locations and the pairs of wires, knowing that a twisted pair is needed for data transmission. The standard also specifies which color of wire goes into pin position 1, 2, and so on.

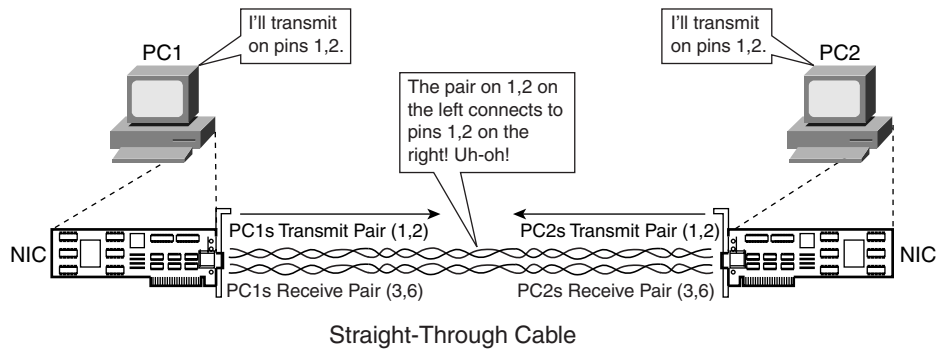
The RJ-45 connectors have clips on the side that allow you to easily insert the connector into the plug (hole) in the networking card in a computer. At that point, the connector is secure and should stay put. If you push the clip close to the rest of the connector, it releases the connector from the card. So, the clip helps keep the connector secure and allows you to pull the connector out when you are ready.

So, why does all this matter? In the next few pages, it will all come together as you see how NICs try to use the cables and wires.

Driving in the Right Lane (Pair) on the Road

Refer to the simple network of Figure 4-1 early in this chapter. That network consisted of two PCs, each with a networking card, and a cable between them. So far, you've learned the basics about wires, cables, and connectors. However, there's one last important thing about basic LAN data transmission that you need to know, and it relates to which wires are used for actual data transmission by Ethernet LANs.

Ethernet NICs in PCs try to send data over the twisted pair that uses pins 1 and 2 of an RJ-45 connector. These same NICs expect to receive data on the twisted pair that uses pins 3 and 6. However, without the right kind of cabling, two Ethernet NICs cannot communicate. Figure 4-9 depicts the effect when both NICs try to send using the twisted pair that uses pins 1 and 2.


Figure 4-9 Both PCs Using the Same Pair (Lane) to Send Data

(Rather than clutter the figure with more cabling, I just drew the two pairs of wires; the wires do indeed sit inside a single cable, with RJ-45 connectors on each end.)

The cable in Figure 4-9 puts one end of a wire in pin 1 of one connector, and the other end into pin 1 of the other connector. Pin 2 on one end of the cable connects to pin 2 on the other side; and so on, for all eight wires. A cable with the wires connected in this manner is called a *straight-through cable*.

Okay, back to the problem illustrated in Figure 4-9. As you can see from the bubbles in the figure, both PC's NICs send on the twisted pair at pins 1 and 2. That electricity goes over the wires and enters the other NIC on pins 1 and 2. But, the NICs aren't receiving data on pins 1 and 2! That's because Ethernet NICs try to send on pins 1 and 2, and they receive data on the pair at pins 3 and 6. In this case, both PCs send, but neither receives data.

The solution is to use a cross-over cable. *Cross-over cables* connect the wire at pin 1 on one end of the cable with pin 3 on the other end; the wire at pin 2 with pin 6 on the other end; the wire at pin 3 with pin 1 at the other side; and the wire at pin 6 with pin 2 at the other side. The result: The PCs can receive the data sent by the other device! Figure 4-10 shows the basic idea.

Figure 4-10 PCs Using Different Pairs (Lanes) to Send Data

Now you know the basics of how you can allow two PCs to attach to a cable and transfer bits between each other. This chapter focuses on how to build the local network roadway—essentially, the networking components that allow bits to be transferred. Next, you'll read about how to connect several devices on an Ethernet LAN, using a device called a hub.

Sharing the Local Roadway: Ethernet Hubs


An old friend of mine is a native of San Francisco and still lives there. The traffic is horrible there, in part due to the need for lots of people to commute across the San Francisco Bay. There simply aren't enough bridges for the number of cars that need to cross the Bay. His solution: Fill in the Bay with dirt and pave the whole thing. Then there would be plenty of roads to allow people to cross the Bay.

If you tried to connect 10 PCs in a network, using a cross-over cable between NICs (as shown in figure 4-10), you would begin to do the equivalent of paving the San Francisco Bay. You know how two computers can use Ethernet NICs with a cross-over cable to communicate. However, to connect to other PCs, you would need more Ethernet NICs, and your PC probably does not have enough room for

all the cards. Also, you would need to run cables between your PC and all the other PCs, or at least get the electrician to run the cables. If you tried to do this for 100 PCs on the same floor of the building, and every PC wanted to connect to every other, you would have 99 cables connected to 99 NICs inside each PC!


The alternative to running a cable to every other PC is to run a cable from each PC to a wiring closet and connect the cables to a networking device, called an Ethernet hub. An *Ethernet hub* provides several functions, but mainly it allows the electrician to cable each device to the hub using only a single NIC and single cable, eliminating the cabling problem. The hub simply listens for incoming electrical signals, and when received, the hub repeats the same electrical signal to every other device that's connected to the hub. Figure 4-11 shows the basic operation.

Figure 4-11 Ethernet Hub Repeats Everything It Hears



As seen in the figure, Larry sends data to the hub, and the hub repeats what Larry sent out the cables to Archie and Bob. It's that simple!

Figure 4-11 shows cabling, but it does not show which pairs of wires are used. Interestingly, the hub expects straight-through Ethernet cabling between itself and the PCs. Why is that? Well, the people who make hubs, knowing that Ethernet NICs in PCs send on the twisted pair that uses pins 1 and 2 and receive on twisted pair that uses pins 3 and 6, do the opposite. Therefore, a straight-through cable works between a PC NIC and a hub. Figure 4-12 illustrates the concept.

Figure 4-12 Hubs Use Straight-Through Cabling to PCs

The hub's logic is simple:

1. Receive traffic on pins 1 and 2 on each physical interface.
2. When received, repeat the same electrical signal out all other ports, except the one in which the data was received.
3. When repeating out other ports, repeat the traffic out pins 3 and 6 so that the PCs will be listening.


Dirt Roads Versus the DOT

You can spend a long time working around or in the networking arena and not ever need to worry about how structured cabling works. This section gives you a brief glimpse into the world of structured cabling, just so you know what people are talking about if they bring it up.

You can go to the store and buy premade Ethernet UTP cables with RJ-45 connectors on them. You can even buy a hub—and for not a lot of money, typically less than \$20. You could then connect a bunch of computers to the hub using the straight-through Ethernet cables, and voilà, you have a network, or at least the part that allows the computers to send and receive data. Running your own cabling is quick and easy.

In a real network in a real building, you will not typically get away with stringing cables on top of the carpet, over people's cubicles, and so on. Instead, you should allow the electricians to do their job right, which means that the cables will run either under the floor or inside the ceiling. Also, instead of a single cable from a PC to the hub, the equivalent will be created. First, you install a short cable from the PC to a wall plate. The electrician runs another cable from the back of the wall plate to the wiring closet into a patch panel. Finally, either the electrician or the network engineer connects a cable from the patch panel in the wiring closet to the hub. Together, these cables provide two twisted pair between the PC NIC and the hub. Figure 4-13 shows the major components.

Figure 4-13 Major Components of a Structured Wiring Plan



Structured wiring allows the electrician to take care of the difficult part of cabling, with a minimum of cost, effort, and clutter, while still making sure that the cabling works correctly. Even though multiple cables are used, the net result of the cabling simply needs to ensure that the correct twisted pairs end up at the right place in the connectors at the endpoints. And the only cables that the people in the offices can see are the short ones between the PC and the wall plate; the rest of the cables are hidden.

Imagine that if instead of structured wiring, the electrician simply ran a single cable from the computer, under the floor, and straight into the hub in the wiring closet. Later, the person in the cubicle decides that she wants her PC on the left side of her desk, and she might discover that the cable is too short. So what does she do? She calls the electrician and asks for another cable to be run.



Of course, the electrician probably isn't going to want to run a new cable from the cubicle to the wiring closet. If he had used structured wiring, he could have run a cable from the wall plate to the patch panel in the wiring closet. A **patch panel**, sometimes called a **wiring panel**, provides the electrician a place to connect the actual wires in the cable on one side of the panel. On the other side, receptacles, much like the ones in the wall plate in the cubicle, allow the electrician or network engineer to use a short cable to connect those pairs of wires to some other devices, such as the hub shown in Figure 4-13.

Because the patch panel is located in a well-known place, and because the wall plate in the cubicles does not move, the electrician can do the hard part (running the cables under the floor) once. If the PC in the cubicle needs to be moved farther away from the wall, the PC user can get a slightly longer patch cable. A **patch cable** is simply a short LAN cable. Likewise, inside the wiring closet, short patch cables can be used to connect from the patch panel to the networking device, such as the hub shown in Figure 4-13.

You can think of structured wiring as the equivalent of having the DOT build your roads versus just using dirt roads. It takes time, planning, effort, and more cost, but in the end, you have a much better road system. With structured wiring, you end up with much better wiring and far less clutter. Furthermore, you can make changes without a lot of effort.

The benefits of structured cabling can be summarized as follows:

- Helps minimize the need for running new cables because the distance from each wall plate to the wiring panel can be determined easily
- Allows PC installers to run a short patch cable from the wall plate to the PC, without requiring help from the electrician
- Allows network engineers to run a short patch cable from the wiring panel to the networking devices, such as a hub
- Helps keep the wiring closet more organized

Chapter Summary

Computers communicate by transmitting an electrical signal over wires, with differing electrical characteristics meaning binary 0 and binary 1. For effective communication, both transmit and receive paths are needed; with a single path, the electrical signals would overlap, confusing the meaning of both electrical signals. Also, a twisted pair of wires is used for transmission in each direction because a twisted pair reduces the impact of electromagnetic interference on the wires, which in turn reduces transmission errors.

LAN UTP cabling includes up to four twisted pairs of unshielded wires, typically terminated with RJ-45 connectors. Although four pairs are supported, only two pairs are needed for Ethernet LAN communications—one pair for each direction of transmission. Depending on which devices are being connected, you might need a cross-over cable or a straight-through cable, ensuring that the pair used for transmission by one device connects to the pair used for receiving data on the other device.

Ethernet hubs allow multiple devices to connect to it, with the result that all devices receive a copy of whatever each device sends to the hub. By doing so, each device needs a single NIC and a single cable connecting it to the LAN hub, while allowing all devices to communicate with all other devices.

Finally, a structured wiring plan helps prevent unnecessary cabling runs inside the difficult-to-reach parts of the physical location (under the floor or in the ceiling). It also allows other personnel to change cabling later, without involving the electrician.



Chapter Review Questions

You can find the answers to the following questions in Appendix A, “Answers to Chapter Review Questions.”

1. Give an example of how a computer might transmit binary 0s and 1s using variations in voltage on a wire.
2. Imagine that two computers are sending and receiving data over a wire. How often would the receiving device need to sample the electrical signal on the wire if the standard called for 20 bps transmission speed?
3. As explained in the chapter, do computers usually transmit data over the same wires they use to receive data? Why or why not?
4. As explained in the chapter, do computers usually transmit data over a single wire, or do they use two wires?
5. How many wires can fit inside the end of an RJ-45 connector?
6. How many wires are needed for a PC to successfully send and receive data to another PC using an Ethernet LAN cable with RJ-45 connectors?
7. Explain why a straight-through LAN cable does not work when connecting two PCs directly, just using the cable.
8. Explain why a cross-over LAN cable *does* work when connecting two PCs directly, just using the cable.
9. Explain why a straight-through LAN cable works when connecting a PC to a hub.

10. What is the main difference between UTP and STP cabling? Which is more popular today?
11. List some of the key benefits of a structured cabling plan.
12. Imagine a hub with ten physical ports, and cables connecting each port to ten different PCs. The hub receives an electrical signal from the PC on port 1. Where does the hub forward the electrical signal?
13. Define the term “pin” in relation to an RJ-45 connector.