CHAPTER 11

The Data Link Layer

Power consists in one's capacity to link his will with the purpose of others, to lead by reason and a gift of cooperation.

— Woodrow Wilson

The Data Link layer is Layer 2 of the OSI reference model. This layer allows for direct communication between nodes over the physical channel provided at the lower layer. The communication can be point-to-point (one-to-one communication between two nodes) or point-to-multipoint (one-to-many communication, from one node to many nodes), depending on the nature and configuration of the network.

LAN technology exists primarily at the Data Link and Physical layers of the architecture. The functions performed by a network bridge or switch occur mainly at the Data Link layer. Network switches are able to tremendously enhance the capabilities provided by the Data Link layer. This is true to the point where you have to be careful that the implementation of the features doesn't affect the operations of some protocols within the upper layers.

The generic operation performed at the Data Link layer is the movement of data between nodes within a network over a physical connection. Once the Data Link layer has ensured that a connection is set up, the layer divides data into frames and transmits them to other nodes within a network. The receiving node sends acknowledgments and ensures that the data is received by keeping track of bit patterns in the received frames.

In this chapter, we discuss the Data Link layer. We cover concerns that are experienced in a LAN, as well as some of the mechanisms that are in place to recover from problems. In addition to the operations of this layer, we will discuss the use of switches and bridges in a LAN.

RANDOM BONUS DEFINITION

Layer 2 switch — Synonymous with bridge.

11.1 Concerns of the LAN

Most typical network users do not care about all the protocols and mechanisms that are in use to get their data; they just care that they get it. Because you are not the typical network user, however, you should care how this data gets there. Networks of all sizes produce conditions that are less than optimal, so actions have to be taken to address these needs. If there were no way to control the flow of data, the networking world would be a mess. If you worked in an organization that only used several thousand 10/100 Mbps Ethernet hubs, you would find that the end users would be less than satisfied (especially if you consider the types of data that are flowing in a normal LAN).

What should you concern yourself with in relation to operations at the Data Link layer within a network? This question is the reason this section is in the book. There are a lot of considerations to be aware of if you want to have an understanding of good old Layer 2.

POP QUIZ

The Data Link layer is what layer of the OSI reference model?

In general, the Data Link layer must provide mechanisms for framing, addressing, and detecting errors in data that is being sent to and fro over the physical link. The framing mechanisms provide a way for the frames to be delimited. Node addressing identifies the source and destination for communication on the LAN. Error detection ensures that only good data is received at the destination and then delivered to the upper layers. In some cases, the Data Link layer discards any errors it discovers, or it may employ a recovery mechanism — it all depends on which action it was designed to do.

11.1.1 It Just Is

It's really hard to compare one LAN to another. There really isn't anything typical about any LAN. There are similarities in both design and functionality, but each LAN is unique. This LAN is your LAN, if you will. The concerns of the LAN are fairly typical. There are commonalities that exist in any LAN.[1] The purpose and the expected outcome are the same in any LAN. The purpose of the LAN is to provide the avenue needed for communication of data. The expected outcome is for the LAN to live up to its configured expectations.

No matter what you do, there are some things about a LAN that are a fact.[2] These have remained (and probably always will remain) a constant throughout the lifetime of LANs.

- A LAN consists of multiple nodes that are attached to a single shared medium.
- There are geographical distance limitations.
- Every LAN will have a ceiling on the number of nodes that it supports.
- A LAN cannot survive without error detection, correction, and recovery.
- A LAN needs to support broadcasting and multicasting.
- Like nodes are peers to one another.
- A LAN is administered locally and is not subject to the same rules that are maintained by networks outside of the LAN.

ACRONYM ALERT

DA — Destination address

These are only a few of the things that most LANs have in common. This is an important list for this chapter because these are what the Data Link layer is all about.

11.1.2 Highs and Lows

Another concern for any LAN is to ensure that the highs and lows are met. What do we mean by this? The LAN is there to provide the best possible methods to deliver data over a shared link. This means that the LAN should meet the following expectations:

- **Highs** — This is the portion of data communication that you want more of.
  - **High throughput**[3] — The data throughput is simply the rate of error-free delivery of messages within a network. This includes data

[1] Believe it or not, there are some network administrators who still do not understand that point.

[2] An important thing to remember is that there is technology coming out all the time that not only pushes the limits of the facts of the LAN, but also gives reason for upgrades.

[3] There are other terms that mean the same thing, but some of those have multiple definitions. For instance, when we were determining exactly which term to use, we had originally considered using the term ''data rate.'' Although this would have been perfectly appropriate, it may have been a bit confusing. Data rate is a term that is used to define signaling rate, bit rate, transfer rate, etc. Throughout, we determined, is more specific in this case.

that is transmitted over a physical or wireless channel, switched through a node, or passed through the portals on both sides of the link. The expectation of the LAN is that the data throughput stays at a level as close as possible to the maximum allowable throughput. This is determined by the configuration and design of the network.

  - **High total bandwidth**[4] — Bandwidth is the available capacity of the physical or wireless channel, and network nodes provide for the delivery of data messages in the LAN.

- **Lows** — Things that you know will happen, but you don't want to happen.
  - **Low delays** — No delays is optimal, but unlikely. There will be peaks and there will be lulls. You can take action to try to stagger network chores (for instance, you can transfer large amounts of data at night so that it does not affect the times when users are all at work). Delays will occur, but the goal is to have as few as possible.
  - **Low error rate** — The number of errors in the network needs to stay as low as possible. You can take actions to detect and recover from errors, but you want to be as proactive as possible to prevent them from occurring in the first place.

POP QUIZ

Multiple nodes attached to a single shared medium can define what?

AN UNRELATED MOMENT OF PAUSE: FUN TECHNICAL TRIVIA!

1. **The Macintosh computer was launched by Apple Computer in 1984, with an ad that played during the Super Bowl. (The Raiders beat the Redskins, 38 to 9.)**
2. **How many approximate lines of code did the following Microsoft OS original releases have?**
   - **Windows 3.1 had over 3 million lines of code.**

(continued)

[4] Consider bandwidth as the amount available, and throughput as the actual amount of successful data messages that are transmitted. The throughput normally does not match the bandwidth, as there is other chatter that consumes some of the capacity of the communication channel (Hellos, Acks, etc.). For instance, if the link is a 10 GB Ethernet link, the bandwidth is going to be 10,000 Mbps. The throughput would be the rate of successful messages sent over the link. Of course, this is based on the performance of the network and is variable.

AN UNRELATED MOMENT OF PAUSE: FUN TECHNICAL TRIVIA! *(continued)*

- **Windows 95 had over 15 million lines of code.**
- **Windows 98 had over 18 million lines of code.**
- **Windows 2000 had over 35 million lines of code.**

3. **The computer mouse was invented in 1963 by Dr. Douglas C. Engelbart.**
4. **The term "computer" was first used to describe a mechanical calculating device in 1897.**
5. **We all know that 8 bits is called a byte, but did you know that 4 bits is called a nibble?**
6. **The type of keyboard that we are all familiar with is known as the QWERTY keyboard. This name is derived from the first six letters on the top line.**
7. **Netscape was the most popular Internet browser until Microsoft released Internet Explorer 4.**

11.2 Accessing the Medium

We all know that the LAN is made up of nodes connected to one another over a shared medium. We also know that it is called ''shared medium'' because everyone shares it for transmitting data. It's important to cover a few of the rules that must be upheld in a LAN as far as actually connecting to the network. Sure, we have discussed some of this before, but now is a good time for a refresher!

POP QUIZ

Define *throughput*.

11.2.1 Rules of Accessing the Medium

The previous section talked about some of the facts of a typical[5] LAN. When dealing with the shared medium, there are some facts as well.

- Within a shared medium, only one node can successfully transmit data at any given time.

[5] Typical is used loosely.

- Bandwidth is allocated to support the nodes that are sharing the medium so that each node gets a fair amount of bandwidth, with little to none left over.[6]
- The shared medium should support as much throughput as it is intended to handle.
- The network administrator should ensure that delays are kept to a minimum for data that is transported over a shared medium. A reasonable amount of waste, overhead, and delay should be taken into account when setting up and maintaining the network, and network monitoring will help you ensure that you meet the goals that you set.

RANDOM BONUS DEFINITION

Gigabit Ethernet — 1000 Mbps Ethernet.

11.2.2 From Tokens to Contention

So, how exactly do you go about ensuring the bandwidth is distributed fairly to the nodes using the shared medium? There are two methods that can be used in a LAN: tokens and contention. When using the token method, a token is passed from node to node. The nodes then pass data among one another in a round-robin fashion. When the contention method is used, the nodes transmit data when they want to. Therefore, it is entirely possible that two stations send data at the same time, causing a collision) to occur (see Figure 11-1).

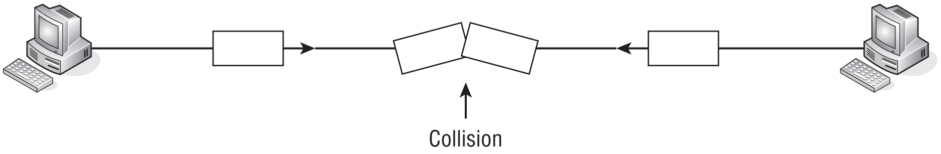

Figure 11-1 A collision

A collision causes datagrams to be dropped, but it doesn't necessarily mean that the data can't be recovered in some way. There are mechanisms that can

POP QUIZ

Name two methods of ensuring bandwidth is distributed fairly to the nodes that share connectivity within a LAN.

[6] When allocating bandwidth, it is important to use as much as possible. Some will be used by other processes, so a small amount of waste is possible.

be configured to recover from data loss and even prevent conditions that may cause it. Even so, there is a potential for data loss, so you can consider the token method a guarantee, whereas the contention method is more of a probability.

11.2.2.1 Using the Token Method

In Chapter 1 you learned that the token-passing topology consists of a single frame, known as a *token*, that is passed from one station to the next. When a node wants to pass data, it must wait until it receives an empty token. The node can then add its data to the token and pass it along the way. IEEE 802.5 is the official standard for Token Ring, which is the most common LAN token method in use.

A Token Ring topology can be set up physically in either a token ring (Figure 11-2) or a token bus (Figure 11-3) configuration. There is no logical difference between the two methods, as both operate in a token-passing manner.

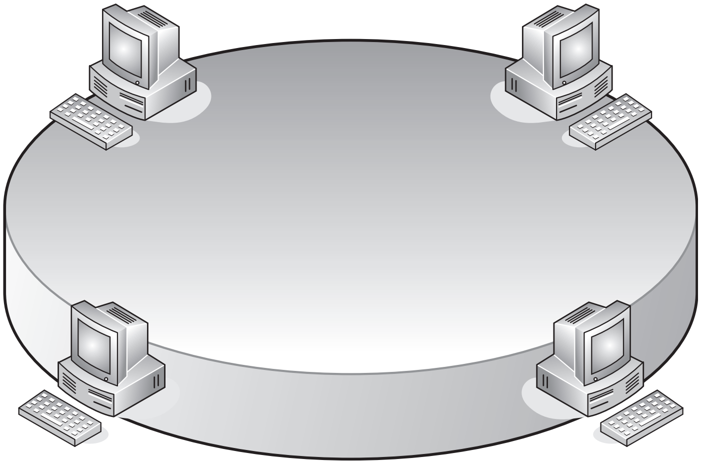

Figure 11-2 A token ring

In a token bus configuration, there is a central node called a media access unit (MAU) or a multistation access unit (MSAU). This device is similar to an Ethernet hub, but it has a computer chip that provides the logical ring that the end nodes are concerned with. The benefit of the token bus is that when a node goes down, the ring can be adjusted so that the other nodes will continue to operate on the network. In a physical ring, if a node goes down, the communication for all nodes goes down as well.

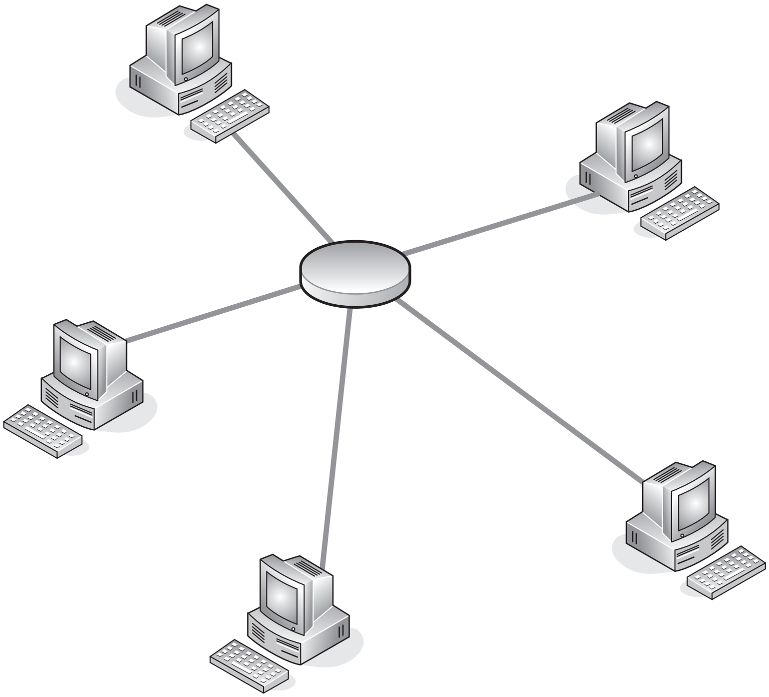

Figure 11-3 A token bus

11.2.2.2 Using the Contention Method

Nodes that use the contention method transmit their data at any time. The first node to get data on the line gets served first. When two nodes transmit at the same time, a collision occurs and the data will be resent. If the network is experiencing a high rate of data at any particular time, there will most likely be a lot of collisions, which will continue until bandwidth availability is restored.

Fortunately, some enterprising individuals came up with a way to sense when there is data being transmitted, thus reducing the number of collisions that can occur. Following are the protocols that are used to ensure that data flow in a contention method environment passes as smoothly as possible:

- **Carrier Sense Multiple Access (CSMA)** — Allows multiple nodes to be attached to a shared network. Prior to transmission, the nodes listen to see if the shared channel is busy and will transmit when they sense that the channel is not busy. ''Carrier sense'' simply means that a node is listening to see if it can detect an unused channel. If the node senses that there is a busy channel, it will defer transmission of its data until the channel is idle. ''Multiple access'' defines the fact that there are multiple nodes accessing the shared medium to transmit data.

- **Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)** — This is an enhanced version of the CSMA protocol in that it adds collision avoidance as a function. In this type of network, collisions are avoided because the station will not transmit data when it senses the channel is busy. The node will listen to the channel for a defined amount of time, and when the node is ready to send data, it will send a jam signal,[7] which lets all the other nodes know that the node is ready to transmit data.
- **Carrier Sense Multiple Access with Collision Detection (CSMA/CD)** — This is an enhanced version of the CSMA protocol in that it adds collision detection as a function. This function allows the transmitting node to monitor the channel for other transmissions. If while transmitting a frame, the node detects a signal coming from another node, it will terminate the transmission, send out a jam signal, and then try to send the frame again.[8] There are different ways for collisions to be detected, depending on the shared medium that is being used. The most popular and most often used CSMA/CD protocol is Ethernet.

11.3 Meet the Sublayers

In order to handle service requests from the network, the Data Link layer is broken into two sublayers: the Logical Link Control (LLC) sublayer and the Media Access Control (MAC) sublayer (see Figure 11-4).

ACRONYM ALERT

DEC — Digital Equipment Corporation

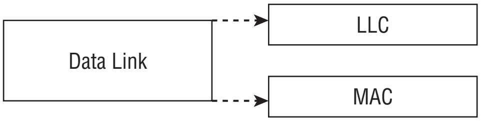

Figure 11-4 The Data Link layer's sublayers

LLC is the upper sublayer and is responsible for flow control, error control, and multiplexing and demultiplexing data transmitted over the MAC sublayer. The LLC sublayer is the sublayer that serves the higher layer client. LLC does

[7] A jam signal in CSMA/CD is a message to all other nodes that a collision has occurred and that they should stop transmitting.

[8] A random time interval is set that will determine when a station will try to transmit a frame again.

not have to worry about the design and functions of the LAN, which allows it to buffer these functions so that the higher layer protocols need not worry about the details but can focus on the tasks at hand.

The MAC sublayer is responsible for framing formats and determining which frame is going to be the next to access the shared medium.

11.3.1 Logical Link Control

LLC is a protocol developed by the IEEE 802.2 working group and provides three different types of service:

- **LLC Type 1 (LLC-1)** — Used for connectionless services.
- **LLC Type 2 (LLC-2)** — Used for connection-oriented services.
- **LLC Type 3 (LLC-3)** — Used for acknowledgments in conjunction with connectionless services.

LLC-1 is used for connectionless services. It is a best-effort delivery, providing none of the bells and whistles (for instance, flow control). LLC-1 provides multiplexing services to the upper layers. LLC-2[9] is used for connection-oriented services. Because it serves the connection-oriented operations, it does support the bells and whistles (for example, flow control, error control and recovery, call setup, call management, and call termination). LLC-3, which is seldom used, acknowledges frame delivery in a connectionless environment.[10]

POP QUIZ

The most popular and most often used CSMA/CD protocol is ______________.

11.3.1.1 LLC Framing

LAN source and destination addresses are determined by the MAC sublayer and will be in the MAC header portion of the frame. The LLC PDU[11] contains the following fields:

- **Destination Service Access Point (DSAP)** — This is used to identify the LLC that is supposed to receive the PDU.

[9] Note that nodes that support LLC-2 must also support LLC-1. This is because LLC-2 connections are established from an LLC-1 connectionless session.

[10] LLC-3 is used over LLC-1. This provides you with a bit of reliability without having the overhead of LLC-2. In most LANs you will usually only see LLC-1 and LLC-2. This is because many upper-layer protocols provide for recovery and don't need more than best-effort delivery.

[11] In case you forgot, PDU stands for protocol data unit. The PDU is the entity that all information is transferred in within a network.

- **Source Service Access Point (SSAP)** — This is used to identify the LLC that is supposed to send the PDU.
- **Control** — The Control field provides sequencing data, command information, and responses to requests. Note that any or all of these can be used in any combination.

Figure 11-5 shows the format of the LLC header. The LLC header is either 3 bytes or 4 bytes in length, depending on the version of LLC service type that you are using. The DSAP and SSAP are always 8 bytes in length each, which leaves the control field as the variable length field in the PDU.

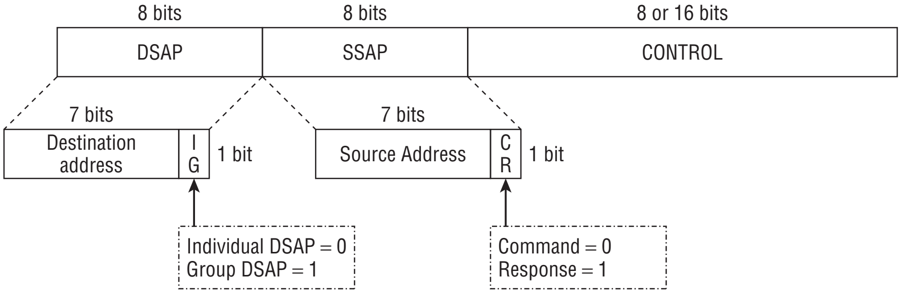

Figure 11-5 An LLC PDU (LLC header)

The DSAP field and the SSAP field are pretty straightforward. The DSAP field is an 8-bit (or 1-byte) field that contains 7 bits for the destination address portion of the field; the additional bit identifies whether it is destined for an individual or group DSAP. The SSAP field is an 8-bit (or 1-byte) field that contains 7 bits for the source address portion of the field; the additional bit identifies whether it is a request or a response to a request.

The Control field is a variable length, depending on what type of LLC you are using, and the type of the frame. The three frame formats you will see are:

- **Informational frame (I-frame)** — Used with LLC-2 only. This type uses a 2-byte (16-bit) field. Its purpose is to send numbered information transfer in LLC-2. Figure 11-6 shows an example of the format of the I-frame format.

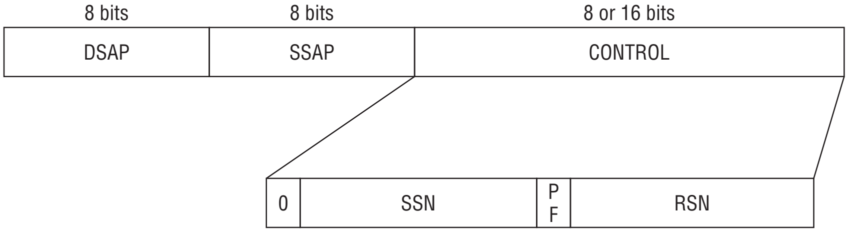

Figure 11-6 The format of the I-frame

- SSN — Sender sequence number
- RSN — Receiver sequence number
- PF — Poll on command frames or Final on response frames

- **Supervisory frame (S-frame)** — Used with LLC-2 only. This type uses a 2-byte (16-bit) field. It is responsible for handling acknowledgments, retransmitting requests, and terminating requests of the I-frames in LLC-2. Figure 11-7 is an example of the format of the I-frame format.
    - S — Supervisory function bits
    - PF — Poll on command frames or Final on response frames
    - RSN — Receiver sequence number

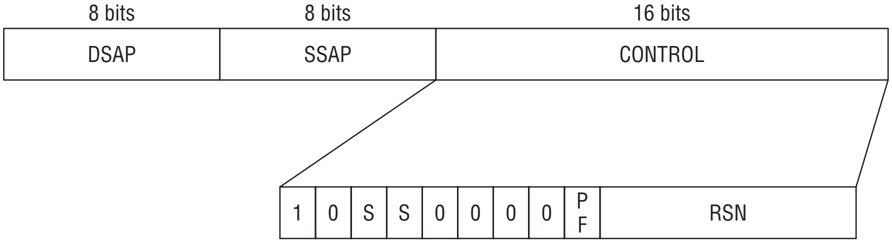

Figure 11-7 The format of the S-frame

- **Unnumbered frame (U-frame)** — Can be used with all LLC types. This type uses a 1-byte (8-bit) field. It is responsible for unsequenced data transfer and may handle some control functions as well (see Figure 11-8).
    - M — Modifier bits
    - PF — Poll on command frames or Final on response frames

RANDOM BONUS DEFINITION

frame — The Data Link layer encapsulation of transmitted or received information.

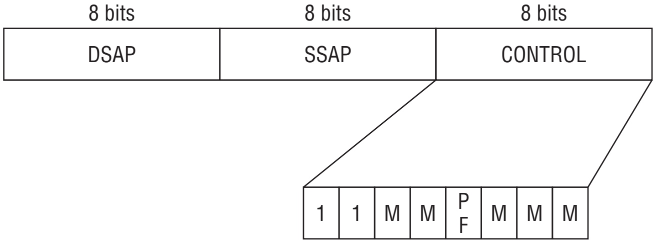

Figure 11-8 The format of the U-frame

11.3.1.2 Subnetwork Access Protocol

The Subnetwork[12] Access Protocol (SNAP) is used in conjunction with LLC-1 for the purpose of upward multiplexing to more upper-layer protocols than what is available with the standard LLC 8-bit SAP fields. When SNAP is not in use, the LLC DSAP's 8-bit field provides support of multiplexing to a maximum of 256 clients. Because the DSAP field reserves half its space for group SAPs, you actually can only multiplex to 128 clients.

POP QUIZ

What does DSAP stand for?

As far as the PDU goes, the SNAP header is placed directly behind the LLC header in the PDU. If SNAP encapsulation is being used, the DSAP and SSAP fields will be set to 0xAA, which indicates that SNAP is being used and that there is a SNAP header in the PDU. See Figure 11-9 for an example of SNAP encapsulation.

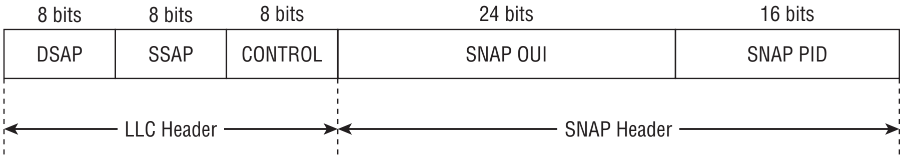

Figure 11-9 SNAP encapsulation

The fields in the SNAP header are as follows:

- **SNAP OUI** — This is a 24-bit field that contains the organizationally unique identifier (OUI). The OUI identifies the organization that the PID is assigned to.
- **SNAP PID** — This is a 16-bit field that contains the protocol identifier (PID), which identifies the upper-layer protocol that the PDU is destined for.

ACRONYM ALERT

EIA — Electronic Industries Association

Here's the clincher: SNAP encapsulation allows you[13] to have up to 65,536 upper-layer protocol identifiers.[14]

[12] It's important to note that the term ''subnetwork,'' in the SNAP sense, does not have anything to do with a subnetwork in a TCP/IP sense. This is one of those acronyms that may have actually come before the term. It's nothing more than a way to make SNAP have that fancy ring that we ''catenet'' lovers like.

[13] In saying you, we are referencing the applicable organization.

[14] This simply blows the 256 (if you are lucky) identifiers out of the water.

A LITTLE MORE ABOUT THE OUI

The OUI is a 24-bit number that can be purchased from the IEEE. The number is unique to an organization (vendor, company, etc.) and serves several purposes. Many protocols reference the OUI (SNAP, for instance). Some even append a few bits to increase the functionality of the OUI. There are a lot of other terms that are used for the OUI. It is also known as a MAC address (more on this in the following section), vendor ID, NIC address, and many more.

Here is a list of a few OUIs that are assigned today. Note that these are globally assigned, which is why they are unique for that particular organization. Also note that often the same company can be assigned multiple OUIs, regardless of the location they are registered to (as in the case of Nortel Networks).

00-00-C0 (hex)	**Western Digital Corporation**
0000C0 (base 16)	**Western Digital Corporation** **8105 Irvine Center Drive** **Irvine, CA 92718** **United States**
00-0C-41 (hex)	**Cisco-Linksys**
000C41 (base 16)	**Cisco-Linksys** **121 Theory Drive** **Irvine, CA 92612** **United States**
00-0D-54 (hex)	**3Com Ltd.**
000D54 (base 16)	**3Com Ltd.** **Peoplebuilding 2** **Peoplebuilding Estate** **Maylands Avenue** **Hemel Hempstead** **Hertfordshire HP2 4NW** **United Kingdom**
00-0D-56 (hex)	**Dell PCBA Test**
000D56 (base 16)	**Dell PCBA Test** **One Dell Way** **RR5 MS-8545**

(continued)

A LITTLE MORE ABOUT THE OUI *(continued)*

	Round Rock, TX 78682 **United States**
00-0E-40 (hex)	**Nortel Networks**
000E40 (base 16)	**Nortel Networks** **8200 Dixie Road** **Suite 100** **Brampton, Ontario L6T 5P6** **Canada**
00-1F-9A (hex)	**Nortel Networks**
001F9A (base 16)	**Nortel Networks** **2221 Lakeside Boulevard** **Richardson, TX 75082-4399** **United States**
00-23-0D (hex)	**Nortel Networks**
00230D (base 16)	**Nortel Networks** **2221 Lakeside Boulevard** **Richardson, TX 75082-4399** **United States**

This list is an example and is only a short list compared to all the OUIs that are registered. If you want to see a complete list, you can view it on the IEEE website (`http://standards.ieee.org`).

11.3.2 The MAC Sublayer

The MAC sublayer is responsible for interfacing between the LLC sublayer and Layer 1, the Physical layer. The MAC sublayer provides access control as well as addressing for the PDU. This sublayer is what makes multipoint communication with a LAN/WAN a reality. This sublayer is also able to operate as a full-duplex logical channel in a LAN. This logical channel supports unicast (point-to-point) services, multicast services (point-to-multipoint), and broadcast (point-to-multipoint) services. All these services are discussed in Section 11.4.

The MAC sublayer uses a MAC address,[15] the address assigned to the node's network adaptor (commonly, a NIC). For channel access, the MAC sublayer employs some control functions that allow multiple nodes to use the same physical medium. We discuss both the MAC address and the channel access control functions next.

POP QUIZ

What does SNAP stand for?

11.3.2.1 The MAC Address

The IEEE 802 MAC address is a 48-bit address that is used to identify the network adaptor for a particular node or interface in the network. The MAC address was originally designed as a permanent address that is unique to the adaptor it is assigned to. Most hardware today allows a MAC address manipulation method known as *MAC spoofing*. That tidbit of trivia is informational, but we won't be going into the details, as it is beyond the scope of this book.

The format of the IEEE 802 MAC address is set up to make it as easy as possible to understand. It consists of six groups of two hexadecimal digits. The groups are separated by either a colon (:) or a hyphen (-). Following is an example of each method:

- 01:00:23:00:bf:00
- 01-00-23-00-bf-00

In these examples, the OUI would be 01:00:23, with the remainder being the NIC-specific identifier. Combined, they make up the MAC address.

MAC addresses can be administered both universally and locally. When the address is administered universally, the MAC address is assigned to the interface by the device's manufacturer. Locally administered addresses are manipulated by a network administrator for purposes that serve the needs of the LAN.

RANDOM BONUS DEFINITION

error control — A procedure used to recover from detected errors.

11.3.2.2 Access Control for the Channel

The MAC sublayer is responsible for ensuring that multiple nodes are able to connect to and share the same physical medium. The groups of protocols that operate and perform this function are known as *multiple access protocols (MAP)*.

[15] The MAC address is also referred to as a *physical address*.

These protocols detect and avoid collisions in contention environments and ensure that there are enough resources to set up a logical connection when needed. Remember, earlier in this chapter we said that the most popular contention method used is Ethernet.

11.4 The "ings" — Casting, Detecting, and Addressing

LAN traffic flow is a fairly simple process. There are a lot of standards and configuration options in a LAN that provide a lot of freedom to configure and maintain in an attempt to reach the maximum highs and minimum lows that we discussed in Section 11.1.2. None of this would really mean anything if we didn't have a way of getting the data from the upper layers and making sure it reaches the appropriate process on the other end of the link. Well, we do have a way to do all of this in one very helpful and handy layer — Layer 2!

It is important to keep in mind that some of the ''ing'' operations occur at other layers of the OSI reference model (upward/downward multiplexing, data link multicasting vs. IP multicasting, etc.). Unless otherwise stated, this section pertains only to the processes at the Data Link layer.

This section covers MAC addressing and end-to-end delivery of data to a single node as well as multiple nodes.

11.4.1 Data Link Addressing

> Main Entry: ad.dress[16]
> Function: noun, verb, -dresses or -drest, -dress.ing.
> 1. a direction as to the intended recipient, written on or attached
> to a piece of mail.
> 2. the place or the name of the place where a person, organization, or
> the like is located or may be reached: What is your address when you're
> in Des Moines?
> 3. to direct (data) to a specified location in an electronic computer.

Directing data is what addressing is all about. At the Data Link layer, this is done by pointing PDUs to the destination MAC address for delivery of a frame within a LAN. The MAC address is the number that is assigned by the manufacturer of a NIC or a network interface. In Figure 11-10, you

ACRONYM ALERT

IC — Integrated circuit

[16] *Dictionary.com Unabridged (v 1.1)*. Random House, Inc. April 18, 2008.

can see a group of individuals sharing a physical medium. If Bob needs to send anything to Larry, he simply enters the MAC address (01:bb:04:af:00:1f) that is assigned to the NIC card on Larry's PC in the frame and sends it toward Larry's PC.

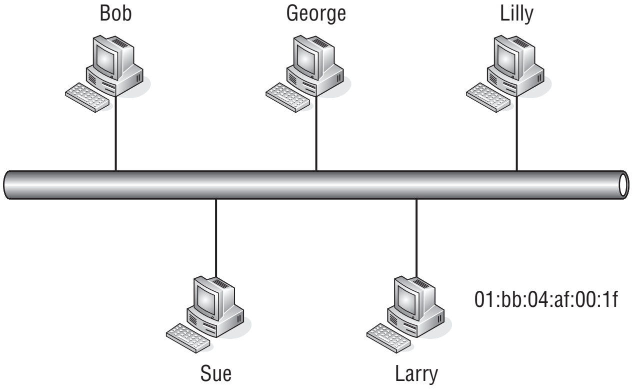

Figure 11-10 Data Link layer frame delivery

That sounds simple, doesn't it? But what we haven't really discussed is how Bob's PC learned the MAC address of Larry's PC. We also need to cover how Larry's PC knows how to get back in touch with Bob's.

POP QUIZ

What does SSAP stand for?

It wasn't until the early 1980s' PC boom that there was really a need to formulate addressing that could be learned in a dynamic fashion and could support several hundred nodes. Prior to the PC boom, there were not more than a few nodes in a network, and addressing was locally assigned and administered. In a network of only a few nodes, it was easy to maintain networks in this manner. Now, however, with hundreds of nodes communicating with hundreds of networks with hundreds of nodes, there is a real need to have a way to bridge traffic that is easy to administer.

For now, and way into the future, LANs will continue to evolve and expand geographically as well as in technical achievements that are, and will probably remain, a constant.

11.4.1.1 The MAC Address Format

The MAC header of a frame contains the destination and source MAC addresses for the interfaces involved in the communication stream. Figure 11-11 shows the 48 bits that make up the MAC address.

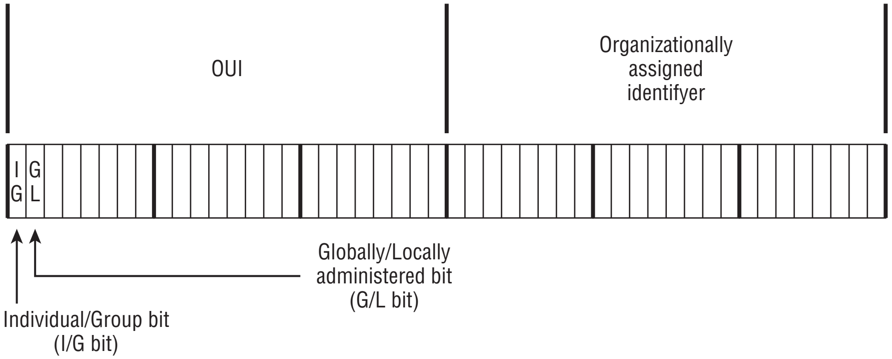
OUI
Organizationally assigned identifyer
I G
G L
Globally/Locally administered bit (G/L bit)
Individual/Group bit (I/G bit)

Figure 11-11 The MAC address format

Regardless of whether it is the source or destination address, the format is the same for all but the first bit. When referring to the destination address field, the first bit (the I/G bit) identifies whether the destination target is an individual (unicast) or a group (multicast). The source address field only uses the first bit when using Token Ring or FDDI. When used, it identifies if there is any source routing data in the frame.

The second bit in the source and destination address field indicates whether the address is globally or locally unique. This bit is called the *G/L bit* and it identifies whether the organizational assigned identifier[17] is globally unique (G/L bit set to 0) or locally unique (G/L bit set to 1). If it is a locally unique identifier, then the address is unique only to the LAN.

In any given LAN, there can be a mix of both globally and locally unique addresses. The nodes within that LAN do not have to worry about whether an identifier matches theirs that is in another LAN. This is because LAN-to-LAN communication is handled at the Network layer and the IP addressing scheme negates this concern. Nodes within a LAN cannot directly communicate at the Data Link layer with nodes in other LANs. Therefore, it is possible to have a duplicate locally assigned MAC, but they will not be aware of one another.

RANDOM BONUS DEFINITION

edge switch — A switch that is implemented at the boundary of a VLAN-unaware segment and a VLAN-aware segment of a LAN.

11.4.1.2 Unicast Addressing

A *unicast address* is simply the address of a particular node's interface within the LAN. The unicast address is the MAC address that is assigned to a device

[17] The last 24 bits of the MAC address.

or an interface within the LAN. Unicasting is the act of sending a frame from one source node to a single destination node. Figure 11-12 shows an example of unicasting.

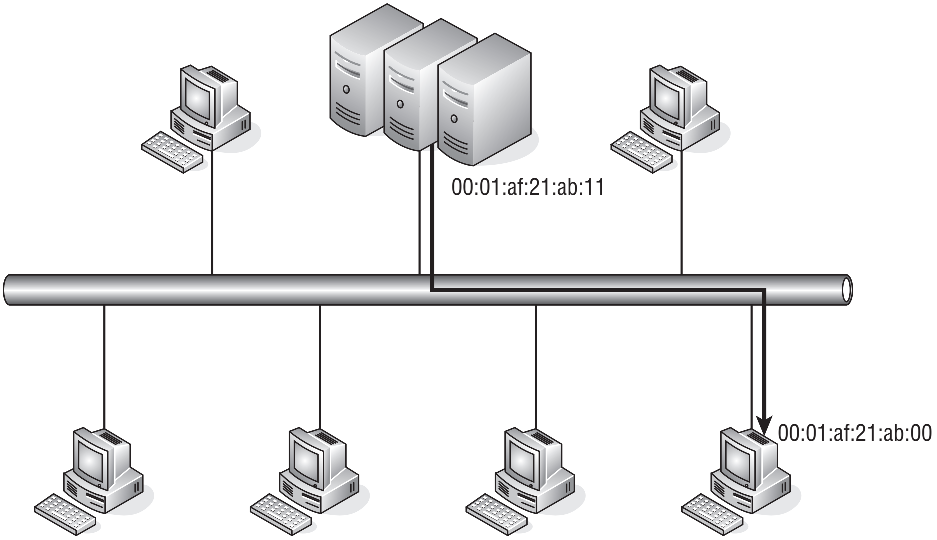

Figure 11-12 Unicasting

The figure shows the server farm sending data to a single node on the LAN. The unicast address of the source is 00:01:af:21:ab:11, which is the MAC address of the interface on the source side of the transmission. The destination unicast address is the MAC address of the interface used by the destination node — in this case, 00:01:af:21:ab:00. All transmitted frames during the session will use the same destination and source unicast addresses.

11.4.1.3 Multicast Addressing

Multicasting[18] is the act of sending a message to multiple nodes. Multicasting can be handled at the Layer 3 level (IP multicasting) or at Layer 2 (Ethernet multicasting). This section will focus on Layer 2 multicasting; the Layer 3 multicasting was discussed in Chapter 10.

[18] A type of multicast that you might come across in a LAN is the broadcast, which is destined for everyone in the network. Often called the ‘‘all F’s’’ MAC address, the broadcast address is always ff:ff:ff:ff:ff:ff. Table 11.2 shows how this address maps to various protocols over Ethernet.

Nodes that participate in a multicast group will be related in some logical fashion.[19] Multicast addresses are group addresses of nodes within a shared internetwork. Multicasting provides the ability for multiple nodes to receive data sent from a single transmission. Figure 11-13 shows an example of multicasting.

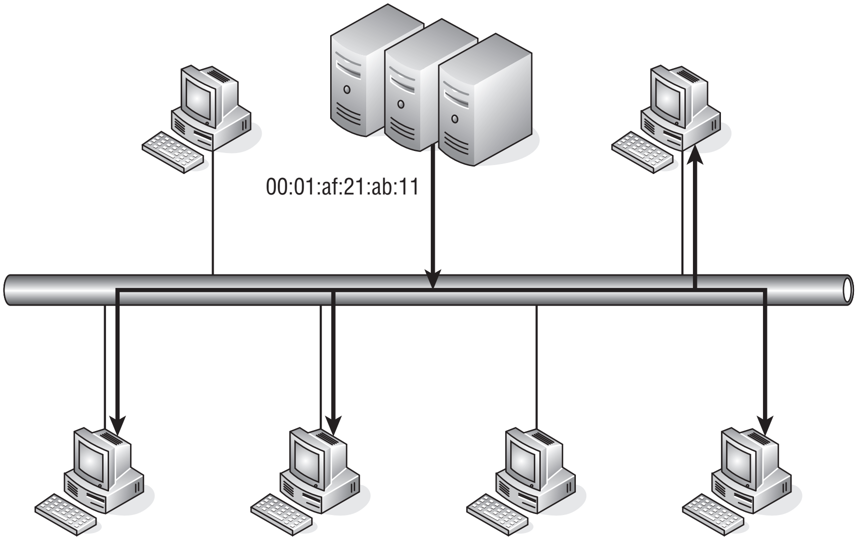

Figure 11-13 Multicasting

Notice that in the figure not all nodes are receiving the transmission that is being sourced from the server farm. This is because not all of the nodes are in the same multicast group. Also notice that the source address for the originator will be a unicast address.[20]

When a node decides to join a multicast group, it needs to determine if a received frame is a unicast or a multicast frame. NIC cards are configured to recognize when a frame is unicast and when it is not. How is this done? Remember the I/G bit that we discussed in Section 11.4.1.2? This is the bit that identifies if the frame is a unicast (I/G bit set to 0) or multicast (I/G bit set to 1).

ACRONYM ALERT

LACP — Link Aggregation Control Protocol

[19] The multicast will be sent only to those stations that share the function that requires them to receive the message. The stations that are not applicable won't be bothered.

[20] This is because there is only one source node involved.

Table 11-1 shows some well-known multicast MAC addresses that are used by Ethernet.

Table 11-1 Ethernet Multicast MAC Addresses

ADDRESS	TYPE	FUNCTION
01:80:C2:00:00:00	Length field	Spanning tree BPDU
09:00:07:FF:FF:FF	Length field	AppleTalk Multicast
09:00:07:00:00:FC	Length field	AppleTalk Zone Multicast
09:00:2B:00:00:03	8038	DEC LanBridge Hello packet
09:00:2B:00:00:0F	6004	DEC LAT
09:00:2B:00:00:00	8038	DEC LanBridge copy packet
09:00:2B:00:00:01	8038	DEC LanBridge Hello packet
09:00:4EL00:00:02	8137	Novell IPX
AB:00:04:04:00:00	6003	DECnet Phase IV router Hello packets
AB:00:00:03:00:00	6003	DECnet Phase IV end node Hello packets
CF:00:00:00:00:00	0900	Ethernet configuration test

Broadcasting is really nothing more than multicasting to everyone in the LAN. Table 11-2 shows some of the various types and functions performed in the broadcast message.

POP QUIZ

A ____________ address is simply the address of a particular node's interface within the LAN.

Table 11-2 Ethernet Broadcast MAC Addresses

ADDRESS	TYPE	FUNCTION
FF:FF:FF:FF:FF:FF	0600	XNS hello packets
FF:FF:FF:FF:FF:FF	0800	IP
FF:FF:FF:FF:FF:FF	0806	ARP
FF:FF:FF:FF:FF:FF	8035	Reverse ARP
FF:FF:FF:FF:FF:FF	809B	Ethertalk

11.4.2 Error Detection

Frames are either fixed-length PDUs (ATM uses a fixed-length PDU) or bit-oriented, which is more common and is what we discuss in this book. Regardless of the frame type, errors can occur in the LAN, and frames can disappear, duplicate, and even become corrupted on their way to a destination.

An error in a length-type frame can cause the frame to terminate and skew the beginning of a new frame. Likewise, a bit can be set incorrectly in a bit-oriented type frame, which can cause duplication and even deletion of the frame. Errors can be caused by numerous reasons, environmental as well as traffic-related. Electrical interference can cause noise on the physical medium, which can corrupt the bits in the frame. Other causes of transmission errors include:

- Signal distortion
- Synchronization issues
- Crosstalk

Errors will occur and there are acceptable error rates that are figured into any LAN design. Excessive errors are not good. Depending on the protocols in use, errors can cause transmission delays, and if not handled correctly, the problem can propagate itself, causing sluggishness and possible outages in the LAN. This is why you need a way to detect and possibly correct errors at the Data Link layer level (as well as some protocols within other layers).

ACRONYM ALERT

TB — Transparent bridge

There are two methods of error detection used at Layer 2, parity check and cyclic redundancy check (CRC):

- **Parity check** — The simplest of the error-checking methods. This method adds a bit to a string of bits to ensure that the total number of 1s in the string is equal to an even or an odd number. For example:
  - **Odd parity —** 01010101 + 1 parity bit = 010101011. Notice that the total number of 1s is an odd number. An

POP QUIZ

What is the Ethernet standard broadcast MAC address?

odd parity bit is always set to 1 if the total number of 1s in the string (before the parity bit is considered) is an even number. By adding 1 to the even number, it ensures that the number is odd, which matches the type of parity in use in this case.

Figure 11-14 shows an example of data transmission using odd parity.

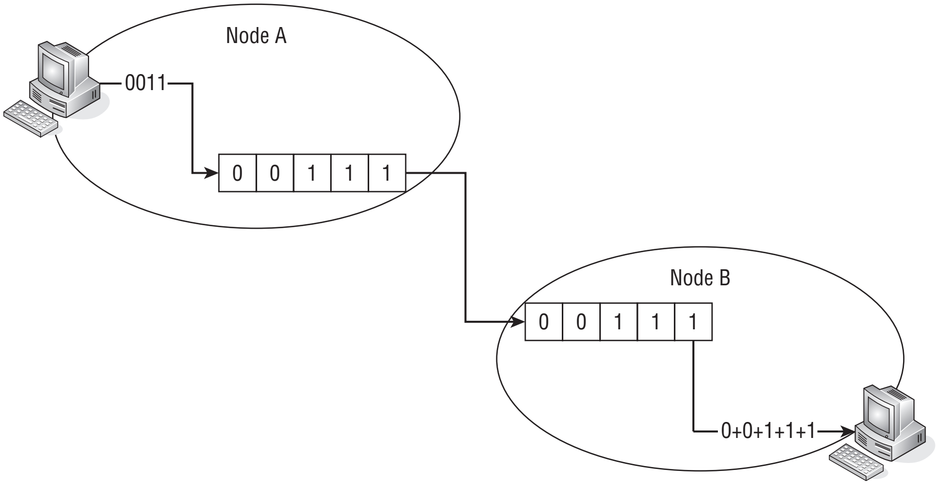

Figure 11-14 Odd parity

In Figure 11-14, node A wants to send the data stream 0011 to node B. Node A computes the value of the data stream (0+0+1+1)[21] and because odd parity checking is being used, node A turns the parity bit on to 1 before it transmits the data.

Node B then receives the data and computes the overall value (0+0+1+1+1), which is an odd value. Odd Parity is in use, so node B reports a good frame received.

- **Even parity —** 01010100 + 1 parity bit = 010101001. Notice that the total number of 1s is an even number. An even parity bit is always set to 1 if the total number of 1s in the string (before the parity bit is considered) is an odd number. By adding 1 to the odd number, it ensures that the number is even, which matches the type of parity in use in this case.

  Figure 11-15 shows an example of data transmission using even parity.

[21] When does 1 + 1 = 1? When dealing with binary, when you are either on (1) or off (1).

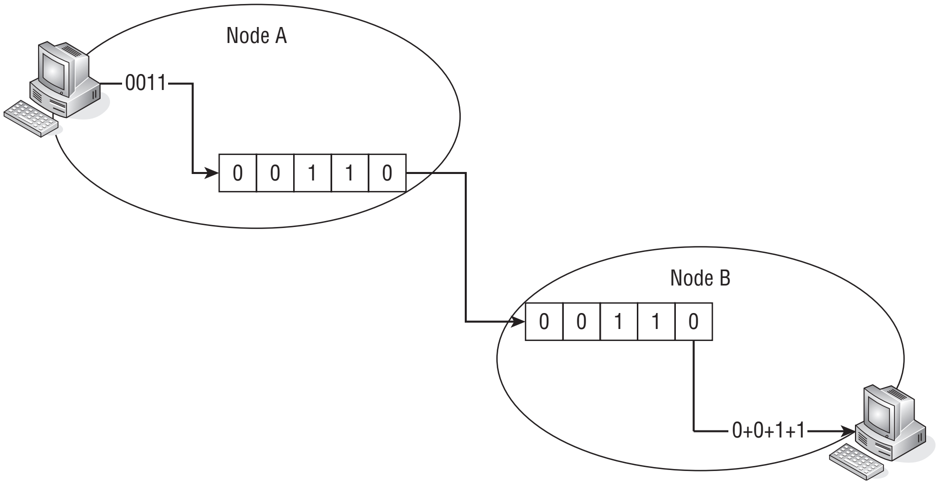

Figure 11-15 Even parity

In Figure 11-15, node A wants to send the data stream 0011 to node B. Node A computes the value of the data stream (0+0+1+1) and because even parity checking is being used, node A does not turn on the parity bit before it transmits the data.

Node B then receives the data and computes the overall value (0+0+1+1+0), which is an even value. Even parity is in use, so node B reports a good frame received.

Finally, let's take a look at the parity check when an error has occurred. Figure 11-16 shows an example of a data stream that is being sent using even parity. Notice that an error occurs before the stream reached the destination. When node B receives the data, it counts the number of 1s and notices that there is an odd number, therefore realizing that an error has occurred.

- **Cyclic redundancy check (CRC)** — Also known as the *frame check sequence (FCS)*. The CRC is a function used to detect common errors that may occur during data transmission. CRC is a much more complex method of error checking than the parity check method, but it isn't necessarily complicated.

  The way the CRC method works is that the node that is transmitting the frame adds a value, known as a *checksum*, to the message that is being transmitted, The receiver uses the CRC method to calculate the checksum on its end and compares it with the checksum that was added by the transmitting node to determine if there was any corruption along the way. Figure 11-17 shows an example of a simple checksum.

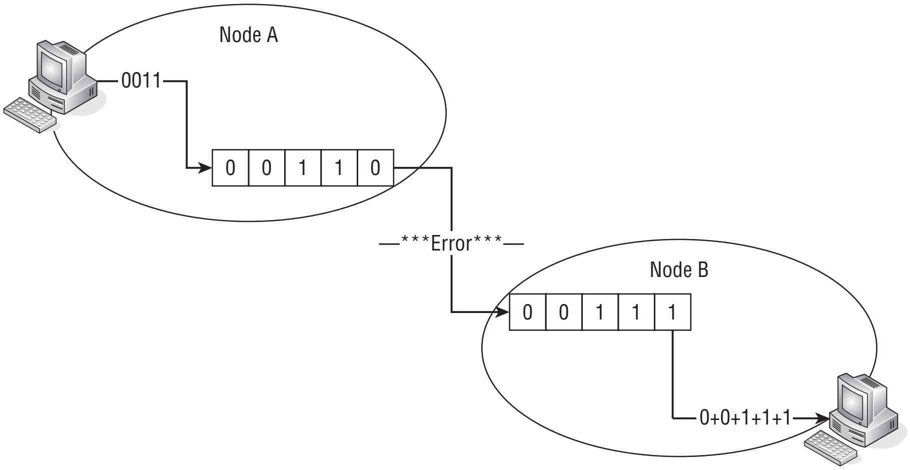

Figure 11-16 A parity error

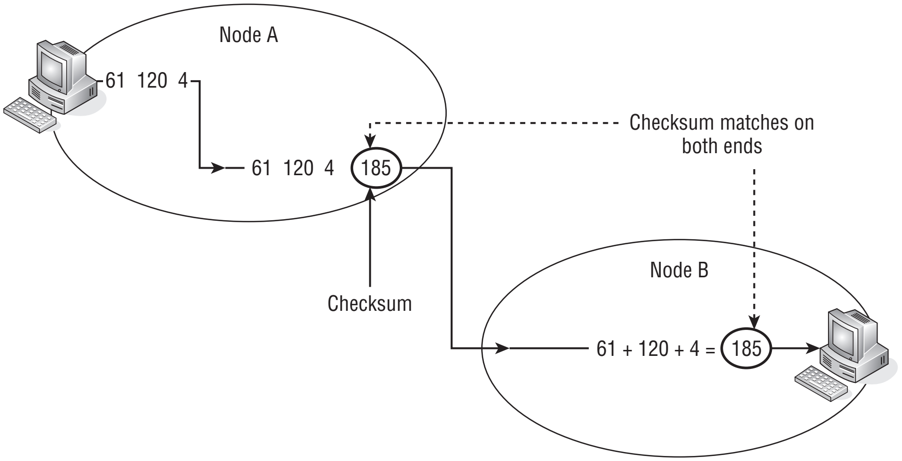

Figure 11-17 A simple checksum

For simplicity sake, this example uses decimal notation. Each decimal number represents a byte of data in a message. This means that there are 256 possibilities in each byte. The checksum algorithm in use simply adds

ACRONYM ALERT

POP — Point of presence or Post Office Protocol the value of all the bytes and uses the combined value as the checksum. In Figure 11-17, node A is sending the message `61----120----4`. Node A adds the total number of bytes and appends the checksum to the message (61 + 120 + 4 = the checksum of 185).

Node B receives the message and adds the value of the bytes in the message (61 + 120 + 4). By adding the total numerical value, node B determines that the checksum should be 185. Once node B determines the value that it thinks it should be, it compares its checksum with the checksum of node A. If there is a match, it knows the message was received intact.

Next, let's take a look at how the simple checksum works when an error occurs. Figure 11-18 shows such an example.

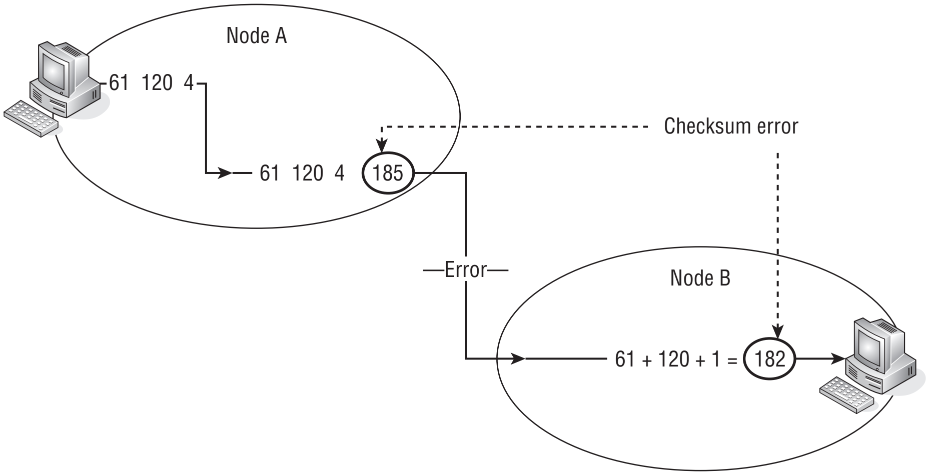

Figure 11-18 Checksum failure

> **RANDOM BONUS DEFINITION**
>
> access domain — The collection of nodes that share a network segment among which MAC arbitration can occur.

In this example, node A is sending the message `61----120----4`. Node A adds the total number of bytes and appends the checksum to the message (61 + 120 + 4 = the checksum of 185). Notice that somewhere between node A and node B, there is an error that causes the last digit of the message to change from a 4 to a 1.

Node B receives the message and adds the value of the bytes in the message (61 + 120 + 1). By adding the total numerical value, node B determines that the checksum should be 182. Once node B determines the value that it thinks it should be, it compares its checksum with the checksum of node A. In this example, node B recognizes that the message was corrupted (185 does not equal 182), so an error occurs.

The simple checksum used in the example above would not be that reliable. There are too many possibilities of errors occurring with the checksum still intact at the opposite end. For instance:

61 + 120 + 4 = 185

51 + 130 + 4 = 185

60 + 120 + 5 = 185

CRC computes the checksum by using an algorithm that is basically long division for binary. Additionally, the CRC uses the first 16 bits in the calculation, creating 65,536 possibilities. (The chances of an erroneous calculation is far less than with the simple example above.) Taking it a step further, the remainder (not the quotient) is what is used as the checksum.

POP QUIZ

What is the simplest of all error-checking methods?

Let's assume that an originator wants to send the first 2 bytes of data used in the example above (61 and 120). Also assume that the CRC divider will be a constant 1-byte divider whose value would equal the decimal number 7. So, now we want to convert these numbers to binary:

61 = 00111101

120 = 01111000

7 = 00000111

Take a look at Figure 11-19, which shows that node A is sending a message to node B.

The message is binary 00111101 01111000. The CRC constant divisor is set to binary 00000111. You want to take the message and divide it by the divider. The remainder is your checksum value.[22]

[22] Calculated out (we cheated and used a scientific calculator), the remainder in this case would be binary 00000110 and that is what will be used as the checksum value.

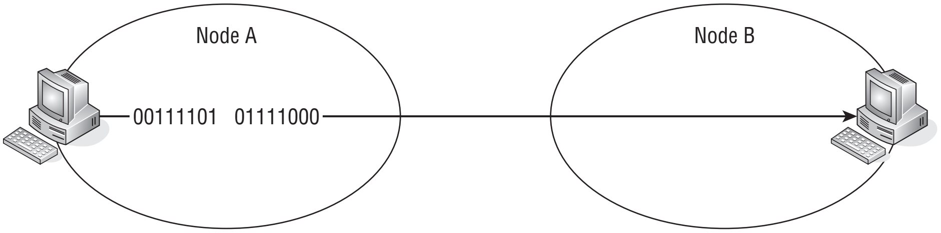

Figure 11-19 The CRC function

RANDOM BONUS DEFINITION

disabled state — A state used in the Spanning Tree Protocol that identifies a bridge port that has been set to not receive or transmit any frames.

CRC algorithms more commonly use a method that closely resembles polynomial arithmetic. Instead of a simple divisor, message value, quotient, and remainder, as you saw in the last example, these integers are actually seen as polynomials with a binary coefficient. There are many ways to take this even further but are beyond the scope of this book. Really, we could fill pages with the different algorithms that can be used. Therefore, we bring this section to a close. What you really need to understand is that the CRC uses a checksum that can be complex and is used to validate data integrity in a LAN.

Now, all this may sound complicated, but it is fairly simple. To show how simple it is, we want you to take a moment away from the book and relax. Just clear your thoughts and take a break from all of this technical mumbo jumbo and relax. And what better way of relaxing than eating a pizza?

AN UNRELATED MOMENT OF PAUSE — DENISE'S PESTO CHICKEN PIZZA

If you are a fan of interesting foods, you are bound to love this pizza recipe. It is a super-easy meal to make and well worth the time it takes to make it. Just be careful if you have an allergy to nuts, as pesto contains pine nuts and you may have a reaction.

Note: Delivery is also an option, but they won't have this recipe.

Ingredients:

- **Refrigerated pizza dough**
- **Pesto sauce**
- **Cooked chicken, cubed or shredded**

(continued)

AN UNRELATED MOMENT OF PAUSE – DENISE'S PESTO CHICKEN PIZZA *(continued)*

- **Fresh mozzarella (that doesn't mean shredded!)**
- **Sun-dried tomatoes in oil (rinsed)**
- **Rosemary**
- **Olive oil**

Directions:

1. **Preheat oven to 425°.**
2. **Roll out the pizza dough, brush on olive oil, and sprinkle with rosemary.**
3. **Bake for about 5 minutes.**
4. **Remove the pizza and spread on the pesto, chicken, and sun-dried tomatoes. Top with mozzarella.**
5. **Bake 10 to 15 minutes, or until the cheese is melted.**

11.4.3 Control of the Flow

Flow control is used to prevent the sender of data from sending more data than the receiver can handle. Without flow control, the sender would not be aware that the receiver can't accept any more data and would continue to send the data, only to have to send it again once they are aware there is a problem.

There are different methods of flow control that can be used. Sometimes it is medium dependant, but there are options that work with higher-layer protocols. The receiving node does not necessarily have to provide feedback when it can or cannot accept more data. Ethernet uses what are known as *PAUSE frames* for flow control.

A PAUSE frame is a message sent by a receiver to the sender, letting the sender know that the receiving node can no longer receive data and that the transmission needs to be paused for a specified period of time. The PAUSE function only works within full-duplex environments.[23]

The PAUSE function has a reserved multicast MAC address of 01-80-C2-00-00-01. This is a MAC address that was set up by the IEEE and is used for the MAC PAUSE frame function.

[23] Because this is the most common standard in today's networks, we decided to focus on the PAUSE function in our discussion of flow control. Understand that, from a data link perspective, flow control is a function that prevents a sender from overloading a receiver.

11.5 “Knode” the LAN

We assume that you are all thinking, ‘‘What in the heck is *knode*?’’ A knode is fictional, simply a term that we created as a combination of ‘‘know’’ and ‘‘*node*.’’ This may be a bit silly, but it is also good food for thought. Knowing your LAN is every bit as important as having the nodes you need to do what you want in the LAN.

In Chapter 3, ‘‘Network Hardware and Transmission Media,’’ we introduced bridges and switches, two types of hardware that operate at the Data Link layer. We are going to finish this chapter by talking about bridge/switch deployment within the LANs. Although only an overview, this section should be a great lead-in to Part III of the book, which deals with network design and implementation. If you are interested in getting a good reference book, Jim’s last book,[24] *The All-New Switch Book: The Complete Guide to LAN Switching Technology*, is a comprehensive reference to everything network switching.

So what is different between a bridge and a Layer 2 switch? The answer to this question may surprise you. There is no functional difference between a bridge and a switch. That’s right! None! Nada! Zero! Zip! Switch is nothing more than a marketing term that came out in the 1990s. The change was brought about due to the ever-growing LANs. Original bridges could not offer wire speed transmission rates on more than two ports within the bridge. Bridges that could handle the higher rates were still not that reliable and carried a very high price tag.

This all changed when the application-specific integrated circuit (ASIC)[25] was developed. Along with improvements in system memory and higher processor speeds, the ASIC allowed the bridge to be developed, supporting a lot of ports that were capable of concurrent wire speed transmissions. The best part was that the cost was less than traditional bridges with the same number of ports per area, but not transmission speeds.

RANDOM BONUS DEFINITION

chassis switch — A switch that is designed in a modular fashion. This type of switch consists of a chassis and multiple plug-in modules.

[24] This is more than just a shameless plug – it’s a really good book.

[25] Pronounced ‘‘a sick.’’

These new devices were introduced to the world and the salespeople of the world decided to call them ''switches,'' and that was where the switch was born.[26] For simplicity sake, we will use the term ''bridge'' for the remainder of this section. Feel free to substitute the word ''switch'' if you are so inclined.

Two methods[27] of address-to-port mapping are used:

- **Source route bridging** — This type of bridging is used in a source-routed internetwork. The path to a destination is determined by the end nodes, not by the bridge itself. An example of an environment that uses source route bridging is Token Ring.
- **Transparent bridging** — This is the type of bridging that is used in Ethernet (and others).[28] In a transparent bridging environment, the bridge makes the path determinations and the end nodes are not aware of decisions that are being made. They simply throw the data to the bridge and leave the decision making up to it.

ACRONYM ALERT

RISC — Reduced instruction-set computer

Let's take a moment to look at a bridge in a network segment and how the bridge learns and gets the data to and from a set of endpoint nodes. Don't be disappointed if there is not enough meat and potatoes in this section; we will discuss node implementation in further depth in the upcoming parts of this book. As a matter of fact, network design and implementation are up next.

11.5.1 Diary of a Network Bridge

A bridge is a device that operates much like a repeater or a hub, but it makes data forwarding decisions that bridges traffic from one network segment to another. A network segment is simply a group of nodes that are connected to one another via a shared medium (see Figure 11-20).

The only limitations to the number of network segments the bridge can connect would be the number of ports the bridge physically has. In order for a bridge to operate correctly, each node that connects to a segment that is connected to the bridge must have a globally unique MAC address.[29] The bridge will have at least one interface that connects to the network segment

[26] The term ''switch'' is often used when a new product is marketed. Some examples of a node that is called a switch but is nothing like a bridge include Layer 3 switching, routing switch, and application switch.

[27] You may need to run the two types together if you are running a mixed environment. This is supported (thank goodness).

[28] This is also the type that we will focus on (in keeping in line with our focus on Ethernet).

[29] You wouldn't want to introduce confusion when you have the same locally assigned MAC address in a node in two different network segments.

that it knows about, and it will build a table that maps the globally assigned MAC addresses to the port or interface that the bridge has determined the MAC can be reached through.

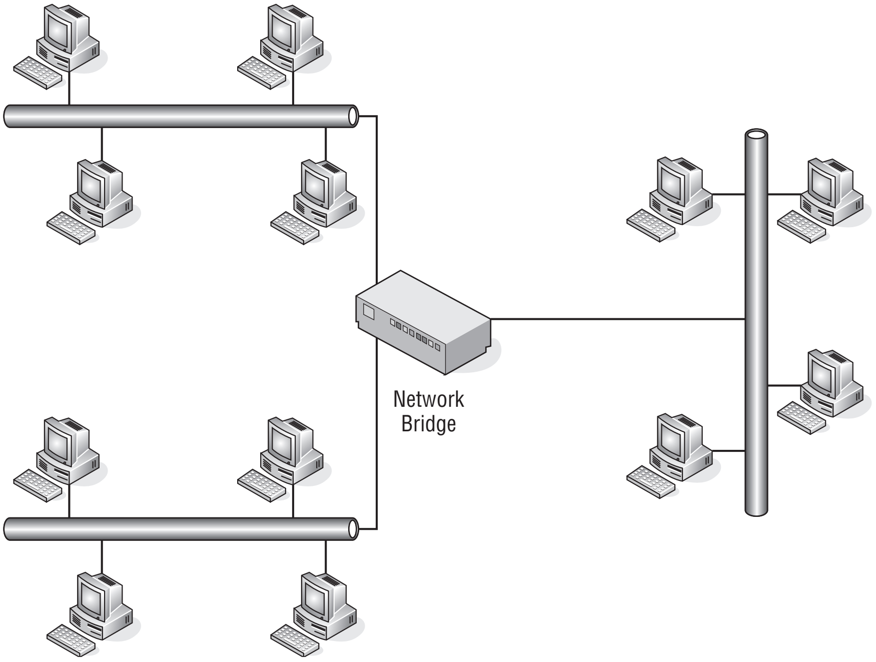

Figure 11-20 A bridge connecting three network segments

The bridge operates in *promiscuous mode*, which means it will take each frame that it receives (regardless of the destination MAC address). The switch then will use information in the frame to make a decision on which segment a MAC address belongs to. In Ethernet, the source and destination MAC address is the information the bridge uses. Figure 11-21 shows an example of two network segments that are able to communicate via an Ethernet bridge. Notice that the bridge has learned the MAC address of each of the nodes and has logged it in the MAC address table, along with the port that it knows it has to go through to reach the MAC.

11.5.1.1 Unicast Operation

Earlier in this chapter, we said that unicasting is the act of sending a frame from one source node to a single destination node. Now let's take a deeper look into what happens in multicast operations. The bridge receives frames on any active interface. The bridge then reviews the frame, looking at the destination and the source MAC addresses. It checks the MAC table to see if it knows the destination and forwards the frame to the destination.[30] The bridge follows the rules of the network protocols that are in use (for instance, the rules of CSMA/CD, flow control, congestion control, waiting for the token, etc.). Another thing that a bridge does when forwarding the frame is to use the source node's MAC address as the outbound interface address, instead of its own. This keeps the bridge transparent to the end nodes and reduces the computations necessary when receiving and retransmitting a frame. Figure 11-22 shows an example of frame forwarding.

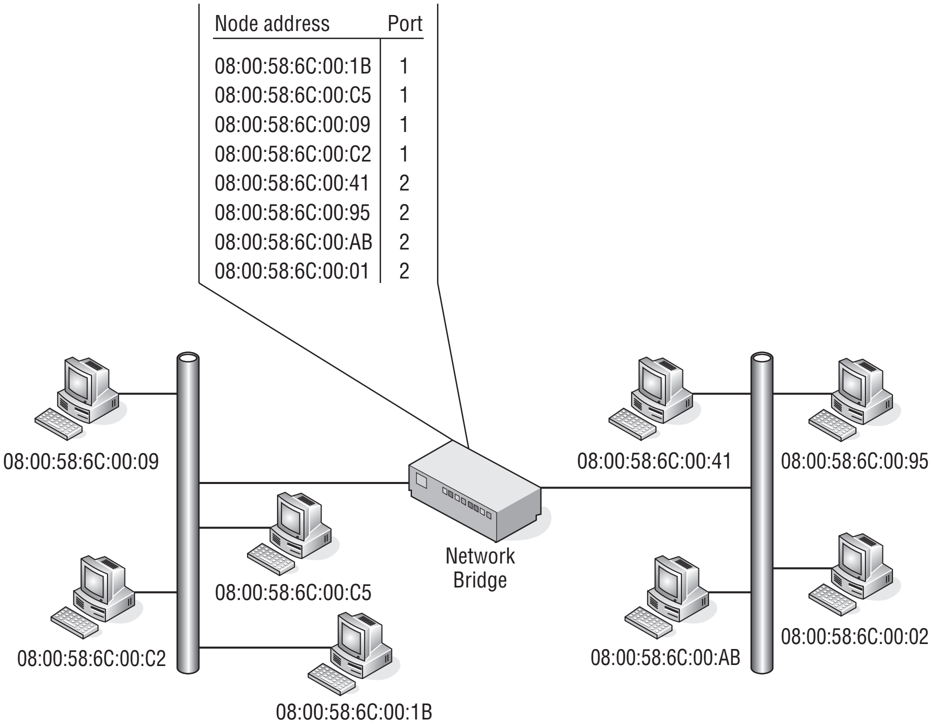

Node address	Port
08:00:58:6C:00:1B	1
08:00:58:6C:00:C5	1
08:00:58:6C:00:09	1
08:00:58:6C:00:C2	1
08:00:58:6C:00:41	2
08:00:58:6C:00:95	2
08:00:58:6C:00:AB	2
08:00:58:6C:00:01	2

Figure 11-21 The operation of a bridge — mapping the addresses to the interface they belong on

You can see that the source node contains a MAC address of 08:00:58:6C:00:09, and it is sending a frame to MAC address 08:00:58:6C:00:AB. The bridge receives the frame, noting that the destination MAC address is 08:00:58:6C:00:AB. Looking at the address table, the bridge knows that MAC address is reachable via port 3, and the bridge forwards the frame on, leaving itself transparent by identifying itself with the source MAC of the destination.

[30] When the bridge sees a source MAC address that it does not currently have in its address table, it will add the information at that point.

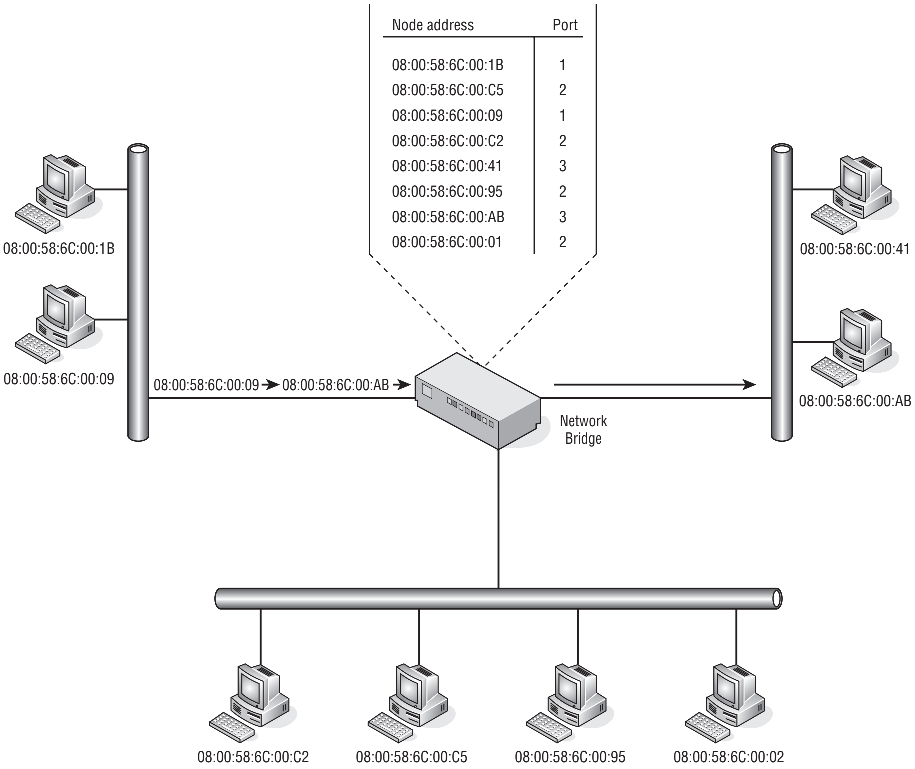

Node address	Port
08:00:58:6C:00:1B	1
08:00:58:6C:00:C5	2
08:00:58:6C:00:09	1
08:00:58:6C:00:C2	2
08:00:58:6C:00:41	3
08:00:58:6C:00:95	2
08:00:58:6C:00:AB	3
08:00:58:6C:00:01	2

Figure 11-22 Unicast frame forwarding

11.5.1.2 Multicast Operation

When a bridge receives a frame that is destined for a multicast address, the bridge forwards the frame to all the ports except the port on which it is received.[31] In Figure 11-23, the source node that has a MAC address of 08:00:58:6C:00:09 sends a frame to all members of the multicast group. The bridge recognizes that this is a multicast frame and forwards it out to all ports except the port that it received the frame on (in this case, port 1).

11.5.1.3 When the Bridge Just Does Not Know

Sometimes a bridge receives a frame that is destined for a node it does not know about. The bridge is limited in what it can do in these cases. It can

[31] This is known as flooding. A couple of things that can be done to cut down on the amount of ports that are flooded to are multicast pruning and virtual LANs (VLANS). These allow the ports in the switch to be separated into groups so that not all are affected when a frame is flooded. This ensures that the multicast traffic only goes out the ports that are part of that multicast group.

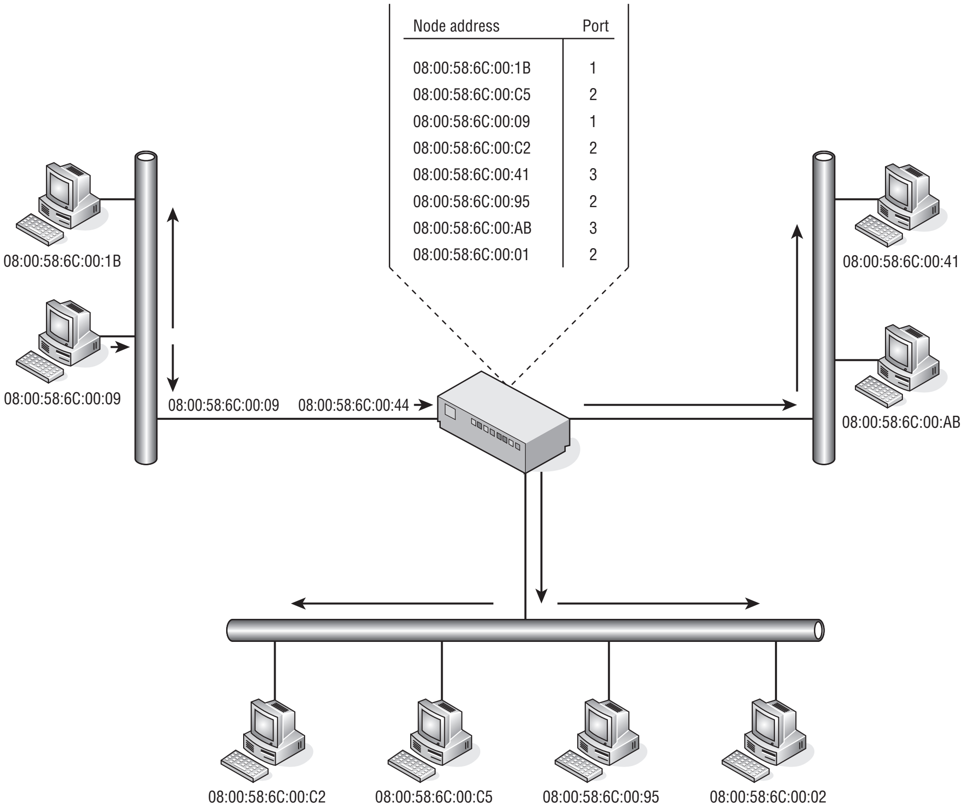

Node address	Port
08:00:58:6C:00:1B	1
08:00:58:6C:00:C5	2
08:00:58:6C:00:09	1
08:00:58:6C:00:C2	2
08:00:58:6C:00:41	3
08:00:58:6C:00:95	2
08:00:58:6C:00:AB	3
08:00:58:6C:00:01	2

Figure 11-23 Multicast frame forwarding

forward the frame to all ports (except the one it received the frame on), or it can discard the frame.

Figure 11-24 shows node 08:00:58:6C:00:09 sending a frame to node 08:00:58:6C:00:44. The bridge can't find that MAC address in its table, so it floods the frame out and eventually the frame will arrive at the node via port 3.

RANDOM BONUS DEFINITION

ARP cache — A data structure that provides the current mapping of 32-bit IP addresses to 48-bit MAC addresses.

11.5.2 The Address Table

The bridge would be nothing more than a bulkier rendition of a network hub if it were not for its ability to direct traffic to a proper port for data delivery. The address table is the backbone for the proper operation of a bridge. The

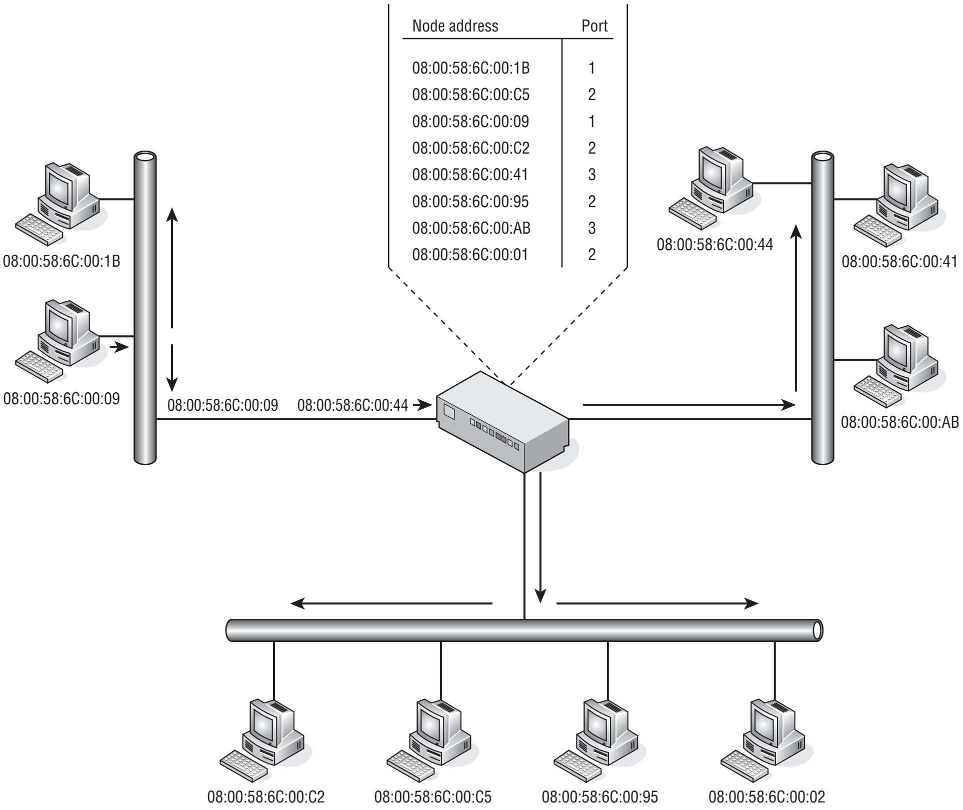

Node address	Port
08:00:58:6C:00:1B	1
08:00:58:6C:00:C5	2
08:00:58:6C:00:09	1
08:00:58:6C:00:C2	2
08:00:58:6C:00:41	3
08:00:58:6C:00:95	2
08:00:58:6C:00:AB	3
08:00:58:6C:00:01	2

Figure 11-24 Unknown destination frame forwarding

address table is built based on the source address of received frames. As we discussed previously, one of the functions of the bridge is to forward and flood frames based on the information in the address table. Another important function is to see if a source address contained in the frame is in the address table, and if not, to add it. If it is, then the port mapping is updated so that the latest destination information is synchronized. Eventually, the bridge will know about every bridge[32] that connects to a shared segment that it interfaces with.

Another important process that needs to occur pertaining to the address table is that the address entries must expire after a period of time. Imagine how big an address table would become if entries were never removed. Additionally, the performance of the bridge could suffer, as the list could become cumbersome to review if too large. When the bridge receives a frame, it checks the address table to see if the source MAC is present. If it is, it flags

[32] And have updated and accurate forwarding information.

the address so that it realizes that the MAC is still active and the information needs to be retained until the address finally does expire.

11.6 Chapter Exercises

1. How is a jam signal used in a CSMA/CD environment?
2. How is a jam signal used in a CSMA/CA environment?
3. An unnumbered frame type is used with which type of LLC?
4. Find the MAC address of your PC's NIC card. Once you have found it, take the OUI and look it up on the IEEE website. What is the information that is listed for that particular OUI?
5. What are the three fields in an LLC PDU, and what do they do?
6. How many bits are in an IEEE 802 MAC address?
7. What are the two error-checking methods used at the Data Link layer?
8. What does full-duplex Ethernet use for flow control?
9. What is the functional difference between a bridge and a Layer 2 switch?

11.7 Pop Quiz Answers

1. The Data Link layer is what layer of the OSI reference model?

   Layer 2
2. Multiple nodes attached to a single shared medium can define what?

   A LAN
3. Define *throughput*.

   Throughput is the average rate of successful messages transmitted over a channel.
4. Name two methods of ensuring bandwidth is distributed fairly to the nodes that share connectivity within a LAN.
   - Token
   - Contention
5. The most popular and most often used CSMA/CD protocol is *Ethernet*.
6. What does DSAP stand for?

   Destination Service Access Point

7. What does SNAP stand for?

   Subnetwork Access Protocol

8. What does SSAP stand for?

   Source Service Access Point

9. A *unicast* address is simply the address of a particular node's interface within the LAN.

10. What is the Ethernet standard broadcast MAC address?

    FF:FF:FF:FF:FF:FF

11. What is the simplest of all error-checking methods?

    Parity check