

6

THE TRANSPORT LAYER

The transport layer is not just another layer. It is the heart of the whole protocol hierarchy. Its task is to provide reliable, cost-effective data transport from the source machine to the destination machine, independent of the physical network or networks currently in use. Without the transport layer, the whole concept of layered protocols would make little sense. In this chapter we will study the transport layer in detail, including its services, design, protocols, and performance.

6.1. THE TRANSPORT SERVICE

In the following sections we will provide an introduction to the transport service. We look at what kind of service is provided to the application layer (or session layer, if one exists), and especially how one can characterize the quality of service. Then we will look at how applications access the transport service, that is, what the interface is like.

6.1.1. Services Provided to the Upper Layers

The ultimate goal of the transport layer is to provide efficient, reliable, and cost-effective service to its users, normally processes in the application layer. To achieve this goal, the transport layer makes use of the services provided

by the network layer. The hardware and/or software within the transport layer that does the work is called the **transport entity**. The transport entity can be in the operating system kernel, in a separate user process, in a library package bound into network applications, or on the network interface card. In some cases, the carrier may even provide reliable transport service, in which case the transport entity lives on special interface machines at the edge of the subnet to which hosts connect. The (logical) relationship of the network, transport, and application layers is illustrated in Fig. 6-1.




Fig. 6-1. The network, transport, and application layers.

Just as there are two types of network service, connection-oriented and connectionless, there are also the same two types of transport service. The connection-oriented transport service is similar to the connection-oriented network service in many ways. In both cases, connections have three phases: establishment, data transfer, and release. Addressing and flow control are also similar in both layers. Furthermore, the connectionless transport service is also very similar to the connectionless network service.

The obvious question is then: If the transport layer service is so similar to the network layer service, why are there two distinct layers? Why is one layer not adequate? The answer is subtle, but crucial, and goes back to Fig. 1-16. In this figure we can see that the network layer is part of the communication subnet and is run by the carrier (at least for WANs). What happens if the network layer offers connection-oriented service but is unreliable? Suppose that it frequently loses packets? What happens if routers crash from time to time?

Problems occur, that's what. The users have no control over the subnet, so they cannot solve the problem of poor service by using better routers or putting more error handling in the data link layer. The only possibility is to put another

layer on top of the network layer that improves the quality of the service. If a transport entity is informed halfway through a long transmission that its network connection has been abruptly terminated, with no indication of what has happened to the data currently in transit, it can set up a new network connection to the remote transport entity. Using this new network connection, it can send a query to its peer asking which data arrived and which did not, and then pick up from where it left off.

In essence, the existence of the transport layer makes it possible for the transport service to be more reliable than the underlying network service. Lost packets and mangled data can be detected and compensated for by the transport layer. Furthermore, the transport service primitives can be designed to be independent of the network service primitives which may vary considerably from network to network (e.g., connectionless LAN service may be quite different than connection-oriented WAN service).

Thanks to the transport layer, it is possible for application programs to be written using a standard set of primitives, and to have these programs work on a wide variety of networks, without having to worry about dealing with different subnet interfaces and unreliable transmission. If all real networks were flawless and all had the same service primitives, the transport layer would probably not be needed. However, in the real world it fulfills the key function of isolating the upper layers from the technology, design, and imperfections of the subnet.

For this reason, many people have made a distinction between layers 1 through 4 on the one hand, and layer(s) above 4 on the other. The bottom four layers can be seen as the **transport service provider**, whereas the upper layer(s) are the **transport service user**. This distinction of provider versus user has a considerable impact on the design of the layers and puts the transport layer in a key position, since it forms the major boundary between the provider and user of the reliable data transmission service.

6.1.2. Quality of Service

Another way of looking at the transport layer is to regard its primary function as enhancing the **QoS (Quality of Service)** provided by the network layer. If the network service is impeccable, the transport layer has an easy job. If, however, the network service is poor, the transport layer has to bridge the gap between what the transport users want and what the network layer provides.

While at first glance, quality of service might seem like a vague concept (getting everyone to agree what constitutes “good” service is a nontrivial exercise), QoS can be characterized by a number of specific parameters, as we saw in Chap. 5. The transport service may allow the user to specify preferred, acceptable, and minimum values for various service parameters at the time a connection is set up. Some of the parameters also apply to connectionless transport. It is up to the transport layer to examine these parameters, and depending on the kind of

network service or services available to it, determine whether it can provide the required service. In the remainder of this section we will discuss some possible QoS parameters. They are summarized in Fig. 6-2. Note that few networks or protocols provide all of these parameters. Many just try their best to reduce the residual error rate and leave it at that. Others have elaborate QoS architectures (Campbell et al., 1994).

Connection establishment delay
Connection establishment failure probability
Throughput
Transit delay
Residual error ratio
Protection
Priority
Resilience

Fig. 6-2. Typical transport layer quality of service parameters.

The *Connection establishment delay* is the amount of time elapsing between a transport connection being requested and the confirmation being received by the user of the transport service. It includes the processing delay in the remote transport entity. As with all parameters measuring a delay, the shorter the delay, the better the service.

The *Connection establishment failure probability* is the chance of a connection not being established within the maximum establishment delay time, for example, due to network congestion, lack of table space somewhere, or other internal problems.

The *Throughput* parameter measures the number of bytes of user data transferred per second, measured over some time interval. The throughput is measured separately for each direction.

The *Transit delay* measures the time between a message being sent by the transport user on the source machine and its being received by the transport user on the destination machine. As with throughput, each direction is handled separately.

The *Residual error ratio* measures the number of lost or garbled messages as a fraction of the total sent. In theory, the residual error rate should be zero, since it is the job of the transport layer to hide all network layer errors. In practice it may have some (small) finite value.

The *Protection* parameter provides a way for the transport user to specify interest in having the transport layer provide protection against unauthorized third parties (wiretappers) reading or modifying the transmitted data.

The *Priority* parameter provides a way for a transport user to indicate that some of its connections are more important than other ones, and in the event of congestion, to make sure that the high-priority connections get serviced before the low-priority ones.

Finally, the *Resilience* parameter gives the probability of the transport layer itself spontaneously terminating a connection due to internal problems or congestion.

The QoS parameters are specified by the transport user when a connection is requested. Both the desired and minimum acceptable values can be given. In some cases, upon seeing the QoS parameters, the transport layer may immediately realize that some of them are unachievable, in which case it tells the caller that the connection attempt failed, without even bothering to contact the destination. The failure report specifies the reason for the failure.

In other cases, the transport layer knows it cannot achieve the desired goal (e.g., 600 Mbps throughput), but it can achieve a lower, but still acceptable rate (e.g., 150 Mbps). It then sends the lower rate and the minimum acceptable rate to the remote machine, asking to establish a connection. If the remote machine cannot handle the proposed value, but it can handle a value above the minimum, it may make a counteroffer. If it cannot handle any value above the minimum, it rejects the connection attempt. Finally, the originating transport user is informed of whether the connection was established or rejected, and if it was established, the values of the parameters agreed upon.

This process is called **option negotiation**. Once the options have been negotiated, they remain that way throughout the life of the connection. To keep customers from being too greedy, most carriers have the tendency to charge more money for better quality service.

6.1.3. Transport Service Primitives

The transport service primitives allow transport users (e.g., application programs) to access the transport service. Each transport service has its own access primitives. In this section, we will first examine a simple (hypothetical) transport service and then look at a real example.

The transport service is similar to the network service, but there are also some important differences. The main difference is that the network service is intended to model the service offered by real networks, warts and all. Real networks can lose packets, so the network service is generally unreliable.

The (connection-oriented) transport service, in contrast, is reliable. Of course, real networks are not error-free, but that is precisely the purpose of the transport layer—to provide a reliable service on top of an unreliable network.

As an example, consider two processes connected by pipes in UNIX. They assume the connection between them is perfect. They do not want to know about acknowledgements, lost packets, congestion, or anything like that. What they

want is a 100 percent reliable connection. Process *A* puts data into one end of the pipe, and process *B* takes it out of the other. This is what the connection-oriented transport service is all about—hiding the imperfections of the network service so that user processes can just assume the existence of an error-free bit stream.

As an aside, the transport layer can also provide unreliable (datagram) service, but there is relatively little to say about that, so we will concentrate on the connection-oriented transport service in this chapter.

A second difference between the network service and transport service is whom the services are intended for. The network service is used only by the transport entities. Few users write their own transport entities, and thus few users or programs ever see the bare network service. In contrast, many programs (and thus programmers) see the transport primitives. Consequently, the transport service must be convenient and easy to use.

To get an idea of what a transport service might be like, consider the five primitives listed in Fig. 6-3. This transport interface is truly bare bones but it gives the essential flavor of what a connection-oriented transport interface has to do. It allows application programs to establish, use, and release connections, which is sufficient for many applications.

Primitive	TPDU sent	Meaning
LISTEN	(none)	Block until some process tries to connect
CONNECT	CONNECTION REQ.	Actively attempt to establish a connection
SEND	DATA	Send information
RECEIVE	(none)	Block until a DATA TPDU arrives
DISCONNECT	DISCONNECTION REQ.	This side wants to release the connection

Fig. 6-3. The primitives for a simple transport service.

To see how these primitives might be used, consider an application with a server and a number of remote clients. To start with, the server executes a LISTEN primitive, typically by calling a library procedure that makes a system call to block the server until a client turns up. When a client wants to talk to the server, it executes a CONNECT primitive. The transport entity carries out this primitive by blocking the caller and sending a packet to the server. Encapsulated in the payload of this packet is a transport layer message for the server's transport entity.

A quick note on terminology is now in order. For lack of a better term, we will reluctantly use the somewhat ungainly acronym **TPDU (Transport Protocol Data Unit)** for messages sent from transport entity to transport entity. Thus TPDU's (exchanged by the transport layer) are contained in packets (exchanged by the network layer). In turn, packets are contained in frames (exchanged by the data link layer). When a frame arrives, the data link layer processes the frame header and passes the contents of the frame payload field up to the network entity.

The network entity processes the packet header and passes the contents of the packet payload up to the transport entity. This nesting is illustrated in Fig. 6-4.




Fig. 6-4. Nesting of TPDU, packets, and frames.

Getting back to our client-server example, the client's CONNECT call causes a CONNECTION REQUEST TPDU to be sent to the server. When it arrives, the transport entity checks to see that the server is blocked on a LISTEN (i.e., is interested in handling requests). It then unblocks the server and sends a CONNECTION ACCEPTED TPDU back to the client. When this TPDU arrives, the client is unblocked and the connection is established.

Data can now be exchanged using the SEND and RECEIVE primitives. In the simplest form, either party can do a (blocking) RECEIVE to wait for the other party to do a SEND. When the TPDU arrives, the receiver is unblocked. It can then process the TPDU and send a reply. As long as both sides can keep track of whose turn it is to send, this scheme works fine.

Note that at the network layer, even a simple unidirectional data exchange is more complicated than at the transport layer. Every data packet sent will also be acknowledged (eventually). The packets bearing control TPDU's are also acknowledged, implicitly or explicitly. These acknowledgements are managed by the transport entities using the network layer protocol and are not visible to the transport users. Similarly, the transport entities will need to worry about timers and retransmissions. None of this machinery is seen by the transport users. To the transport users, a connection is a reliable bit pipe: one user stuffs bits in and they magically appear at the other end. This ability to hide complexity is the reason that layered protocols are such a powerful tool.

When a connection is no longer needed, it must be released to free up table space within the two transport entities. Disconnection has two variants: asymmetric and symmetric. In the asymmetric variant, either transport user can issue a DISCONNECT primitive, which results in a DISCONNECT TPDU being sent to the remote transport entity. Upon arrival, the connection is released.

In the symmetric variant, each direction is closed separately, independently of the other one. When one side does a DISCONNECT, that means it has no more data

to send, but it is still willing to accept data from its partner. In this model, a connection is released when both sides have done a DISCONNECT.

A state diagram for connection establishment and release for these simple primitives is given in Fig. 6-5. Each transition is triggered by some event, either a primitive executed by the local transport user or an incoming packet. For simplicity, we assume here that each TPDU is separately acknowledged. We also assume that a symmetric disconnection model is used, with the client going first. Please note that this model is quite unsophisticated. We will look at more realistic models later on.




Fig. 6-5. A state diagram for a simple connection management scheme. Transitions labeled in italics are caused by packet arrivals. The solid lines show the client's state sequence. The dashed lines show the server's state sequence.

Berkeley Sockets

Let us now briefly inspect another set of transport primitives, the socket primitives used in Berkeley UNIX for TCP. They are listed in Fig. 6-6. Roughly speaking, they follow the model of our first example but offer more features and flexibility. We will not look at the corresponding TPDU's here. That discussion will have to wait until we study TCP later in this chapter.

The first four primitives in the list are executed in that order by servers. The SOCKET primitive creates a new end point and allocates table space for it within

Primitive	Meaning
SOCKET	Create a new communication end point
BIND	Attach a local address to a socket
LISTEN	Announce willingness to accept connections; give queue size
ACCEPT	Block the caller until a connection attempt arrives
CONNECT	Actively attempt to establish a connection
SEND	Send some data over the connection
RECEIVE	Receive some data from the connection
CLOSE	Release the connection

Fig. 6-6. The socket primitives for TCP.

the transport entity. The parameters of the call specify the addressing format to be used, the type of service desired (e.g., reliable byte stream), and the protocol. A successful `SOCKET` call returns an ordinary file descriptor for use in succeeding calls, the same way an `OPEN` call does.

Newly created sockets do not have addresses. These are assigned using the `BIND` primitive. Once a server has bound an address to a socket, remote clients can connect to it. The reason for not having the `SOCKET` call create an address directly is that some processes care about their address (e.g., they have been using the same address for years and everyone knows this address), whereas others do not care.

Next comes the `LISTEN` call, which allocates space to queue incoming calls for the case that several clients try to connect at the same time. In contrast to `LISTEN` in our first example, in the socket model `LISTEN` is not a blocking call.

To block waiting for an incoming connection, the server executes an `ACCEPT` primitive. When a TPDU asking for a connection arrives, the transport entity creates a new socket with the same properties as the original one and returns a file descriptor for it. The server can then fork off a process or thread to handle the connection on the new socket and go back to waiting for the next connection on the original socket.

Now let us look at the client side. Here, too, a socket must first be created using the `SOCKET` primitive, but `BIND` is not required since the address used does not matter to the server. The `CONNECT` primitive blocks the caller and actively starts the connection process. When it completes (i.e., when the appropriate TPDU is received from the server), the client process is unblocked and the connection is established. Both sides can now use `SEND` and `RECEIVE` to transmit and receive data over the full-duplex connection.

Connection release with sockets is symmetric. When both sides have executed a `CLOSE` primitive, the connection is released.

6.2. ELEMENTS OF TRANSPORT PROTOCOLS

The transport service is implemented by a **transport protocol** used between the two transport entities. In some ways, transport protocols resemble the data link protocols we studied in detail in Chap. 3. Both have to deal with error control, sequencing, and flow control, among other issues.

However, significant differences between the two also exist. These differences are due to major dissimilarities between the environments in which the two protocols operate, as shown in Fig. 6-7. At the data link layer, two routers communicate directly via a physical channel, whereas at the transport layer, this physical channel is replaced by the entire subnet. This difference has many important implications for the protocols.




Fig. 6-7. (a) Environment of the data link layer. (b) Environment of the transport layer.

For one thing, in the data link layer, it is not necessary for a router to specify which router it wants to talk to—each outgoing line uniquely specifies a particular router. In the transport layer, explicit addressing of destinations is required.

For another thing, the process of establishing a connection over the wire of Fig. 6-7(a) is simple: the other end is always there (unless it has crashed, in which case it is not there). Either way, there is not much to do. In the transport layer, initial connection establishment is more complicated, as we will see.

Another, exceedingly annoying, difference between the data link layer and the transport layer is the potential existence of storage capacity in the subnet. When a router sends a frame, it may arrive or be lost, but it cannot bounce around for a while, go into hiding in a far corner of the world, and then suddenly emerge at an inopportune moment 30 sec later. If the subnet uses datagrams and adaptive routing inside, there is a nonnegligible probability that a packet may be stored for a number of seconds and then delivered later. The consequences of this ability of the subnet to store packets can sometimes be disastrous and require the use of special protocols.

A final difference between the data link and transport layers is one of amount rather than of kind. Buffering and flow control are needed in both layers, but the presence of a large and dynamically varying number of connections in the

transport layer may require a different approach than we used in the data link layer. In Chap. 3, some of the protocols allocate a fixed number of buffers to each line, so that when a frame arrives there is always a buffer available. In the transport layer, the larger number of connections that must be managed make the idea of dedicating many buffers to each one less attractive. In the following sections, we will examine all of these important issues and others.

6.2.1. Addressing

When an application process wishes to set up a connection to a remote application process, it must specify which one to connect to. (Connectionless transport has the same problem: To whom should each message be sent?) The method normally used is to define transport addresses to which processes can listen for connection requests. In the Internet, these end points are (IP address, local port) pairs. In ATM networks, they are AAL-SAPs. We will use the neutral term **TSAP (Transport Service Access Point)**. The analogous end points in the network layer (i.e., network layer addresses) are then called **NSAPs**. IP addresses are examples of NSAPs.

Figure 6-8 illustrates the relationship between the NSAP, TSAP, network connection, and transport connection for a connection-oriented subnet (e.g., ATM). Note that a transport entity normally supports multiple TSAPs. On some networks, multiple NSAPs also exist, but on others each machine has only one NSAP (e.g., one IP address). A possible connection scenario for a transport connection over a connection-oriented network layer is as follows.

1. A time-of-day server process on host 2 attaches itself to TSAP 122 to wait for an incoming call. How a process attaches itself to a TSAP is outside the networking model and depends entirely on the local operating system. A call such as our LISTEN might be used, for example.
2. An application process on host 1 wants to find out the time-of-day, so it issues a CONNECT request specifying TSAP 6 as the source and TSAP 122 as the destination.
3. The transport entity on host 1 selects a network address on its machine (if it has more than one) and sets up a network connection between them. (With a connectionless subnet, establishing this network layer connection would not be done.) Using this network connection, host 1's transport entity can talk to the transport entity on host 2.
4. The first thing the transport entity on 1 says to its peer on 2 is: "Good morning. I would like to establish a transport connection between my TSAP 6 and your TSAP 122. What do you say?"

5. The transport entity on 2 then asks the time-of-day server at TSAP 122 if it is willing to accept a new connection. If it agrees, the transport connection is established.

Note that the transport connection goes from TSAP to TSAP, whereas the network connection only goes part way, from NSAP to NSAP.




Fig. 6-8. TSAPs, NSAPs, and connections.

The picture painted above is fine, except we have swept one little problem under the rug: How does the user process on host 1 know that the time-of-day server is attached to TSAP 122? One possibility is that the time-of-day server has been attaching itself to TSAP 122 for years, and gradually all the network users have learned this. In this model, services have stable TSAP addresses which can be printed on paper and given to new users when they join the network.

While stable TSAP addresses might work for a small number of key services that never change, in general, user processes often want to talk to other user processes that only exist for a short time and do not have a TSAP address that is known in advance. Furthermore, if there are potentially many server processes, most of which are rarely used, it is wasteful to have each of them active and listening to a stable TSAP address all day long. In short, a better scheme is needed.

One such scheme, used by UNIX hosts on the Internet, is shown in Fig. 6-9 in a simplified form. It is known as the **initial connection protocol**. Instead of every conceivable server listening at a well-known TSAP, each machine that wishes to

offer service to remote users has a special **process server** that acts as a proxy for less-heavily used servers. It listens to a set of ports at the same time, waiting for a TCP connection request. Potential users of a service begin by doing a **CONNECT** request, specifying the TSAP address (TCP port) of the service they want. If no server is waiting for them, they get a connection to the process server, as shown in Fig. 6-9(a).




Fig. 6-9. How a user process in host 1 establishes a connection with a time-of-day server in host 2.

After it gets the incoming request, the process server spawns off the requested server, allowing it to inherit the existing connection with the user. The new server then does the requested work, while the process server goes back to listening for new requests, as shown in Fig. 6-9(b).

While the initial connection protocol works fine for those servers that can be created as they are needed, there are many situations in which services do exist independently of the process server. A file server, for example, needs to run on special hardware (a machine with a disk) and cannot just be created on-the-fly when someone wants to talk to it.

To handle this situation, an alternative scheme is often used. In this model, there exists a special process called a **name server** or sometimes a **directory server**. To find the TSAP address corresponding to a given service name, such as "time-of-day," a user sets up a connection to the name server (which listens to a well-known TSAP). The user then sends a message specifying the service name,

and the name server sends back the TSAP address. Then the user releases the connection with the name server and establishes a new one with the desired service.

In this model, when a new service is created, it must register itself with the name server, giving both its service name (typically an ASCII string) and the address of its TSAP. The name server records this information in its internal database, so that when queries come in later, it will know the answers.

The function of the name server is analogous to the directory assistance operator in the telephone system—it provides a mapping of names onto numbers. Just as in the telephone system, it is essential that the address of the well-known TSAP used by the name server (or the process server in the initial connection protocol) is indeed well known. If you do not know the number of the information operator, you cannot call the information operator to find it out. If you think the number you dial for information is obvious, try it in a foreign country some time.

Now let us suppose that the user has successfully located the address of the TSAP to be connected to. Another interesting question is how does the local transport entity know on which machine that TSAP is located? More specifically, how does the transport entity know which network layer address to use to set up a network connection to the remote transport entity that manages the TSAP requested?

The answer depends on the structure of TSAP addresses. One possible structure is that TSAP addresses are **hierarchical addresses**. With hierarchical addresses, the address consists of a sequence of fields used to disjointly partition the address space. For example, a truly universal TSAP address might have the following structure:

address = <galaxy> <star> <planet> <country> <network> <host> <port>

With this scheme, it is straightforward to locate a TSAP anywhere in the known universe. Equivalently, if a TSAP address is a concatenation of an NSAP address and a port (a local identifier specifying one of the local TSAPs), then when a transport entity is given a TSAP address to connect to, it uses the NSAP address contained in the TSAP address to reach the proper remote transport entity.

As a simple example of a hierarchical address, consider the telephone number 19076543210. This number can be parsed as 1-907-654-3210, where 1 is a country code (United States + Canada), 907 is an area code (Alaska), 654 is an end office in Alaska, and 3210 is one of the “ports” (subscriber lines) in that end office.

The alternative to a hierarchical address space is a **flat address space**. If the TSAP addresses are not hierarchical, a second level of mapping is needed to locate the proper machine. There would have to be a name server that took transport addresses as input and returned network addresses as output. Alternatively, in some situations (e.g., on a LAN), it is possible to broadcast a query asking the destination machine to please identify itself by sending a packet.

6.2.2. Establishing a Connection

Establishing a connection sounds easy, but it is actually surprisingly tricky. At first glance, it would seem sufficient for one transport entity to just send a CONNECTION REQUEST TPDU to the destination and wait for a CONNECTION ACCEPTED reply. The problem occurs when the network can lose, store, and duplicate packets.

Imagine a subnet that is so congested that acknowledgements hardly ever get back in time, and each packet times out and is retransmitted two or three times. Suppose that the subnet uses datagrams inside, and every packet follows a different route. Some of the packets might get stuck in a traffic jam inside the subnet and take a long time to arrive, that is, they are stored in the subnet and pop out much later.

The worst possible nightmare is as follows. A user establishes a connection with a bank, sends messages telling the bank to transfer a large amount of money to the account of a not-entirely-trustworthy person, and then releases the connection. Unfortunately, each packet in the scenario is duplicated and stored in the subnet. After the connection has been released, all the packets pop out of the subnet and arrive at the destination in order, asking the bank to establish a new connection, transfer money (again), and release the connection. The bank has no way of telling that these are duplicates. It must assume that this is a second, independent transaction, and transfers the money again. For the remainder of this section we will study the problem of delayed duplicates, with special emphasis on algorithms for establishing connections in a reliable way, so that nightmares like the one above cannot happen.

The crux of the problem is the existence of delayed duplicates. It can be attacked in various ways, none of them very satisfactory. One way is to use throwaway transport addresses. In this approach, each time a transport address is needed, a new one is generated. When a connection is released, the address is discarded. This strategy makes the process server model of Fig. 6-9 impossible.

Another possibility is to give each connection a connection identifier (i.e., a sequence number incremented for each connection established), chosen by the initiating party, and put in each TPDU, including the one requesting the connection. After each connection is released, each transport entity could update a table listing obsolete connections as (peer transport entity, connection identifier) pairs. Whenever a connection request came in, it could be checked against the table, to see if it belonged to a previously released connection.

Unfortunately, this scheme has a basic flaw: it requires each transport entity to maintain a certain amount of history information indefinitely. If a machine crashes and loses its memory, it will no longer know which connection identifiers have already been used.

Instead, we need to take a different tack. Rather than allowing packets to live forever within the subnet, we must devise a mechanism to kill off aged packets

that are still wandering about. If we can ensure that no packet lives longer than some known time, the problem becomes somewhat more manageable.

Packet lifetime can be restricted to a known maximum using one of the following techniques:

1. Restricted subnet design.
2. Putting a hop counter in each packet.
3. Timestamping each packet.

The first method includes any method that prevents packets from looping, combined with some way of bounding congestion delay over the (now known) longest possible path. The second method consists of having the hop count incremented each time the packet is forwarded. The data link protocol simply discards any packet whose hop counter has exceeded a certain value. The third method requires each packet to bear the time it was created, with the routers agreeing to discard any packet older than some agreed upon time. This latter method requires the router clocks to be synchronized, which itself is a nontrivial task unless synchronization is achieved external to the network, for example by listening to WWV or some other radio station that broadcasts the precise time periodically.

In practice, we will need to guarantee not only that a packet is dead, but also that all acknowledgements to it are also dead, so we will now introduce T , which is some small multiple of the true maximum packet lifetime. The multiple is protocol-dependent and simply has the effect of making T longer. If we wait a time T after a packet has been sent, we can be sure that all traces of it are now gone and that neither it nor its acknowledgements will suddenly appear out of the blue to complicate matters.

With packet lifetimes bounded, it is possible to devise a foolproof way to establish connections safely. The method described below is due to Tomlinson (1975). It solves the problem but introduces some peculiarities of its own. The method was further refined by Sunshine and Dalal (1978). Variants of it are widely used in practice.

To get around the problem of a machine losing all memory of where it was after a crash, Tomlinson proposed equipping each host with a time-of-day clock. The clocks at different hosts need not be synchronized. Each clock is assumed to take the form of a binary counter that increments itself at uniform intervals. Furthermore, the number of bits in the counter must equal or exceed the number of bits in the sequence numbers. Last, and most important, the clock is assumed to continue running even if the host goes down.

The basic idea is to ensure that two identically numbered TPDU's are never outstanding at the same time. When a connection is set up, the low-order k bits of the clock are used as the initial sequence number (also k bits). Thus, unlike our protocols of Chap. 3, each connection starts numbering its TPDU's with a different

sequence number. The sequence space should be so large that by the time sequence numbers wrap around, old TPDU's with the same sequence number are long gone. This linear relation between time and initial sequence numbers is shown in Fig. 6-10.




Fig. 6-10. (a) TPDU's may not enter the forbidden region. (b) The resynchronization problem.

Once both transport entities have agreed on the initial sequence number, any sliding window protocol can be used for data flow control. In reality, the initial sequence number curve (shown by the heavy line) is not really linear, but a staircase, since the clock advances in discrete steps. For simplicity we will ignore this detail.

A problem occurs when a host crashes. When it comes up again, its transport entity does not know where it was in the sequence space. One solution is to require transport entities to be idle for T sec after a recovery to let all old TPDU's die off. However, in a complex internetwork, T may be large, so this strategy is unattractive.

To avoid requiring T sec of dead time after a crash, it is necessary to introduce a new restriction on the use of sequence numbers. We can best see the need for this restriction by means of an example. Let T , the maximum packet lifetime, be 60 sec and let the clock tick once per second. As shown in Fig. 6-10, the initial sequence number for a connection opened at time x will be x . Imagine that at $t = 30$ sec, an ordinary data TPDU being sent on (a previously opened) connection 5 is given sequence number 80. Call this TPDU X . Immediately after sending TPDU X , the host crashes and then quickly restarts. At $t = 60$, it begins reopening connections 0 through 4. At $t = 70$, it reopens connection 5, using initial sequence number 70 as required. Within the next 15 sec it sends data TPDU's 70 through 80. Thus at $t = 85$, a new TPDU with sequence number 80 and connection 5 has been injected into the subnet. Unfortunately, TPDU X still exists. If it

should arrive at the receiver before the new TPDU 80, TPDU *X* will be accepted and the correct TPDU 80 will be rejected as a duplicate.

To prevent such problems, we must prevent sequence numbers from being used (i.e., assigned to new TPDU's) for a time *T* before their potential use as initial sequence numbers. The illegal combinations of time and sequence number are shown as the **forbidden region** in Fig. 6-10(a). Before sending any TPDU on any connection, the transport entity must read the clock and check to see that it is not in the forbidden region.

The protocol can get itself into trouble in two different ways. If a host sends too much data too fast on a newly opened connection, the actual sequence number versus time curve may rise more steeply than the initial sequence number versus time curve. This means that the maximum data rate on any connection is one TPDU per clock tick. It also means that the transport entity must wait until the clock ticks before opening a new connection after a crash restart, lest the same number be used twice. Both of these points argue for a short clock tick (a few milliseconds).

Unfortunately, entering the forbidden region from underneath by sending too fast is not the only way to get into trouble. From Fig. 6-10(b), it should be clear that at any data rate less than the clock rate, the curve of actual sequence numbers used versus time will eventually run into the forbidden region from the left. The greater the slope of the actual sequence number curve, the longer this event will be delayed. As we stated above, just before sending every TPDU, the transport entity must check to see if it is about to enter the forbidden region, and if so, either delay the TPDU for *T* sec or resynchronize the sequence numbers.

The clock-based method solves the delayed duplicate problem for data TPDU's, but for this method to be useful, a connection must first be established. Since control TPDU's may also be delayed, there is a potential problem in getting both sides to agree on the initial sequence number. Suppose, for example, that connections are established by having host 1 send a CONNECTION REQUEST TPDU containing the proposed initial sequence number and destination port number to a remote peer, host 2. The receiver, host 2, then acknowledges this request by sending a CONNECTION ACCEPTED TPDU back. If the CONNECTION REQUEST TPDU is lost but a delayed duplicate CONNECTION REQUEST suddenly shows up at host 2, the connection will be established incorrectly.

To solve this problem, Tomlinson (1975) introduced the **three-way handshake**. This establishment protocol does not require both sides to begin sending with the same sequence number, so it can be used with synchronization methods other than the global clock method. The normal setup procedure when host 1 initiates is shown in Fig. 6-11(a). Host 1 chooses a sequence number, *x*, and sends a CONNECTION REQUEST TPDU containing it to host 2. Host 2 replies with a CONNECTION ACCEPTED TPDU acknowledging *x* and announcing its own initial sequence number, *y*. Finally, host 1 acknowledges host 2's choice of an initial sequence number in the first data TPDU that it sends.




Fig. 6-11. Three protocol scenarios for establishing a connection using a three-way handshake. CR and ACC denote CONNECTION REQUEST and CONNECTION ACCEPTED, respectively. (a) Normal operation. (b) Old duplicate CONNECTION REQUEST appearing out of nowhere. (c) Duplicate CONNECTION REQUEST and duplicate ACK.

Now let us see how the three-way handshake works in the presence of delayed duplicate control TPDU. In Fig. 6-12(b), the first TPDU is a delayed duplicate CONNECTION REQUEST from an old connection. This TPDU arrives at host 2 without host 1's knowledge. Host 2 reacts to this TPDU by sending host 1 a CONNECTION ACCEPTED TPDU, in effect asking for verification that host 1 was indeed trying to set up a new connection. When host 1 rejects host 2's attempt to establish, host 2 realizes that it was tricked by a delayed duplicate and abandons the connection. In this way, a delayed duplicate does no damage.

The worst case is when both a delayed CONNECTION REQUEST and an acknowledgement to a CONNECTION ACCEPTED are floating around in the subnet. This case is shown in Fig. 6-11(c). As in the previous example, host 2 gets a delayed CONNECTION REQUEST and replies to it. At this point it is crucial to realize that host 2 has proposed using y as the initial sequence number for host 2 to host 1 traffic, knowing full well that no TPDU's containing sequence number y or acknowledgements to y are still in existence. When the second delayed TPDU arrives at host 2, the fact that z has been acknowledged rather than y tells host 2 that this, too, is an old duplicate. The important thing to realize here is that there is no combination of old CONNECTION REQUEST, CONNECTION ACCEPTED, or other TPDU's that can cause the protocol to fail and have a connection set up by accident when no one wants it.

An alternative scheme for establishing connections reliably in the face of delayed duplicates is described in (Watson, 1981). It uses multiple timers to exclude undesired events.

6.2.3. Releasing a Connection

Releasing a connection is easier than establishing one. Nevertheless, there are more pitfalls than one might expect. As we mentioned earlier, there are two styles of terminating a connection: asymmetric release and symmetric release. Asymmetric release is the way the telephone system works: when one party hangs up, the connection is broken. Symmetric release treats the connection as two separate unidirectional connections and requires each one to be released separately.




Fig. 6-12. Abrupt disconnection with loss of data.

Asymmetric release is abrupt and may result in data loss. Consider the scenario of Fig. 6-12. After the connection is established, host 1 sends a TPDU

that arrives properly at host 2. Then host 1 sends another TPDU. Unfortunately, host 2 issues a DISCONNECT before the second TPDU arrives. The result is that the connection is released and data are lost.

Clearly, a more sophisticated release protocol is required to avoid data loss. One way is to use symmetric release, in which each direction is released independently of the other one. Here, a host can continue to receive data even after it has sent a DISCONNECT TPDU.

Symmetric release does the job when each process has a fixed amount of data to send and clearly knows when it has sent it. In other situations, determining that all the work has been done and the connection should be terminated is not so obvious. One can envision a protocol in which host 1 says: "I am done. Are you done too?" If host 2 responds: "I am done too. Goodbye." the connection can be safely released.

Unfortunately, this protocol does not always work. There is a famous problem that deals with this issue. It is called the **two-army problem**. Imagine that a white army is encamped in a valley, as shown in Fig. 6-13. On both of the surrounding hillsides are blue armies. The white army is larger than either of the blue armies alone, but together they are larger than the white army. If either blue army attacks by itself, it will be defeated, but if the two blue armies attack simultaneously, they will be victorious.




Fig. 6-13. The two-army problem.

The blue armies want to synchronize their attacks. However, their only communication medium is to send messengers on foot down into the valley, where they might be captured and the message lost (i.e., they have to use an unreliable communication channel). The question is: Does a protocol exist that allows the blue armies to win?

Suppose that the commander of blue army #1 sends a message reading: "I propose we attack at dawn on March 29. How about it?" Now suppose that the

message arrives, and the commander of blue army #2 agrees, and that his reply gets safely back to blue army #1. Will the attack happen? Probably not, because commander #2 does not know if his reply got through. If it did not, blue army #1 will not attack, so it would be foolish for him to charge into battle.

Now let us improve the protocol by making it a three-way handshake. The initiator of the original proposal must acknowledge the response. Assuming no messages are lost, blue army #2 will get the acknowledgement, but the commander of blue army #1 will now hesitate. After all, he does not know if his acknowledgement got through, and if it did not, he knows that blue army #2 will not attack. We could now make a four-way handshake protocol, but that does not help either.

In fact, it can be proven that no protocol exists that works. Suppose that some protocol did exist. Either the last message of the protocol is essential or it is not. If it is not, remove it (and any other unessential messages) until we are left with a protocol in which every message is essential. What happens if the final message does not get through? We just said that it was essential, so if it is lost, the attack does not take place. Since the sender of the final message can never be sure of its arrival, he will not risk attacking. Worse yet, the other blue army knows this, so it will not attack either.

To see the relevance of the two-army problem to releasing connections, just substitute “disconnect” for “attack.” If neither side is prepared to disconnect until it is convinced that the other side is prepared to disconnect too, the disconnection will never happen.

In practice, one is usually prepared to take more risks when releasing connections than when attacking white armies, so the situation is not entirely hopeless. Figure 6-14 illustrates four scenarios of releasing using a three-way handshake. While this protocol is not infallible, it is usually adequate.

In Fig. 6-14(a), we see the normal case in which one of the users sends a DR (DISCONNECTION REQUEST) TPDU in order to initiate the connection release. When it arrives, the recipient sends back a DR TPDU, too, and starts a timer, just in case its DR is lost. When this DR arrives, the original sender sends back an ACK TPDU and releases the connection. Finally, when the ACK TPDU arrives, the receiver also releases the connection. Releasing a connection means that the transport entity removes the information about the connection from its table of open connections and signals the connection’s owner (the transport user) somehow. This action is different from a transport user issuing a DISCONNECT primitive.

If the final ACK TPDU is lost, as shown in Fig. 6-14(b), the situation is saved by the timer. When the timer expires, the connection is released anyway.

Now consider the case of the second DR being lost. The user initiating the disconnection will not receive the expected response, will time out, and will start all over again. In Fig. 6-14(c) we see how this works, assuming that the second time no TPDU is lost and all TPDU are delivered correctly and on time.




Fig. 6-14. Four protocol scenarios for releasing a connection. (a) Normal case of three-way handshake. (b) Final ACK lost. (c) Response lost. (d) Response lost and subsequent DRs lost.

Our last scenario, Fig. 6-14(d), is the same as Fig. 6-14(c) except that now we assume all the repeated attempts to retransmit the DR also fail due to lost TPDU's. After N retries, the sender just gives up and releases the connection. Meanwhile, the receiver times out and also exits.

While this protocol usually suffices, in theory it can fail if the initial DR and N retransmissions are all lost. The sender will give up and release the connection, while the other side knows nothing at all about the attempts to disconnect and is still fully active. This situation results in a half-open connection.

We could have avoided this problem by not allowing the sender to give up after N retries but forcing it to go on forever until it gets a response. However, if the other side is allowed to time out, then the sender will indeed go on forever, because no response will ever be forthcoming. If we do not allow the receiving side to time out, then the protocol hangs in Fig. 6-14(b).

One way to kill off half-open connections is to have a rule saying that if no TPDU's have arrived for a certain number of seconds, the connection is automatically disconnected. That way, if one side ever disconnects, the other side will detect the lack of activity and also disconnect. Of course, if this rule is introduced, it is necessary for each transport entity to have a timer that is stopped and then restarted whenever a TPDU is sent. If this timer expires, a dummy TPDU is transmitted, just to keep the other side from disconnecting. On the other hand, if the automatic disconnect rule is used and too many dummy TPDU's in a row are lost on an otherwise idle connection, first one side, then the other side will automatically disconnect.

We will not belabor this point any more, but by now it should be clear that releasing a connection is not nearly as simple as it at first appears.

6.2.4. Flow Control and Buffering

Having examined connection establishment and release in some detail, let us now look at how connections are managed while they are in use. One of the key issues has come up before: flow control. In some ways the flow control problem in the transport layer is the same as in the data link layer, but in other ways it is different. The basic similarity is that in both layers a sliding window or other scheme is needed on each connection to keep a fast transmitter from overrunning a slow receiver. The main difference is that a router usually has relatively few lines whereas a host may have numerous connections. This difference makes it impractical to implement the data link buffering strategy in the transport layer.

In the data link protocols of Chap. 3, frames were buffered at both the sending router and at the receiving router. In protocol 6, for example, both sender and receiver are required to dedicate $MaxSeq + 1$ buffers to each line, half for input and half for output. For a host with a maximum of, say, 64 connections, and a 4-bit sequence number, this protocol would require 1024 buffers.

In the data link layer, the sending side must buffer outgoing frames because they might have to be retransmitted. If the subnet provides datagram service, the sending transport entity must also buffer, and for the same reason. If the receiver knows that the sender buffers all TPDU's until they are acknowledged, the receiver may or may not dedicate specific buffers to specific connections, as it sees fit. The receiver may, for example, maintain a single buffer pool shared by all connections. When a TPDU comes in, an attempt is made to dynamically acquire a new buffer. If one is available, the TPDU is accepted; otherwise, it is discarded. Since the sender is prepared to retransmit TPDU's lost by the subnet, no harm is

done by having the receiver drop TPDU's, although some resources are wasted. The sender just keeps trying until it gets an acknowledgement.

In summary, if the network service is unreliable, the sender must buffer all TPDU's sent, just as in the data link layer. However, with reliable network service, other trade-offs become possible. In particular, if the sender knows that the receiver always has buffer space, it need not retain copies of the TPDU's it sends. However, if the receiver cannot guarantee that every incoming TPDU will be accepted, the sender will have to buffer anyway. In the latter case, the sender cannot trust the network layer's acknowledgement, because the acknowledgement means only that the TPDU arrived, not that it was accepted. We will come back to this important point later.

Even if the receiver has agreed to do the buffering, there still remains the question of the buffer size. If most TPDU's are nearly the same size, it is natural to organize the buffers as a pool of identical size buffers, with one TPDU per buffer, as in Fig. 6-15(a). However, if there is wide variation in TPDU size, from a few characters typed at a terminal to thousands of characters from file transfers, a pool of fixed-sized buffers presents problems. If the buffer size is chosen equal to the largest possible TPDU, space will be wasted whenever a short TPDU arrives. If the buffer size is chosen less than the maximum TPDU size, multiple buffers will be needed for long TPDU's, with the attendant complexity.




Fig. 6-15. (a) Chained fixed-size buffers. (b) Chained variable-size buffers.
 (c) One large circular buffer per connection.

Another approach to the buffer size problem is to use variable-size buffers, as in Fig. 6-15(b). The advantage here is better memory utilization, at the price of

more complicated buffer management. A third possibility is to dedicate a single large circular buffer per connection, as in Fig. 6-15(c). This system also makes good use of memory, provided that all connections are heavily loaded but is poor if some connections are lightly loaded.

The optimum trade-off between source buffering and destination buffering depends on the type of traffic carried by the connection. For low-bandwidth bursty traffic, such as that produced by an interactive terminal, it is better not to dedicate any buffers, but rather to acquire them dynamically at both ends. Since the sender cannot be sure the receiver will be able to acquire a buffer, the sender must retain a copy of the TPDU until it is acknowledged. On the other hand, for file transfer and other high-bandwidth traffic, it is better if the receiver does dedicate a full window of buffers, to allow the data to flow at maximum speed. Thus for low-bandwidth bursty traffic, it is better to buffer at the sender, and for high-bandwidth, smooth traffic, it is better to buffer at the receiver.

As connections are opened and closed, and as the traffic pattern changes, the sender and receiver need to dynamically adjust their buffer allocations. Consequently, the transport protocol should allow a sending host to request buffer space at the other end. Buffers could be allocated per connection, or collectively, for all the connections running between the two hosts. Alternatively, the receiver, knowing its buffer situation (but not knowing the offered traffic) could tell the sender "I have reserved X buffers for you." If the number of open connections should increase, it may be necessary for an allocation to be reduced, so the protocol should provide for this possibility.

A reasonably general way to manage dynamic buffer allocation is to decouple the buffering from the acknowledgements, in contrast to the sliding window protocols of Chap. 3. Dynamic buffer management means, in effect, a variable-sized window. Initially, the sender requests a certain number of buffers, based on its perceived needs. The receiver then grants as many of these as it can afford. Every time the sender transmits a TPDU, it must decrement its allocation, stopping altogether when the allocation reaches zero. The receiver then separately piggybacks both acknowledgements and buffer allocations onto the reverse traffic.

Figure 6-16 shows an example of how dynamic window management might work in a datagram subnet with 4-bit sequence numbers. Assume that buffer allocation information travels in separate TPDU's, as shown, and is not piggybacked onto reverse traffic. Initially, A wants eight buffers, but is granted only four of these. It then sends three TPDU's, of which the third is lost. TPDU 6 acknowledges receipt of all TPDU's up to and including sequence number 1, thus allowing A to release those buffers, and furthermore informs A that it has permission to send three more TPDU's starting beyond 1 (i.e., TPDU's 2, 3, and 4). A knows that it has already sent number 2, so it thinks that it may send TPDU's 3 and 4, which it proceeds to do. At this point it is blocked and must wait for more buffer allocation. Timeout induced retransmissions (line 9), however, may occur while blocked, since they use buffers that have already been allocated. In line 10, B

acknowledges receipt of all TPDU's up to and including 4, but refuses to let *A* continue. Such a situation is impossible with the fixed window protocols of Chap. 3. The next TPDU from *B* to *A* allocates another buffer and allows *A* to continue.

	<u>A</u>	<u>Message</u>	<u>B</u>	<u>Comments</u>
1	→	< request 8 buffers >	→	A wants 8 buffers
2	←	<ack = 15, buf = 4>	←	B grants messages 0-3 only
3	→	<seq = 0, data = m0>	→	A has 3 buffers left now
4	→	<seq = 1, data = m1>	→	A has 2 buffers left now
5	→	<seq = 2, data = m2>	•••	Message lost but A thinks it has 1 left
6	←	<ack = 1, buf = 3>	←	B acknowledges 0 and 1, permits 2-4
7	→	<seq = 3, data = m3>	→	A has buffer left
8	→	<seq = 4, data = m4>	→	A has 0 buffers left, and must stop
9	→	<seq = 2, data = m2>	→	A times out and retransmits
10	←	<ack = 4, buf = 0>	←	Everything acknowledged, but A still blocked
11	←	<ack = 4, buf = 1>	←	A may now send 5
12	←	<ack = 4, buf = 2>	←	B found a new buffer somewhere
13	→	<seq = 5, data = m5>	→	A has 1 buffer left
14	→	<seq = 6, data = m6>	→	A is now blocked again
15	←	<ack = 6, buf = 0>	←	A is still blocked
16	•••	<ack = 6, buf = 4>	←	Potential deadlock

Fig. 6-16. Dynamic buffer allocation. The arrows show the direction of transmission. An ellipsis (...) indicates a lost TPDU.

Potential problems with buffer allocation schemes of this kind can arise in datagram networks if control TPDU's can get lost. Look at line 16. *B* has now allocated more buffers to *A*, but the allocation TPDU was lost. Since control TPDU's are not sequenced or timed out, *A* is now deadlocked. To prevent this situation, each host should periodically send control TPDU's giving the acknowledgement and buffer status on each connection. That way, the deadlock will be broken, sooner or later.

Up until now we have tacitly assumed that the only limit imposed on the sender's data rate is the amount of buffer space available in the receiver. As memory prices continue to fall dramatically, it may become feasible to equip hosts with so much memory that lack of buffers is rarely, if ever, a problem.

When buffer space no longer limits the maximum flow, another bottleneck will appear: the carrying capacity of the subnet. If adjacent routers can exchange at most x frames/sec and there are k disjoint paths between a pair of hosts, there is no way that those hosts can exchange more than kx TPDU's/sec, no matter how much buffer space is available at each end. If the sender pushes too hard (i.e., sends more than kx TPDU's/sec), the subnet will become congested, because it will be unable to deliver TPDU's as fast as they are coming in.

What is needed is a mechanism based on the subnet's carrying capacity rather than on the receiver's buffering capacity. Clearly, the flow control mechanism must be applied at the sender to prevent it from having too many unacknowledged TPDU's outstanding at once. Belsnes (1975) proposed using a sliding window flow control scheme in which the sender dynamically adjusts the window size to match the network's carrying capacity. If the network can handle c TPDU's/sec and the cycle time (including transmission, propagation, queueing, processing at the receiver, and return of the acknowledgement) is r , then the sender's window should be cr . With a window of this size the sender normally operates with the pipeline full. Any small decrease in network performance will cause it to block.

In order to adjust the window size periodically, the sender could monitor both parameters and then compute the desired window size. The carrying capacity can be determined by simply counting the number of TPDU's acknowledged during some time period and then dividing by the time period. During the measurement, the sender should send as fast as it can, to make sure that the network's carrying capacity, and not the low input rate, is the factor limiting the acknowledgement rate. The time required for a transmitted TPDU to be acknowledged can be measured exactly and a running mean maintained. Since the capacity of the network depends on the amount of traffic in it, the window size should be adjusted frequently, to track changes in the carrying capacity. As we will see later, the Internet uses a similar scheme.

6.2.5. Multiplexing

Multiplexing several conversations onto connections, virtual circuits, and physical links plays a role in several layers of the network architecture. In the transport layer the need for multiplexing can arise in a number of ways. For example, in networks that use virtual circuits within the subnet, each open connection consumes some table space in the routers for the entire duration of the connection. If buffers are dedicated to the virtual circuit in each router as well, a user who left a terminal logged into a remote machine during a coffee break is nevertheless consuming expensive resources. Although this implementation of packet switching defeats one of the main reasons for having packet switching in the first place—to bill the user based on the amount of data sent, not the connect time—many carriers have chosen this approach because it so closely resembles the circuit switching model to which they have grown accustomed over the decades.

The consequence of a price structure that heavily penalizes installations for having many virtual circuits open for long periods of time is to make multiplexing of different transport connections onto the same network connection attractive. This form of multiplexing, called **upward multiplexing**, is shown in Fig. 6-17(a). In this figure, four distinct transport connections all use the same network connection (e.g., ATM virtual circuit) to the remote host. When connect time forms the

major component of the carrier's bill, it is up to the transport layer to group transport connections according to their destination and map each group onto the minimum number of network connections. If too many transport connections are mapped onto one network connection, the performance will be poor, because the window will usually be full, and users will have to wait their turn to send one message. If too few transport connections are mapped onto one network connection, the service will be expensive. When upward multiplexing is used with ATM, we have the ironic (tragic?) situation of having to identify the connection using a field in the transport header, even though ATM provides more than 4000 virtual circuit numbers per virtual path expressly for that purpose.




Fig. 6-17. (a) Upward multiplexing. (b) Downward multiplexing.

Multiplexing can also be useful in the transport layer for another reason, related to carrier technical decisions rather than carrier pricing decisions. Suppose, for example, that a certain important user needs a high-bandwidth connection from time to time. If the subnet enforces a sliding window flow control with an n -bit sequence number, the user must stop sending as soon as $2^n - 1$ packets are outstanding and must wait for the packets to propagate to the remote host and be acknowledged. If the physical connection is via a satellite, the user is effectively limited to $2^n - 1$ packets every 540 msec. With, for example, $n = 8$ and 128-byte packets, the usable bandwidth is about 484 kbps, even though the physical channel bandwidth is more than 100 times higher.

One possible solution is to have the transport layer open multiple network connections and distribute the traffic among them on a round-robin basis, as indicated in Fig. 6-17(b). This modus operandi is called **downward multiplexing**. With k network connections open, the effective bandwidth is increased by a factor of k . With 4095 virtual circuits, 128-byte packets, and an 8-bit sequence number,

it is theoretically possible to achieve data rates in excess of 1.6 Gbps. Of course, this performance can be achieved only if the output line can support 1.6 Gbps, because all 4095 virtual circuits are still being sent out over one physical line, at least in Fig. 6-17(b). If multiple output lines are available, downward multiplexing can also be used to increase the performance even more.

6.2.6. Crash Recovery

If hosts and routers are subject to crashes, recovery from these crashes becomes an issue. If the transport entity is entirely within the hosts, recovery from network and router crashes is straightforward. If the network layer provides datagram service, the transport entities expect lost TPDU's all the time and know how to cope with them. If the network layer provides connection-oriented service, then loss of a virtual circuit is handled by establishing a new one and then probing the remote transport entity to ask it which TPDU's it has received and which ones it has not received. The latter ones can be retransmitted.

A more troublesome problem is how to recover from host crashes. In particular, it may be desirable for clients to be able to continue working when servers crash and then quickly reboot. To illustrate the difficulty, let us assume that one host, the client, is sending a long file to another host, the file server, using a simple stop-and-wait protocol. The transport layer on the server simply passes the incoming TPDU's to the transport user, one by one. Part way through the transmission, the server crashes. When it comes back up, its tables are reinitialized, so it no longer knows precisely where it was.

In an attempt to recover its previous status, the server might send a broadcast TPDU to all other hosts, announcing that it had just crashed and requesting that its clients inform it of the status of all open connections. Each client can be in one of two states: one TPDU outstanding, *S1*, or no TPDU's outstanding, *S0*. Based on only this state information, the client must decide whether or not to retransmit the most recent TPDU.

At first glance it would seem obvious: the client should retransmit only if it has an unacknowledged TPDU outstanding (i.e., is in state *S1*) when it learns of the crash. However, a closer inspection reveals difficulties with this naive approach. Consider, for example, the situation when the server's transport entity first sends an acknowledgement, and then, when the acknowledgement has been sent, performs the write up to the application process. Writing a TPDU onto the output stream and sending an acknowledgement are two distinct indivisible events that cannot be done simultaneously. If a crash occurs after the acknowledgement has been sent but before the write has been done, the client will receive the acknowledgement and thus be in state *S0* when the crash recovery announcement arrives. The client will therefore not retransmit, (incorrectly) thinking that the TPDU has arrived. This decision by the client leads to a missing TPDU.

At this point you may be thinking: "That problem can be solved easily. All you have to do is reprogram the transport entity to first do the write and then send the acknowledgement." Try again. Imagine that the write has been done but the crash occurs before the acknowledgement can be sent. The client will be in state *S1* and thus retransmit, leading to an undetected duplicate TPDU in the output stream to the server application process.

No matter how the sender and receiver are programmed, there are always situations where the protocol fails to recover properly. The server can be programmed in one of two ways: acknowledge first or write first. The client can be programmed in one of four ways: always retransmit the last TPDU, never retransmit the last TPDU, retransmit only in state *S0*, or retransmit only in state *S1*. This gives eight combinations, but as we shall see, for each combination there is some set of events that makes the protocol fail.

Three events are possible at the server: sending an acknowledgement (*A*), writing to the output process (*W*), and crashing (*C*). The three events can occur in six different orderings: *AC(W)*, *AWC*, *C(AW)*, *C(WA)*, *WAC*, and *WC(A)*, where the parentheses are used to indicate that neither *A* nor *W* may follow *C* (i.e., once it has crashed, it has crashed). Figure 6-18 shows all eight combinations of client and server strategy and the valid event sequences for each one. Notice that for each strategy there is some sequence of events that causes the protocol to fail. For example, if the client always retransmits, the *AWC* event will generate an undetected duplicate, even though the other two events work properly.

Strategy used by sending host	Strategy used by receiving host					
	First ACK, then write			First write, then ACK		
	<i>AC(W)</i>	<i>AWC</i>	<i>C(AW)</i>	<i>C(WA)</i>	<i>WAC</i>	<i>WC(A)</i>
Always retransmit	OK	DUP	OK	OK	DUP	DUP
Never retransmit	LOST	OK	LOST	LOST	OK	OK
Retransmit in <i>S0</i>	OK	DUP	LOST	LOST	DUP	OK
Retransmit in <i>S1</i>	LOST	OK	OK	OK	OK	DUP

OK = Protocol functions correctly
 DUP = Protocol generates a duplicate message
 LOST = Protocol loses a message

Fig. 6-18. Different combinations of client and server strategy.

Making the protocol more elaborate does not help. Even if the client and server exchange several TPDU's before the server attempts to write, so that the client knows exactly what is about to happen, the client has no way of knowing whether a crash occurred just before or just after the write. The conclusion is

inescapable: under our ground rules of no simultaneous events, host crash and recovery cannot be made transparent to higher layers.

Put in more general terms, this result can be restated as recovery from a layer N crash can only be done by layer $N + 1$, and then only if the higher layer retains enough status information. As mentioned above, the transport layer can recover from failures in the network layer, provided that each end of a connection keeps track of where it is.

This problem gets us into the issue of what a so-called end-to-end acknowledgement really means. In principle, the transport protocol is end-to-end and not chained like the lower layers. Now consider the case of a user entering requests for transactions against a remote database. Suppose that the remote transport entity is programmed to first pass TPDU's to the next layer up and then acknowledge. Even in this case, the receipt of an acknowledgement back at the user's machine does not necessarily mean that the remote host stayed up long enough to actually update the database. A truly end-to-end acknowledgement, whose receipt means that the work has actually been done, and lack thereof means that it has not, is probably impossible to achieve. This point is discussed in more detail by Saltzer et al. (1984).

6.3. A SIMPLE TRANSPORT PROTOCOL

To make the ideas discussed so far more concrete, in this section we will study an example transport layer in detail. The example has been carefully chosen to be reasonably realistic, yet still simple enough to be easy to understand. The abstract service primitives we will use are the connection-oriented primitives of Fig. 6-3.

6.3.1. The Example Service Primitives

Our first problem is how to express these transport primitives concretely. `CONNECT` is easy: we will just have a library procedure *connect* that can be called with the appropriate parameters necessary to establish a connection. The parameters are the local and remote TSAPs. During the call, the caller is blocked (i.e., suspended) while the transport entity tries to set up the connection. If the connection succeeds, the caller is unblocked, and can start transmitting data.

When a process wants to be able to accept incoming calls, it calls *listen*, specifying a particular TSAP to listen to. The process then blocks until some remote process attempts to establish a connection to its TSAP.

Note that this model is highly asymmetric. One side is passive, executing a *listen* and waiting until something happens. The other side is active and initiates the connection. An interesting question arises of what to do if the active side

begins first. One strategy is to have the connection attempt fail if there is no listener at the remote TSAP. Another strategy is to have the initiator block (possibly forever) until a listener appears.

A compromise, used in our example, is to hold the connection request at the receiving end for a certain time interval. If a process on that host calls *listen* before the timer goes off, the connection is established; otherwise, it is rejected and the caller is unblocked and given an error return.

To release a connection, we will use a procedure *disconnect*. When both sides have disconnected, the connection is released. In other words, we are using a symmetric disconnection model.

Data transmission has precisely the same problem as connection establishment: sending is active but receiving is passive. We will use the same solution for data transmission as for connection establishment, an active call *send* that transmits data, and a passive call *receive* that blocks until a TPDU arrives.

Our concrete service definition thus consists of five primitives: CONNECT, LISTEN, DISCONNECT, SEND, and RECEIVE. Each primitive corresponds exactly with a library procedure that executes the primitive. The parameters for the service primitives and library procedures are as follows:

```
connum = LISTEN(local)
connum = CONNECT(local, remote)
status = SEND(connum, buffer, bytes)
status = RECEIVE(connum, buffer, bytes)
status = DISCONNECT(connum)
```

The LISTEN primitive announces the caller's willingness to accept connection requests directed at the indicated TSAP. The user of the primitive is blocked until an attempt is made to connect to it. There is no timeout.

The CONNECT primitive takes two parameters, a local TSAP (i.e., transport address), *local*, and a remote TSAP, *remote*, and tries to establish a transport connection between the two. If it succeeds, it returns in *connum* a nonnegative number used to identify the connection on subsequent calls. If it fails, the reason for failure is put in *connum* as a negative number. In our simple model, each TSAP may participate in only one transport connection, so a possible reason for failure is that one of the transport addresses is currently in use. Some other reasons are: remote host down, illegal local address, and illegal remote address.

The SEND primitive transmits the contents of the buffer as a message on the indicated transport connection, possibly in several units if it is too big. Possible errors, returned in *status*, are no connection, illegal buffer address, or negative count.

The RECEIVE primitive indicates the caller's desire to accept data. The size of the incoming message is placed in *bytes*. If the remote process has released the connection or the buffer address is illegal (e.g., outside the user's program), *status* is set to an error code indicating the nature of the problem.

The DISCONNECT primitive terminates a transport connection. The parameter *connum* tells which one. Possible errors are *connum* belongs to another process, or *connum* is not a valid connection identifier. The error code, or 0 for success, is returned in *status*.

6.3.2. The Example Transport Entity

Before looking at the code of the example transport entity, please be sure you realize that this example is analogous to the early examples presented in Chap. 3: it is more for pedagogical purposes than a serious proposal. Many of the technical details (such as extensive error checking) that would be needed in a production system have been omitted here for the sake of simplicity.

The transport layer makes use of the network service primitives to send and receive TPDUs. For this example, we need to choose network service primitives to use. One choice would have been unreliable datagram service. We have not made that choice to keep the example simple. With unreliable datagram service, the transport code would have been large and complex, mostly dealing with lost and delayed packets. Furthermore, most of these ideas have already been discussed at length in Chap. 3.

Instead, we have chosen to use a connection-oriented reliable network service. This way we can focus on transport issues that do not occur in the lower layers. These include connection establishment, connection release, and credit management, among others. A simple transport service built on top of an ATM network might look something like this.

In general, the transport entity may be part of the host's operating system or it may be a package of library routines running within the user's address space. It may also be contained on a coprocessor chip or network board plugged into the host's backplane. For simplicity, our example has been programmed as though it were a library package, but the changes needed to make it part of the operating system are minimal (primarily how user buffers are accessed).

It is worth noting, however, that in this example, the "transport entity" is not really a separate entity at all, but part of the user process. In particular, when the user executes a primitive that blocks, such as LISTEN, the entire transport entity blocks as well. While this design is fine for a host with only a single user process, on a host with multiple users, it would be more natural to have the transport entity be a separate process, distinct from all the user processes.

The interface to the network layer is via the procedures *to_net* and *from_net* (not shown). Each has six parameters. First comes the connection identifier, which maps one-to-one onto network virtual circuits. Next come the *Q* and *M* bits, which, when set to 1, indicate control message and more data from this message follows in the next packet, respectively. After that we have the packet type, chosen from the set of six packet types listed in Fig. 6-19. Finally, we have a pointer to the data itself, and an integer giving the number of bytes of data.

Network packet	Meaning
CALL REQUEST	Sent to establish a connection
CALL ACCEPTED	Response to CALL REQUEST
CLEAR REQUEST	Sent to release a connection
CLEAR CONFIRMATION	Response to CLEAR REQUEST
DATA	Used to transport data
CREDIT	Control packet for managing the window

Fig. 6-19. The network layer packets used in our example.

On calls to *to_net*, the transport entity fills in all the parameters for the network layer to read; on calls to *from_net*, the network layer dismembers an incoming packet for the transport entity. By passing information as procedure parameters rather than passing the actual outgoing or incoming packet itself, the transport layer is shielded from the details of the network layer protocol. If the transport entity should attempt to send a packet when the underlying virtual circuit's sliding window is full, it is suspended within *to_net* until there is room in the window. This mechanism is transparent to the transport entity and is controlled by the network layer using commands like *enable_transport_layer* and *disable_transport_layer* analogous to those described in the protocols of Chap. 3. The management of the packet layer window is also done by the network layer.

In addition to this transparent suspension mechanism, there are also explicit *sleep* and *wakeup* procedures (not shown) called by the transport entity. The procedure *sleep* is called when the transport entity is logically blocked waiting for an external event to happen, generally the arrival of a packet. After *sleep* has been called, the transport entity (and the user process, of course) stop executing.

The actual code of the transport entity is shown in Fig. 6-20. Each connection is always in one of seven states, as follows:

1. IDLE—Connection not established yet.
2. WAITING—CONNECT has been executed and CALL REQUEST sent.
3. QUEUED—A CALL REQUEST has arrived; no LISTEN yet.
4. ESTABLISHED—The connection has been established.
5. SENDING—The user is waiting for permission to send a packet.
6. RECEIVING—A RECEIVE has been done.
7. DISCONNECTING—A DISCONNECT has been done locally.

Transitions between states can occur when any of the following events occur: a primitive is executed, a packet arrives, or the timer expires.

```

#define MAX_CONN 32                /* maximum number of simultaneous connections */
#define MAX_MSG_SIZE 8192          /* largest message in bytes */
#define MAX_PKT_SIZE 512          /* largest packet in bytes */
#define TIMEOUT 20
#define CRED 1
#define OK 0

#define ERR_FULL -1
#define ERR_REJECT -2
#define ERR_CLOSED -3
#define LOW_ERR -3

typedef int transport_address;
typedef enum {CALL_REQ,CALL_ACC,CLEAR_REQ,CLEAR_CONF,DATA_PKT,CREDIT} pkt_type;
typedef enum {IDLE,WAITING,QUEUED,ESTABLISHED,SENDING,RECEIVING,DISCONN} cstate;

/* Global variables. */
transport_address listen_address;    /* local address being listened to */
int listen_conn;                    /* connection identifier for listen */
unsigned char data[MAX_PKT_SIZE];    /* scratch area for packet data */

struct conn {
    transport_address local_address, remote_address;
    cstate state;                    /* state of this connection */
    unsigned char *user_buf_addr;    /* pointer to receive buffer */
    int byte_count;                  /* send/receive count */
    int clr_req_received;             /* set when CLEAR_REQ packet received */
    int timer;                        /* used to time out CALL_REQ packets */
    int credits;                      /* number of messages that may be sent */
} conn[MAX_CONN];

/* prototypes */
void sleep(void);
void wakeup(void);
void to_net(int cid, int q, int m, pkt_type pt, unsigned char *p, int bytes);
void from_net(int *cid, int *q, int *m, pkt_type *pt, unsigned char *p, int *bytes);

int listen(transport_address t)
{ /* User wants to listen for a connection. See if CALL_REQ has already arrived. */
    int i = 1, found = 0;

    for (i = 1; i <= MAX_CONN; i++) /* search the table for CALL_REQ */
        if (conn[i].state == QUEUED && conn[i].local_address == t) {
            found = i;
            break;
        }

    if (found == 0) {
        /* No CALL_REQ is waiting. Go to sleep until arrival or timeout. */
        listen_address = t; sleep(); i = listen_conn;
    }
    conn[i].state = ESTABLISHED; /* connection is ESTABLISHED */
    conn[i].timer = 0;          /* timer is not used */
}

```

```

listen_conn = 0; /* 0 is assumed to be an invalid address */
to_net(i, 0, 0, CALL_ACC, data, 0); /* tell net to accept connection */
return(i); /* return connection identifier */
}

int connect(transport_address l, transport_address r)
{ /* User wants to connect to a remote process; send CALL_REQ packet. */
  int i;
  struct conn *cptr;

  data[0] = r; data[1] = l; /* CALL_REQ packet needs these */
  i = MAX_CONN; /* search table backward */
  while (conn[i].state != IDLE && i > 1) i = i - 1;
  if (conn[i].state == IDLE) {
    /* Make a table entry that CALL_REQ has been sent. */
    cptr = &conn[i];
    cptr->local_address = l; cptr->remote_address = r;
    cptr->state = WAITING; cptr->clr_req_received = 0;
    cptr->credits = 0; cptr->timer = 0;
    to_net(i, 0, 0, CALL_REQ, data, 2);
    sleep(); /* wait for CALL_ACC or CLEAR_REQ */
    if (cptr->state == ESTABLISHED) return(i);
    if (cptr->clr_req_received) {
      /* Other side refused call. */
      cptr->state = IDLE; /* back to IDLE state */
      to_net(i, 0, 0, CLEAR_CONF, data, 0);
      return(ERR_REJECT);
    }
  }
  else return(ERR_FULL); /* reject CONNECT: no table space */
}

int send(int cid, unsigned char bufptr[], int bytes)
{ /* User wants to send a message. */
  int i, count, m;
  struct conn *cptr = &conn[cid];

  /* Enter SENDING state. */
  cptr->state = SENDING;
  cptr->byte_count = 0; /* # bytes sent so far this message */
  if (cptr->clr_req_received == 0 && cptr->credits == 0) sleep();
  if (cptr->clr_req_received == 0) {
    /* Credit available; split message into packets if need be. */
    do {
      if (bytes - cptr->byte_count > MAX_PKT_SIZE) { /* multipacket message */
        count = MAX_PKT_SIZE; m = 1; /* more packets later */
      } else { /* single packet message */
        count = bytes - cptr->byte_count; m = 0; /* last pkt of this message */
      }
      for (i = 0; i < count; i++) data[i] = bufptr[cptr->byte_count + i];
      to_net(cid, 0, m, DATA_PKT, data, count); /* send 1 packet */
      cptr->byte_count = cptr->byte_count + count; /* increment bytes sent so far */
    } while (cptr->byte_count < bytes); /* loop until whole message sent */
  }
}

```

```

    cptr->credits--;          /* each message uses up one credit */
    cptr->state = ESTABLISHED;
    return(OK);
} else {
    cptr->state = ESTABLISHED;
    return(ERR_CLOSED);      /* send failed: peer wants to disconnect */
}
}

int receive(int cid, unsigned char bufptr[], int *bytes)
{ /* User is prepared to receive a message. */
    struct conn *cptr = &conn[cid];

    if (cptr->clr_req_received == 0) {
        /* Connection still established; try to receive. */
        cptr->state = RECEIVING;
        cptr->user_buf_addr = bufptr;
        cptr->byte_count = 0;
        data[0] = CRED;
        data[1] = 1;
        to_net(cid, 1, 0, CREDIT, data, 2);    /* send credit */
        sleep();                               /* block awaiting data */
        *bytes = cptr->byte_count;
    }
    cptr->state = ESTABLISHED;
    return(cptr->clr_req_received ? ERR_CLOSED : OK);
}

int disconnect(int cid)
{ /* User wants to release a connection. */
    struct conn *cptr = &conn[cid];

    if (cptr->clr_req_received) {             /* other side initiated termination */
        cptr->state = IDLE;                   /* connection is now released */
        to_net(cid, 0, 0, CLEAR_CONF, data, 0);
    } else {                                  /* we initiated termination */
        cptr->state = DISCONN;               /* not released until other side agrees */
        to_net(cid, 0, 0, CLEAR_REQ, data, 0);
    }
    return(OK);
}

void packet_arrival(void)
{ /* A packet has arrived, get and process it. */
    int cid;                                 /* connection on which packet arrived */
    int count, i, q, m;
    pkt_type ptype; /* CALL_REQ, CALL_ACC, CLEAR_REQ, CLEAR_CONF, DATA_PKT, CREDIT */
    unsigned char data[MAX_PKT_SIZE]; /* data portion of the incoming packet */
    struct conn *cptr;

    from_net(&cid, &q, &m, &ptype, data, &count); /* go get it */
    cptr = &conn[cid];

```

```

switch (ptype) {
case CALL_REQ:                /* remote user wants to establish connection */
    cptr->local_address = data[0]; cptr->remote_address = data[1];
    if (cptr->local_address == listen_address) {
        listen_conn = cid; cptr->state = ESTABLISHED; wakeup();
    } else {
        cptr->state = QUEUED; cptr->timer = TIMEOUT;
    }
    cptr->clr_req_received = 0; cptr->credits = 0;
    break;
case CALL_ACC:                /* remote user has accepted our CALL_REQ */
    cptr->state = ESTABLISHED;
    wakeup();
    break;
case CLEAR_REQ:               /* remote user wants to disconnect or reject call */
    cptr->clr_req_received = 1;
    if (cptr->state == DISCONN) cptr->state = IDLE; /* clear collision */
    if (cptr->state == WAITING || cptr->state == RECEIVING || cptr->state == SENDING) wakeup();
    break;
case CLEAR_CONF:             /* remote user agrees to disconnect */
    cptr->state = IDLE;
    break;
case CREDIT:                 /* remote user is waiting for data */
    cptr->credits += data[1];
    if (cptr->state == SENDING) wakeup();
    break;
case DATA_PKT:              /* remote user has sent data */
    for (i = 0; i < count; i++) cptr->user_buf_addr[cptr->byte_count + i] = data[i];
    cptr->byte_count += count;
    if (m == 0) wakeup();
}
}

void clock(void)
{ /* The clock has ticked, check for timeouts of queued connect requests. */
    int i;
    struct conn *cptr;
    for (i = 1; i <= MAX_CONN; i++) {
        cptr = &conn[i];
        if (cptr->timer > 0) {                /* timer was running */
            cptr->timer--;
            if (cptr->timer == 0) {          /* timer has now expired */
                cptr->state = IDLE;
                to_net(i, 0, 0, CLEAR_REQ, data, 0);
            }
        }
    }
}
}

```

Fig. 6-20. An example transport entity.

The procedures shown in Fig. 6-20 are of two types. Most are directly callable by user programs. *packet_arrival* and *clock* are different, however. They are spontaneously triggered by external events: the arrival of a packet and the clock ticking, respectively. In effect, they are interrupt routines. We will assume that they are never invoked while a transport entity procedure is running. Only when the user process is sleeping or executing outside the transport entity may they be called. This property is crucial to the correct functioning of the transport entity.

The existence of the *Q* (Qualifier) bit in the packet header allows us to avoid the overhead of a transport protocol header. Ordinary data messages are sent as data packets with *Q* = 0. Transport protocol control messages, of which there is only one (CREDIT) in our example, are sent as data packets with *Q* = 1. These control messages are detected and processed by the receiving transport entity.

The main data structure used by the transport entity is the array *conn*, which has one record for each potential connection. The record maintains the state of the connection, including the transport addresses at either end, the number of messages sent and received on the connection, the current state, the user buffer pointer, the number of bytes of the current messages sent or received so far, a bit indicating that the remote user has issued a DISCONNECT, a timer, and a permission counter used to enable sending of messages. Not all of these fields are used in our simple example, but a complete transport entity would need all of them, and perhaps more. Each *conn* entry is assumed initialized to the *IDLE* state.

When the user calls CONNECT, the network layer is instructed to send a CALL REQUEST packet to the remote machine, and the user is put to sleep. When the CALL REQUEST packet arrives at the other side, the transport entity is interrupted to run *packet_arrival* to check if the local user is listening on the specified address. If so, a CALL ACCEPTED packet is sent back and the remote user is awakened; if not, the CALL REQUEST is queued for *TIMEOUT* clock ticks. If a LISTEN is done within this period, the connection is established; otherwise, it times out and is rejected with a CLEAR REQUEST packet. This mechanism is needed to prevent the initiator from blocking forever in the event that the remote process does not want to connect to it.

Although we have eliminated the transport protocol header, we still need a way to keep track of which packet belongs to which transport connection, since multiple connections may exist simultaneously. The simplest approach is to use the network layer virtual circuit number as the transport connection number as well. Furthermore, the virtual circuit number can also be used as the index into the *conn* array. When a packet comes in on network layer virtual circuit *k*, it belongs to transport connection *k*, whose state is in the record *conn[k]*. For connections initiated at a host, the connection number is chosen by the originating transport entity. For incoming calls, the network layer makes the choice, choosing any unused virtual circuit number.

To avoid having to provide and manage buffers within the transport entity, a flow control mechanism different from the traditional sliding window is used

here. Instead, when a user calls RECEIVE, a special **credit message** is sent to the transport entity on the sending machine and is recorded in the *conn* array. When SEND is called, the transport entity checks to see if a credit has arrived on the specified connection. If so, the message is sent (in multiple packets if need be) and the credit decremented; if not, the transport entity puts itself to sleep until a credit arrives. This mechanism guarantees that no message is ever sent unless the other side has already done a RECEIVE. As a result, whenever a message arrives there is guaranteed to be a buffer available into which it can be put. The scheme can easily be generalized to allow receivers to provide multiple buffers and request multiple messages.

You should keep the simplicity of Fig. 6-20 in mind. A realistic transport entity would normally check all user supplied parameters for validity, handle recovery from network layer crashes, deal with call collisions, and support a more general transport service including such facilities as interrupts, datagrams, and nonblocking versions of the SEND and RECEIVE primitives.

6.3.3. The Example as a Finite State Machine

Writing a transport entity is difficult and exacting work, especially for more realistic protocols. To reduce the chance of making an error, it is often useful to represent the state of the protocol as a finite state machine.

We have already seen that our example protocol has seven states per connection. It is also possible to isolate 12 events that can happen to move a connection from one state to another. Five of these events are the five service primitives. Another six are the arrivals of the six legal packet types. The last one is the expiration of the timer. Figure 6-21 shows the main protocol actions in matrix form. The columns are the states and the rows are the 12 events.

Each entry in the matrix (i.e., the finite state machine) of Fig. 6-21 has up to three fields: a predicate, an action, and a new state. The predicate indicates under what conditions the action is taken. For example, in the upper left-hand entry, if a LISTEN is executed and there is no more table space (predicate *P1*), the LISTEN fails and the state does not change. On the other hand, if a CALL REQUEST packet has already arrived for the transport address being listened to (predicate *P2*), the connection is established immediately. Another possibility is that *P2* is false, that is, no CALL REQUEST has come in, in which case the connection remains in the *IDLE* state, awaiting a CALL REQUEST packet.

It is worth pointing out that the choice of states to use in the matrix is not entirely fixed by the protocol itself. In this example, there is no state *LISTENING*, which might have been a reasonable thing to have following a LISTEN. There is no *LISTENING* state because a state is associated with a connection record entry, and no connection record is created by LISTEN. Why not? Because we have decided to use the network layer virtual circuit numbers as the connection

		State						
		Idle	Waiting	Queued	Established	Sending	Receiving	Dis-connecting
Primitives	LISTEN	P1: ~Idle P2: A1/Estab P2: A2/Idle		~/Estab				
	CONNECT	P1: ~Idle P1: A3/Wait						
	DISCONNECT				P4: A5/Idle P4: A6/Disc			
	SEND				P5: A7/Estab P5: A8/Send			
	RECEIVE				A9/Receiving			
Incoming packets	Call_req	P3: A1/Estab P3: A4/Queue'd						
	Call_acc		~/Estab					
	Clear_req		~/Idle		A10/Estab	A10/Estab	A10/Estab	~/Idle
	Clear_conf							~/Idle
	DataPkt						A12/Estab	
Clock	Credit				A11/Estab	A7/Estab		
	Timeout			~/Idle				

Predicates	Actions
P1: Connection table full	A1: Send Call_acc A7: Send message
P2: Call_req pending	A2: Wait for Call_req A8: Wait for credit
P3: LISTEN pending	A3: Send Call_req A9: Send credit
P4: Clear_req pending	A4: Start timer A10: Set Clr_req_received flag
P5: Credit available	A5: Send Clear_conf A11: Record credit
	A6: Send Clear_req A12: Accept message

Fig. 6-21. The example protocol as a finite state machine. Each entry has an optional predicate, an optional action, and the new state. The tilde indicates that no major action is taken. An overbar above a predicate indicates the negation of the predicate. Blank entries correspond to impossible or invalid events.

identifiers, and for a LISTEN, the virtual circuit number is ultimately chosen by the network layer when the CALL REQUEST packet arrives.

The actions A1 through A12 are the major actions, such as sending packets and starting timers. Not all the minor actions, such as initializing the fields of a connection record, are listed. If an action involves waking up a sleeping process,

the actions following the wakeup also count. For example, if a CALL REQUEST packet comes in and a process was asleep waiting for it, the transmission of the CALL ACCEPT packet following the wakeup counts as part of the action for CALL REQUEST. After each action is performed, the connection may move to a new state, as shown in Fig. 6-21.

The advantage of representing the protocol as a matrix is threefold. First, in this form it is much easier for the programmer to systematically check each combination of state and event to see if an action is required. In production implementations, some of the combinations would be used for error handling. In Fig. 6-21 no distinction is made between impossible situations and illegal ones. For example, if a connection is in *waiting* state, the DISCONNECT event is impossible because the user is blocked and cannot execute any primitives at all. On the other hand, in *sending* state, data packets are not expected because no credit has been issued. The arrival of a data packet is a protocol error.

The second advantage of the matrix representation of the protocol is in implementing it. One could envision a two-dimensional array in which element $a[i][j]$ was a pointer or index to the procedure that handled the occurrence of event i when in state j . One possible implementation is to write the transport entity as a short loop, waiting for an event at the top of the loop. When an event happens, the relevant connection is located and its state is extracted. With the event and state now known, the transport entity just indexes into the array a and calls the proper procedure. This approach gives a much more regular and systematic design than our transport entity.

The third advantage of the finite state machine approach is for protocol description. In some standards documents, the protocols are given as finite state machines of the type of Fig. 6-21. Going from this kind of description to a working transport entity is much easier if the transport entity is also driven by a finite state machine based on the one in the standard.

The primary disadvantage of the finite state machine approach is that it may be more difficult to understand than the straight programming example we used initially. However, this problem may be partially solved by drawing the finite state machine as a graph, as is done in Fig. 6-22.

6.4. THE INTERNET TRANSPORT PROTOCOLS (TCP AND UDP)

The Internet has two main protocols in the transport layer, a connection-oriented protocol and a connectionless one. In the following sections we will study both of them. The connection-oriented protocol is TCP. The connectionless protocol is UDP. Because UDP is basically just IP with a short header added, we will focus on TCP.

TCP (Transmission Control Protocol) was specifically designed to provide a reliable end-to-end byte stream over an unreliable internetwork. An




Fig. 6-22. The example protocol in graphical form. Transitions that leave the connection state unchanged have been omitted for simplicity.

internetwork differs from a single network because different parts may have wildly different topologies, bandwidths, delays, packet sizes, and other parameters. TCP was designed to dynamically adapt to properties of the internetwork and to be robust in the face of many kinds of failures.

TCP was formally defined in RFC 793. As time went on, various errors and inconsistencies were detected, and the requirements were changed in some areas. These clarifications and some bug fixes are detailed in RFC 1122. Extensions are given in RFC 1323.

Each machine supporting TCP has a TCP transport entity, either a user process or part of the kernel that manages TCP streams and interfaces to the IP layer. A TCP entity accepts user data streams from local processes, breaks them up into pieces not exceeding 64K bytes (in practice, usually about 1500 bytes), and sends each piece as a separate IP datagram. When IP datagrams containing TCP data arrive at a machine, they are given to the TCP entity, which reconstructs the original byte streams. For simplicity, we will sometimes use just “TCP” to mean the

TCP transport entity (a piece of software) or the TCP protocol (a set of rules). From the context it will be clear which is meant. For example, in "The user gives TCP the data," the TCP transport entity is clearly intended.

The IP layer gives no guarantee that datagrams will be delivered properly, so it is up to TCP to time out and retransmit them as need be. Datagrams that do arrive may well do so in the wrong order; it is also up to TCP to reassemble them into messages in the proper sequence. In short, TCP must furnish the reliability that most users want and that IP does not provide.

6.4.1. The TCP Service Model

TCP service is obtained by having both the sender and receiver create end points, called sockets, as discussed in Sec. 6.1.3. Each socket has a socket number (address) consisting of the IP address of the host and a 16-bit number local to that host, called a **port**. A port is the TCP name for a TSAP. To obtain TCP service, a connection must be explicitly established between a socket on the sending machine and a socket on the receiving machine. The socket calls are listed in Fig. 6-6.

A socket may be used for multiple connections at the same time. In other words, two or more connections may terminate at the same socket. Connections are identified by the socket identifiers at both ends, that is, (*socket1*, *socket2*). No virtual circuit numbers or other identifiers are used.

Port numbers below 256 are called **well-known ports** and are reserved for standard services. For example, any process wishing to establish a connection to a host to transfer a file using FTP can connect to the destination host's port 21 to contact its FTP daemon. Similarly, to establish a remote login session using TELNET, port 23 is used. The list of well-known ports is given in RFC 1700.

All TCP connections are full-duplex and point-to-point. Full duplex means that traffic can go in both directions at the same time. Point-to-point means that each connection has exactly two end points. TCP does not support multicasting or broadcasting.

A TCP connection is a byte stream, not a message stream. Message boundaries are not preserved end to end. For example, if the sending process does four 512-byte writes to a TCP stream, these data may be delivered to the receiving process as four 512-byte chunks, two 1024-byte chunks, one 2048-byte chunk (see Fig. 6-23), or some other way. There is no way for the receiver to detect the unit(s) in which the data were written.

Files in UNIX have this property too. The reader of a file cannot tell whether the file was written a block at a time, a byte at a time, or all in one blow. As with a UNIX file, the TCP software has no idea of what the bytes mean and no interest in finding out. A byte is just a byte.

When an application passes data to TCP, TCP may send it immediately or buffer it (in order to collect a larger amount to send at once), at its discretion.




Fig. 6-23. (a) Four 512-byte segments sent as separate IP datagrams. (b) The 2048 bytes of data delivered to the application in a single READ call.

However, sometimes, the application really wants the data to be sent immediately. For example, suppose a user is logged into a remote machine. After a command line has been finished and the carriage return typed, it is essential that the line be shipped off to the remote machine immediately and not buffered until the next line comes in. To force data out, applications can use the PUSH flag, which tells TCP not to delay the transmission.

Some early applications used the PUSH flag as a kind of marker to delineate messages boundaries. While this trick sometimes works, it sometimes fails since not all implementations of TCP pass the PUSH flag to the application on the receiving side. Furthermore, if additional PUSHes come in before the first one has been transmitted (e.g., because the output line is busy), TCP is free to collect all the PUSHed data into a single IP datagram, with no separation between the various pieces.

One last feature of the TCP service that is worth mentioning here is **urgent data**. When an interactive user hits the DEL or CTRL-C key to break off a remote computation that has already begun, the sending application puts some control information in the data stream and gives it to TCP along with the URGENT flag. This event causes TCP to stop accumulating data and transmit everything it has for that connection immediately.

When the urgent data are received at the destination, the receiving application is interrupted (e.g., given a signal in UNIX terms), so it can stop whatever it was doing and read the data stream to find the urgent data. The end of the urgent data is marked, so the application knows when it is over. The start of the urgent data is not marked. It is up to the application to figure that out. This scheme basically provides a crude signaling mechanism and leaves everything else up to the application.

6.4.2. The TCP Protocol

In this section we will give a general overview of the TCP protocol. In the next one we will go over the protocol header, field by field. Every byte on a TCP connection has its own 32-bit sequence number. For a host blasting away at full

speed on a 10-Mbps LAN, theoretically the sequence numbers could wrap around in an hour, but in practice it takes much longer. The sequence numbers are used both for acknowledgements and for the window mechanism, which use separate 32-bit header fields.

The sending and receiving TCP entities exchange data in the form of segments. A **segment** consists of a fixed 20-byte header (plus an optional part) followed by zero or more data bytes. The TCP software decides how big segments should be. It can accumulate data from several writes into one segment or split data from one write over multiple segments. Two limits restrict the segment size. First, each segment, including the TCP header, must fit in the 65,535 byte IP payload. Second, each network has a **maximum transfer unit** or **MTU**, and each segment must fit in the MTU. In practice, the MTU is generally a few thousand bytes and thus defines the upper bound on segment size. If a segment passes through a sequence of networks without being fragmented and then hits one whose MTU is smaller than the segment, the router at the boundary fragments the segment into two or more smaller segments.

A segment that is too large for a network that it must transit can be broken up into multiple segments by a router. Each new segment gets its own IP header, so fragmentation by routers increases the total overhead (because each additional segment adds 20 bytes of extra header information in the form of an IP header).

The basic protocol used by TCP entities is the sliding window protocol. When a sender transmits a segment, it also starts a timer. When the segment arrives at the destination, the receiving TCP entity sends back a segment (with data if any exists, otherwise without data) bearing an acknowledgement number equal to the next sequence number it expects to receive. If the sender's timer goes off before the acknowledgement is received, the sender transmits the segment again.

Although this protocol sounds simple, there are many ins and outs that we will cover below. For example, since segments can be fragmented, it is possible that part of a transmitted segment arrives and is acknowledged by the receiving TCP entity, but the rest is lost. Segments can also arrive out of order, so bytes 3072–4095 can arrive but cannot be acknowledged because bytes 2048–3071 have not turned up yet. Segments can also be delayed so long in transit that the sender times out and retransmits them. If a retransmitted segment takes a different route than the original, and is fragmented differently, bits and pieces of both the original and the duplicate can arrive sporadically, requiring a careful administration to achieve a reliable byte stream. Finally, with so many networks making up the Internet, it is possible that a segment may occasionally hit a congested (or broken) network along its path.

TCP must be prepared to deal with these problems and solve them in an efficient way. A considerable amount of effort has gone into optimizing the performance of TCP streams, even in the face of network problems. A number of the algorithms used by many TCP implementations will be discussed below.

6.4.3. The TCP Segment Header

Figure 6-24 shows the layout of a TCP segment. Every segment begins with a fixed-format 20-byte header. The fixed header may be followed by header options. After the options, if any, up to $65,535 - 20 - 20 = 65,515$ data bytes may follow, where the first 20 refers to the IP header and the second to the TCP header. Segments without any data are legal and are commonly used for acknowledgements and control messages.




Fig. 6-24. The TCP header.

Let us dissect the TCP header field by field. The *Source port* and *Destination port* fields identify the local end points of the connection. Each host may decide for itself how to allocate its own ports starting at 256. A port plus its host's IP address forms a 48-bit unique TSAP. The source and destination socket numbers together identify the connection.

The *Sequence number* and *Acknowledgement number* fields perform their usual functions. Note that the latter specifies the next byte expected, not the last byte correctly received. Both are 32 bits long because every byte of data is numbered in a TCP stream.

The *TCP header length* tells how many 32-bit words are contained in the TCP header. This information is needed because the *Options* field is of variable length, so the header is too. Technically, this field really indicates the start of the data

within the segment, measured in 32-bit words, but that number is just the header length in words, so the effect is the same.

Next comes a 6-bit field that is not used. The fact that this field has survived intact for over a decade is testimony to how well thought out TCP is. Lesser protocols would have needed it to fix bugs in the original design.

Now come six 1-bit flags. *URG* is set to 1 if the *Urgent pointer* is in use. The *Urgent pointer* is used to indicate a byte offset from the current sequence number at which urgent data are to be found. This facility is in lieu of interrupt messages. As we mentioned above, this facility is a bare bones way of allowing the sender to signal the receiver without getting TCP itself involved in the reason for the interrupt.

The *ACK* bit is set to 1 to indicate that the *Acknowledgement number* is valid. If *ACK* is 0, the segment does not contain an acknowledgement so the *Acknowledgement number* field is ignored.

The *PSH* bit indicates PUSHed data. The receiver is hereby kindly requested to deliver the data to the application upon arrival and not buffer it until a full buffer has been received (which it might otherwise do for efficiency reasons).

The *RST* bit is used to reset a connection that has become confused due to a host crash or some other reason. It is also used to reject an invalid segment or refuse an attempt to open a connection. In general, if you get a segment with the *RST* bit on, you have a problem on your hands.

The *SYN* bit is used to establish connections. The connection request has *SYN* = 1 and *ACK* = 0 to indicate that the piggyback acknowledgement field is not in use. The connection reply does bear an acknowledgement, so it has *SYN* = 1 and *ACK* = 1. In essence the *SYN* bit is used to denote CONNECTION REQUEST and CONNECTION ACCEPTED, with the *ACK* bit used to distinguish between those two possibilities.

The *FIN* bit is used to release a connection. It specifies that the sender has no more data to transmit. However, after closing a connection, a process may continue to receive data indefinitely. Both *SYN* and *FIN* segments have sequence numbers and are thus guaranteed to be processed in the correct order.

Flow control in TCP is handled using a variable-size sliding window. The *Window* field tells how many bytes may be sent starting at the byte acknowledged. A *Window* field of 0 is legal and says that the bytes up to and including *Acknowledgement number* - 1 have been received, but that the receiver is currently badly in need of a rest and would like no more data for the moment, thank you. Permission to send can be granted later by sending a segment with the same *Acknowledgement number* and a nonzero *Window* field.

A *Checksum* is also provided for extreme reliability. It checksums the header, the data, and the conceptual pseudoheader shown in Fig. 6-25. When performing this computation, the TCP *Checksum* field is set to zero, and the data field is padded out with an additional zero byte if its length is an odd number. The checksum algorithm is simply to add up all the 16-bit words in 1's complement and then to

take the 1's complement of the sum. As a consequence, when the receiver performs the calculation on the entire segment, including the *Checksum* field, the result should be 0.




Fig. 6-25. The pseudoheader included in the TCP checksum.

The pseudoheader contains the 32-bit IP addresses of the source and destination machines, the protocol number for TCP (6), and the byte count for the TCP segment (including the header). Including the pseudoheader in the TCP checksum computation helps detect misdelivered packets, but doing so violates the protocol hierarchy since the IP addresses in it belong to the IP layer, not the TCP layer.

The *Options* field was designed to provide a way to add extra facilities not covered by the regular header. The most important option is the one that allows each host to specify the maximum TCP payload it is willing to accept. Using large segments is more efficient than using small ones because the 20-byte header can then be amortized over more data, but small hosts may not be able to handle very large segments. During connection setup, each side can announce its maximum and see its partner's. The smaller of the two numbers wins. If a host does not use this option, it defaults to a 536-byte payload. All Internet hosts are required to accept TCP segments of $536 + 20 = 556$ bytes.

For lines with high bandwidth, high delay, or both, the 64 KB window is often a problem. On a T3 line (44.736 Mbps), it takes only 12 msec to output a full 64 KB window. If the round trip propagation delay is 50 msec (typical for a transcontinental fiber), the sender will be idle 3/4 of the time waiting for acknowledgements. On a satellite connection, the situation is even worse. A larger window size would allow the sender to keep pumping data out, but using the 16-bit *Window size* field, there is no way to express such a size. In RFC 1323, a *Window scale* option was proposed, allowing the sender and receiver to negotiate a window scale factor. This number allows both sides to shift the *Window size* field up to 16 bits to the left, thus allowing windows of up to 2^{32} bytes. Most TCP implementations now support this option.

Another option proposed by RFC 1106 and now widely implemented is the use of the selective repeat instead of go back n protocol. If the receiver gets one bad segment and then a large number of good ones, the normal TCP protocol will

eventually time out and retransmit all the unacknowledged segments, including all those that were received correctly. RFC 1106 introduced NAKs, to allow the receiver to ask for a specific segment (or segments). After it gets these, it can acknowledge all the buffered data, thus reducing the amount of data retransmitted.

6.4.4. TCP Connection Management

Connections are established in TCP using the three-way handshake discussed in Sec. 6.2.2. To establish a connection, one side, say the server, passively waits for an incoming connection by executing the LISTEN and ACCEPT primitives, either specifying a specific source or nobody in particular.

The other side, say the client, executes a CONNECT primitive, specifying the IP address and port to which it wants to connect, the maximum TCP segment size it is willing to accept, and optionally some user data (e.g., a password). The CONNECT primitive sends a TCP segment with the *SYN* bit on and *ACK* bit off and waits for a response.

When this segment arrives at the destination, the TCP entity there checks to see if there is a process that has done a LISTEN on the port given in the *Destination port* field. If not, it sends a reply with the *RST* bit on to reject the connection.




Fig. 6-26. (a) TCP connection establishment in the normal case. (b) Call collision.

If some process is listening to the port, that process is given the incoming TCP segment. It can then either accept or reject the connection. If it accepts, an acknowledgement segment is sent back. The sequence of TCP segments sent in the normal case is shown in Fig. 6-26(a). Note that a *SYN* segment consumes 1 byte of sequence space so it can be acknowledged unambiguously.

In the event that two hosts simultaneously attempt to establish a connection between the same two sockets, the sequence of events is as illustrated in Fig. 6-26(b). The result of these events is that just one connection is established, not two because connections are identified by their end points. If the first setup results in a connection identified by (x, y) and the second one does too, only one table entry is made, namely, for (x, y) .

The initial sequence number on a connection is not 0 for the reasons we discussed earlier. A clock-based scheme is used, with a clock tick every 4 μ sec. For additional safety, when a host crashes, it may not reboot for the maximum packet lifetime (120 sec) to make sure that no packets from previous connections are still roaming around the Internet somewhere.

Although TCP connections are full duplex, to understand how connections are released it is best to think of them as a pair of simplex connections. Each simplex connection is released independently of its sibling. To release a connection, either party can send a TCP segment with the *FIN* bit set, which means that it has no more data to transmit. When the *FIN* is acknowledged, that direction is shut down for new data. Data may continue to flow indefinitely in the other direction, however. When both directions have been shut down, the connection is released. Normally, four TCP segments are needed to release a connection, one *FIN* and one *ACK* for each direction. However, it is possible for the first *ACK* and the second *FIN* to be contained in the same segment, reducing the total count to three.

Just as with telephone calls in which both people say goodbye and hang up the phone simultaneously, both ends of a TCP connection may send *FIN* segments at the same time. These are each acknowledged in the usual way, and the connection shut down. There is, in fact, no essential difference between the two hosts releasing sequentially or simultaneously.

To avoid the two-army problem, timers are used. If a response to a *FIN* is not forthcoming within two maximum packet lifetimes, the sender of the *FIN* releases the connection. The other side will eventually notice that nobody seems to be listening to it any more, and time out as well. While this solution is not perfect, given the fact that a perfect solution is theoretically impossible, it will have to do. In practice, problems rarely arise.

The steps required to establish and release connections can be represented in a finite state machine with the 11 states listed in Fig. 6-27. In each state, certain events are legal. When a legal event happens, some action may be taken. If some other event happens, an error is reported.

Each connection starts in the *CLOSED* state. It leaves that state when it does either a passive open (*LISTEN*), or an active open (*CONNECT*). If the other side does the opposite one, a connection is established and the state becomes *ESTABLISHED*. Connection release can be initiated by either side. When it is complete, the state returns to *CLOSED*.

The finite state machine itself is shown in Fig. 6-28. The common case of a client actively connecting to a passive server is shown with heavy lines—solid for

State	Description
CLOSED	No connection is active or pending
LISTEN	The server is waiting for an incoming call
SYN RCVD	A connection request has arrived; wait for ACK
SYN SENT	The application has started to open a connection
ESTABLISHED	The normal data transfer state
FIN WAIT 1	The application has said it is finished
FIN WAIT 2	The other side has agreed to release
TIMED WAIT	Wait for all packets to die off
CLOSING	Both sides have tried to close simultaneously
CLOSE WAIT	The other side has initiated a release
LAST ACK	Wait for all packets to die off

Fig. 6-27. The states used in the TCP connection management finite state machine.

the client, dotted for the server. The lightface lines are unusual event sequences. Each line in Fig. 6-28 is marked by an *event/action* pair. The event can either be a user-initiated system call (CONNECT, LISTEN, SEND, or CLOSE), a segment arrival (*SYN*, *FIN*, *ACK*, or *RST*), or in one case, a timeout of twice the maximum packet lifetime. The action is the sending of a control segment (*SYN*, *FIN*, or *RST*) or nothing, indicated by —. Comments are shown in parentheses.

The diagram can best be understood by first following the path of a client (the heavy solid line) then later the path of a server (the heavy dashed line). When an application on the client machine issues a CONNECT request, the local TCP entity creates a connection record, marks it as being in the *SYN SENT* state, and sends a *SYN* segment. Note that many connections may be open (or being opened) at the same time on behalf of multiple applications, so the state is per connection and recorded in the connection record. When the *SYN+ACK* arrives, TCP sends the final *ACK* of the three-way handshake and switches into the *ESTABLISHED* state. Data can now be sent and received.

When an application is finished, it executes a CLOSE primitive, which causes the local TCP entity to send a *FIN* segment and wait for the corresponding *ACK* (dashed box marked active close). When the *ACK* arrives, a transition is made to state *FIN WAIT 2* and one direction of the connection is now closed. When the other side closes, too, a *FIN* comes in, which is acknowledged. Now both sides are closed, but TCP waits a time equal to the maximum packet lifetime to guarantee that all packets from the connection have died off, just in case the acknowledgement was lost. When the timer goes off, TCP deletes the connection record.




Fig. 6-28. TCP connection management finite state machine. The heavy solid line is the normal path for a client. The heavy dashed line is the normal path for a server. The light lines are unusual events.

Now let us examine connection management from the server's viewpoint. The server does a LISTEN and settles down to see who turns up. When a SYN comes in, it is acknowledged and the server goes to the SYN RCVD state. When the server's SYN is itself acknowledged, the three-way handshake is complete and the server goes to the ESTABLISHED state. Data transfer can now occur.

When the client has had enough, it does a CLOSE, which causes a FIN to arrive at the server (dashed box marked passive close). The server is then

signaled. When it, too, does a CLOSE, a *FIN* is sent to the client. When the client's acknowledgement shows up, the server releases the connection and deletes the connection record.

6.4.5. TCP Transmission Policy

Window management in TCP is not directly tied to acknowledgements as it is in most data link protocols. For example, suppose the receiver has a 4096-byte buffer as shown in Fig. 6-29. If the sender transmits a 2048-byte segment that is correctly received, the receiver will acknowledge the segment. However, since it now has only 2048 of buffer space (until the application removes some data from the buffer), it will advertise a window of 2048 starting at the next byte expected.




Fig. 6-29. Window management in TCP.

Now the sender transmits another 2048 bytes, which are acknowledged, but the advertised window is 0. The sender must stop until the application process on

the receiving host has removed some data from the buffer, at which time TCP can advertise a larger window.

When the window is 0, the sender may not normally send segments, with two exceptions. First, urgent data may be sent, for example, to allow the user to kill the process running on the remote machine. Second, the sender may send a 1-byte segment to make the receiver reannounce the next byte expected and window size. The TCP standard explicitly provides this option to prevent deadlock if a window announcement ever gets lost.

Senders are not required to transmit data as soon as they come in from the application. Neither are receivers required to send acknowledgements as soon as possible. For example, in Fig. 6-29, When the first 2 KB of data came in, TCP, knowing that it had a 4-KB window available, would have been completely correct in just buffering the data until another 2 KB came in, to be able to transmit a segment with a 4-KB payload. This freedom can be exploited to improve performance.

Consider a TELNET connection to an interactive editor that reacts on every keystroke. In the worst case, when a character arrives at the sending TCP entity, TCP creates a 21-byte TCP segment, which it gives to IP to send as a 41-byte IP datagram. At the receiving side, TCP immediately sends a 40-byte acknowledgement (20 bytes of TCP header and 20 bytes of IP header). Later, when the editor has read the byte, TCP sends a window update, moving the window 1 byte to the right. This packet is also 40 bytes. Finally, when the editor has processed the character, it echoes it as a 41-byte packet. In all, 162 bytes of bandwidth are used and four segments are sent for each character typed. When bandwidth is scarce, this method of doing business is not desirable.

One approach that many TCP implementations use to optimize this situation is to delay acknowledgements and window updates for 500 msec in the hope of acquiring some data on which to hitch a free ride. Assuming the editor echoes within 500 msec, only one 41-byte packet now need be sent back to the remote user, cutting the packet count and bandwidth usage in half.

Although this rule reduces the load placed on the network by the receiver, the sender is still operating inefficiently by sending 41-byte packets containing 1 byte of data. A way to reduce this usage is known as **Nagle's algorithm** (Nagle, 1984). What Nagle suggested is simple: when data come into the sender one byte at a time, just send the first byte and buffer all the rest until the outstanding byte is acknowledged. Then send all the buffered characters in one TCP segment and start buffering again until they are all acknowledged. If the user is typing quickly and the network is slow, a substantial number of characters may go in each segment, greatly reducing the bandwidth used. The algorithm additionally allows a new packet to be sent if enough data have trickled in to fill half the window or a maximum segment.

Nagle's algorithm is widely used by TCP implementations, but there are times when it is better to disable it. In particular, when an X-Windows application is

being run over the Internet, mouse movements have to be sent to the remote computer. Gathering them up to send in bursts makes the mouse cursor move erratically, which makes for unhappy users.

Another problem that can ruin TCP performance is the **silly window syndrome** (Clark, 1982). This problem occurs when data are passed to the sending TCP entity in large blocks, but an interactive application on the receiving side reads data 1 byte at a time. To see the problem, look at Fig. 6-30. Initially, the TCP buffer on the receiving side is full and the sender knows this (i.e., has a window of size 0). Then the interactive application reads one character from the TCP stream. This action makes the receiving TCP happy, so it sends a window update to the sender saying that it is all right to send 1 byte. The sender obliges and sends 1 byte. The buffer is now full, so the receiver acknowledges the 1-byte segment but sets the window to 0. This behavior can go on forever.




Fig. 6-30. Silly window syndrome.

Clark's solution is to prevent the receiver from sending a window update for 1 byte. Instead it is forced to wait until it has a decent amount of space available and advertise that instead. Specifically, the receiver should not send a window update until it can handle the maximum segment size it advertised when the connection was established, or its buffer is half empty, whichever is smaller.

Furthermore, the sender can also help by not sending tiny segments. Instead, it should try to wait until it has accumulated enough space in the window to send a full segment or at least one containing half of the receiver's buffer size (which it must estimate from the pattern of window updates it has received in the past).

Nagle's algorithm and Clark's solution to the silly window syndrome are complementary. Nagle was trying to solve the problem caused by the sending application delivering data to TCP a byte at a time. Clark was trying to solve the problem of the receiving application sucking the data up from TCP a byte at a time. Both solutions are valid and can work together. The goal is for the sender not to send small segments and the receiver not to ask for them.

The receiving TCP can go further in improving performance than just doing window updates in large units. Like the sending TCP, it also has the ability to buffer data, so it can block a READ request from the application until it has a large chunk of data to provide. Doing this reduces the number of calls to TCP, and hence the overhead. Of course, it also increases the response time, but for noninteractive applications like file transfer, efficiency may outweigh response time to individual requests.

Another receiver issue is what to do with out of order segments. They can be kept or discarded, at the receiver's discretion. Of course, acknowledgements can be sent only when all the data up to the byte acknowledged have been received. If the receiver gets segments 0, 1, 2, 4, 5, 6, and 7, it can acknowledge everything up to and including the last byte in segment 2. When the sender times out, it then retransmits segment 3. If the receiver has buffered segments 4 through 7, upon receipt of segment 3 it can acknowledge all bytes up to the end of segment 7.

6.4.6. TCP Congestion Control

When the load offered to any network is more than it can handle, congestion builds up. The Internet is no exception. In this section we will discuss algorithms that have been developed over the past decade to deal with congestion. Although the network layer also tries to manage congestion, most of the heavy lifting is done by TCP because the real solution to congestion is to slow down the data rate.

In theory, congestion can be dealt with by employing a principle borrowed from physics: the law of conservation of packets. The idea is not to inject a new packet into the network until an old one leaves (i.e., is delivered). TCP attempts to achieve this goal by dynamically manipulating the window size.

The first step in managing congestion is detecting it. In the old days, detecting congestion was difficult. A timeout caused by a lost packet could have been caused by either (1) noise on a transmission line or (2) packet discard at a congested router. Telling the difference was difficult.

Nowadays, packet loss due to transmission errors is relatively rare because most long-haul trunks are fiber (although wireless networks are a different story). Consequently, most transmission timeouts on the Internet are due to congestion. All the Internet TCP algorithms assume that timeouts are caused by congestion and monitor timeouts for signs of trouble the way miners watch their canaries.

Before discussing how TCP reacts to congestion, let us first describe what it does to try to prevent it from occurring in the first place. When a connection is

established, a suitable window size has to be chosen. The receiver can specify a window based on its buffer size. If the sender sticks to this window size, problems will not occur due to buffer overflow at the receiving end, but they may still occur due to internal congestion within the network.

In Fig. 6-31, we see this problem illustrated hydraulically. In Fig. 6-31(a), we see a thick pipe leading to a small-capacity receiver. As long as the sender does not send more water than the bucket can contain, no water will be lost. In Fig. 6-31(b), the limiting factor is not the bucket capacity, but the internal carrying capacity of the network. If too much water comes in too fast, it will back up and some will be lost (in this case by overflowing the funnel).




Fig. 6-31. (a) A fast network feeding a low-capacity receiver. (b) A slow network feeding a high-capacity receiver.

The Internet solution is to realize that two potential problems exist—network capacity and receiver capacity—and to deal with each of them separately. To do so, each sender maintains two windows: the window the receiver has granted and a second window, the **congestion window**. Each reflects the number of bytes the sender may transmit. The number of bytes that may be sent is the minimum of the two windows. Thus the effective window is the minimum of what the sender

thinks is all right and what the receiver thinks is all right. If the receiver says “Send 8K” but the sender knows that bursts of more than 4K clog the network up, it sends 4K. On the other hand, if the receiver says “Send 8K” and the sender knows that bursts of up to 32K get through effortlessly, it sends the full 8K requested.

When a connection is established, the sender initializes the congestion window to the size of the maximum segment in use on the connection. It then sends one maximum segment. If this segment is acknowledged before the timer goes off, it adds one segment’s worth of bytes to the congestion window to make it two maximum size segments and sends two segments. As each of these segments is acknowledged, the congestion window is increased by one maximum segment size. When the congestion window is n segments, if all n are acknowledged on time, the congestion window is increased by the byte count corresponding to n segments. In effect, each burst successfully acknowledged doubles the congestion window.

The congestion window keeps growing exponentially until either a timeout occurs or the receiver’s window is reached. The idea is that if bursts of size, say, 1024, 2048, and 4096 bytes work fine, but a burst of 8192 bytes gives a timeout, the congestion window should be set to 4096 to avoid congestion. As long as the congestion window remains at 4096, no bursts longer than that will be sent, no matter how much window space the receiver grants. This algorithm is called **slow start**, but it is not slow at all (Jacobson, 1988). It is exponential. All TCP implementations are required to support it.

Now let us look at the Internet congestion control algorithm. It uses a third parameter, the **threshold**, initially 64K, in addition to the receiver and congestion windows. When a timeout occurs, the threshold is set to half of the current congestion window, and the congestion window is reset to one maximum segment. Slow start is then used to determine what the network can handle, except that exponential growth stops when the threshold is hit. From that point on, successful transmissions grow the congestion window linearly (by one maximum segment for each burst) instead of one per segment. In effect, this algorithm is guessing that it is probably acceptable to cut the congestion window in half, and then it gradually works its way up from there.

As an illustration of how the congestion algorithm works, see Fig. 6-32. The maximum segment size here is 1024 bytes. Initially the congestion window was 64K, but a timeout occurred, so the threshold is set to 32K and the congestion window to 1K for transmission 0 here. The congestion window then grows exponentially until it hits the threshold (32K). Starting then it grows linearly.

Transmission 13 is unlucky (it should have known) and a timeout occurs. The threshold is set to half the current window (by now 40K, so half is 20K) and slow start initiated all over again. When the acknowledgements from transmission 18 start coming in, the first four each increment the congestion window by one segment, but after that, growth becomes linear again.




Fig. 6-32. An example of the Internet congestion algorithm.

If no more timeouts occur, the congestion window will continue to grow up to the size of the receiver's window. At that point, it will stop growing and remain constant as long as there are no more timeouts and the receiver's window does not change size. As an aside, if an ICMP SOURCE QUENCH packet comes in and is passed to TCP, this event is treated the same way as a timeout.

Work on improving the congestion control mechanism is continuing. For example, Brakmo et al. (1994) have reported improving TCP throughput by 40 percent to 70 percent by managing the clock more accurately, predicting congestion before timeouts occur, and using this early warning system to improve the slow start algorithm.

6.4.7. TCP Timer Management

TCP uses multiple timers (at least conceptually) to do its work. The most important of these is the **retransmission timer**. When a segment is sent, a retransmission timer is started. If the segment is acknowledged before the timer expires, the timer is stopped. If, on the other hand, the timer goes off before the acknowledgement comes in, the segment is retransmitted (and the timer started again). The question that arises is: How long should the timeout interval be?

This problem is much more difficult in the Internet transport layer than in the generic data link protocols of Chap. 3. In the latter case, the expected delay is highly predictable (i.e., has a low variance), so the timer can be set to go off just slightly after the acknowledgement is expected, as shown in Fig. 6-33(a). Since acknowledgements are rarely delayed in the data link layer, the absence of an acknowledgement at the expected time generally means the frame or the acknowledgement has been lost.




Fig. 6-33. (a) Probability density of acknowledgement arrival times in the data link layer. (b) Probability density of acknowledgement arrival times for TCP.

TCP is faced with a radically different environment. The probability density function for the time it takes for a TCP acknowledgement to come back looks more like Fig. 6-33(b) than Fig. 6-33(a). Determining the round-trip time to the destination is tricky. Even when it is known, deciding on the timeout interval is also difficult. If the timeout is set too short, say T_1 in Fig. 6-33(b), unnecessary retransmissions will occur, clogging the Internet with useless packets. If it is set too long, (T_2), performance will suffer due to the long retransmission delay whenever a packet is lost. Furthermore, the mean and variance of the acknowledgement arrival distribution can change rapidly within a few seconds as congestion builds up or is resolved.

The solution is to use a highly dynamic algorithm that constantly adjusts the timeout interval, based on continuous measurements of network performance. The algorithm generally used by TCP is due to Jacobson (1988) and works as follows. For each connection, TCP maintains a variable, RTT , that is the best current estimate of the round-trip time to the destination in question. When a segment is sent, a timer is started, both to see how long the acknowledgement takes and to

trigger a retransmission if it takes too long. If the acknowledgement gets back before the timer expires, TCP measures how long the acknowledgement took, say, M . It then updates RTT according to the formula

$$RTT = \alpha RTT + (1 - \alpha)M$$

where α is a smoothing factor that determines how much weight is given to the old value. Typically $\alpha = 7/8$.

Even given a good value of RTT , choosing a suitable retransmission timeout is a nontrivial matter. Normally, TCP uses βRTT , but the trick is choosing β . In the initial implementations, β was always 2, but experience showed that a constant value was inflexible because it failed to respond when the variance went up.

In 1988, Jacobson proposed making β roughly proportional to the standard deviation of the acknowledgement arrival time probability density function so a large variance means a large β and vice versa. In particular, he suggested using the *mean deviation* as a cheap estimator of the *standard deviation*. His algorithm requires keeping track of another smoothed variable, D , the deviation. Whenever an acknowledgement comes in, the difference between the expected and observed values, $|RTT - M|$ is computed. A smoothed value of this is maintained in D by the formula

$$D = \alpha D + (1 - \alpha) |RTT - M|$$

where α may or may not be the same value used to smooth RTT . While D is not exactly the same as the standard deviation, it is good enough and Jacobson showed how it could be computed using only integer adds, subtracts, and shifts, a big plus. Most TCP implementations now use this algorithm and set the timeout interval to

$$\text{Timeout} = RTT + 4 * D$$

The choice of the factor 4 is somewhat arbitrary, but it has two advantages. First, multiplication by 4 can be done with a single shift. Second, it minimizes unnecessary timeouts and retransmissions because less than one percent of all packets come in more than four standard deviations late. (Actually, Jacobson initially said to use 2, but later work has shown that 4 gives better performance.)

One problem that occurs with the dynamic estimation of RTT is what to do when a segment times out and is sent again. When the acknowledgement comes in, it is unclear whether the acknowledgement refers to the first transmission or a later one. Guessing wrong can seriously contaminate the estimate of RTT . Phil Karn discovered this problem the hard way. He is an amateur radio enthusiast interested in transmitting TCP/IP packets by ham radio, a notoriously unreliable medium (on a good day, half the packets get through). He made a simple proposal: do not update RTT on any segments that have been retransmitted. Instead, the timeout is doubled on each failure until the segments get through the first time. This fix is called **Karn's algorithm**. Most TCP implementations use it.

The retransmission timer is not the only one TCP uses. A second timer is the **persistence timer**. It is designed to prevent the following deadlock. The receiver sends an acknowledgement with a window size of 0, telling the sender to wait. Later, the receiver updates the window, but the packet with the update is lost. Now both the sender and the receiver are waiting for each other to do something. When the persistence timer goes off, the sender transmits a probe to the receiver. The response to the probe gives the window size. If it is still zero, the persistence timer is set again and the cycle repeats. If it is nonzero, data can now be sent.

A third timer that some implementations use is the **keepalive timer**. When a connection has been idle for a long time, the keepalive timer may go off to cause one side to check if the other side is still there. If it fails to respond, the connection is terminated. This feature is controversial because it adds overhead and may terminate an otherwise healthy connection due to a transient network partition.

The last timer used on each TCP connection is the one used in the *TIMED WAIT* state while closing. It runs for twice the maximum packet lifetime to make sure that when a connection is closed, all packets created by it have died off.

6.4.8. UDP

The Internet protocol suite also supports a connectionless transport protocol, **UDP (User Data Protocol)**. UDP provides a way for applications to send encapsulated raw IP datagrams and send them without having to establish a connection. Many client-server applications that have one request and one response use UDP rather than go to the trouble of establishing and later releasing a connection. UDP is described in RFC 768.




Fig. 6-34. The UDP header.

A UDP segment consists of an 8-byte header followed by the data. The header is shown in Fig. 6-34. The two ports serve the same function as they do in TCP: to identify the end points within the source and destination machines. The *UDP length* field includes the 8-byte header and the data. The *UDP checksum* includes the same format pseudoheader shown in Fig. 6-25, the UDP header, and the UDP data, padded out to an even number of bytes if need be. It is optional and stored as 0 if not computed (a true computed 0 is stored as all 1s, which is the same in 1's complement). Turning it off is foolish unless the quality of the data does not matter (e.g., digitized speech).

6.4.9. Wireless TCP and UDP

In theory, transport protocols should be independent of the technology of the underlying network layer. In particular, TCP should not care whether IP is running over fiber or over radio. In practice, it does matter because most TCP implementations have been carefully optimized based on assumptions that are true for wired networks but which fail for wireless networks. Ignoring the properties of wireless transmission can lead to a TCP implementation that is logically correct but has horrendous performance.

The principal problem is the congestion control algorithm. Nearly all TCP implementations nowadays assume that timeouts are caused by congestion, not by lost packets. Consequently, when a timer goes off, TCP slows down and sends less vigorously (e.g., Jacobson's slow start algorithm). The idea behind this approach is to reduce the network load and thus alleviate the congestion.

Unfortunately, wireless transmission links are highly unreliable. They lose packets all the time. The proper approach to dealing with lost packets is to send them again, and as quickly as possible. Slowing down just makes matters worse. If, say, 20 percent of all packets are lost, then when the sender transmits 100 packets/sec, the throughput is 80 packets/sec. If the sender slows down to 50 packets/sec, the throughput drops to 40 packets/sec.

In effect, when a packet is lost on a wired network, the sender should slow down. When one is lost on a wireless network, the sender should try harder. When the sender does not know what the network is, it is difficult to make the correct decision.

Frequently, the path from sender to receiver is inhomogeneous. The first 1000 km might be over a wired network, but the last 1 km might be wireless. Now making the correct decision on a timeout is even harder, since it matters where the problem occurred. A solution proposed by Bakne and Badrinath (1995), **indirect TCP**, is to split the TCP connection into two separate connections, as shown in Fig. 6-35. The first connection goes from the sender to the base station. The second one goes from the base station to the receiver. The base station simply copies packets between the connections in both directions.




Fig. 6-35. Splitting a TCP connection into two connections.

The advantage of this scheme is that both connections are now homogeneous. Timeouts on the first connection can slow the sender down, whereas timeouts on the second one can speed it up. Other parameters can also be tuned separately for the two connections. The disadvantage is that it violates the semantics of TCP. Since each part of the connection is a full TCP connection, the base station acknowledges each TCP segment in the usual way. Only now, receipt of an acknowledgement by the sender does not mean that the receiver got the segment, only that the base station got it.

A different solution, due to Balakrishnan et al. (1995), does not break the semantics of TCP. It works by making several small modifications to the network layer code in the base station. One of the changes is the addition of a snooping agent that observes and caches TCP segments going out to the mobile host, and acknowledgements coming back from it. When the snooping agent sees a TCP segment going out to the mobile host but does not see an acknowledgement coming back before its (relatively short) timer goes off, it just retransmits that segment, without telling the source that it is doing so. It also generates a retransmission when it sees duplicate acknowledgements from the mobile host go by, invariably meaning that the mobile host has missed something. Duplicate acknowledgements are discarded on the spot, to avoid having the source misinterpret them as a sign of congestion.

One disadvantage of this transparency, however, is that if the wireless link is very lossy, the source may time out waiting for an acknowledgement and invoke the congestion control algorithm. With indirect TCP, the congestion control algorithm will never be started unless there really is congestion in the wired part of the network.

The Balakrishnan et al. paper also has a solution to the problem of lost segments originating at the mobile host. When the base station notices a gap in the inbound sequence numbers, it generates a request for a selective repeat of the missing bytes using a TCP option. Using these two fixes, the wireless link is made more reliable in both directions, without the source knowing about it, and without changing the semantics of TCP.

While UDP does not suffer from the same problems as TCP, wireless communication also introduces difficulties for it. The main trouble is that programs use UDP expecting it to be highly reliable. They know that no guarantees are given, but they still expect it to be near perfect. In a wireless environment, it will be far from perfect. For programs that are able to recover from lost UDP messages, but only at considerable cost, suddenly going from an environment where messages theoretically can be lost but rarely are, to one in which they are constantly being lost can result in a performance disaster.

Wireless communication also affects areas other than just performance. For example, how does a mobile host find a local printer to connect to, rather than use its home printer? Somewhat related to this is how to get the WWW page for the local cell, even if its name is not known. Also, WWW page designers tend to

assume lots of bandwidth is available. Putting a large logo on every page becomes counterproductive if it is going to take 30 sec to transmit at 9600 bps every time the page is referenced, irritating the users no end.

6.5. THE ATM AAL LAYER PROTOCOLS

It is not really clear whether or not ATM has a transport layer. On the one hand, the ATM layer has the functionality of a network layer, and there is another layer on top of it (AAL), which sort of makes AAL a transport layer. Some experts agree with this view (e.g., De Prycker, 1993, page 112). One of the protocols used here (AAL 5) is functionally similar to UDP, which is unquestionably a transport protocol.

On the other hand, none of the AAL protocols provide a reliable end-to-end connection, as TCP does (although with only very minor changes they could). Also, in most applications another transport layer is used on top of AAL. Rather than split hairs, we will discuss the AAL layer and its protocols in this chapter without making a claim that it is a true transport layer.

The AAL layer in ATM networks is radically different than TCP, largely because the designers were primarily interested in transmitting voice and video streams, in which rapid delivery is more important than accurate delivery. Remember that the ATM layer just outputs 53-byte cells one after another. It has no error control, no flow control, and no other control. Consequently, it is not well matched to the requirements that most applications need.

To bridge this gap, in Recommendation I.363, ITU has defined an end-to-end layer on top of the ATM layer. This layer, called **AAL (ATM Adaptation Layer)** has a tortuous history, full of mistakes, revisions, and unfinished business. In the following sections we will look at it and its design.

The goal of AAL is to provide useful services to application programs and to shield them from the mechanics of chopping data up into cells at the source and reassembling them at the destination. When ITU began defining AAL, it realized that different applications had different requirements, so it organized the service space along three axes:

1. Real-time service versus nonreal-time service.
2. Constant bit rate service versus variable bit rate service.
3. Connection-oriented service versus connectionless service.

In principle, with three axes and two values on each axis, eight distinct services can be defined, as shown in Fig. 6-36. ITU felt that only four of these were of any use, and named them classes A, B, C, and D, as noted. The others were not supported. Starting with ATM 4.0, Fig. 6-36 is somewhat obsolete, so it has been presented here mostly as background information to help understand why the

AAL protocols have been designed as they have been. Instead of these service classes, the major distinction now is between the traffic classes we studied in Chap. 5 (ABR, CBR, NRT-VBR, RT-VBR, and UBR).

	A		B		C		D	
Timing	Real time	None	Real time	None	Real time	None	Real time	None
Bit rate	Constant		Variable		Constant		Variable	
Mode	Connection orientated				Connectionless			

Fig. 6-36. Original service classes supported by AAL (now obsolete).

To handle these four classes of service, ITU defined four protocols, AAL 1 through AAL 4, respectively. However, later it discovered that the technical requirements for classes C and D were so similar that AAL 3 and AAL 4 were combined into AAL 3/4. Then the computer industry, which had been asleep at the switch, realized that none of them were any good. It solved this problem by the simple expedient of defining another protocol, AAL 5. We will look at all four of these shortly. We will also look at an interesting control protocol used on ATM systems.

6.5.1. Structure of the ATM Adaptation Layer

The ATM adaptation layer is divided into two major parts, one of which is often further subdivided, as illustrated in Fig. 6-37.

The upper part of the ATM adaptation layer is called the **convergence sublayer**. Its job is to provide the interface to the application. It consists of a subpart that is common to all applications (for a given AAL protocol) and an application specific subpart. The functions of each of these parts are protocol dependent but can include message framing and error detection.

In addition, at the source, the convergence sublayer is responsible for accepting bit streams or arbitrary length messages from the applications and breaking them up into units of 44 to 48 bytes for transmission. The exact size is protocol dependent, since some protocols use part of the 48-byte ATM payload for their own headers. At the destination, this sublayer reassembles the cells into the original messages. In general, message boundaries are preserved, when present. In other words, if the source sends four 512-byte messages, they will arrive as four 512-byte messages, not one 2048-byte message. For data streams, no message boundaries exist, so they are not preserved.

The lower part of the AAL is called the **SAR (Segmentation And Reassembly)** sublayer. It can add headers and trailers to the data units given to it by the




Fig. 6-37. The ATM model showing the ATM adaptation layer and its sublayers.

convergence sublayer to form cell payloads. These payloads are then given to the ATM layer for transmission. At the destination, the SAR sublayer reassembles the cells into messages. The SAR sublayer is basically concerned with cells, whereas the convergence sublayer is concerned with messages.

The generic operation of the convergence and SAR sublayers is shown in Fig. 6-38. When a message comes in to the AAL from the application, the convergence sublayer may give it a header and/or trailer. The message is then broken up into 44- to 48-byte units, which are passed to the SAR sublayer. The SAR sublayer may add its own header and trailer to each piece and pass them down to the ATM layer for transmission as independent cells. Note that the figure shows the most general case because some of the AAL protocols have null headers and/or trailers.

The SAR sublayer also has some additional functions for some (but not all) service classes. In particular, it sometimes handles error detection and multiplexing. The SAR sublayer is present for all service classes but does more or less work, depending on the specific protocol.

The communication between the application and AAL layer uses the standard OSI *request* and *indication* primitives that we discussed in Chap. 1. The communication between the sublayers uses different primitives.

6.5.2. AAL 1

AAL 1 is the protocol used for transmitting class A traffic, that is, real-time, constant bit rate, connection-oriented traffic, such as uncompressed audio and video. Bits are fed in by the application at a constant rate and must be delivered at the far end at the same constant rate, with a minimum of delay, jitter, and overhead. The input is a stream of bits, with no message boundaries. For this traffic, error detecting protocols such as stop-and-wait are not used because the delays that are introduced by timeouts and retransmissions are unacceptable. However,




Fig. 6-38. The headers and trailers that can be added to a message in an ATM network.

missing cells are reported to the application, which must then take its own action (if any) to recover from them.

AAL 1 uses a convergence sublayer and a SAR sublayer. The convergence sublayer detects lost and misinserted cells. (A misinserted cell is one that is delivered to the wrong destination as a result of an undetected error in its virtual circuit or virtual path identifiers.) It also smoothes out incoming traffic to provide delivery of cells at a constant rate. Finally, the convergence sublayer breaks up the input messages or stream into 46- or 47-byte units that are given to the SAR sublayer for transmission. At the other end it extracts these and reconstructs the original input. The AAL 1 convergence sublayer does not have any protocol headers of its own.

In contrast, the AAL 1 SAR sublayer does have a protocol. The formats of its cells are given in Fig. 6-39. Both formats begin with a 1-byte header containing a 3-bit cell sequence number, *SN*, (to detect missing or misinserted cells). This field is followed by a 3-bit sequence number protection, *SNP*, (i.e., checksum) over the sequence number to allow correction of single errors and detection of double errors in the sequence field. It uses a cyclic redundancy check with the polynomial $x^3 + x + 1$. An even parity bit covering the header byte further reduces the likelihood of a bad sequence number sneaking in unnoticed. AAL 1 cells need not be filled with a full 47 bytes. For example, to transmit digitized voice arriving at a rate of 1 byte every 125 μ sec, filling a cell with 47 bytes means collecting samples for 5.875 msec. If this delay before transmission is




Fig. 6-39. The AAL 1 cell format.

unacceptable, partial cells can be sent. In this case, the number of actual data bytes per cell is the same for all cells and agreed on in advance.

The *P* cells are used when message boundaries must be preserved. The *Pointer* field is used to give the offset of the start of the next message. Only cells with an even sequence number may be *P* cells, so the pointer is in the range 0 to 92, to put it within the payload of either its own cell or the one following it. Note that this scheme allows messages to be an arbitrary number of bytes long, so messages can be run continuously and need not align on cell boundaries.

The high-order bit of the *Pointer* field is reserved for future use. The initial header bit of all the odd-numbered cells forms a data stream used for clock synchronization.

6.5.3. AAL 2

AAL 1 is designed for simple, connection-oriented, real-time data streams without error detection, except for missing and misinserted cells. For pure uncompressed audio or video, or any other data stream in which having a few garbled bits once in a while is not a problem, AAL 1 is adequate.

For compressed audio or video, the rate can vary strongly in time. For example, many compression schemes transmit a full video frame periodically and then send only the differences between subsequent frames and the last full frame for several frames. When the camera is stationary and nothing is moving, the difference frames are small, but when the camera is panning rapidly, they are large. Also, message boundaries must be preserved so that the start of the next full frame can be recognized, even in the presence of lost cells or bad data. For these reasons, a fancier protocol is needed. AAL 2 has been designed for this purpose.

As in AAL 1, the CS sublayer does not have a protocol but the SAR sublayer does. The SAR cell format is shown in Fig. 6-40. It has a 1-byte header and a 2-byte trailer, leaving room for up to 45 data bytes per cell.

The *SN* field (*Sequence Number*) is used for numbering cells in order to detect missing or misinserted cells. The *IT* field (*Information Type*) is used to indicate




Fig. 6-40. The AAL 2 cell format.

that the cell is the start, middle, or end of a message. The *LI* (*Length indicator*) field tells how big the payload is, in bytes (it might be less than 45 bytes). Finally, the *CRC* field is a checksum over the entire cell, so errors can be detected.

Strange as it may sound, the field sizes are not included in the standard. According to one insider, at the very end of the standardization process the committee realized that AAL 2 had so many problems that it should not be used. Unfortunately, it was too late to stop the standardization process. They had a deadline to meet. In a last ditch effort, the committee removed all the field sizes so that the formal standard could be issued on time, but in such a way that nobody could actually use it. Such is life in the world of standardization.

6.5.4. AAL 3/4

Originally, ITU had different protocols for classes C and D, connection-oriented service and connectionless service for data transport that is sensitive to loss or errors but is not time dependent. Then ITU discovered that there was no real need for two protocols, so they were combined into a single protocol, AAL 3/4.

AAL 3/4 can operate in two modes: stream or message. In message mode, each call from the application to AAL 3/4 injects one message into the network. The message is delivered as such, that is, message boundaries are preserved. In stream mode the boundaries are not preserved. The discussion below will concentrate on message mode. Reliable and unreliable (i.e., no guarantee) transport are available in each mode.

A feature of AAL 3/4 not present in any of the other protocols is multiplexing. This aspect of AAL 3/4 allows multiple sessions (e.g., remote logins) from a single host to travel along the same virtual circuit and be separated at the destination, as illustrated in Fig. 6-41.

The reason that this facility is desirable is that carriers often charge for each connection setup and for each second that a connection is open. If a pair of hosts have several sessions open simultaneously, giving each one its own virtual circuit will be more expensive than multiplexing all of them onto the same virtual circuit. If one virtual circuit has sufficient bandwidth to handle the job, there is no need




Fig. 6-41. Multiplexing of several sessions onto one virtual circuit.

for more than one. All sessions using a single virtual circuit get the same quality of service, since this is negotiated per virtual circuit.

This issue is the real reason that there were originally separate AAL 3 and AAL 4 formats: the Americans wanted multiplexing and the Europeans did not. So each group went off and made its own standard. Eventually, the Europeans decided that saving 10 bits in the header was not worth the price of having the United States and Europe not be able to communicate. For the same money, they could have stuck to their guns and we would have had four incompatible AAL standards (of which one is broken) instead of three.

Unlike AAL 1 and AAL 2, AAL 3/4 has both a convergence sublayer protocol and a SAR sublayer protocol. Messages as large as 65,535 bytes come into the convergence sublayer from the application. These are first padded out to a multiple of 4 bytes. Then a header and a trailer are attached, as shown in Fig. 6-42.




Fig. 6-42. AAL 3/4 convergence sublayer message format.

The *CPI* field (*Common Part Indicator*) gives the message type and the counting unit for the *BA size* and *Length* fields. The *Btag* and *Etag* fields are used to frame messages. The two bytes must be the same and are incremented by one on every new message sent. This mechanism checks for lost or misinserted cells. The *BA size* field is used for buffer allocation. It tells the receiver how much buffer space to allocate for the message in advance of its arrival. The *Length* field gives the payload length again. In message mode, it must be equal to *BA size*, but in stream mode it may be different. The trailer also contains 1 unused byte.

After the convergence sublayer has constructed and added a header and trailer to the message, as shown in Fig. 6-42, it passes the message to the SAR sublayer,

which chops the message up into 44-byte chunks. Note that to support multiplexing, the convergence sublayer may have several messages constructed internally at once and may pass 44-byte chunks to the SAR sublayer first from one message, then from another, in any order.

The SAR sublayer inserts each 44-byte chunk into the payload of a cell whose format is shown in Fig. 6-43. These cells are then transmitted to the destination for reassembly, after which checksum verification is performed and action taken if need be.




Fig. 6-43. The AAL 3/4 cell format.

The fields in the AAL 3/4 cell are as follows. The *ST* (*Segment Type*) field is used for message framing. It indicates whether the cell begins a message, is in the middle of a message, is the last cell of a message, or is a small (i.e., single cell) message. Next comes a 4-bit sequence number, *SN*, for detecting missing and misinserted cells. The *MID* (*Multiplexing ID*) field is used to keep track of which cell belongs to which session. Remember that the convergence sublayer may have several messages, belonging to different sessions, buffered at once, and it may send pieces of these messages in whatever order it wishes. All the pieces from messages belonging to session *i* carry *i* in the *MID* field, so they can be correctly reassembled at the destination. The trailer contains the payload length and cell checksum.

Notice that AAL 3/4 has two layers of protocol overhead: 8 bytes are added to every message and 4 bytes are added to every cell. All in all, it is a heavyweight mechanism, especially for short messages.

6.5.5. AAL 5

The AAL 1 through AAL 3/4 protocols were largely designed by the telecommunications industry and standardized by ITU without a lot of input from the computer industry. When the computer industry finally woke up and began to understand the implications of Fig. 6-43, a sense of panic set in. The complexity and inefficiency generated by two layers of protocol, coupled with the surprisingly short checksum (only 10 bits), caused some researchers to invent a new

adaptation protocol. It was called **SEAL (Simple Efficient Adaptation Layer)**, which suggests what the designers thought of the old ones. After some discussion, the ATM Forum accepted SEAL and assigned it the name AAL 5. For more information about AAL 5 and how it differs from AAL 3/4, see (Suzuki, 1994).

AAL 5 offers several kinds of service to its applications. One choice is reliable service (i.e., guaranteed delivery with flow control to prevent overruns). Another choice is unreliable service (i.e., no guaranteed delivery), with options to have cells with checksum errors either discarded or passed to the application anyway (but reported as bad). Both unicast and multicast are supported, but multicast does not provide guaranteed delivery.

Like AAL 3/4, AAL 5 supports both message mode and stream mode. In message mode, an application can pass a datagram of length 1 to 65,535 bytes to the AAL layer and have it delivered to the destination, either on a guaranteed or a best efforts basis. Upon arrival in the convergence sublayer, a message is padded out and a trailer added, as shown in Fig. 6-44. The amount of padding (0 to 47 bytes) is chosen to make the entire message, including the padding and trailer, be a multiple of 48 bytes. AAL 5 does not have a convergence sublayer header, just an 8-byte trailer.




Fig. 6-44. AAL 5 convergence sublayer message format.

The *UU (User to User)* field is not used by the AAL layer itself. Instead, it is available for a higher layer for its own purposes, for example, sequencing or multiplexing. The higher layer in question may be the service-specific subpart of the convergence sublayer. The *Length* field tells how long the true payload is, in bytes, not counting the padding. A value of 0 is used to abort the current message in midstream. The *CRC* field is the standard 32-bit checksum over the entire message, including the padding and the trailer (with the *CRC* field set to 0). One 8-bit field in the trailer is reserved for future use.

The message is transmitted by passing it to the SAR sublayer, which does not add any headers or trailers. Instead, it breaks the message into 48-byte units and passes each of these to the ATM layer for transmission. It also tells the ATM layer to set a bit in the *PTI* field on the last cell, so message boundaries are preserved. A case can be made that this is an incorrect mixing of protocol layers because the AAL layer should not be using bits in the ATM layer's header. Doing so violates the most basic principle of protocol engineering, and suggests the layering should have perhaps been done differently.

The principal advantage of AAL 5 over AAL 3/4 is the much greater efficiency. While AAL 3/4 adds only 4 bytes per message, it also adds 4 bytes per

cell, reducing the payload capacity to 44 bytes, a loss of 8 percent on long messages. AAL 5 has a slightly large trailer per message (8 bytes) but has no overhead in each cell. The lack of sequence numbers in the cells is compensated for by the longer checksum, which can detect lost, misinserted, or missing cells without using sequence numbers.

Within the Internet community, it is expected that the normal way of interfacing to ATM networks will be to transport IP packets with the AAL 5 payload field. Various issues relating to this approach are discussed in RFC 1483 and RFC 1577.

6.5.6. Comparison of AAL Protocols

The reader is hereby forgiven if he or she thinks that the various AAL protocols seem unnecessarily similar to one another and poorly thought out. The value of having distinct convergence and SAR sublayers is also questionable, especially since AAL 5 does not have anything in the SAR sublayer. A slightly enhanced ATM layer header could have provided for sequencing, multiplexing, and framing quite adequately.

Some of the differences between the various AAL protocols are summarized in Fig. 6-45. These relate to efficiency, error handling, multiplexing, and the relation between the AAL sublayers.

Item	AAL 1	AAL 2	AAL 3/4	AAL 5
Service class	A	B	C/D	C/D
Multiplexing	No	No	Yes	No
Message delimiting	None	None	Btag/Etag	Bit in PTI
Advance buffer allocation	No	No	Yes	No
User bytes available	0	0	0	1
CS padding	0	0	32-Bit word	0-47 bytes
CS protocol overhead (bytes)	0	0	8	8
CS checksum	None	None	None	32 Bits
SAR payload bytes	46-47	45	44	48
SAR protocol overhead (bytes)	1-2	3	4	0
SAR checksum	None	None	10 Bits	None

Fig. 6-45. Some differences between the various AAL protocols.

The overall impression that AAL gives is of too many variants with too many minor differences and a job half done. The original four service classes, A, B, C, D, have been effectively abandoned. AAL 1 is probably not really necessary;

AAL 2 is broken; AAL 3 and AAL 4 never saw the light of day; and AAL 3/4 is inefficient and has too short a checksum.

The future lies with AAL 5, but even here there is room for improvement. AAL 5 messages should have had a sequence number and a bit to distinguish data from control messages, so it could have been used as a reliable transport protocol. Unused space in the trailer was even available for them. As it stands, for reliable transport, the additional overhead of a transport layer is required on top of it, when it could have been avoided. If the full AAL committee had turned its work in as a class project, the professor would probably have given it back with instructions to fix it and turn it in again when it was finished. More criticism of ATM can be found in (Sterbenz et al., 1995).

6.5.7. SSCOP—Service Specific Connection-Oriented Protocol

Despite all these different AAL protocols, none of them provides for simple end-to-end reliable transport connections. For applications where that is required, another AAL protocol exists: **SSCOP (Service Specific Connection Oriented Protocol)**. However, SSCOP is only used for control, not for data transmission.

SSCOP users send messages, each of which is assigned a 24-bit sequence number. Messages can be up to 64K bytes and are not fragmented. They must be delivered in order. Unlike some other reliable transport protocols, missing messages are always retransmitted using selective repeat rather than go back n.

SSCOP is fundamentally a dynamic sliding window protocol. For each connection, the receiver maintains a window of message sequence numbers that it is prepared to receive, and a bit map marking the ones it already has. This window can change size during protocol operation.

What makes SSCOP unusual is the way acknowledgements are handled: there is no piggybacking. Instead, periodically, the sender polls the receiver and asks it to send back the bit map giving the window status. Based on the result, the sender discards messages that have been accepted and updates its window. SSCOP is described in detail in (Henderson, 1995).

6.6. PERFORMANCE ISSUES

Performance issues are very important in computer networks. When hundreds or thousands of computers are connected together, complex interactions, with unforeseen consequences, are common. Frequently, this complexity leads to poor performance and no one knows why. In the following sections, we will examine many issues related to network performance to see what kinds of problems exist and what can be done about them.

Unfortunately, understanding network performance is more of an art than a science. There is little underlying theory that is actually of any use in practice.

The best we can do is give rules of thumb gained from hard experience and present examples taken from the real world. We have intentionally delayed this discussion until after studying the transport layer in TCP and ATM networks in order to be able to point out places where they have done things right or done things wrong.

The transport layer is not the only place performance issues arise. We saw some of them in the network layer in the previous chapter. Nevertheless, the network layer tends to be largely concerned with routing and congestion control. The broader, system-oriented issues tend to be transport related, so this chapter is an appropriate place to examine them.

In the next five sections, we will look at five aspects of network performance:

1. Performance problems.
2. Measuring network performance.
3. System design for better performance.
4. Fast TPDU processing.
5. Protocols for future high-performance networks.

As an aside, we need a name for the units exchanged by transport entities. The TCP term, segment, is confusing at best and is never used outside the TCP world in this context. The proper ATM terms, CS-PDU, SAR-PDU, and CPCS-PDU, are specific to ATM. Packets clearly refer to the network layer and messages belong to the application layer. For lack of a standard term, we will go back to calling the units exchanged by transport entities TPDU. When we mean both TPDU and packet together, we will use packet as the collective term, as in “The CPU must be fast enough to process incoming packets in real time.” By this we mean both the network layer packet and the TPDU encapsulated in it.

6.6.1. Performance Problems in Computer Networks

Some performance problems, such as congestion, are caused by temporary resource overloads. If more traffic suddenly arrives at a router than the router can handle, congestion will build up and performance will suffer. We studied congestion in detail in the previous chapter.

Performance also degrades when there is a structural resource imbalance. For example, if a gigabit communication line is attached to a low-end PC, the poor CPU will not be able to process the incoming packets fast enough, and some will be lost. These packets will eventually be retransmitted, adding delay, wasting bandwidth, and generally reducing performance.

Overloads can also be synchronously triggered. For example, if a TPDU contains a bad parameter (e.g., the port or process for which it is destined), in many

cases the receiver will thoughtfully send back an error notification. Now consider what could happen if a bad TPDU is broadcast to 10,000 machines: each one might send back an error message. The resulting **broadcast storm** could cripple the network. UDP suffered from this problem until the protocol was changed to cause hosts to refrain from responding to errors in UDP TPDU's sent to broadcast addresses.

A second example of synchronous overload is what happens after an electrical power failure. When the power comes back on, all the machines simultaneously jump to their ROMs to start rebooting. A typical reboot sequence might require first going to some (RARP) server to learn one's true identity, and then to some file server to get a copy of the operating system. If hundreds of machines all do this at once, the server will probably collapse under the load.

Even in the absence of synchronous overloads and when there are sufficient resources available, poor performance can occur due to lack of system tuning. For example, if a machine has plenty of CPU power and memory, but not enough of the memory has been allocated for buffer space, overruns will occur and TPDU's will be lost. Similarly, if the scheduling algorithm does not give a high enough priority to processing incoming TPDU's, some of them may be lost.

Another tuning issue is setting timeouts correctly. When a TPDU is sent, a timer is typically set to guard against its loss. If the timeout is set too short, unnecessary retransmissions will occur, clogging the wires. If the timeout is set too long, unnecessary delays will occur after a TPDU is lost. Other tunable parameters include how long to wait for data to piggyback onto before sending a separate acknowledgement and the number of retransmissions before giving up.

Gigabit networks bring with them new performance problems. Consider, for example, sending data from San Diego to Boston when the receiver's buffer is 64K bytes. Suppose that the link is 1 Gbps and the one-way speed-of-light-in-fiber delay is 20 msec. Initially, at $t = 0$, the pipe is empty, as illustrated in Fig. 6-46(a). Only 500 μ sec later, in Fig. 6-46(b), all the TPDU's are out on the fiber. The lead TPDU will now be somewhere in the vicinity of Brawley, still deep in Southern California. However, the transmitter must stop until it gets a window update.

After 20 msec, the lead TPDU hits Boston, as shown in Fig. 6-46(c) and is acknowledged. Finally, 40 msec after starting, the first acknowledgement gets back to the sender and the second burst can be transmitted. Since the transmission line was used for 0.5 msec out of 40, the efficiency is about 1.25 percent. This situation is typical of running older protocols over gigabit lines.

A useful quantity to keep in mind when analyzing network performance is the **bandwidth-delay product**. It is obtained by multiplying the bandwidth (in bits/sec) by the round-trip delay time (in sec). The product is the capacity of the pipe from the sender to the receiver and back (in bits).

For the example of Fig. 6-46 the bandwidth-delay product is 40 million bits. In other words, the sender would have to transmit a burst of 40 million bits to be




Fig. 6-46. The state of transmitting one megabit from San Diego to Boston. (a) At $t = 0$. (b) After $500 \mu\text{sec}$. (c) After 20 msec . (d) After 40 msec .

able to keep going full speed until the first acknowledgement came back. It takes this many bits to fill the pipe (in both directions). This is why a burst of half a million bits only achieves a 1.25 percent efficiency: it is only 1.25 percent of the pipe capacity.

The conclusion to be drawn here is that to achieve good performance, the receiver's window must be at least as large as the bandwidth-delay product, preferably somewhat larger since the receiver may not respond instantly. For a transcontinental gigabit line, at least 5 megabytes are required for each connection.

If the efficiency is terrible for sending a megabit, imagine what it is like when sending a few hundred bytes for a remote procedure call. Unless some other use can be found for the line while the first client is waiting for its reply, a gigabit line is no better than a megabit line, just more expensive.

Another performance problem that occurs with time-critical applications like audio and video is jitter. Having a short mean transmission time is not enough. A small standard deviation is also required. Achieving a short mean transmission time along with a small standard deviation demands a serious engineering effort.

6.6.2. Measuring Network Performance

When a network performs poorly, its users often complain to the folks running it, demanding improvements. To improve the performance, the operators must first determine exactly what is going on. To find out what is really happening, the operators must make measurements. In this section we will look at network performance measurements. The discussion below is based on the work of Mogul (1993). For a more thorough discussion of the measurement process, see (Jain, 1991; and Villamizan and Song, 1995).

The basic loop used to improve network performance contains the following steps:

1. Measure the relevant network parameters and performance.
2. Try to understand what is going on.
3. Change one parameter.

These steps are repeated until the performance is good enough or it is clear that the last drop of improvement has been squeezed out.

Measurements can be made in many ways and at many locations (both physically and in the protocol stack). The most basic kind of measurement is to start a timer when beginning some activity and use it to see how long that activity takes. For example, knowing how long it takes for a TPDU to be acknowledged is a key measurement. Other measurements are made with counters that record how often some event has happened (e.g., number of lost TPDU's). Finally, one is often interested in knowing the amount of something, such as the number of bytes processed in a certain time interval.

Measuring network performance and parameters has many potential pitfalls. Below we list a few of them. Any systematic attempt to measure network performance should be careful to avoid these.

Make Sure that the Sample Size Is Large Enough

Do not measure the time to send one TPDU, but repeat the measurement, say, one million times and take the average. Having a large sample will reduce the uncertainty in the measured mean and standard deviation. This uncertainty can be computed using standard statistical formulas.

Make Sure that the Samples Are Representative

Ideally, the whole sequence of one million measurements should be repeated at different times of the day and the week to see the effect of different system loads on the measured quantity. Measurements of congestion, for example, are of

little use if they are made at a moment when there is no congestion. Sometimes the results may be counterintuitive at first, such as heavy congestion at 10, 11, 1, and 2 o'clock, but no congestion at noon (when all the users are away at lunch).

Be Careful When Using a Coarse-Grained Clock

Computer clocks work by adding one to some counter at regular intervals. For example, a millisecond timer adds one to a counter every 1 msec. Using such a timer to measure an event that takes less than 1 msec is not impossible, but requires some care.

To measure the time to send a TPDU, for example, the system clock (say, in milliseconds) should be read out when the transport layer code is entered, and again when it is exited. If the true TPDU send time is 300 μ sec, the difference between the two readings will be either 0 or 1, both wrong. However, if the measurement is repeated one million times and the total of all measurements added up and divided by one million, the mean time will be accurate to better than 1 μ sec.

Be Sure that Nothing Unexpected Is Going On during Your Tests

Making measurements on a university system the day some major lab project has to be turned in may give different results than if made the next day. Likewise, if some researcher has decided to run a video conference over your network during your tests, you may get a biased result. It is best to run tests on an idle system and create the entire workload yourself. Even this approach has pitfalls though. While you might think nobody will be using the network at 3 A.M., that might be precisely when the automatic backup program begins copying all the disks to videotape. Furthermore, there might be heavy traffic for your wonderful World Wide Web pages from distant time zones.

Caching Can Wreak Havoc with Measurements

To measure file transfer times, the obvious way to do it is to open a large file, read the whole thing, close it, and see how long it takes. Then repeat the measurement many more times to get a good average. The trouble is, the system may cache the file, so that only the first measurement actually involves network traffic. The rest are just reads from the local cache. The results from such a measurement are essentially worthless (unless you want to measure cache performance).

Often you can get around caching by simply overflowing the cache. For example, if the cache is 10 MB, the test loop could open, read, and close two 10-MB files on each pass, in an attempt to force the cache hit rate to 0. Still, caution is advised unless you are absolutely sure you understand the caching algorithm.

Buffering can have a similar effect. One popular TCP/IP performance utility program has been known to report that UDP can achieve a performance

substantially higher than the physical line allows. How does this occur? A call to UDP normally returns control as soon as the message has been accepted by the kernel and added to the transmission queue. If there is sufficient buffer space, timing 1000 UDP calls does not mean that all the data have been sent. Most of them may still be in the kernel, but the performance utility thinks they have all been transmitted.

Understand What You Are Measuring

When you measure the time to read a remote file, your measurements depend on the network, the operating systems on both the client and server, the particular hardware interface boards used, their drivers, and other factors. If done carefully, you will ultimately discover the file transfer time for the configuration you are using. If your goal is to tune this particular configuration, these measurements are fine.

However, if you are making similar measurements on three different systems in order to choose which network interface board to buy, your results could be thrown off completely by the fact that one of the network drivers is truly awful and is only getting 10 percent of the performance of the board.

Be Careful about Extrapolating the Results

Suppose that you make measurements of something with simulated network loads running from 0 (idle) to 0.4 (40 percent of capacity), as shown by the data points and solid line through them in Fig. 6-47. It may be tempting to extrapolate linearly, as shown by the dotted line. However, many queueing results involve a factor of $1/(1 - \rho)$, where ρ is the load, so the true values may look more like the dashed line.

6.6.3. System Design for Better Performance

Measuring and tinkering can often improve performance considerably, but they cannot substitute for good design in the first place. A poorly designed network can be improved only so much. Beyond that, it has to be redone from scratch.

In this section, we will present some rules of thumb based on experience with many networks. These rules relate to system design, not just network design, since the software and operating system are often more important than the routers and interface boards. Most of these ideas have been common knowledge to network designers for years and have been passed on from generation to generation by word of mouth. They were first stated explicitly by Mogul (1993); our treatment largely parallels his. Another relevant source is (Metcalf, 1993).




Fig. 6-47. Response as a function of load.

Rule #1: CPU Speed Is More Important than Network Speed

Long experience has shown that in nearly all networks, operating system and protocol overhead dominates actual time on the wire. For example, in theory, the minimum RPC time on an Ethernet is 102 μsec , corresponding to a minimum (64-byte) request followed by a minimum (64-byte) reply. In practice, getting the RPC time down to 1500 μsec is a considerable achievement (Van Renesse et al., 1988). Note that 1500 μsec is 15 times worse than the theoretical minimum. Nearly all the overhead is in the software.

Similarly, the biggest problem in running at 1 Gbps is getting the bits from the user's buffer out onto the fiber fast enough and having the receiving CPU process them as fast as they come in. In short, if you double the CPU speed, you often can come close to doubling the throughput. Doubling the network capacity often has no effect since the bottleneck is generally in the hosts.

Rule #2: Reduce Packet Count to Reduce Software Overhead

Processing a TPDU has a certain amount of overhead per TPDU (e.g., header processing) and a certain amount of processing per byte (e.g., doing the checksum). When sending 1 million bytes, the per-byte overhead is the same no matter what the TPDU size is. However, using 128-byte TPDU's means 32 times as much per-TPDU overhead as using 4K TPDU's. This overhead adds up fast.

In addition to the TPDU overhead, there is overhead in the lower layers to consider. Each arriving packet causes an interrupt. On a modern RISC processor, each interrupt breaks the CPU pipeline, interferes with the cache, requires a change to the memory management context, and forces a substantial number of CPU registers to be saved. An n -fold reduction in TPDU's sent thus reduces the interrupt and packet overhead by a factor of n .

This observation argues for collecting a substantial amount of data before transmission in order to reduce interrupts at the other side. Nagle's algorithm and Clark's solution to the silly window syndrome are attempts to do precisely this.

Rule #3: Minimize Context Switches

Context switches (e.g., from kernel mode to user mode) are deadly. They have the same bad properties as interrupts, the worst being a long series of initial cache misses. Context switches can be reduced by having the library procedure that sends data do internal buffering until it has a substantial amount of them. Similarly, on the receiving side, small incoming TPDU's should be collected together and passed to the user in one fell swoop instead of individually to minimize context switches.

In the best case, an incoming packet causes a context switch from the current user to the kernel, and then a switch to the receiving process to give it the newly-arrived data. Unfortunately, with many operating systems, additional context switches happen. For example, if the network manager runs as a special process in user space, a packet arrival is likely to cause a context switch from the current user to the kernel, then another one from the kernel to the network manager followed by another one back to the kernel, and finally one from the kernel to the receiving process. This sequence is shown in Fig. 6-48. All these context switches on each packet are very wasteful of CPU time and will have a devastating effect on network performance.




Fig. 6-48. Four context switches to handle one packet with a user-space network manager.

Rule #4: Minimize Copying

Even worse than multiple context switches is making multiple copies. It is not unusual for an incoming packet to be copied three or four times before the TPDU enclosed in it is delivered. After a packet is received by the network interface in a special on-board hardware buffer, it is typically copied to a kernel buffer. From there it is copied to a network layer buffer, then to a transport layer buffer, and finally to the receiving application process.

A clever operating system will copy a word at a time, but it is not unusual to require about five instructions per word (a load, a store, incrementing an index register, a test for end-of-data, and a conditional branch). On a 50-MIPS machine, making three copies of each packet at five instructions per 32-bit word copied requires 75 nsec per incoming byte. Such a machine can thus accept data at a maximum rate of about 107 Mbps. When overhead for header processing, interrupt handling, and context switches is factored in, 50 Mbps might be achievable, and we have not even considered the actual processing of the data. Clearly, handling a 1-Gbps line is out of the question.

In fact, probably a 50-Mbps line is out of the question, too. In the computation above, we have assumed that a 50-MIPS machine can execute any 50 million instructions/sec. In reality, machines can only run at such speeds if they are not referencing memory. Memory operations are often a factor of three slower than register-register instructions, so actually getting 16 Mbps out of the 1 Gbps line might be considered pretty good. Note that hardware assistance will not help here. The problem is too much copying by the operating system.

Rule #5: You Can Buy More Bandwidth but Not Lower Delay

The next three rules deal with communication, rather than protocol processing. The first rule states that if you want more bandwidth, you can just buy it. Putting a second fiber next to the first one doubles the bandwidth but does nothing to reduce the delay. Making the delay shorter requires improving the protocol software, the operating system, or the network interface. Even if all of these are done, the delay will not be reduced if the bottleneck is the transmission time.

Rule #6: Avoiding Congestion Is Better than Recovering from It

The old maxim that an ounce of prevention is worth a pound of cure certainly holds for network congestion. When a network is congested, packets are lost, bandwidth is wasted, useless delays are introduced, and more. Recovering from it takes time and patience. Not having it occur in the first place is better. Congestion avoidance is like getting your DTP vaccination: it hurts a little at the time you get it, but it prevents something that would hurt a lot more.

Rule #7: Avoid Timeouts

Timers are necessary in networks, but they should be used sparingly and timeouts should be minimized. When a timer goes off, some action is generally repeated. If it is truly necessary to repeat the action, so be it, but repeating it unnecessarily is wasteful.

The way to avoid extra work is to be careful that timers are set a little bit on the conservative side. A timer that takes too long to expire adds a small amount of extra delay to one connection in the (unlikely) event of a TPDU being lost. A timer that goes off when it should not have uses up scarce CPU time, wastes bandwidth, and puts extra load on perhaps dozens of routers for no good reason.

6.6.4. Fast TPDU Processing

The moral of the story above is that the main obstacle to fast networking is protocol software. In this section we will look at some ways to speed up this software. For more information, see (Clark et al., 1989; Edwards and Muir, 1995; and Chandranmenon and Varghese, 1995).

TPDU processing overhead has two components: overhead per TPDU and overhead per byte. Both must be attacked. The key to fast TPDU processing is to separate out the normal case (one-way data transfer) and handle it specially. Although a sequence of special TDUs are needed to get into the *ESTABLISHED* state, once there, TPDU processing is straightforward until one side starts to close the connection.

Let us begin by examining the sending side in the *ESTABLISHED* state when there are data to be transmitted. For the sake of clarity, we assume here that the transport entity is in the kernel, although the same ideas apply if it is a user-space process or a library inside the sending process. In Fig. 6-49, the sending process traps into the kernel to do the SEND. The first thing the transport entity does is make a test to see if this is the normal case: the state is *ESTABLISHED*, neither side is trying to close the connection, a regular (i.e., not an out-of-band) full TPDU is being sent, and there is enough window space available at the receiver. If all conditions are met, no further tests are needed and the fast path through the sending transport entity can be taken.

In the normal case, the headers of consecutive data TDUs are almost the same. To take advantage of this fact, a prototype header is stored within the transport entity. At the start of the fast path, it is copied as fast as possible to a scratch buffer, word by word. Those fields that change from TPDU to TPDU are then overwritten in the buffer. Frequently, these fields are easily derived from state variables, such as the next sequence number. A pointer to the full TPDU header plus a pointer to the user data are then passed to the network layer. Here the same strategy can be followed (not shown in Fig. 6-49). Finally, the network layer gives the resulting packet to the data link layer for transmission.




Fig. 6-49. The fast path from sender to receiver is shown with a heavy line. The processing steps on this path are shaded.

As an example of how this principle works in practice, let us consider TCP/IP. Fig. 6-50(a) shows the TCP header. The fields that are the same between consecutive TPDU's on a one-way flow are shaded. All the sending transport entity has to do is copy the five words from the prototype header into the output buffer, fill in the next sequence number (by copying it from a word in memory), compute the checksum, and increment the sequence number in memory. It can then hand the header and data to a special IP procedure for sending a regular, maximum TPDU. IP then copies its five-word prototype header [see Fig. 6-50(b)] into the buffer, fills in the *Identification* field, and computes its checksum. The packet is now ready for transmission.




Fig. 6-50. (a) TCP header. (b) IP header. In both cases, the shaded fields are taken from the prototype without change.

Now let us look at fast path processing on the receiving side of Fig. 6-49. Step 1 is locating the connection record for the incoming TPDU. For ATM,

finding the connection record is easy: the *VPI* field can be used as an index into the path table to find the virtual circuit table for that path and the *VCI* can be used as an index to find the connection record. For TCP, the connection record can be stored in a hash table for which some simple function of the two IP addresses and two ports is the key. Once the connection record has been located, both addresses and both ports must be compared to verify that the correct record has been found.

An optimization that often speeds up connection record lookup even more is just to maintain a pointer to the last one used and try that one first. Clark et al. (1989) tried this and observed a hit rate exceeding 90 percent. Other lookup heuristics are described in (McKenney and Dove, 1992).

The TPDU is then checked to see if it is a normal one: the state is *ESTABLISHED*, neither side is trying to close the connection, the TPDU is a full one, no special flags are set, and the sequence number is the one expected. These tests take just a handful of instructions. If all conditions are met, a special fast path TCP procedure is called.

The fast path updates the connection record and copies the data to the user. While it is copying, it also computes the checksum, eliminating an extra pass over the data. If the checksum is correct, the connection record is updated and an acknowledgement is sent back. The general scheme of first making a quick check to see if the header is what is expected, and having a special procedure to handle that case, is called **header prediction**. Many TCP implementations use it. When this optimization and all the other ones discussed in this chapter are used together, it is possible to get TCP to run at 90 percent of the speed of a local memory-to-memory copy, assuming the network itself is fast enough.

Two other areas where major performance gains are possible are buffer management and timer management. The issue in buffer management is avoiding unnecessary copying, as we mentioned above. Timer management is important because nearly all timers set do not expire. They are set to guard against TPDU loss, but most TPDU's arrive correctly and their acknowledgements also arrive correctly. Hence it is important to optimize timer management for the case of timers rarely expiring.

A common scheme is to use a linked list of timer events sorted by expiry time. The head entry contains a counter telling how many ticks away from expiry it is. Each successive entry contains a counter telling how many ticks after the previous entry it is. Thus if timers expire in 3, 10, and 12 ticks, respectively, the three counters are 3, 7, and 2, respectively.

At every clock tick, the counter in the head entry is decremented. When it hits zero, its event is processed and the next item on the list becomes the head. Its counter does not have to be changed. In this scheme, inserting and deleting timers are expensive operations, with execution times proportional to the length of the list.

A more efficient approach can be used if the maximum timer interval is bounded and known in advance. Here an array, called a **timing wheel**, can be

used, as shown in Fig. 6-51. Each slot corresponds to one clock tick. The current time shown is $T = 4$. Timers are scheduled to expire at 3, 10, and 12 ticks from now. If a new timer suddenly is set to expire in seven ticks, an entry is just made in slot 11. Similarly, if the timer set for $T + 10$ has to be canceled, the list starting in slot 14 has to be searched and the required entry removed. Note that the array of Fig. 6-51 cannot accommodate timers beyond $T + 15$.




Fig. 6-51. A timing wheel.

When the clock ticks, the current time pointer is advanced by one slot (circularly). If the entry now pointed to is nonzero, all of its timers are processed. Many variations on the basic idea are discussed in (Varghese and Lauck, 1987).

6.6.5. Protocols for Gigabit Networks

At the start of the 1990s, gigabit networks began to appear. People's first reaction was to use the old protocols on them, but various problems quickly arose. In this section we will discuss some of these problems and the directions new protocols are taking to solve them. Other information can be found in (Baransel et al., 1995; and Partridge, 1994).

The first problem is that many protocols use 16-bit or 32-bit sequence numbers. In the old days, 2^{32} was a pretty good approximation to infinity. It no longer is. At a data rate of 1 Gbps, it takes about 32 sec to send 2^{32} bytes. If sequence numbers refer to bytes, as they do in TCP, then a sender can start transmitting byte 0, blast away, and 32 sec later be back at byte 0. Even assuming that all bytes have been acknowledged, the sender cannot safely transmit new data

labeled starting at 0 because the old packets may still be floating around somewhere. In the Internet, for example, packets can live for 120 sec. If packets are numbered instead of bytes, the problem is less severe, unless the sequence numbers are 16 bits, in which case the problem is even worse.

The problem is that many protocol designers simply assumed, without stating it, that the time to use up the entire sequence space would greatly exceed the maximum packet lifetime. Consequently there was no need to even worry about the problem of old duplicates still existing when the sequence numbers wrapped around. At gigabit speeds, that unstated assumption fails.

A second problem is that communication speeds have improved much faster than computing speeds. (Note to computer engineers: Go out and beat those communication engineers! We are counting on you.) In the 1970s, the ARPANET ran at 56 kbps and had computers that ran at about 1 MIPS. Packets were 1008 bits, so the ARPANET was capable of delivering about 56 packets/sec. With almost 18 msec available per packet, a host could afford to spend 18,000 instructions processing a packet. Of course, doing so would soak up the entire CPU, but it could devote 9000 instructions per packet and still have half the CPU left over to do real work.

Compare these numbers to modern 100-MIPS computers exchanging 4-KB packets over a gigabit line. Packets can flow in at a rate of over 30,000 per second, so packet processing must be completed in 15 μ sec if we want to reserve half the CPU for applications. In 15 μ sec, a 100-MIPS computer can execute 1500 instructions, only 1/6 of what the ARPANET hosts had available. Furthermore, modern RISC instructions do less per instruction than the old CISC instructions did, so the problem is even worse than it appears. The conclusion is: there is less time available for protocol processing than there used to be, so protocols must become simpler.

A third problem is that the go back n protocol performs poorly on lines with a large bandwidth-delay product. Consider, for example, a 4000-km line operating at 1 Gbps. The round-trip transmission time is 40 msec, in which time a sender can transmit 5 megabytes. If an error is detected, it will be 40 msec before the sender is told about it. If go back n is used, the sender will have to retransmit not just the bad packet, but also the 5 megabytes worth of packets that came afterward. Clearly, this is a massive waste of resources.

A fourth problem is that gigabit lines are fundamentally different from megabit lines in that long ones are delay limited rather than bandwidth limited. In Fig. 6-52 we show the time it takes to transfer a 1-megabit file 4000 km at various transmission speeds. At speeds up to 1 Mbps, the transmission time is dominated by the rate at which the bits can be sent. By 1 Gbps, the 40-msec round-trip delay dominates the 1 msec it takes to put the bits on the fiber. Further increases in bandwidth have hardly any effect at all.

Figure 6-52 has unfortunate implications for network protocols. It says that stop-and-wait protocols, such as RPC, have an inherent upper bound on their




Fig. 6-52. Time to transfer and acknowledge a 1-megabit file over a 4000-km line.

performance. This limit is dictated by the speed of light. No amount of technological progress in optics will improve matters (new laws of physics would help, though).

A fifth problem that is worth mentioning is not a technological or protocol one like the others, but a result of new applications. Simply stated, it is that for many gigabit applications, such as multimedia, the variance in the packet arrival times is as important as the mean delay itself. A slow-but-uniform delivery rate, is often preferable to a fast-but-jumpy one.

Let us now turn from the problems to ways of dealing with them. We will first make some general remarks, then look at protocol mechanisms, packet layout, and protocol software.

The basic principle that all gigabit network designers should learn by heart is:

Design for speed, not for bandwidth optimization.

Old protocols were often designed to minimize the number of bits on the wire, frequently by using small fields and packing them together into bytes and words. Nowadays, there is plenty of bandwidth. Protocol processing is the problem, so protocols should be designed to minimize it.

A tempting way to go fast is to build fast network interfaces in hardware. The difficulty with this strategy is that unless the protocol is exceedingly simple, hardware just means a plug-in board with a second CPU and its own program. To avoid having the network coprocessor be as expensive as the main CPU, it is often a slower chip. The consequence of this design is that much of the time the main

(fast) CPU is idle waiting for the second (slow) CPU to do the critical work. It is a myth to think that the main CPU has other work to do while waiting. Furthermore, when two general-purpose CPUs communicate, race conditions can occur, so elaborate protocols are needed between the two processors to synchronize them correctly. Usually, the best approach is to make the protocols simple and have the main CPU do the work.

Let us now look at the issue of feedback in high-speed protocols. Due to the (relatively) long delay loop, feedback should be avoided: it takes too long for the receiver to signal the sender. One example of feedback is governing the transmission rate using a sliding window protocol. To avoid the (long) delays inherent in the receiver sending window updates to the sender, it is better to use a rate-based protocol. In such a protocol, the sender can send all it wants to, provided it does not send faster than some rate the sender and receiver have agreed upon in advance.

A second example of feedback is Jacobson's slow start algorithm. This algorithm makes multiple probes to see how much the network can handle. With high-speed networks, making half a dozen or so small probes to see how the network responds wastes a huge amount of bandwidth. A more efficient scheme is to have the sender, receiver, and network all reserve the necessary resources at connection setup time. Reserving resources in advance also has the advantage of making it easier to reduce jitter. In short, going to high speeds inexorably pushes the design toward connection-oriented operation, or something fairly close to it.

Packet layout is an important consideration in gigabit networks. The header should contain as few fields as possible, to reduce processing time, and these fields should be big enough to do the job and be word aligned for ease of processing. In this context, "big enough" means that problems such as sequence numbers wrapping around while old packets still exist, receivers being unable to advertise enough window space because the window field is too small, and so on, do not occur.

The header and data should be separately checksummed, for two reasons. First, to make it possible to checksum the header but not the data. Second, to verify that the header is correct before starting to copy the data into user space. It is desirable to do the data checksum at the time the data are copied to user space, but if the header is incorrect, the copy may be to the wrong process. To avoid an incorrect copy but to allow the data checksum to be done during copying, it is essential that the two checksums be separate.

The maximum data size should be large, to permit efficient operation even in the face of long delays. Also, the larger the data block, the smaller the fraction of the total bandwidth devoted to headers.

Another valuable feature is the ability to send a normal amount of data along with the connection request. In this way, one round-trip time can be saved.

Finally, a few words about the protocol software are appropriate. A key thought is concentrating on the successful case. Many older protocols tend to

emphasize what to do when something goes wrong (e.g., a packet getting lost). To make the protocols run fast, the designer should aim for minimizing processing time when everything goes right. Minimizing processing time when an error occurs is secondary.

A second software issue is minimizing copying time. As we saw earlier, copying data is often the main source of overhead. Ideally, the hardware should dump each incoming packet into memory as a contiguous block of data. The software should then copy this packet to the user buffer with a single block copy. Depending on how the cache works, it may even be desirable to avoid a copy loop. In other words, to copy 1024 words, the fastest way may be to have 1024 back-to-back MOVE instructions (or 1024 load-store pairs). The copy routine is so critical it should be carefully handcrafted in assembly code, unless there is a way to trick the compiler into producing precisely the optimal code.

In the late 1980s, there was a brief flurry of interest in fast special-purpose protocols such as NETBLT (Clark et al., 1987), VTMP (Cheriton and Williamson, 1989), and XTP (Chesson, 1989). A survey is given in (Doeringer et al., 1990). However, the trend now is toward simplifying general-purpose protocols to make them fast, too. ATM exhibits many of the features discussed above, and IPv6 does too.

6.7. SUMMARY

The transport layer is the key to understanding layered protocols. It provides various services, the most important of which is an end-to-end, reliable, connection-oriented byte stream from sender to receiver. It is accessed through service primitives that permit the establishment, use and release of connections.

Transport protocols must be able to do connection management over unreliable networks. Connection establishment is complicated by the existence of delayed duplicate packets that can reappear at inopportune moments. To deal with them, three-way handshakes are needed to establish connections. Releasing a connection is easier than establishing one but is still far from trivial due to the two-army problem.

Even when the network layer is completely reliable, the transport layer has plenty of work to do, as we saw in our example. It must handle all the service primitives, manage connections and timers, and allocate and utilize credits.

The main Internet transport protocol is TCP. It uses a 20-byte header on all segments. Segments can be fragmented by routers within the Internet, so hosts must be prepared to do reassembly. A great deal of work has gone into optimizing TCP performance, using algorithms from Nagle, Clark, Jacobson, Karn, and others.

ATM has four protocols in the AAL layer. All of them break messages into cells at the source and reassemble the cells into messages at the destination. The

CS and SAR sublayers add their own headers and trailers in various ways, leaving from 44 to 48 bytes of cell payload.

Network performance is typically dominated by protocol and TPDU processing overhead, and this situation gets worse at higher speeds. Protocols should be designed to minimize the number of TPDU's, context switches, and times each TPDU is copied. For gigabit networks, simple protocols using rate, rather than credit, flow control are called for.

PROBLEMS

1. In our example transport primitives of Fig. 6-3, LISTEN is a blocking call. Is this strictly necessary? If not, explain how a nonblocking primitive could be used. What advantage would this have over the scheme described in the text?
2. In the model underlying Fig. 6-5, it is assumed that packets may be lost by the network layer and thus must be individually acknowledged. Suppose that the network layer is 100 percent reliable and never loses packets. What changes, if any, are needed to Fig. 6-5?
3. Imagine a generalized n -army problem, in which the agreement of any two of the armies is sufficient for victory. Does a protocol exist that allows blue to win?
4. Suppose that the clock-driven scheme for generating initial sequence numbers is used with a 15-bit wide clock counter. The clock ticks once every 100 msec, and the maximum packet lifetime is 60 sec. How often need resynchronization take place
 - (a) in the worst case?
 - (b) when the data consumes 240 sequence numbers/min?
5. Why does the maximum packet lifetime, T , have to be large enough to ensure that not only the packet, but also its acknowledgements, have vanished?
6. Imagine that a two-way handshake rather than a three-way handshake were used to set up connections. In other words, the third message was not required. Are deadlocks now possible? Give an example or show that none exist.
7. Consider the problem of recovering from host crashes (i.e., Fig. 6-18). If the interval between writing and sending an acknowledgement, or vice versa, can be made relatively small, what are the two best sender-receiver strategies for minimizing the chance of a protocol failure?
8. Are deadlocks possible with the transport entity described in the text?
9. Out of curiosity, the implementer of the transport entity of Fig. 6-20 has decided to put counters inside the *sleep* procedure to collect statistics about the *conn* array. Among these are the number of connections in each of the seven possible states, n_i ($i = 1, \dots, 7$). After writing a massive FORTRAN program to analyze the data, our implementer discovered that the relation $\sum n_i = MAX_CONN$ appears to always be true. Are there any other invariants involving only these seven variables?

10. What happens when the user of the transport entity given in Fig. 6-20 sends a zero length message? Discuss the significance of your answer.
11. For each event that can potentially occur in the transport entity of Fig. 6-20, tell whether it is legal or not when the user is sleeping in *sending* state.
12. Discuss the advantages and disadvantages of credits versus sliding window protocols.
13. Datagram fragmentation and reassembly are handled by IP and are invisible to TCP. Does this mean that TCP does not have to worry about data arriving in the wrong order?
14. A process on host 1 has been assigned port p and a process on host 2 has been assigned port q . Is it possible for there to be two or more TCP connections between these two ports at the same time?
15. The maximum payload of a TCP segment is 65,515 bytes. Why was such a strange number chosen?
16. Describe two ways to get into the *SYN RCVD* state of Fig. 6-28.
17. Give a potential disadvantage when Nagle's algorithm is used on a badly congested network.
18. Consider the effect of using slow start on a line with a 10-msec round-trip time and no congestion. The receive window is 24 KB and the maximum segment size is 2 KB. How long does it take before the first full window can be sent?
19. Suppose that the TCP congestion window is set to 18K bytes and a timeout occurs. How big will the window be if the next four transmission bursts are all successful? Assume that the maximum segment size is 1 KB.
20. If the TCP round-trip time, RTT , is currently 30 msec and the following acknowledgements come in after 26, 32, and 24 msec, respectively, what is the new RTT estimate? Use $\alpha = 0.9$.
21. A TCP machine is sending windows of 65,535 bytes over a 1-Gbps channel that has a 10-msec one-way delay. What is the maximum throughput achievable? What is the line efficiency?
22. In a network that has a maximum TPDU size of 128 bytes, a maximum TPDU lifetime of 30 sec, and an 8-bit sequence number, what is the maximum data rate per connection?
23. Why does UDP exist? Would it not have been enough to just let user processes send raw IP packets?
24. A group of N users located in the same building are all using the same remote computer via an ATM network. The average user generates L lines of traffic (input + output) per hour, on the average, with the mean line length being P bytes, excluding the ATM headers. The packet carrier charges C cents per byte of user data transported, plus X cents per hour for each ATM virtual circuit open. Under what conditions is it cost effective to multiplex all N transport connections onto the same ATM virtual circuit, if such multiplexing adds 2 bytes of data to each packet? Assume that even one ATM virtual circuit has enough bandwidth for all the users.

25. Can AAL 1 handle messages shorter than 40 bytes using the scheme with the *Pointer* field? Explain your answer.
26. Make a guess at what the field sizes for AAL 2 were before they were pulled from the standard.
27. AAL 3/4 allows multiple sessions to be multiplexed onto a single virtual circuit. Give an example of a situation in which that has no value. Assume that one virtual circuit has sufficient bandwidth to carry all the traffic. *Hint:* Think about virtual paths.
28. What is the payload size of the maximum length message that fits in a single AAL 3/4 cell?
29. When a 1024-byte message is sent with AAL 3/4, what is the efficiency obtained? In other words, what fraction of the bits transmitted are useful data bits? Repeat the problem for AAL 5.
30. An ATM device is transmitting single-cell messages at 600 Mbps. One cell in 100 is totally scrambled due to random noise. How many undetected errors per week can be expected with the 32-bit AAL 5 checksum?
31. A client sends a 128-byte request to a server located 100 km away over a 1-gigabit optical fiber. What is the efficiency of the line during the remote procedure call?
32. Consider the situation of the previous problem again. Compute the minimum possible response time both for the given 1-Gbps line and for a 1-Mbps line. What conclusion can you draw?
33. Suppose that you are measuring the time to receive a TPDU. When an interrupt occurs, you read out the system clock in milliseconds. When the TPDU is fully processed, you read out the clock again. You measure 0 msec 270,000 times and 1 msec 730,000 times. How long does it take to receive a TPDU?
34. A CPU executes instructions at the rate of 100 MIPS. Data can be copied 64 bits at a time, with each word copied costing six instructions. If an coming packet has to be copied twice, can this system handle a 1-Gbps line? For simplicity, assume that all instructions, even those instructions that read or write memory, run at the full 100-MIPS rate.
35. To get around the problem of sequence numbers wrapping around while old packets still exist, one could use 64-bit sequence numbers. However, theoretically, an optical fiber can run at 75 Tbps. What maximum packet lifetime is required to make sure that future 75 Tbps networks do not have wraparound problems even with 64-bit sequence numbers? Assume that each byte has its own sequence number, as TCP does.
36. In the text we calculated that a gigabit line dumps 30,000 packets/sec on the host, giving it only 1500 instructions to process it and leaving half the CPU time for applications. This calculation assumed a 4-KB packet. Redo the calculation for an ARPANET-sized packet (128 bytes).
37. For a 1-Gbps network operating over 4000 km, the delay is the limiting factor, not the bandwidth. Consider a MAN with the average source and destination 20 km apart. At what data rate does the round-trip delay due to the speed of light equal the transmission delay for a 1-KB packet?

38. Modify the program of Fig. 6-20 to do error recovery. Add a new packet type, *reset*, that can arrive after a connection has been opened by both sides but closed by neither. This event, which happens simultaneously on both ends of the connection, means that any packets that were in transit have either been delivered or destroyed, but in either case are no longer in the subnet.
39. Write a program that simulates buffer management in a transport entity using a sliding window for flow control rather than the credit system of Fig. 6-20. Let higher-layer processes randomly open connections, send data, and close connections. To keep it simple, have all the data travel from machine *A* to machine *B*, and none the other way. Experiment with different buffer allocation strategies at *B*, such as dedicating buffers to specific connections versus a common buffer pool, and measure the total throughput achieved by each one.

7

THE APPLICATION LAYER

Having finished all the preliminaries, we now come to the application layer, where all the interesting applications can be found. The layers below the application layer are there to provide reliable transport, but they do not do any real work for users. In this chapter we will study some real applications.

However, even in the application layer there is a need for support protocols to allow the real applications to function. Accordingly, we will look at three of these before starting with the applications themselves. The first area is security, which is not a single protocol, but a large number of concepts and protocols that can be used to ensure privacy where needed. The second is DNS, which handles naming within the Internet. The third support protocol is for network management. After that, we will examine four real applications: electronic mail, USENET (net news), the World Wide Web, and finally, multimedia.

7.1. NETWORK SECURITY

For the first few decades of their existence, computer networks were primarily used by university researchers for sending email, and by corporate employees for sharing printers. Under these conditions, security did not get a lot of attention. But now, as millions of ordinary citizens are using networks for banking, shopping, and filing their tax returns, network security is looming on the horizon as a

potentially massive problem. In the following sections, we will study network security from several angles, point out numerous pitfalls, and discuss many algorithms and protocols for making networks more secure.

Security is a broad topic and covers a multitude of sins. In its simplest form, it is concerned with making sure that nosy people cannot read, or worse yet, modify messages intended for other recipients. It is concerned with people trying to access remote services that they are not authorized to use. It also deals with how to tell whether that message purportedly from the IRS saying: "Pay by Friday or else" is really from the IRS or from the Mafia. Security also deals with the problems of legitimate messages being captured and replayed, and with people trying to deny that they sent certain messages.

Most security problems are intentionally caused by malicious people trying to gain some benefit or harm someone. A few of the most common perpetrators are listed in Fig. 7-1. It should be clear from this list that making a network secure involves a lot more than just keeping it free of programming errors. It involves outsmarting often intelligent, dedicated, and sometimes well-funded adversaries. It should also be clear that measures that will stop casual adversaries will have little impact on the serious ones.

Adversary	Goal
Student	To have fun snooping on people's email
Hacker	To test out someone's security system; steal data
Sales rep	To claim to represent all of Europe, not just Andorra
Businessman	To discover a competitor's strategic marketing plan
Ex-employee	To get revenge for being fired
Accountant	To embezzle money from a company
Stockbroker	To deny a promise made to a customer by email
Con man	To steal credit card numbers for sale
Spy	To learn an enemy's military strength
Terrorist	To steal germ warfare secrets

Fig. 7-1. Some people who cause security problems and why.

Network security problems can be divided roughly into four intertwined areas: secrecy, authentication, nonrepudiation, and integrity control. Secrecy has to do with keeping information out of the hands of unauthorized users. This is what usually comes to mind when people think about network security. Authentication deals with determining whom you are talking to before revealing sensitive information or entering into a business deal. Nonrepudiation deals with signatures:

How do you prove that your customer really placed an electronic order for ten million left-handed doobickeys at 89 cents each when he later claims the price was 69 cents? Finally, how can you be sure that a message you received was really the one sent and not something that a malicious adversary modified in transit or concocted?

All these issues (secrecy, authentication, nonrepudiation, and integrity control) occur in traditional systems, too, but with some significant differences. Secrecy and integrity are achieved by using registered mail and locking documents up. Robbing the mail train is harder than it was in Jesse James' day.

Also, people can usually tell the difference between an original paper document and a photocopy, and it often matters to them. As a test, make a photocopy of a valid check. Try cashing the original check at your bank on Monday. Now try cashing the photocopy of the check on Tuesday. Observe the difference in the bank's behavior. With electronic checks, the original and the copy are indistinguishable. It may take a while for banks to get used to this.

People authenticate other people by recognizing their faces, voices, and handwriting. Proof of signing is handled by signatures on letterhead paper, raised seals, and so on. Tampering can usually be detected by handwriting, paper, and ink experts. None of these options are available electronically. Clearly, other solutions are needed.

Before getting into the solutions themselves, it is worth spending a few moments considering where in the protocol stack network security belongs. There is probably no one single place. Every layer has something to contribute. In the physical layer, wiretapping can be foiled by enclosing transmission lines in sealed tubes containing argon gas at high pressure. Any attempt to drill into a tube will release some gas, reducing the pressure and triggering an alarm. Some military systems use this technique.

In the data link layer, packets on a point-to-point line can be encoded as they leave one machine and decoded as they enter another. All the details can be handled in the data link layer, with higher layers oblivious to what is going on. This solution breaks down when packets have to traverse multiple routers, however, because packets have to be decrypted at each router, leaving them vulnerable to attacks from within the router. Also, it does not allow some sessions to be protected (e.g., those involving on-line purchases by credit card) and others not. Nevertheless, **link encryption**, as this method is called, can be added to any network easily and is often useful.

In the network layer, firewalls can be installed to keep packets in or keep packets out. We looked at firewalls in Chap. 5. In the transport layer, entire connections can be encrypted, end to end, that is, process to process. Although these solutions help with secrecy issues and many people are working hard to improve them, none of them solve the authentication or nonrepudiation problem in a sufficiently general way. To tackle these problems, the solutions must be in the application layer, which is why they are being studied in this chapter.

7.1.1. Traditional Cryptography

Cryptography has a long and colorful history. In this section we will just sketch some of the highlights, as background information for what follows. For a complete history, Kahn's (1967) book is still recommended reading. For a comprehensive treatment of the current state-of-the-art, see (Kaufman et al., 1995; Schneier, 1996; and Stinson, 1995).

Historically, four groups of people have used and contributed to the art of cryptography: the military, the diplomatic corps, diarists, and lovers. Of these, the military has had the most important role and has shaped the field. Within military organizations, the messages to be encrypted have traditionally been given to poorly paid code clerks for encryption and transmission. The sheer volume of messages prevented this work from being done by a few elite specialists.

Until the advent of computers, one of the main constraints on cryptography had been the ability of the code clerk to perform the necessary transformations, often on a battlefield with little equipment. An additional constraint has been the difficulty in switching over quickly from one cryptographic method to another one, since this entails retraining a large number of people. However, the danger of a code clerk being captured by the enemy has made it essential to be able to change the cryptographic method instantly, if need be. These conflicting requirements have given rise to the model of Fig. 7-2.




Fig. 7-2. The encryption model.

The messages to be encrypted, known as the **plaintext**, are transformed by a function that is parametrized by a **key**. The output of the encryption process, known as the **ciphertext**, is then transmitted, often by messenger or radio. We assume that the enemy, or **intruder**, hears and accurately copies down the complete ciphertext. However, unlike the intended recipient, he does not know what the decryption key is and so cannot decrypt the ciphertext easily. Sometimes the

intruder can not only listen to the communication channel (passive intruder) but can also record messages and play them back later, inject his own messages, or modify legitimate messages before they get to the receiver (active intruder). The art of breaking ciphers is called **cryptanalysis**. The art of devising ciphers (cryptography) and breaking them (cryptanalysis) is collectively known as **cryptology**.

It will often be useful to have a notation for relating plaintext, ciphertext, and keys. We will use $C = E_K(P)$ to mean that the encryption of the plaintext P using key K gives the ciphertext C . Similarly, $P = D_K(C)$ represents of decryption of C to get the plaintext again. It then follows that

$$D_K(E_K(P)) = P$$

This notation suggests that E and D are just mathematical functions, which they are. The only tricky part is that both are functions of two parameters, and we have written one of the parameters (the key) as a subscript, rather than as an argument, to distinguish it from the message.

A fundamental rule of cryptography is that one must assume that the cryptanalyst knows the general method of encryption used. In other words, the cryptanalyst knows how the encryption method, E , of Fig. 7-2 works. The amount of effort necessary to invent, test, and install a new method every time the old method is compromised or thought to be compromised has always made it impractical to keep this secret, and thinking it is secret when it is not does more harm than good.

This is where the key enters. The key consists of a (relatively) short string that selects one of many potential encryptions. In contrast to the general method, which may only be changed every few years, the key can be changed as often as required. Thus our basic model is a stable and publicly known general method parametrized by a secret and easily changed key.

The nonsecrecy of the algorithm cannot be emphasized enough. By publicizing the algorithm, the cryptographer gets free consulting from a large number of academic cryptologists eager to break the system so they can publish papers demonstrating how smart they are. If many experts have tried to break the algorithm for 5 years after its publication and no one has succeeded, it is probably pretty solid.

The real secrecy is in the key, and its length is a major design issue. Consider a simple combination lock. The general principle is that you enter digits in sequence. Everyone knows this, but the key is secret. A key length of two digits means that there are 100 possibilities. A key length of three digits means 1000 possibilities, and a key length of six digits means a million. The longer the key, the higher the **work factor** the cryptanalyst has to deal with. The work factor for breaking the system by exhaustive search of the key space is exponential in the key length. Secrecy comes from having a strong (but public) algorithm and a long key. To prevent your kid brother from reading your email, 64-bit keys will do. To keep major governments at bay, keys of at least 256 bits are needed.

From the cryptanalyst's point of view, the cryptanalysis problem has three principal variations. When he has a quantity of ciphertext and no plaintext, he is confronted with the **ciphertext only** problem. The cryptograms that appear in the puzzle section of newspapers pose this kind of problem. When he has some matched ciphertext and plaintext, the problem becomes known as the **known plaintext** problem. Finally, when the cryptanalyst has the ability to encrypt pieces of plaintext of his own choosing, we have the **chosen plaintext** problem. Newspaper cryptograms could be broken trivially if the cryptanalyst were allowed to ask such questions as: What is the encryption of ABCDE?

Novices in the cryptography business often assume that if a cipher can withstand a ciphertext only attack, it is secure. This assumption is very naive. In many cases the cryptanalyst can make a good guess at parts of the plaintext. For example, the first thing many timesharing systems say when you call them up is "PLEASE LOGIN." Equipped with some matched plaintext-ciphertext pairs, the cryptanalyst's job becomes much easier. To achieve security, the cryptographer should be conservative and make sure that the system is unbreakable even if his opponent can encrypt arbitrary amounts of chosen plaintext.

Encryption methods have historically been divided into two categories: substitution ciphers and transposition ciphers. We will now deal with each of these briefly as background information for modern cryptography.

Substitution Ciphers

In a **substitution cipher** each letter or group of letters is replaced by another letter or group of letters to disguise it. One of the oldest known ciphers is the **Caesar cipher**, attributed to Julius Caesar. In this method, *a* becomes *D*, *b* becomes *E*, *c* becomes *F*, . . . , and *z* becomes *C*. For example, *attack* becomes *DWWDFN*. In examples, plaintext will be given in lowercase letters, and ciphertext in uppercase letters.

A slight generalization of the Caesar cipher allows the ciphertext alphabet to be shifted by *k* letters, instead of always 3. In this case *k* becomes a key to the general method of circularly shifted alphabets. The Caesar cipher may have fooled the Carthaginians, but it has not fooled anyone since.

The next improvement is to have each of the symbols in the plaintext, say the 26 letters for simplicity, map onto some other letter. For example,

```
plaintext:  a b c d e f g h i j k l m n o p q r s t u v w x y z
ciphertext: Q W E R T Y U I O P A S D F G H J K L Z X C V B N M
```

This general system is called a **monoalphabetic substitution**, with the key being the 26-letter string corresponding to the full alphabet. For the key above, the plaintext *attack* would be transformed into the ciphertext *QZZQEA*.

At first glance this might appear to be a safe system because although the cryptanalyst knows the general system (letter for letter substitution), he does not know which of the $26! \approx 4 \times 10^{26}$ possible keys is in use. In contrast with the Caesar cipher, trying all of them is not a promising approach. Even at 1 μ sec per solution, a computer would take 10^{13} years to try all the keys.

Nevertheless, given a surprisingly small amount of ciphertext, the cipher can be broken easily. The basic attack takes advantage of the statistical properties of natural languages. In English, for example, *e* is the most common letter, followed by *t*, *o*, *a*, *n*, *i*, etc. The most common two letter combinations, or **digrams**, are *th*, *in*, *er*, *re*, and *an*. The most common three letter combinations, or **trigrams**, are *the*, *ing*, *and*, and *ion*.

A cryptanalyst trying to break a monoalphabetic cipher would start out by counting the relative frequencies of all letters in the ciphertext. Then he might tentatively assign the most common one to *e* and the next most common one to *t*. He would then look at trigrams to find a common one of the form *tXe*, which strongly suggests that *X* is *h*. Similarly, if the pattern *thYt* occurs frequently, the *Y* probably stands for *a*. With this information, he can look for a frequently occurring trigram of the form *aZW*, which is most likely *and*. By making guesses at common letters, digrams, and trigrams, and knowing about likely patterns of vowels and consonants, the cryptanalyst builds up a tentative plaintext, letter by letter.

Another approach is to guess a probable word or phrase. For example, consider the following ciphertext from an accounting firm (blocked into groups of five characters):

```
CTBMN BYCTC BTJDS QXBNS GSTJC BSWX CTQTZ CQVUJ
QJSGS TJQZZ MNQJS VLNSX VSZJU JDSTS JQUUS JUBXJ
DSKSU JSNTK BGAQJ ZBGYQ TLCTZ BNYBN QJSW
```

A likely word in a message from an accounting firm is *financial*. Using our knowledge that *financial* has a repeated letter (*i*), with four other letters between their occurrences, we look for repeated letters in the ciphertext at this spacing. We find 12 hits, at positions 6, 15, 27, 31, 42, 48, 56, 66, 70, 71, 76, and 82. However, only two of these, 31 and 42, have the next letter (corresponding to *n* in the plaintext) repeated in the proper place. Of these two, only 31 also has the *a* correctly positioned, so we know that *financial* begins at position 30. From this point on, deducing the key is easy by using the frequency statistics for English text.

Transposition Ciphers

Substitution ciphers preserve the order of the plaintext symbols but disguise them. **Transposition ciphers**, in contrast, reorder the letters but do not disguise them. Figure 7-3 depicts a common transposition cipher, the columnar

transposition. The cipher is keyed by a word or phrase not containing any repeated letters. In this example, MEGABUCK is the key. The purpose of the key is to number the columns, column 1 being under the key letter closest to the start of the alphabet, and so on. The plaintext is written horizontally, in rows. The ciphertext is read out by columns, starting with the column whose key letter is the lowest.

<u>M</u>	<u>E</u>	<u>G</u>	<u>A</u>	<u>B</u>	<u>U</u>	<u>C</u>	<u>K</u>	
7	4	5	1	2	8	3	6	
p	l	e	a	s	e	t	r	Plaintext
a	n	s	f	e	r	o	n	pleasetransferonemilliondollarsto
e	m	i	l	l	i	o	n	myswissbankaccountsixtwo
d	o	l	l	a	r	s	t	Ciphertext
o	m	y	s	w	i	s	s	AFLLSKSOSELAWAIATOSSCTCLNMOMANT
b	a	n	k	a	c	c	o	ESILYNTWRNNTSOWDPAEDOBUEIRICXB
u	n	t	s	i	x	t	w	
o	t	w	o	a	b	c	d	

Fig. 7-3. A transposition cipher.

To break a transposition cipher, the cryptanalyst must first be aware that he is dealing with a transposition cipher. By looking at the frequency of *E*, *T*, *A*, *O*, *I*, *N*, etc., it is easy to see if they fit the normal pattern for plaintext. If so, the cipher is clearly a transposition cipher, because in such a cipher every letter represents itself.

The next step is to make a guess at the number of columns. In many cases a probable word or phrase may be guessed at from the context of the message. For example, suppose that our cryptanalyst suspected the plaintext phrase *milliondollars* to occur somewhere in the message. Observe that digrams *MO*, *IL*, *LL*, *LA*, *IR* and *OS* occur in the ciphertext as a result of this phrase wrapping around. The ciphertext letter *O* follows the ciphertext letter *M* (i.e., they are vertically adjacent in column 4) because they are separated in the probable phrase by a distance equal to the key length. If a key of length seven had been used, the digrams *MD*, *IO*, *LL*, *LL*, *IA*, *OR*, and *NS* would have occurred instead. In fact, for each key length, a different set of digrams is produced in the ciphertext. By hunting for the various possibilities, the cryptanalyst can often easily determine the key length.

The remaining step is to order the columns. When the number of columns, k , is small, each of the $k(k-1)$ column pairs can be examined to see if its digram frequencies match those for English plaintext. The pair with the best match is assumed to be correctly positioned. Now each remaining column is tentatively tried as the successor to this pair. The column whose digram and trigram frequencies give the best match is tentatively assumed to be correct. The predecessor

column is found in the same way. The entire process is continued until a potential ordering is found. Chances are that the plaintext will be recognizable at this point (e.g., if *milloin* occurs, it is clear what the error is).

Some transposition ciphers accept a fixed-length block of input and produce a fixed-length block of output. These ciphers can be completely described by just giving a list telling the order in which the characters are to be output. For example, the cipher of Fig. 7-3 can be seen as a 64 character block cipher. Its output is 4, 12, 20, 28, 36, 44, 52, 60, 5, 13, . . . , 62. In other words, the fourth input character, *a*, is the first to be output, followed by the twelfth, *f*, and so on.

One-Time Pads

Constructing an unbreakable cipher is actually quite easy; the technique has been known for decades. First choose a random bit string as the key. Then convert the plaintext into a bit string, for example by using its ASCII representation. Finally, compute the EXCLUSIVE OR of these two strings, bit by bit. The resulting ciphertext cannot be broken, because every possible plaintext is an equally probable candidate. The ciphertext gives the cryptanalyst no information at all. In a sufficiently large sample of ciphertext, each letter will occur equally often, as will every digram and every trigram.

This method, known as the **one-time pad**, has a number of practical disadvantages, unfortunately. To start with, the key cannot be memorized, so both sender and receiver must carry a written copy with them. If either one is subject to capture, written keys are clearly undesirable. Additionally, the total amount of data that can be transmitted is limited by the amount of key available. If the spy strikes it rich and discovers a wealth of data, he may find himself unable to transmit it back to headquarters because the key has been used up. Another problem is the sensitivity of the method to lost or inserted characters. If the sender and receiver get out of synchronization, all data from then on will appear garbled.

With the advent of computers, the one-time pad might potentially become practical for some applications. The source of the key could be a special CD that contains several gigabits of information, and if transported in a music CD box and prefixed by a few songs, would not even be suspicious. Of course, at gigabit network speeds, having to insert a new CD every 5 sec could become tedious. For this reason, we will now start looking at modern encryption algorithms that can process arbitrarily large amounts of plaintext.

7.1.2. Two Fundamental Cryptographic Principles

Although we will study many different cryptographic systems in the pages ahead, there are two principles underlying all of them that are important to understand. The first principle is that all encrypted messages must contain some

redundancy, that is, information not needed to understand the message. An example may make it clear why this is needed. Consider a mail-order company, The Couch Potato (TCP), with 60,000 products. Thinking they are being very efficient, TCP's programmers decide that ordering messages should consist of a 16-byte customer name followed by a 3-byte data field (1 byte for the quantity and 2 bytes for the product number). The last 3 bytes are to be encrypted using a very long key known only by the customer and TCP.

At first this might seem secure, and in a sense it is because passive intruders cannot decrypt the messages. Unfortunately, it also has a fatal flaw that renders it useless. Suppose that a recently-fired employee wants to punish TCP for firing her. Just before leaving, she takes (part of) the customer list with her. She works through the night writing a program to generate fictitious orders using real customer names. Since she does not have the list of keys, she just puts random numbers in the last 3 bytes, and sends hundreds of orders off to TCP.

When these messages arrive, TCP's computer uses the customer's name to locate the key and decrypt the message. Unfortunately for TCP, almost every 3-byte message is valid, so the computer begins printing out shipping instructions. While it might seem odd for a customer to order 137 sets of children's swings, or 240 sandboxes, for all the computer knows, the customer might be planning to open a chain of franchised playgrounds. In this way an active intruder (the ex-employee) can cause a massive amount of trouble, even though she cannot understand the messages her computer is generating.

This problem can be solved by adding redundancy to all messages. For example, if order messages are extended to 12 bytes, the first 9 of which must be zeros, then this attack no longer works because the ex-employee no longer can generate a large stream of valid messages. The moral of the story is that all messages must contain considerable redundancy so that active intruders cannot send random junk and have it be interpreted as a valid message.

However, adding redundancy also makes it much easier for cryptanalysts to break messages. Suppose that the mail order business is highly competitive, and The Couch Potato's main competitor, The Sofa Tuber, would dearly love to know how many sandboxes TCP is selling. Consequently, they have tapped TCP's telephone line. In the original scheme with 3-byte messages, cryptanalysis was nearly impossible, because after guessing a key, the cryptanalyst had no way of telling whether the guess was right. After all, almost every message is technically legal. With the new 12-byte scheme, it is easy for the cryptanalyst to tell a valid message from an invalid one.

Thus cryptographic principle number one is that all messages must contain redundancy to prevent active intruders from tricking the receiver into acting on a false message. However, this same redundancy makes it much easier for passive intruders to break the system, so there is some tension here. Furthermore, the redundancy should never be in the form of n zeros at the start or end of a message, since running such messages through some cryptographic algorithms gives more

predictable results, making the cryptanalysts' job easier. A random string of English words would be a much better choice for the redundancy.

The second cryptographic principle is that some measures must be taken to prevent active intruders from playing back old messages. If no such measures were taken, our ex-employee could tap TCP's phone line and just keep repeating previously sent valid messages. One such measure is including in every message a timestamp valid only for, say, 5 minutes. The receiver can then just keep messages around for 5 minutes, to compare newly arrived messages to previous ones to filter out duplicates. Messages older than 5 minutes can be thrown out, since any replays sent more than 5 minutes later will be rejected as too old. Measures other than timestamps will be discussed later.

7.1.3. Secret-Key Algorithms

Modern cryptography uses the same basic ideas as traditional cryptography, transposition and substitution, but its emphasis is different. Traditionally, cryptographers have used simple algorithms and relied on very long keys for their security. Nowadays the reverse is true: the object is to make the encryption algorithm so complex and involuted that even if the cryptanalyst acquires vast mounds of enciphered text of his own choosing, he will not be able to make any sense of it at all.

Transpositions and substitutions can be implemented with simple circuits. Figure 7-4(a) shows a device, known as a **P-box** (P stands for permutation), used to effect a transposition on an 8-bit input. If the 8 bits are designated from top to bottom as 01234567, the output of this particular P-box is 36071245. By appropriate internal wiring, a P-box can be made to perform any transposition, and do it at practically the speed of light.




Fig. 7-4. Basic elements of product ciphers. (a) P-box. (b) S-box. (c) Product.

Substitutions are performed by **S-boxes**, as shown in Fig. 7-4(b). In this example a 3-bit plaintext is entered and a 3-bit ciphertext is output. The 3-bit input selects one of the eight lines exiting from the first stage and sets it to 1; all the other lines are 0. The second stage is a P-box. The third stage encodes the

selected input line in binary again. With the wiring shown, if the eight octal numbers 01234567 were input one after another, the output sequence would be 24506713. In other words, 0 has been replaced by 2, 1 has been replaced by 4, etc. Again, by appropriate wiring of the P-box inside the S-box, any substitution can be accomplished.

The real power of these basic elements only becomes apparent when we cascade a whole series of boxes to form a **product cipher**, as shown in Fig. 7-4(c). In this example, 12 input lines are transposed by the first stage. Theoretically, it would be possible to have the second stage be an S-box that mapped a 12-bit number onto another 12-bit number. However, such a device would need $2^{12} = 4096$ crossed wires in its middle stage. Instead, the input is broken up into four groups of 3 bits, each of which is substituted independently of the others. Although this method is less general, it is still powerful. By including a sufficiently large number of stages in the product cipher, the output can be made to be an exceedingly complicated function of the input.

DES

In January 1977, the U.S. government adopted a product cipher developed by IBM as its official standard for unclassified information. This cipher, **DES (Data Encryption Standard)**, was widely adopted by the industry for use in security products. It is no longer secure in its original form (Wayner, 1995), but in a modified form it is still useful. We will now explain how DES works.

An outline of DES is shown in Fig. 7-5(a). Plaintext is encrypted in blocks of 64 bits, yielding 64 bits of ciphertext. The algorithm, which is parametrized by a 56-bit key, has 19 distinct stages. The first stage is a key independent transposition on the 64-bit plaintext. The last stage is the exact inverse of this transposition. The stage prior to the last one exchanges the leftmost 32 bits with the rightmost 32 bits. The remaining 16 stages are functionally identical but are parametrized by different functions of the key. The algorithm has been designed to allow decryption to be done with the same key as encryption. The steps are just run in the reverse order.

The operation of one of these intermediate stages is illustrated in Fig. 7-5(b). Each stage takes two 32-bit inputs and produces two 32-bit outputs. The left output is simply a copy of the right input. The right output is the bitwise EXCLUSIVE OR of the left input and a function of the right input and the key for this stage, K_i . All the complexity lies in this function.

The function consists of four steps, carried out in sequence. First, a 48-bit number, E , is constructed by expanding the 32-bit R_{i-1} according to a fixed transposition and duplication rule. Second, E and K_i are EXCLUSIVE ORed together. This output is then partitioned into eight groups of 6 bits each, each of which is fed into a different S-box. Each of the 64 possible inputs to an S-box is mapped onto a 4-bit output. Finally, these 8×4 bits are passed through a P-box.




Fig. 7-5. The data encryption standard. (a) General outline. (b) Detail of one iteration.

In each of the 16 iterations, a different key is used. Before the algorithm starts, a 56-bit transposition is applied to the key. Just before each iteration, the key is partitioned into two 28-bit units, each of which is rotated left by a number of bits dependent on the iteration number. K_i is derived from this rotated key by applying yet another 56-bit transposition to it. A different 48-bit subset of the 56 bits is extracted and permuted on each round.

DES Chaining

Despite all this complexity, DES is basically a monoalphabetic substitution cipher using a 64-bit character. Whenever the same 64-bit plaintext block goes in the front end, the same 64-bit ciphertext block comes out the back end. A cryptanalyst can exploit this property to help break DES.

To see how this monoalphabetic substitution cipher property can be used to subvert DES, let us consider encrypting a long message the obvious way: by breaking it up into consecutive 8-byte (64-bit) blocks and encrypting them one

after another with the same key. The last block is padded out to 64 bits, if need be. This technique is known as **electronic code book mode**.

In Fig. 7-6 we have the start of a computer file listing the annual bonuses a company has decided to award to its employees. This file consists of consecutive 32-byte records, one per employee, in the format shown: 16 bytes for the name, 8 bytes for the position, and 8 bytes for the bonus. Each of the sixteen 8-byte blocks (numbered from 0 to 15) is encrypted by DES.

Name	Position	Bonus
A d a m s . L e s l i e .	C l e r k .	\$ 1 0
B l a c k . R o b i n .	B o s s .	\$ 5 0 0 . 0 0
C o l l i n s . K i m .	M a n a g e r .	\$ 1 0 0 . 0 0
D a v i s . B o b b i e .	J a n i t o r .	\$ 5

Bytes ← 16 → 8 → 8 →

Fig. 7-6. The plaintext of a file encrypted as 16 DES blocks.

Leslie just had a fight with the boss and is not expecting much of a bonus. Kim, in contrast is the boss' favorite, and everyone knows this. Leslie can get access to the file after it is encrypted but before it is sent to the bank. Can Leslie rectify this unfair situation, given only the encrypted file?

No problem at all. All Leslie has to do is make a copy of ciphertext block 11 (which contains Kim's bonus) and use it to replace ciphertext block 3 (which contains Leslie's bonus). Even without knowing what block 11 says, Leslie can expect to have a much merrier Christmas this year. (Copying ciphertext block 7 is also a possibility, but is more likely to be detected; besides, Leslie is not a greedy person.)

To thwart this type of attack, DES (and all block ciphers) can be chained in various ways so that replacing a block the way Leslie did will cause the plaintext decrypted starting at the replaced block to be garbage. One way of chaining is **cipher block chaining**. In this method, shown in Fig. 7-7, each plaintext block is EXCLUSIVE ORed (#) with the previous ciphertext block before being encrypted. Consequently, the same plaintext block no longer maps onto the same ciphertext block, and the encryption is no longer a big monoalphabetic substitution cipher. The first block is EXCLUSIVE ORed with a randomly chosen **initialization vector, IV**, that is transmitted along with the ciphertext.

We can see how cipher block chaining works by examining the example of Fig. 7-7. We start out by computing $C_0 = E(P_0 \text{ XOR } IV)$. Then we compute $C_1 = E(P_1 \text{ XOR } C_0)$, and so on. Decryption works the other way, with $P_0 = IV \text{ XOR } D(C_0)$, and so on. Note that the encryption of block i is a function




Fig. 7-7. Cipher block chaining

of all the plaintext in blocks 0 through $i - 1$, so the same plaintext generates different ciphertext depending on where it occurs. A transformation of the type Leslie made will result in nonsense for two blocks starting at Leslie's bonus field. To an astute security officer, this peculiarity might suggest where to start the ensuing investigation.

Cipher block chaining also has the advantage that the same plaintext block will not result in the same ciphertext block, making cryptanalysis more difficult. In fact, this is the main reason it is used.

However, cipher block chaining has the disadvantage of requiring an entire 64-bit block to arrive before decryption can begin. For use with interactive terminals, where people can type lines shorter than eight characters and then stop, waiting for a response, this mode is unsuitable. For byte-by-byte encryption, **cipher feedback mode**, shown in Fig. 7-8, can be used. In this figure, the state of the encryption machine is shown after bytes 0 through 9 have been encrypted and sent. When plaintext byte 10 arrives, as illustrated in Fig. 7-8(a), the DES algorithm operates on the 64-bit shift register to generate a 64-bit ciphertext. The left-most byte of that ciphertext is extracted and EXCLUSIVE Ored with P_{10} . That byte is transmitted on the transmission line. In addition, the shift register is shifted left 8 bits, causing C_2 to fall off the left end, and C_{10} is inserted in the position just vacated at the right end by C_9 . Note that the contents of the shift register depend on the entire previous history of the plaintext, so a pattern that repeats multiple times in the plaintext will be encrypted differently each time in the ciphertext. As with cipher block chaining, an initialization vector is needed to start the ball rolling.

Decryption with cipher feedback mode just does the same thing as encryption. In particular, the contents of the shift register is *encrypted*, not *decrypted*, so the selected byte that is EXCLUSIVE Ored with C_{10} to get P_{10} is the same one that was EXCLUSIVE Ored with P_{10} to generate C_{10} in the first place. As long as the two shift registers remain identical, decryption works correctly.




Fig. 7-8. Cipher feedback mode.

As an aside, it should be noted that if one bit of the ciphertext is accidentally inverted during transmission, the 8 bytes that are decrypted while the bad byte is in the shift register will be corrupted. Once the bad byte is pushed out of the shift register, correct plaintext will once again be generated. Thus the effects of a single inverted bit are relatively localized and do not ruin the rest of the message.

Nevertheless, there exist applications in which having a 1-bit transmission error mess up 64 bits of plaintext is too large an effect. For these applications, a fourth option, **output feedback mode**, exists. It is identical to cipher feedback mode, except that the byte fed back into the right end of the shift register is taken from just before the EXCLUSIVE OR box, not just after it.

Output feedback mode has the property that a 1-bit error in the ciphertext causes only a 1-bit error in the resulting plaintext. On the other hand, it is less secure than the other modes, and should be avoided for general-purpose use. Electronic code book mode should also be avoided except under special circumstances (e.g., encrypting a single random number, such as a session key). For normal operation, cipher block chaining should be used when the input arrives in 8-byte units (e.g., for encrypting disk files) and cipher feedback mode should be used for irregular input streams, such as keyboard input.

Breaking DES

DES has been enveloped in controversy from the day it was launched. It was based on a cipher developed and patented by IBM, called Lucifer, except that IBM's cipher used a 128-bit key instead of a 56-bit key. When the U.S. federal government wanted to standardize on one cipher for unclassified use, it "invited"

IBM to “discuss” the matter with NSA, the government’s code-breaking arm, which is the world’s largest employer of mathematicians and cryptologists. NSA is so secret that an industry joke goes:

Q: What does NSA stand for?

A: No Such Agency.

Actually, NSA stands for National Security Agency.

After these discussions took place, IBM reduced the key from 128 bits to 56 bits and decided to keep secret the process by which DES was designed. Many people suspected that the key length was reduced to make sure that NSA could just break DES, but no organization with a smaller budget could. The point of the secret design was supposedly to hide a trapdoor that could make it even easier for NSA to break DES. When an NSA employee discreetly told IEEE to cancel a planned conference on cryptography, that did not make people any more comfortable.

In 1977, two Stanford cryptography researchers, Diffie and Hellman (1977), designed a machine to break DES and estimated that it could be built for 20 million dollars. Given a small piece of plaintext and matched ciphertext, this machine could find the key by exhaustive search of the 2^{56} -entry key space in under 1 day. Nowadays, such a machine would cost perhaps 1 million dollars. A detailed design for a machine that can break DES by exhaustive search in about four hours is presented in (Wiener, 1994).

Here is another strategy. Although software encryption is 1000 times slower than hardware encryption, a high-end home computer can still do about 250,000 encryptions/sec in software and is probably idle 2 million seconds/month. This idle time could be put to use breaking DES. If someone posted a message to one of the popular Internet newsgroups, it should not be hard to sign up the necessary 140,000 people to check all 7×10^{16} keys in a month.

Probably the most innovative idea for breaking DES is the **Chinese Lottery** (Quisquater and Girault, 1991). In this design, every radio and television has to be equipped with a cheap DES chip capable of performing 1 million encryptions/sec in hardware. Assuming that every one of the 1.2 billion people in China owns a radio or television, whenever the Chinese government wants to decrypt a message encrypted by DES, it just broadcasts the plaintext/ciphertext pair, and each of the 1.2 billion chips begins searching its preassigned section of the key space. Within 60 seconds, one (or more) hits will be found. To ensure that they are reported, the chips could be programmed to display or announce the message:

CONGRATULATIONS! YOU HAVE JUST WON THE CHINESE LOTTERY.
TO COLLECT, PLEASE CALL 1-800-BIG-PRIZE

The conclusion that one can draw from these arguments is that DES should no longer be used for anything important. However, although 2^{56} is a paltry

7×10^{16} , 2^{112} is a magnificent 5×10^{33} . Even with a billion DES chips doing a billion operations per second, it would take 100 million years to exhaustively search a 112-bit key space. Thus the thought arises of just running DES twice, with two different 56-bit keys.

Unfortunately, Merkle and Hellman (1981) have developed a method that makes double encryption suspect. It is called the **meet-in-the-middle** attack and works like this (Hellman, 1980). Suppose that someone has doubly encrypted a series of plaintext blocks, using electronic code book mode. For a few values of i , the cryptanalyst has matched pairs (P_i, C_i) where

$$C_i = E_{K_2}(E_{K_1}(P_i))$$

If we now apply the decryption function, D_{K_2} to each side of this equation, we get

$$D_{K_2}(C_i) = E_{K_1}(P_i) \quad (7-1)$$

because encrypting x and then decrypting it with the same key gives back x .

The meet-in-the-middle attack uses this equation to find the DES keys, K_1 and K_2 , as follows:

1. Compute $R_i = E_i(P_1)$ for all 2^{56} values of i , where E is the DES encryption function. Sort this table in ascending order of R_i .
2. Compute $S_j = D_j(C_1)$ for all 2^{56} values of j , where D is the DES decryption function. Sort this table in ascending order of S_j .
3. Scan the first table looking for an R_i that matches some S_j in the second table. When a match is found, we then have a key pair (i, j) such that $D_j(C_1) = E_i(P_1)$. Potentially, i is K_1 and j is K_2 .
4. Check to see if $E_j(E_i(P_2))$ is equal to C_2 . If it is, try all the other (plaintext, ciphertext) pairs. If it is not, continue searching the two tables looking for matches.

Many false alarms will certainly occur before the real keys are located, but eventually they will be found. This attack requires only 2^{57} encryption or decryption operations (to construct the two tables), far less than 2^{112} . However it also requires a total of 2^{60} bytes of storage for the two tables, so it is not currently feasible in this basic form, but Merkle and Hellman have shown various optimizations and trade-offs that permit less storage at the expense of more computing. All in all, double encryption using DES is probably not much more secure than single encryption.

Triple encryption is another matter. As early as 1979, IBM realized that the DES key length was too short and devised a way to effectively increase it using triple encryption (Tuchman, 1979). The method chosen, which has since been incorporated in International Standard 8732, is illustrated in Fig. 7-9. Here two

keys and three stages are used. In the first stage, the plaintext is encrypted with K_1 . In the second stage, DES is run in decryption mode, using K_2 as the key. Finally, another encryption is done with K_1 .




Fig. 7-9. Triple encryption using DES.

This design immediately gives rise to two questions. First, why are only two keys used, instead of three? Second, why is EDE used, instead of EEE? The reason that two keys are used is that even the most paranoid cryptographers concede that 112 bits is sufficient for commercial applications for the time being. Going to 168 bits would just add the unnecessary overhead of managing and transporting another key.

The reason for encrypting, decrypting, and then encrypting again is backward compatibility with existing single-key DES systems. Both the encryption and decryption functions are mappings between sets of 64-bit numbers. From a cryptographic point of view, the two mappings are equally strong. By using EDE, however, instead of EEE, a computer using triple encryption can speak to one using single encryption by just setting $K_1 = K_2$. This property allows triple encryption to be phased in gradually, something of no concern to academic cryptographers, but of considerable import to IBM and its customers.

No method is known for breaking triple DES in EDE mode. Van Oorschot and Wiener (1988) have presented a method to speed up the search of EDE by a factor of 16, but even with their attack, EDE is highly secure. For anyone wishing nothing less than the very best, EEE with three distinct 56-bit keys (168 bits in all) is recommended.

Before leaving the subject of DES, it is worth at least mentioning two recent developments in cryptanalysis. The first development is **differential cryptanalysis** (Biham and Shamir, 1993). This technique can be used to attack any block cipher. It works by beginning with a pair of plaintext blocks that differ in only a small number of bits and watching carefully what happens on each internal iteration as the encryption proceeds. In many cases, some patterns are much more common than other patterns, and this observation leads to a probabilistic attack.

The other development worth noting is **linear cryptanalysis** (Matsui, 1994). It can break DES with only 2^{43} known plaintexts. It works by EXCLUSIVE ORing certain plaintext and ciphertext bits together to generate 1 bit. When done repeatedly, half the bits should be 0s and half should be 1s. Often, however, ciphers introduce a bias in one direction or the other, and this bias, however small, can be exploited to reduce the work factor. For the details, see Matsui's paper.

IDEA

Perhaps all this hammering on why DES is insecure is like beating a dead horse, but the reality is that singly-encrypted DES is still widely used for secure applications, such as banking using automated teller machines. While this choice was probably appropriate when it was made, a decade or more ago, it is no longer adequate.

At this point, the reader is probably legitimately wondering: "If DES is so weak, why hasn't anyone invented a better block cipher?" The fact is, many other block ciphers have been proposed, including BLOWFISH (Schneier, 1994), Crab (Kaliski and Robshaw, 1994), FEAL (Shimizu and Miyaguchi, 1988), KHAFRE (Merkle, 1991), LOKI91 (Brown et al., 1991), NEWDES (Scott, 1985), REDOC-II (Cusick and Wood, 1991), and SAFER K64 (Massey, 1994). Schneier (1996) discusses all of these and innumerable others. Probably the most interesting and important of the post-DES block ciphers is **IDEA** the (**I**nternational **D**ata **E**ncryption **A**lgorithm) (Lai and Massey, 1990; and Lai, 1992). Let us now study IDEA in more detail.

IDEA was designed by two researchers in Switzerland, so it is probably free of any NSA "guidance" that might have introduced a secret trapdoor. It uses a 128-bit key, which will make it immune to brute force, Chinese lottery, and meet-in-the-middle attacks for decades to come. It was also designed to withstand differential cryptanalysis. No currently known technique or machine is thought to be able to break IDEA.

The basic structure of the algorithm resembles DES in that 64-bit plaintext input blocks are mangled in a sequence of parameterized iterations to produce 64-bit ciphertext output blocks, as shown in Fig. 7-10(a). Given the extensive bit mangling (for every iteration, every output bit depends on every input bit), eight iterations are sufficient. As with all block ciphers, IDEA can also be used in cipher feedback mode and the other DES modes.

The details of one iteration are depicted in Fig. 7-10(b). Three operations are used, all on unsigned 16-bit numbers. These operations are EXCLUSIVE OR, addition modulo 2^{16} , and multiplication modulo $2^{16} + 1$. All three of these can easily be done on a 16-bit microcomputer by ignoring the high-order parts of results. The operations have the property that no two pairs obey the associative law or distributive law, making cryptanalysis more difficult. The 128-bit key is used to generate 52 subkeys of 16 bits each, 6 for each of eight iterations and 4 for the final transformation. Decryption uses the same algorithm as encryption, only with different subkeys.

Both software and hardware implementations of IDEA have been constructed. The first software implementation ran on a 33-MHz 386 and achieved an encryption rate of 0.88 Mbps. On a modern machine running ten times as fast, 9 Mbps should be achievable in software. An experimental 25-MHz VLSI chip was built at ETH Zurich and encrypted at a rate of 177 Mbps.




Fig. 7-10. (a) IDEA. (b) Detail of one iteration.

7.1.4. Public-Key Algorithms

Historically the key distribution problem has always been the weak link in most cryptosystems. No matter how strong a cryptosystem was, if an intruder could steal the key, the system was worthless. Since all cryptologists always took for granted that the encryption key and decryption key were the same (or easily derived from one another) and the key had to be distributed to all users of the system, it seemed as if there was an inherent built-in problem: keys had to be protected from theft, but they also had to be distributed, so they could not just be locked up in a bank vault.

In 1976, two researchers at Stanford University, Diffie and Hellman (1976), proposed a radically new kind of cryptosystem, one in which the encryption and decryption keys were different, and the decryption key could not be derived from the encryption key. In their proposal, the (keyed) encryption algorithm, E , and the

(keyed) decryption algorithm, D , had to meet the following three requirements. These requirements can be stated simply as follows:

1. $D(E(P)) = P$.
2. It is exceedingly difficult to deduce D from E .
3. E cannot be broken by a chosen plaintext attack.

The first requirement says that if we apply D to an encrypted message, $E(P)$, we get the original plaintext message, P , back. The second requirement speaks for itself. The third requirement is needed because, as we shall see in a moment, intruders may experiment with the algorithm to their hearts' content. Under these conditions, there is no reason that the encryption key cannot be made public.

The method works like this. A person, say, Alice, wanting to receive secret messages, first devises two algorithms, E_A and D_A , meeting the above requirements. The encryption algorithm and key, E_A , is then made public, hence the name **public-key cryptography** (to contrast it with traditional secret-key cryptography). This might be done by putting it in a file that anyone who wanted to could read. Alice publishes the decryption algorithm (to get the free consulting), but keeps the decryption key secret. Thus, E_A is public, but D_A is private.

Now let us see if we can solve the problem of establishing a secure channel between Alice and Bob, who have never had any previous contact. Both Alice's encryption key, E_A , and Bob's encryption key, E_B , are assumed to be in a publicly readable file. (Basically, all users of the network are expected to publish their encryption keys as soon as they become network users.) Now Alice takes her first message, P , computes $E_B(P)$, and sends it to Bob. Bob then decrypts it by applying his secret key D_B [i.e., he computes $D_B(E_B(P)) = P$]. No one else can read the encrypted message, $E_B(P)$, because the encryption system is assumed strong and because it is too difficult to derive D_B from the publicly known E_B . Alice and Bob can now communicate securely.

A note on terminology is perhaps useful here. Public-key cryptography requires each user to have two keys: a public key, used by the entire world for encrypting messages to be sent to that user, and a private key, which the user needs for decrypting messages. We will consistently refer to these keys as the *public* and *private* keys, respectively, and distinguish them from the *secret* keys used for both encryption and decryption in conventional (also called **symmetric key**) cryptography.

The RSA Algorithm

The only catch is that we need to find algorithms that indeed satisfy all three requirements. Due to the potential advantages of public-key cryptography, many researchers are hard at work, and some algorithms have already been published. One good method was discovered by a group at M.I.T. (Rivest et al., 1978). It is

known by the initials of the three discoverers (Rivest, Shamir, Adleman): **RSA**. Their method is based on some principles from number theory. We will now summarize how to use the method below; for details, consult the paper.

1. Choose two large primes, p and q , (typically greater than 10^{100}).
2. Compute $n = p \times q$ and $z = (p - 1) \times (q - 1)$.
3. Choose a number relatively prime to z and call it d .
4. Find e such that $e \times d = 1 \pmod{z}$.

With these parameters computed in advance, we are ready to begin encryption. Divide the plaintext (regarded as a bit string) into blocks, so that each plaintext message, P , falls in the interval $0 \leq P < n$. This can be done by grouping the plaintext into blocks of k bits, where k is the largest integer for which $2^k < n$ is true.

To encrypt a message, P , compute $C = P^e \pmod{n}$. To decrypt C , compute $P = C^d \pmod{n}$. It can be proven that for all P in the specified range, the encryption and decryption functions are inverses. To perform the encryption, you need e and n . To perform the decryption, you need d and n . Therefore, the public key consists of the pair (e, n) and the private key consists of (d, n) .

The security of the method is based on the difficulty of factoring large numbers. If the cryptanalyst could factor the (publicly known) n , he could then find p and q , and from these z . Equipped with knowledge of z and e , d can be found using Euclid's algorithm. Fortunately, mathematicians have been trying to factor large numbers for at least 300 years, and the accumulated evidence suggests that it is an exceedingly difficult problem.

According to Rivest and colleagues, factoring a 200-digit number requires 4 billion years of computer time; factoring a 500-digit number requires 10^{25} years. In both cases, they assume the best known algorithm and a computer with a 1- μ sec instruction time. Even if computers continue to get faster by an order of magnitude per decade, it will be centuries before factoring a 500-digit number becomes feasible, at which time our descendants can simply choose p and q still larger.

A trivial pedagogical example of the RSA algorithm is given in Fig. 7-11. For this example we have chosen $p = 3$ and $q = 11$, giving $n = 33$ and $z = 20$. A suitable value for d is $d = 7$, since 7 and 20 have no common factors. With these choices, e can be found by solving the equation $7e = 1 \pmod{20}$, which yields $e = 3$. The ciphertext, C , for a plaintext message, P , is given by $C = P^3 \pmod{33}$. The ciphertext is decrypted by the receiver according to the rule $P = C^7 \pmod{33}$. The figure shows the encryption of the plaintext "SUZANNE" as an example.

Because the primes chosen for this example are so small, P must be less than 33, so each plaintext block can contain only a single character. The result is a

Plaintext (P)		Ciphertext (C)			After decryption	
Symbolic	Numeric	P^3	$P^3 \pmod{33}$	C^7	$C^7 \pmod{33}$	Symbolic
S	19	6859	28	13492928512	19	S
U	21	9261	21	1801088541	21	U
Z	26	17576	20	1280000000	26	Z
A	01	1	1	1	1	A
N	14	2744	5	78125	14	N
N	14	2744	5	78125	14	N
E	05	125	26	8031810176	5	E

Sender's computation
Receiver's computation

Fig. 7-11. An example of the RSA algorithm.

monoalphabetic substitution cipher, not very impressive. If instead we had chosen p and $q \approx 10^{100}$, we would have $n \approx 10^{200}$, so each block could be up to 664 bits ($2^{664} \approx 10^{200}$) or 83 8-bit characters, versus 8 characters for DES.

It should be pointed out that using RSA as we have described is similar to using DES in ECB mode—the same input block gives the same output block. Therefore some form of chaining is needed for data encryption. However, in practice, most RSA-based systems use public-key cryptography primarily for distributing one-time session keys for use with DES, IDEA, or similar algorithms. RSA is too slow for actually encrypting large volumes of data.

Other Public-Key Algorithms

Although RSA is widely used, it is by no means the only public-key algorithm known. The first public-key algorithm was the knapsack algorithm (Merkle and Hellman, 1978). The idea here is that someone owns a large number of objects, each with a different weight. The owner encodes the message by secretly selecting a subset of the objects and placing them in the knapsack. The total weight of the objects in the knapsack is made public, as is the list of all possible objects. The list of objects in the knapsack is kept secret. With certain additional restrictions, the problem of figuring out a possible list of objects with the given weight was thought to be computationally infeasible, and formed the basis of the public-key algorithm.

The algorithm's inventor, Ralph Merkle, was quite sure that this algorithm could not be broken, so he offered a 100-dollar reward to anyone who could break it. Adi Shamir (the "S" in RSA) promptly broke it and collected the reward. Undeterred, Merkle strengthened the algorithm and offered a 1000-dollar reward to anyone who could break the new one. Ron Rivest (the "R" in RSA) promptly broke the new one and collected the reward. Merkle did not dare offer 10,000

dollars for the next version, so “A” (Leonard Adleman) was out of luck. Although it has been patched up again, the knapsack algorithm is not considered secure and is rarely used.

Other public-key schemes are based on the difficulty of computing discrete logarithms (Rabin, 1979). Algorithms that use this principle have been invented by El Gamal (1985) and Schnorr (1991).

A few other schemes exist, such as those based on elliptic curves (Menezes and Vanstone, 1993), but the three major categories are those based on the difficulty of factoring large numbers, computing discrete logarithms, and determining the contents of a knapsack from its weight. These problems are thought to be genuinely difficult to solve because mathematicians have been working on them for many years without any great breakthroughs.

7.1.5. Authentication Protocols

Authentication is the technique by which a process verifies that its communication partner is who it is supposed to be and not an imposter. Verifying the identity of a remote process in the face of a malicious, active intruder is surprisingly difficult and requires complex protocols based on cryptography. In this section, we will study some of the many authentication protocols that are used on insecure computer networks.

As an aside, some people confuse authorization with authentication. Authentication deals with the question of whether or not you are actually communicating with a specific process. Authorization is concerned with what that process is permitted to do. For example, a client process contacts a file server and says: “I am Scott’s process and I want to delete the file *cookbook.old*.” From the file server’s point of view, two questions must be answered:

1. Is this actually Scott’s process (authentication)?
2. Is Scott allowed to delete *cookbook.old* (authorization)?

Only after both questions have been unambiguously answered in the affirmative can the requested action take place. The former question is really the key one. Once the file server knows whom it is talking to, checking authorization is just a matter of looking up entries in local tables. For this reason, we will concentrate on authentication in this section.

The general model that all authentication protocols use is this. An initiating user (really a process), say, Alice, wants to establish a secure connection with a second user, Bob. Alice and Bob are sometimes called **principals**, the main characters in our story. Bob is a banker with whom Alice would like to do business. Alice starts out by sending a message either to Bob, or to a trusted **key distribution center (KDC)**, which is always honest. Several other message exchanges

follow in various directions. As these messages are being sent, a nasty intruder, Trudy,[†] may intercept, modify, or replay them in order to trick Alice and Bob or just to gum up the works.

Nevertheless, when the protocol has been completed, Alice is sure she is talking to Bob and Bob is sure he is talking to Alice. Furthermore, in most of the protocols, the two of them will also have established a secret **session key** for use in the upcoming conversation. In practice, for performance reasons, all data traffic is encrypted using secret-key cryptography, although public-key cryptography is widely used for the authentication protocols themselves and for establishing the session key.

The point of using a new, randomly-chosen session key for each new connection is to minimize the amount of traffic that gets sent with the users' secret keys or public keys, to reduce the amount of ciphertext an intruder can obtain, and to minimize the damage done if a process crashes and its core dump falls into the wrong hands. Hopefully, the only key present then will be the session key. All the permanent keys should have been carefully zeroed out after the session was established.

Authentication Based on a Shared Secret Key

For our first authentication protocol, we will assume that Alice and Bob already share a secret key, K_{AB} (In the formal protocols, we will abbreviate Alice as A and Bob as B , respectively.) This shared key might have been agreed upon on the telephone, or in person, but, in any event, not on the (insecure) network.

This protocol is based on a principle found in many authentication protocols: one party sends a random number to the other, who then transforms it in a special way and then returns the result. Such protocols are called **challenge-response** protocols. In this and subsequent authentication protocols, the following notation will be used:

A, B are the identities of Alice and Bob

R_i 's are the challenges, where the subscript identifies the challenger

K_i are keys, where i indicates the owner; K_S is the session key

The message sequence for our first shared-key authentication protocol is shown in Fig. 7-12. In message 1, Alice sends her identity, A , to Bob in a way that Bob understands. Bob, of course, has no way of knowing whether this message came from Alice or from Trudy, so he chooses a challenge, a large random number, R_B , and sends it back to "Alice" as message 2, in plaintext. Alice then encrypts the message with the key she shares with Bob and sends the ciphertext, $K_{AB}(R_B)$, back in message 3. When Bob sees this message, he immediately knows that it came from Alice because Trudy does not know K_{AB} and thus could

[†] I thank Kaufman₁ et al.₂₃ (1995) for revealing her name.

not have generated it. Furthermore, since R_B was chosen randomly from a large space (say, 128-bit random numbers), it is very unlikely that Trudy would have seen R_B and its response from an earlier session.




Fig. 7-12. Two-way authentication using a challenge-response protocol.

At this point, Bob is sure he is talking to Alice, but Alice is not sure of anything. For all Alice knows, Trudy might have intercepted message 1 and sent back R_B in response. Maybe Bob died last night. To find out whom she is talking to, Alice picks a random number, R_A and sends it to Bob as plaintext, in message 4. When Bob responds with $K_{AB}(R_A)$, Alice knows she is talking to Bob. If they wish to establish a session key now, Alice can pick one, K_S , and send it to Bob encrypted with K_{AB} .

Although the protocol of Fig. 7-12 works, it contains extra messages. These can be eliminated by combining information, as illustrated in Fig. 7-13. Here Alice initiates the challenge-response protocol instead of waiting for Bob to do it. Similarly, while he is responding to Alice's challenge, Bob sends his own. The entire protocol can be reduced to three messages instead of five.




Fig. 7-13. A shortened two-way authentication protocol.

Is this new protocol an improvement over the original one? In one sense it is: it is shorter. Unfortunately, it is also wrong. Under certain circumstances, Trudy can defeat this protocol by using what is known as a **reflection attack**. In particular, Trudy can break it if it is possible to open multiple sessions with Bob at

once. This situation would be true, for example, if Bob is a bank and is prepared to accept many simultaneous connections from teller machines at once.

Trudy's reflection attack is shown in Fig. 7-14. It starts out with Trudy claiming she is Alice and sending R_T . Bob responds, as usual, with his own challenge, R_B . Now Trudy is stuck. What can she do? She does not know $K_{AB}(R_B)$.




Fig. 7-14. The reflection attack.

She can open a second session with message 3, supplying the R_B taken from message 2 as her challenge. Bob calmly encrypts it and sends back $K_{AB}(R_B)$ in message 4. Now Trudy has the missing information, so she can complete the first session and abort the second one. Bob is now convinced that Trudy is Alice, so when she asks for her bank account balance, he gives it to her without question. Then when she asks him to transfer it all to a secret bank account in Switzerland, he does so without a moment's hesitation.

The moral of this story is:

Designing a correct authentication protocol is harder than it looks.

Three general rules that often help are as follows:

1. Have the initiator prove who she is before the responder has to. In this case, Bob gives away valuable information before Trudy has to give any evidence of who she is.
2. Have the initiator and responder use different keys for proof, even if this means having two shared keys, K_{AB} and K'_{AB} .
3. Have the initiator and responder draw their challenges from different sets. For example, the initiator must use even numbers and the responder must use odd numbers.

All three rules were violated here, with disastrous results. Note that our first (five-message) authentication protocol requires Alice to prove her identity first, so that protocol is not subject to the reflection attack.

Establishing a Shared Key: The Diffie-Hellman Key Exchange

So far we have assumed that Alice and Bob share a secret key. Suppose that they do not? How can they establish one? One way would be for Alice to call Bob and give him her key on the phone, but he would probably start out by saying: "How do I know you are Alice and not Trudy?" They could try to arrange a meeting, with each one bringing a passport, a drivers' license, and three major credit cards, but being busy people, they might not be able to find a mutually acceptable date for months. Fortunately, incredible as it may sound, there is a way for total strangers to establish a shared secret key in broad daylight, even with Trudy carefully recording every message.

The protocol that allows strangers to establish a shared secret key is called the **Diffie-Hellman key exchange** (Diffie and Hellman, 1976) and works as follows. Alice and Bob have to agree on two large prime numbers, n , and g , where $(n - 1)/2$ is also a prime and certain conditions apply to g . These numbers may be public, so either one of them can just pick n and g and tell the other openly. Now Alice picks a large (say, 512-bit) number, x , and keeps it secret. Similarly, Bob picks a large secret number, y .

Alice initiates the key exchange protocol by sending Bob a message containing $(n, g, g^x \bmod n)$, as shown in Fig. 7-15. Bob responds by sending Alice a message containing $g^y \bmod n$. Now Alice takes the number Bob sent her and raises it to the x th power to get $(g^y \bmod n)^x$. Bob performs a similar operation to get $(g^x \bmod n)^y$. By the laws of modular arithmetic, both calculations yield $g^{xy} \bmod n$. Lo and behold, Alice and Bob now share a secret key, $g^{xy} \bmod n$.




Fig. 7-15. The Diffie-Hellman key exchange.

Trudy, of course, has seen both messages. She knows g and n from message 1. If she could compute x and y , she could figure out the secret key. The trouble is, given only $g^x \bmod n$, she cannot find x . No practical algorithm for computing discrete logarithms modulo a very large prime number is known.

To make the above example more concrete, we will use the (completely unrealistic) values of $n = 47$ and $g = 3$. Alice picks $x = 8$ and Bob picks $y = 10$.

Both of these are kept secret. Alice's message to Bob is $(47, 3, 28)$ because $3^8 \bmod 47$ is 28. Bob's message to Alice is (17). Alice computes $17^8 \bmod 47$, which is 4. Bob computes $28^{10} \bmod 47$, which is 4. Alice and Bob have independently determined that the secret key is now 4. Trudy has to solve the equation $3^x \bmod 47 = 28$, which can be done by exhaustive search for small numbers like this, but not when all the numbers are hundreds of bits long. All currently-known algorithms simply take too long, even using a massively parallel supercomputer.

Despite the elegance of the Diffie-Hellman algorithm, there is a problem: when Bob gets the triple $(47, 3, 28)$, how does he know it is from Alice and not from Trudy? There is no way he can know. Unfortunately, Trudy can exploit this fact to deceive both Alice and Bob, as illustrated in Fig. 7-16. Here, while Alice and Bob are choosing x and y , respectively, Trudy picks her own random number, z . Alice sends message 1 intended for Bob. Trudy intercepts it and sends message 2 to Bob, using the correct g and n (which are public anyway) but with her own z instead of x . She also sends message 3 back to Alice. Later Bob sends message 4 to Alice, which Trudy again intercepts and keeps.




Fig. 7-16. The bucket brigade attack.

Now everybody does the modular arithmetic. Alice computes the secret key as $g^{xz} \bmod n$, and so does Trudy (for messages to Alice). Bob computes $g^{yz} \bmod n$ and so does Trudy (for messages to Bob). Alice thinks she is talking to Bob so she establishes a session key (with Trudy). So does Bob. Every message that Alice sends on the encrypted session is captured by Trudy, stored, modified if desired, and then (optionally) passed on to Bob. Similarly in the other direction. Trudy sees everything and can modify all messages at will, while both Alice and Bob are under the illusion that they have a secure channel to one another. This attack is known as the **bucket brigade attack**, because it vaguely resembles an old-time volunteer fire department passing buckets along the line from the fire truck to the fire. It is also called the **(wo)man-in-the-middle attack**, which should not be confused with the meet-in-the-middle attack on block ciphers. Fortunately, more complex algorithms can defeat this attack.

Authentication Using a Key Distribution Center

Setting up a shared secret with a stranger almost worked, but not quite. On the other hand, it probably was not worth doing in the first place (sour grapes attack). To talk to n people this way, you would need n keys. For popular people, key management would become a real burden, especially if each key had to be stored on a separate plastic chip card.

A different approach is to introduce a trusted key distribution center (KDC). In this model, each user has a single key shared with the KDC. Authentication and session key management now goes through the KDC. The simplest known KDC authentication protocol involving two parties and a trusted KDC is depicted in Fig. 7-17.




Fig. 7-17. A first attempt at an authentication protocol using a KDC.

The idea behind this protocol is simple: Alice picks a session key, K_S , and tells the KDC that she wants to talk to Bob using K_S . This message is encrypted with the secret key Alice shares (only) with the KDC, K_A . The KDC decrypts this message, extracting Bob's identity and the session key. It then constructs a new message containing Alice's identity and the session key and sends this message to Bob. This encryption is done with K_B , the secret key Bob shares with the KDC. When Bob decrypts the message, he learns that Alice wants to talk to him, and which key she wants to use.

The authentication here happens for free. The KDC knows that message 1 must have come from Alice, since no one else would have been able to encrypt it with Alice's secret key. Similarly, Bob knows that message 2 must have come from the KDC, whom he trusts, since no one else knows his secret key.

Unfortunately, this protocol has a serious flaw. Trudy needs some money, so she figures out some legitimate service she can perform for Alice, makes an attractive offer, and gets the job. After doing the work, Trudy then politely requests Alice to pay by bank transfer. Alice then establishes a session key with her banker, Bob. Then she sends Bob a message requesting money to be transferred to Trudy's account.

Meanwhile, Trudy is back to her old ways, snooping on the network. She copies both message 2 in Fig. 7-17, and the money-transfer request that follows it.

Later, she replays both of them to Bob. Bob gets them and thinks: "Alice must have hired Trudy again. She clearly does good work." Bob then transfers an equal amount of money from Alice's account to Trudy's. Some time after the 50th message pair, Bob runs out of the office to find Trudy to offer her a big loan so she can expand her obviously successful business. This problem is called the **replay attack**.

Several solutions to the replay attack are possible. The first one is to include a timestamp in each message. Then if anyone receives an obsolete message, it can be discarded. The trouble with this approach is that clocks are never exactly synchronized over a network, so there has to be some interval during which a timestamp is valid. Trudy can replay the message during this interval and get away with it.

The second solution is to put a one-time, unique message number, usually called a **nonce**, in each message. Each party then has to remember all previous nonces and reject any message containing a previously used nonce. But nonces have to be remembered forever, lest Trudy try replaying a 5-year-old message. Also, if some machine crashes and it loses its nonce list, it is again vulnerable to a replay attack. Timestamps and nonces can be combined to limit how long nonces have to be remembered, but clearly the protocol is going to get a lot more complicated.

A more sophisticated approach to authentication is to use a multiway challenge-response protocol. A well-known example of such a protocol is the **Needham-Schroeder authentication** protocol (Needham and Schroeder, 1978), one variant of which is shown in Fig. 7-18.




Fig. 7-18. The Needham-Schroeder authentication protocol.

The protocol begins with Alice telling the KDC that she wants to talk to Bob. This message contains a large random number, R_A , as a nonce. The KDC sends back message 2 containing Alice's random number, a session key, and a ticket that she can send to Bob. The point of the random number, R_A , is to assure Alice that message 2 is fresh, and not a replay. Bob's identity is also enclosed in case Trudy gets any funny ideas about replacing B in message 1 with her own identity

so the KDC will encrypt the ticket at the end of message 2 with K_T instead of K_B . The ticket encrypted with K_B is included inside the encrypted message to prevent Trudy from replacing it with something else on the way back to Alice.

Alice now sends the ticket to Bob, along with a new random number, R_{A2} , encrypted with the session key, K_S . In message 4, Bob sends back $K_S(R_{A2} - 1)$ to prove to Alice that she is talking to the real Bob. Sending back $K_S(R_{A2})$ would not have worked, since Trudy could just have stolen it from message 3.

After receiving message 4, Alice is now convinced that she is talking to Bob, and that no replays could have been used so far. After all, she just generated R_{A2} a few milliseconds ago. The purpose of message 5 is to convince Bob that it is indeed Alice he is talking to, and no replays are being used here either. By having each party both generate a challenge and respond to one, the possibility of any kind of replay attack is eliminated.

Although this protocol seems pretty solid, it does have a slight weakness. If Trudy ever manages to obtain an old session key in plaintext, she can initiate a new session with Bob replaying the message 3 corresponding to the compromised key and convince him that she is Alice (Denning and Sacco, 1981). This time she can plunder Alice's bank account without having to perform the legitimate service even once.

Needham and Schroeder later published a protocol that corrects this problem (Needham and Schroeder, 1987). In the same issue of the same journal, Otway and Rees (1987) also published a protocol that solves the problem in a shorter way. Figure 7-19 shows a slightly modified Otway-Rees protocol.




Fig. 7-19. The Otway-Rees authentication protocol (slightly simplified).

In the Otway-Rees protocol, Alice starts out by generating a pair of random numbers, R , which will be used as a common identifier, and R_A which Alice will use to challenge Bob. When Bob gets this message, he constructs a new message from the encrypted part of Alice's message, and an analogous one of his own. Both the parts encrypted with K_A and K_B identify Alice and Bob, contain the common identifier, and contain a challenge.

The KDC checks to see if the R in both parts is the same. It might not be because Trudy tampered with R in message 1 or replaced part of message 2. If

the two R s match, the KDC believes that the request message from Bob is valid. It then generates a session key and encrypts it twice, once for Alice and once for Bob. Each message contains the receiver's random number, as proof that the KDC, and not Trudy, generated the message. At this point both Alice and Bob are in possession of the same session key and can start communicating. The first time they exchange data messages, each one can see that the other one has an identical copy of K_S , so the authentication is then complete.

Authentication Using Kerberos

An authentication protocol used in many real systems is **Kerberos**, which is based on a variant of Needham-Schroeder. It is named for a multiheaded dog in Greek Mythology that used to guard the entrance to Hades (presumably to keep undesirables out). Kerberos was designed at M.I.T. to allow workstation users to access network resources in a secure way. Its biggest difference with Needham-Schroeder is its assumption that all clocks are fairly-well synchronized. The protocol has gone through several iterations. V4 is the version most widely used in industry, so we will describe it. Afterward, we will say a few words about its successor, V5. For more information, see (Neuman and Ts'o, 1994; and Steiner et al., 1988).

Kerberos involves three servers in addition to Alice (a client workstation):

Authentication Server (AS): verifies users during login
Ticket-Granting Server (TGS): issues "proof of identity tickets"
Bob the server: actually does the work Alice wants performed

AS is similar to a KDC in that it shares a secret password with every user. The TGS's job is to issue tickets that can convince the real servers that the bearer of a TGS ticket really is who he or she claims to be.

To start a session, Alice sits down at an arbitrary public workstation and types her name. The workstation sends her name to the AS in plaintext, as shown in Fig. 7-20. What comes back is a session key and a ticket, $K_{TGS}(A, K_S)$, intended for the TGS. These items are packaged together and encrypted using Alice's secret key, so that only Alice can decrypt them. Only when message 2 arrives, does the workstation ask for Alice's password. The password is then used to generate K_A , in order to decrypt message 2 and obtain the session key and TGS ticket inside it. At this point, the workstation overwrites Alice's password, to make sure that it is only inside the workstation for a few milliseconds at most. If Trudy tries logging in as Alice, the password she types will be wrong and the workstation will detect this because the standard part of message 2 will be incorrect.

After she logs in, Alice may tell the workstation that she wants to contact Bob the file server. The workstation then sends message 3 to the TGS asking for a ticket to use with Bob. The key element in this request is $K_{TGS}(A, K_S)$, which is




Fig. 7-20. The operation of Kerberos V4.

encrypted with the TGS's secret key and is used as proof that the sender really is Alice. The TGS responds by creating a session key, K_{AB} , for Alice to use with Bob. Two versions of it are sent back. The first is encrypted with only K_S , so Alice can read it. The second is encrypted with Bob's key, K_B , so Bob can read it.

Trudy can copy message 3 and try to use it again, but she will be foiled by the encrypted timestamp, t , sent along with it. Trudy cannot replace the timestamp with a more recent one, because she does not know K_S , the session key Alice uses to talk to the TGS. Even if Trudy replays message 3 quickly, all she will get is another copy of message 4, which she could not decrypt the first time and will not be able to decrypt the second time either.

Now Alice can send K_{AB} to Bob to establish a session with him. This exchange is also timestamped. The response is proof to Alice that she is actually talking to Bob, not to Trudy.

After this series of exchanges, Alice can communicate with Bob under cover of K_{AB} . If she later decides she needs to talk to another server, Carol, she just repeats message 3 to the TGS, only now specifying C instead of B . The TGS will promptly respond with a ticket encrypted with K_C that Alice can send to Carol and that Carol will accept as proof that it came from Alice.

The point of all this work is that now Alice can access servers all over the network in a secure way, and her password never has to go over the network. In fact, it only had to be in her own workstation for a few milliseconds. However, note that each server does its own authorization. When Alice presents her ticket to Bob, this merely proves to Bob who sent it. Precisely what Alice is allowed to do is up to Bob.

Since the Kerberos designers did not expect the entire world to trust a single authentication server, they made provision for having multiple **realms**, each with its own AS and TGS. To get a ticket for a server in a distant realm, Alice would ask her own TGS for a ticket accepted by the TGS in the distant realm. If the

distant TGS has registered with the local TGS (the same way local servers do), the local TGS will give Alice a ticket valid at the distant TGS. She can then do business over there, such as getting tickets for servers in that realm. Note, however, that for parties in two realms to do business, each one must trust the other's TGS.

Kerberos V5 is fancier than V4 and has more overhead. It also uses OSI ASN.1 (Abstract Syntax Notation 1) for describing data types and has small changes in the protocols. Furthermore, it has longer ticket lifetimes, allows tickets to be renewed, and will issue postdated tickets. In addition, at least in theory, it is not DES dependent, as V4 is, and supports multiple realms.

Authentication Using Public-Key Cryptography

Mutual authentication can also be done using public-key cryptography. To start with, let us assume Alice and Bob already know each other's public keys (a nontrivial issue). They want to establish a session, and then use secret-key cryptography on that session, since it is typically 100 to 1000 times faster than public-key cryptography. The purpose of the initial exchange then is to authenticate each other and agree on a secret shared session key.

This setup can be done in various ways. A typical one is shown in Fig. 7-21. Here Alice starts by encrypting her identity and a random number, R_A , using Bob's public (or encryption) key, E_B . When Bob receives this message, he has no idea of whether it came from Alice or from Trudy, but he plays along and sends Alice back a message containing Alice's R_A , his own random number, R_B , and a proposed session key, K_S .




Fig. 7-21. Mutual authentication using public-key cryptography.

When Alice gets message 2, she decrypts it using her private key. She sees R_A in it, which gives her a warm feeling inside. The message must have come from Bob, since Trudy has no way of determining R_A . Furthermore, it must be fresh and not a replay, since she just sent Bob R_A . Alice agrees to the session by sending back message 3. When Bob sees R_B encrypted with the session key he just generated, he knows Alice got message 2 and verified R_A .

What can Trudy do to try to subvert this protocol? She can fabricate message 1 and trick Bob into probing Alice, but Alice will see an R_A that she did not send and will not proceed further. Trudy cannot forge message 3 convincingly because

she does not know R_B or K_S and cannot determine them without Alice's private key. She is out of luck.

However, the protocol does have a weakness: it assumes that Alice and Bob already know each other's public keys. Suppose that they do not. Alice could just send Bob her public key in the first message and ask Bob to send his back in the next one. The trouble with this approach is that it is subject to a bucket brigade attack. Trudy can capture Alice's message to Bob and send her own public key back to Alice. Alice will think she has a key for talking to Bob, when, in fact, she has a key for talking to Trudy. Now Trudy can read all the messages encrypted with what Alice thinks is Bob's public key.

The initial public-key exchange can be avoided by having all the public keys stored in a public database. Then Alice and Bob can fetch each other's public keys from the database. Unfortunately, Trudy can still pull off the bucket brigade attack by intercepting the requests to the database and sending simulated replies containing her own public key. After all, how do Alice and Bob know that the replies came from the real data base and not from Trudy?

Rivest and Shamir (1984) have devised a protocol that foils Trudy's bucket brigade attack. In their **interlock protocol**, after the public key exchange, Alice sends only half of her message to Bob, say, only the even bits (after encryption). Bob then responds with his even bits. After getting Bob's even bits, Alice sends her odd bits, then Bob does too.

The trick here is that when Trudy gets Alice's even bits, she cannot decrypt the message, even though Trudy has the private key. Consequently, she is unable to reencrypt the even bits using Bob's public key. If she sends junk to Bob, the protocol will continue, but Bob will shortly discover that the fully assembled message makes no sense and realized that he has been spoofed.

7.1.6. Digital Signatures

The authenticity of many legal, financial, and other documents is determined by the presence or absence of an authorized handwritten signature. And photocopies do not count. For computerized message systems to replace the physical transport of paper and ink documents, a solution must be found to these problems.

The problem of devising a replacement for handwritten signatures is a difficult one. Basically, what is needed is a system by which one party can send a "signed" message to another party in such a way that

1. The receiver can verify the claimed identity of the sender.
2. The sender cannot later repudiate the contents of the message.
3. The receiver cannot possibly have concocted the message himself.

The first requirement is needed, for example, in financial systems. When a customer's computer orders a bank's computer to buy a ton of gold, the bank's

computer needs to be able to make sure that the computer giving the order really belongs to the company whose account is to be debited.

The second requirement is needed to protect the bank against fraud. Suppose that the bank buys the ton of gold, and immediately thereafter the price of gold drops sharply. A dishonest customer might sue the bank, claiming that he never issued any order to buy gold. When the bank produces the message in court, the customer denies having sent it.

The third requirement is needed to protect the customer in the event that the price of gold shoots up and the bank tries to construct a signed message in which the customer asked for one bar of gold instead of one ton.

Secret-Key Signatures

One approach to digital signatures is to have a central authority that knows everything and whom everyone trusts, say Big Brother (*BB*). Each user then chooses a secret key and carries it by hand to *BB*'s office. Thus only Alice and *BB* know Alice's secret, K_A , and so on.

When Alice wants to send a signed plaintext message, P , to her banker, Bob, she generates $K_A(B, R_A, t, P)$ and sends it as depicted in Fig. 7-22. *BB* sees that the message is from Alice, decrypts it, and sends a message to Bob as shown. The message to Bob contains the plaintext of Alice's message and also the signed message $K_{BB}(A, t, P)$, where t is a timestamp. Bob now carries out Alice's request.




Fig. 7-22. Digital signatures with Big Brother.

What happens if Alice later denies sending the message? Step 1 is that everyone sues everyone (at least, in the United States). Finally, when the case comes to court and Alice vigorously denies sending Bob the disputed message, the judge will ask Bob how he can be sure that the disputed message came from Alice and not from Trudy. Bob first points out that *BB* will not accept a message from Alice unless it is encrypted with K_A , so there is no possibility of Trudy sending *BB* a false message from Alice.

Bob then dramatically produces Exhibit A, $K_{BB}(A, t, P)$. Bob says that this is a message signed by *BB* which proves Alice sent P to Bob. The judge then asks

BB (whom everyone trusts) to decrypt Exhibit A. When *BB* testifies that Bob is telling the truth, the judge decides in favor of Bob. Case dismissed.

One potential problem with the signature protocol of Fig. 7-22 is Trudy replaying either message. To minimize this problem, timestamps are used throughout. Furthermore, Bob can check all recent messages to see if R_A was used in any of them. If so, the message is discarded as a replay. Note that Bob will reject very old messages based on the timestamp. To guard against instant replay attacks, Bob just checks the R_A of every incoming message to see if such a message has been received from Alice in the past hour. If not, Bob can safely assume this is a new request.

Public-Key Signatures

A structural problem with using secret-key cryptography for digital signatures is that everyone has to agree to trust Big Brother. Furthermore, Big Brother gets to read all signed messages. The most logical candidates for running the Big Brother server are the government, the banks, or the lawyers. These organizations do not inspire total confidence in all citizens. Hence, it would be nice if signing documents did not require a trusted authority.

Fortunately, public-key cryptography can make an important contribution here. Let us assume that the public-key encryption and decryption algorithms have the property that $E(D(P)) = P$ in addition to the usual property that $D(E(P)) = P$. (RSA has this property, so the assumption is not unreasonable.) Assuming that this is the case, Alice can send a signed plaintext message, P , to Bob by transmitting $E_B(D_A(P))$. Note carefully that Alice knows her own (private) decryption key, D_A , as well as Bob's public key, E_B , so constructing this message is something Alice can do.

When Bob receives the message, he transforms it using his private key, as usual, yielding $D_A(P)$, as shown in Fig. 7-23. He stores this text in a safe place and then decrypts it using E_A to get the original plaintext.




Fig. 7-23. Digital signatures using public-key cryptography.

To see how the signature property works, suppose that Alice subsequently denies having sent the message P to Bob. When the case comes up in court, Bob

can produce both P and $D_A(P)$. The judge can easily verify that Bob indeed has a valid message encrypted by D_A by simply applying E_A to it. Since Bob does not know what Alice's private key is, the only way Bob could have acquired a message encrypted by it is if Alice did indeed send it. While in jail for perjury and fraud, Alice will have plenty of time to devise interesting new public-key algorithms.

Although using public-key cryptography for digital signatures is an elegant scheme, there are problems that are related to the environment in which they operate rather than with the basic algorithm. For one thing, Bob can prove that a message was sent by Alice only as long as D_A remains secret. If Alice discloses her secret key, the argument no longer holds, because anyone could have sent the message, including Bob himself.

The problem might arise, for example, if Bob is Alice's stockbroker. Alice tells Bob to buy a certain stock or bond. Immediately thereafter, the price drops sharply. To repudiate her message to Bob, Alice runs to the police claiming that her home was burglarized and her key was stolen. Depending on the laws in her state or country, she may or may not be legally liable, especially if she claims not to have discovered the break-in until getting home from work, several hours later.

Another problem with the signature scheme is what happens if Alice decides to change her key. Doing so is clearly legal, and it is probably a good idea to do so periodically. If a court case later arises, as described above, the judge will apply the *current* E_A to $D_A(P)$ and discover that it does not produce P . Bob will look pretty stupid at this point. Consequently, it appears that some authority is probably needed to record all key changes and their dates.

In principle, any public-key algorithm can be used for digital signatures. The de facto industry standard is the RSA algorithm. Many security products use it. However, in 1991, NIST (National Institute of Standards and Technology) proposed using a variant of the El Gamal public-key algorithm for their new **Digital Signature Standard (DSS)**. El Gamal gets its security from the difficulty of computing discrete logarithms, rather than the difficulty of factoring large numbers.

As usual when the government tries to dictate cryptographic standards, there was an uproar. DSS was criticized for being

1. Too secret (NSA designed the protocol for using El Gamal).
2. Too new (El Gamal has not yet been thoroughly analyzed).
3. Too slow (10 to 40 times slower than RSA for checking signatures).
4. Too insecure (fixed 512-bit key).

In a subsequent revision, the fourth point was rendered moot when keys up to 1024 bits were allowed. It is not yet clear whether DSS will catch on. For more details, see (Kaufman et al., 1995; Schneier, 1996; and Stinson, 1995).

Message Digests

One criticism of signature methods is that they often couple two distinct functions: authentication and secrecy. Often, authentication is needed but secrecy is not. Since cryptography is slow, it is frequently desirable to be able to send signed plaintext documents. Below we will describe an authentication scheme that does not require encrypting the entire message (De Jonge and Chaum, 1987).

This scheme is based on the idea of a one-way hash function that takes an arbitrarily long piece of plaintext and from it computes a fixed-length bit string. This hash function, often called a **message digest**, has three important properties:

1. Given P , it is easy to compute $MD(P)$.
2. Given $MD(P)$, it is effectively impossible to find P .
3. No one can generate two messages that have the same message digest.

To meet criterion 3, the hash should be at least 128 bits long, preferably more.

Computing a message digest from a piece of plaintext is much faster than encrypting that plaintext with a public-key algorithm, so message digests can be used to speed up digital signature algorithms. To see how this works, consider the signature protocol of Fig. 7-22 again. Instead of signing P with $K_{BB}(A, t, P)$, BB now computes the message digest by applying MD to P , yielding $MD(P)$. BB then encloses $K_{BB}(A, t, MD(P))$ as the fifth item in the list encrypted with K_B that is sent to Bob, instead of $K_{BB}(A, t, P)$.

If a dispute arises, Bob can produce both P and $K_{BB}(A, t, MD(P))$. After Big Brother has decrypted it for the judge, Bob has $MD(P)$, which is guaranteed to be genuine, and the alleged P . However, since it is effectively impossible for Bob to find any other message that gives this hash, the judge will easily be convinced that Bob is telling the truth. Using message digests in this way saves both encryption time and message transport and storage costs.

Message digests work in public-key cryptosystems, too, as shown in Fig. 7-24. Here, Alice first computes the message digest of her plaintext. She then signs the message digest and sends both the signed digest and the plaintext to Bob. If Trudy replaces P underway, Bob will see this when he computes $MD(P)$ himself.




Fig. 7-24. Digital signatures using message digests.

A variety of message digest functions have been proposed. The most widely used ones are MD5 (Rivest, 1992) and SHA (NIST, 1993). MD5 is the fifth in a

series of hash functions designed by Ron Rivest. It operates by mangling bits in a sufficiently complicated way that every output bit is affected by every input bit. Very briefly, it starts out by padding the message to a length of 448 bits (modulo 512). Then the original length of the message is appended as a 64-bit integer to give a total input whose length is a multiple of 512 bits. The last precomputation step is initializing a 128-bit buffer to a fixed value.

Now the computation starts. Each round takes a 512-bit block of input and mixes it thoroughly with the 128-bit buffer. For good measure, a table constructed from the sine function is also thrown in. The point of using a known function like the sine is not because it is more random than a random number generator, but to avoid any suspicion that the designer built in a clever trapdoor through which only he can enter. IBM's refusal to disclose the principles behind the design of the S-boxes in DES led to a great deal of speculation about trapdoors. Four rounds are performed per input block. This process continues until all the input blocks have been consumed. The contents of the 128-bit buffer form the message digest. The algorithm has been optimized for software implementation on 32-bit machines. As a consequence, it may not be fast enough for future high-speed networks (Touch, 1995).

The other major message digest function is **SHA (Secure Hash Algorithm)**, developed by NSA and blessed by NIST. Like MD5, it processes input data in 512-bit blocks, only unlike MD5, it generates a 160-bit message digest. It starts out by padding the message, then adding a 64-bit length to get a multiple of 512 bits. Then it initializes its 160-bit output buffer.

For each input block, the output buffer is updated using the 512-bit input block. No table of random numbers (or sine function values) is used, but for each block 80 rounds are computed, resulting in a thorough mixing. Each group of 20 rounds uses different mixing functions.

Since SHA's hash code is 32 bits longer than MD5's, all other things being equal, it is a factor of 2^{32} more secure than MD5. However, it is also slower than MD5, and having a hash code that is not a power of two might sometimes be an inconvenience. Otherwise, the two are roughly similar technically. Politically, MD5 is defined in an RFC and used heavily on the Internet. SHA is a government standard, and used by companies that have to use it because the government tells them to, or by those that want the extra security. A revised version, SHA-1, has been approved as a standard by NIST.

The Birthday Attack

In the world of crypto, nothing is ever what it seems to be. One might think that it would take on the order of 2^m operations to subvert an m -bit message digest. In fact, $2^{m/2}$ operations will often do using the **birthday attack**, an approach published by Yuval (1979) in his now-classic paper "How to Swindle Rabin."

The idea for this attack comes from a technique that math professors often use in their probability courses. The question is: How many students do you need in a class before the probability of having two people with the same birthday exceeds $1/2$? Most students expect the answer to be way over 100. In fact, probability theory says it is just 23. Without giving a rigorous analysis, intuitively, with 23 people, we can form $(23 \times 22)/2 = 253$ different pairs, each of which has a probability of $1/365$ of being a hit. In this light, it is not really so surprising any more.

More generally, if there is some mapping between inputs and outputs with n inputs (people, messages, etc.) and k possible outputs (birthdays, message digests, etc.), there are $n(n-1)/2$ input pairs. If $n(n-1)/2 > k$, the chance of having at least one match is pretty good. Thus, approximately, a match is likely for $n > \sqrt{k}$. This result means that a 64-bit message digest can probably be broken by generating about 2^{32} messages and looking for two with the same message digest.

Let us look at a practical example. The Dept. of Computer Science at State University has one position for a tenured faculty member and two candidates, Tom and Dick. Tom was hired two years before Dick, so he goes up for review first. If he gets it, Dick is out of luck. Tom knows that the department chairperson, Marilyn, thinks highly of his work, so he asks her to write him a letter of recommendation to the Dean, who will decide on Tom's case. Once sent, all letters become confidential.

Marilyn tells her secretary, Ellen, to write the Dean a letter, outlining what she wants in it. When it is ready, Marilyn will review it, compute and sign the 64-bit digest, and send it to the Dean. Ellen can send the letter later by email.

Unfortunately for Tom, Ellen is romantically involved with Dick and would like to do Tom in, so she writes the letter below with the 32 bracketed options.

Dear Dean Smith,

This [*letter* | *message*] is to give my [*honest* | *frank*] opinion of Prof. Tom Wilson, who is [*a candidate* | *up*] for tenure [*now* | *this year*]. I have [*known* | *worked with*] Prof. Wilson for [*about* | *almost*] six years. He is an [*outstanding* | *excellent*] researcher of great [*talent* | *ability*] known [*worldwide* | *internationally*] for his [*brilliant* | *creative*] insights into [*many* | *a wide variety of*] [*difficult* | *challenging*] problems.

He is also a [*highly* | *greatly*] [*respected* | *admired*] [*teacher* | *educator*]. His students give his [*classes* | *courses*] [*rave* | *spectacular*] reviews. He is [*our* | *the Department's*] [*most popular* | *best-loved*] [*teacher* | *instructor*].

[*In addition* | *Additionally*] Prof. Wilson is a [*gifted* | *effective*] fund raiser. His [*grants* | *contracts*] have brought a [*large* | *substantial*] amount of money into [*the* | *our*] Department. [*This money has* | *These funds have*] [*enabled* | *permitted*] us to [*pursue* | *carry out*] many [*special* | *important*] programs, [*such as* | *for example*] your State 2000 program. Without these funds we would [*be unable* | *not be able*] to continue this program, which is so [*important* | *essential*] to both of us. I strongly urge you to grant him tenure.

Unfortunately for Tom, as soon as Ellen finishes composing and typing in this letter, she also writes a second one:

Dear Dean Smith,

This [*letter* | *message*] is to give my [*honest* | *frank*] opinion of Prof. Tom Wilson, who is [*a candidate* | *up*] for tenure [*now* | *this year*]. I have [*known* | *worked with*] Tom for [*about* | *almost*] six years. He is a [*poor* | *weak*] researcher not well known in his [*field* | *area*]. His research [*hardly ever* | *rarely*] shows [*insight in* | *understanding of*] the [*key* | *major*] problems of [*the* | *our*] day.

Furthermore, he is not a [*respected* | *admired*] [*teacher* | *educator*]. His students give his [*classes* | *courses*] [*poor* | *bad*] reviews. He is [*our* | *the Department's*] least popular [*teacher* | *instructor*], known [*mostly* | *primarily*] within [*the* | *our*] Department for his [*tendency* | *propensity*] to [*ridicule* | *embarrass*] students [*foolish* | *imprudent*] enough to ask questions in his classes.

[*In addition* | *Additionally*] Tom is a [*poor* | *marginal*] fund raiser. His [*grants* | *contracts*] have brought only a [*meager* | *insignificant*] amount of money into [*the* | *our*] Department. Unless new [*money is* | *funds are*] quickly located, we may have to cancel some essential programs, such as your State 2000 program. Unfortunately, under these [*conditions* | *circumstances*] I cannot in good [*conscience* | *faith*] recommend him to you for [*tenure* | *a permanent position*].

Now Ellen sets up her computer to compute the 2^{32} message digests of each letter overnight. Chances are, one digest of the first letter will match one digest of the second letter. If not, she can add a few more options and try again during the weekend. Suppose that she finds a match. Call the “good” letter *A* and the “bad” one *B*.

Ellen now emails letter *A* to Marilyn for her approval. Marilyn, of course, approves, computes her 64-bit message digest, signs the digest, and emails the signed digest off to Dean Smith. Independently, Ellen emails letter *B* to the Dean.

After getting the letter and signed message digest, the Dean runs the message digest algorithm on letter *B*, sees that it agrees with what Marilyn sent him, and fires Tom. (Optional ending: Ellen tells Dick what she did. Dick is appalled and breaks off with her. Ellen is furious and confesses to Marilyn. Marilyn calls the Dean. Tom gets tenure after all.) With MD5 the birthday attack is infeasible because even at 1 billion digests per second, it would take over 500 years to compute all 2^{64} digests of two letters with 64 variants each, and even then a match is not guaranteed.

7.1.7. Social Issues

The implications of network security for individual privacy and society in general are staggering. Below we will just mention a few of the salient issues.

Governments do not like citizens keeping secrets from them. In some

countries (e.g., France) all nongovernmental cryptography is simply forbidden unless the government is given all the keys being used. As Kahn (1980) and Selfridge and Schwartz (1980) point out, government eavesdropping has been practiced on a far more massive scale than most people could dream of, and governments want more than just a pile of indecipherable bits for their efforts.

The U.S. government has proposed an encryption scheme for future digital telephones that includes a special feature to allow the police to tap and decrypt all telephone calls made in the United States. The government promises not to use this feature without a court order, but many people still remember how former FBI Director J. Edgar Hoover illegally tapped the telephones of Martin Luther King, Jr. and other people in an attempt to neutralize them. The police say they need this power to catch criminals. The debate on both sides is vehement, to put it mildly. A discussion of the technology involved (Clipper) is given in (Kaufman et al., 1995). A way to circumvent this technology and send messages that the government cannot read is described in (Blaze, 1994; and Schneier, 1996). Position statements on all sides are given in (Hoffman, 1995).

The United States has a law (22 U.S.C. 2778) that prohibits citizens from exporting munitions (war materiel), such as tanks and jet fighters, without authorization from the DoD. For purposes of this law, cryptographic software is classified as a munition. Phil Zimmermann, who wrote PGP (Pretty Good Privacy), an email protection program, has been accused of violating this law, even though the government admits that he did not export it (but he did give it to a friend who put it on the Internet where foreigners could obtain it). Many people regarded this widely-publicized incident as a gross violation of the rights of an American citizen working to enhance people's privacy.

Not being an American does not help. On July 9, 1986, three Israeli researchers working at the Weizmann Institute in Israel filed a U.S. patent application for a new digital signature scheme that they had invented. They spent the next 6 months discussing their research at conferences all over the world. On Jan. 6, 1987, the U.S. patent office told them to notify all Americans who knew about their results that disclosure of the research would subject them to two years in prison, a 10,000-dollar fine, or both. The patent office also wanted a list of all foreign nationals who knew about the research. To find out how this story turned out, see (Landau, 1988).

Patents are another hot topic. Nearly all public-key algorithms are patented. Patent protection lasts for 17 years. The RSA patent, for example, expires on Sept. 20, 2000.

Network security is politicized to an extent few other technical issues are, and rightly so, since it relates to the difference between a democracy and a police state in the digital era. The March 1993 and November 1994 issues of *Communications of the ACM* have long sections on telephone and network security, respectively, with vigorous arguments explaining and defending many points of view. Chapter 25 of Schneier's security book deals with the politics of cryptography

(Schneier, 1996). Chapter 8 of his email book does too (Schneier, 1995). Privacy and computers are also discussed in (Adam, 1995). These references are highly recommended for readers who wish to pursue their study of this subject.

7.2. DNS—Domain Name System

Programs rarely refer to hosts, mailboxes, and other resources by their binary network addresses. Instead of binary numbers, they use ASCII strings, such as *tana@art.ucsb.edu*. Nevertheless, the network itself only understands binary addresses, so some mechanism is required to convert the ASCII strings to network addresses. In the following sections we will study how this mapping is accomplished in the Internet.

Way back in the ARPANET, there was simply a file, *hosts.txt*, that listed all the hosts and their IP addresses. Every night, all the hosts would fetch it from the site at which it was maintained. For a network of a few hundred large timesharing machines, this approach worked reasonably well.

However, when thousands of workstations were connected to the net, everyone realized that this approach could not continue to work forever. For one thing, the size of the file would become too large. However, even more important, host name conflicts would occur constantly unless names were centrally managed, something unthinkable in a huge international network. To solve these problems, **DNS (the Domain Name System)** was invented.

The essence of DNS is the invention of a hierarchical, domain-based naming scheme and a distributed database system for implementing this naming scheme. It is primarily used for mapping host names and email destinations to IP addresses but can also be used for other purposes. DNS is defined in RFCs 1034 and 1035.

Very briefly, the way DNS is used is as follows. To map a name onto an IP address, an application program calls a library procedure called the **resolver**, passing the name as a parameter. The resolver sends a UDP packet to a local DNS server, which then looks up the name and returns the IP address to the resolver, which then returns it to the caller. Armed with the IP address, the program can then establish a TCP connection with the destination, or send it UDP packets.

7.2.1. The DNS Name Space

Managing a large and constantly changing set of names is a nontrivial problem. In the postal system, name management is done by requiring letters to specify (implicitly or explicitly) the country, state or province, city, and street address of the addressee. By using this kind of hierarchical addressing, there is no confusion between the Marvin Anderson on Main St. in White Plains, N.Y. and the Marvin Anderson on Main St. in Austin, Texas. DNS works the same way.

Conceptually, the Internet is divided into several hundred top-level **domains**, where each domain covers many hosts. Each domain is partitioned into subdomains, and these are further partitioned, and so on. All these domains can be represented by a tree, as shown in Fig. 7-25. The leaves of the tree represent domains that have no subdomains (but do contain machines, of course). A leaf domain may contain a single host, or it may represent a company and contains thousands of hosts.




Fig. 7-25. A portion of the Internet domain name space.

The top-level domains come in two flavors: generic and countries. The generic domains are *com* (*commercial*), *edu* (educational institutions), *gov* (the U.S. federal government), *int* (certain international organizations), *mil* (the U.S. armed forces), *net* (network providers), and *org* (nonprofit organizations). The country domains include one entry for every country, as defined in ISO 3166.

Each domain is named by the path upward from it to the (unnamed) root. The components are separated by periods (pronounced “dot”). Thus Sun Microsystems engineering department might be *eng.sun.com.*, rather than a UNIX-style name such as */com/sun/eng*. Notice that this hierarchical naming means that *eng.sun.com.* does not conflict with a potential use of *eng* in *eng.yale.edu.*, which might be used by the Yale English department.

Domain names can be either absolute or relative. An absolute domain name ends with a period (e.g., *eng.sun.com.*), whereas a relative one does not. Relative names have to be interpreted in some context to uniquely determine their true meaning. In both cases, a named domain refers to a specific node in the tree and all the nodes under it.

Domain names are case insensitive, so *edu* and *EDU* mean the same thing. Component names can be up to 63 characters long, and full path names must not exceed 255 characters.

In principle, domains can be inserted into the tree in two different ways. For example, *cs.yale.edu* could equally well be listed under the *us* country domain as

cs.yale.ct.us. In practice, however, nearly all organizations in the United States are under a generic domain, and nearly all outside the United States are under the domain of their country. There is no rule against registering under two top-level domains, but doing so might be confusing, so few organizations do it.

Each domain controls how it allocates the domains under it. For example, Japan has domains *ac.jp* and *co.jp* that mirror *edu* and *com*. The Netherlands does not make this distinction and puts all organizations directly under *nl*. Thus all three of the following are university computer science departments:

1. *cs.yale.edu* (Yale University, in the United States)
2. *cs.vu.nl* (Vrije Universiteit, in The Netherlands)
3. *cs.keio.ac.jp* (Keio University, in Japan)

To create a new domain, permission is required of the domain in which it will be included. For example, if a VLSI group is started at Yale and wants to be known as *vlsi.cs.yale.edu*, it needs permission from whomever manages *cs.yale.edu*. Similarly, if a new university is chartered, say, the University of Northern South Dakota, it must ask the manager of the *edu* domain to assign it *unsd.edu*. In this way, name conflicts are avoided and each domain can keep track of all its subdomains. Once a new domain has been created and registered, it can create subdomains, such as *cs.unsd.edu*, without getting permission from anybody higher up the tree.

Naming follows organizational boundaries, not physical networks. For example, if the computer science and electrical engineering departments are located in the same building and share the same LAN, they can nevertheless have distinct domains. Similarly, even if computer science is split over Babbage Hall and Turing Hall, all the hosts in both buildings will normally belong to the same domain.

7.2.2. Resource Records

Every domain, whether it is a single host or a top-level domain, can have a set of **resource records** associated with it. For a single host, the most common resource record is just its IP address, but many other kinds of resource records also exist. When a resolver gives a domain name to DNS, what it gets back are the resource records associated with that name. Thus the real function of DNS is to map domain names onto resource records.

A resource record is a five-tuple. Although they are encoded in binary for efficiency, in most expositions resource records are presented as ASCII text, one line per resource record. The format we will use is as follows:

```
Domain_name Time_to_live Type Class Value
```

The *Domain_name* tells the domain to which this record applies. Normally, many records exist for each domain and each copy of the database holds information

about multiple domains. This field is thus the primary search key used to satisfy queries. The order of the records in the database is not significant. When a query is made about a domain, all the matching records of the class requested are returned.

The *Time_to_live* field gives an indication of how stable the record is. Information that is highly stable is assigned a large value, such as 86400 (the number of seconds in 1 day). Information that is highly volatile is assigned a small value, such as 60 (1 minute). We will come back to this point later when we have discussed caching.

The *Type* field tells what kind of record this is. The most important types are listed in Fig. 7-26.

Type	Meaning	Value
SOA	Start of Authority	Parameters for this zone
A	IP address of a host	32-Bit integer
MX	Mail exchange	Priority, domain willing to accept email
NS	Name Server	Name of a server for this domain
CNAME	Canonical name	Domain name
PTR	Pointer	Alias for an IP address
HINFO	Host description	CPU and OS in ASCII
TXT	Text	Uninterpreted ASCII text

Fig. 7-26. The principal DNS resource record types.

An *SOA* record provides the name of the primary source of information about the name server's zone (described below), the email address of its administrator, a unique serial number, and various flags and timeouts.

The most important record type is the *A* (Address) record. It holds a 32-bit IP address for some host. Every Internet host must have at least one IP address, so other machines can communicate with it. Some hosts have two or more network connections, in which case they will have one type *A* resource record per network connection (and thus per IP address).

The next most important record type is the *MX* record. It specifies the name of the domain prepared to accept email for the specified domain. A common use of this record is to allow a machine that is not on the Internet to receive email from Internet sites. Delivery is accomplished by having the non-Internet site make an arrangement with some Internet site to accept email for it and forward it using whatever protocol the two of them agree on.

For example, suppose that Cathy is a computer science graduate student at UCLA. After she gets her degree in AI, she sets up a company, Electrobrain

Corporation, to commercialize her ideas. She cannot afford an Internet connection yet, so she makes an arrangement with UCLA to allow her to have her email sent there. A few times a day she will call up and collect it.

Next, she registers her company with the *com* domain and is assigned the domain *electrobrain.com*. She might then ask the administrator of the *com* domain to add an *MX* record to the *com* database as follows:

```
electrobrain.com 86400 IN MX 1 mailserver.cs.ucla.edu
```

In this way, mail will be forwarded to UCLA where she can pick it up by logging in. Alternatively, UCLA could call her and transfer the email by any protocol they mutually agree on.

The *NS* records specify name servers. For example, every DNS database normally has an *NS* record for each of the top-level domains, so email can be sent to distant parts of the naming tree. We will come back to this point later.

CNAME records allow aliases to be created. For example, a person familiar with Internet naming in general wanting to send a message to someone whose login name is *paul* in the computer science department at M.I.T. might guess that *paul@cs.mit.edu* will work. Actually this address will not work, because the domain for M.I.T.'s computer science department is *lcs.mit.edu*. However, as a service to people who do not know this, M.I.T. could create a *CNAME* entry to point people and programs in the right direction. An entry like this one might do the job:

```
cs.mit.edu 86400 IN CNAME lcs.mit.edu
```

Like *CNAME*, *PTR* points to another name. However, unlike *CNAME*, which is really just a macro definition, *PTR* is a regular DNS datatype whose interpretation depends on the context in which it is found. In practice, it is nearly always used to associate a name with an IP address to allow lookups of the IP address and return the name of the corresponding machine.

HINFO records allow people to find out what kind of machine and operating system a domain corresponds to. Finally, *TXT* records allow domains to identify themselves in arbitrary ways. Both of these record types are for user convenience. Neither is required, so programs cannot count on getting them (and probably cannot deal with them if they do get them).

Getting back to the general structure of resource records, the fourth field of every resource record is the *Class*. For Internet information, it is always *IN*. For non-Internet information, other codes can be used.

Finally, we come to the *Value* field. This field can be a number, a domain name, or an ASCII string. The semantics depend on the record type. A short description of the *Value* fields for each of the principal records types is given in Fig. 7-26.

As an example of the kind of information one might find in the DNS database of a domain, see Fig. 7-27. This figure depicts part of a (semihypothetical)

database for the *cs.vu.nl* domain shown in Fig. 7-25. The database contains seven types of resource records.

```
; Authoritative data for cs.vu.nl
cs.vu.nl.      86400  IN  SOA   star boss (952771,7200,7200,2419200,86400)
cs.vu.nl.      86400  IN  TXT   "Faculteit Wiskunde en Informatica."
cs.vu.nl.      86400  IN  TXT   "Vrije Universiteit Amsterdam."
cs.vu.nl.      86400  IN  MX    1 zephyr.cs.vu.nl.
cs.vu.nl.      86400  IN  MX    2 top.cs.vu.nl.

flits.cs.vu.nl. 86400  IN  HINFO Sun Unix
flits.cs.vu.nl. 86400  IN  A     130.37.16.112
flits.cs.vu.nl. 86400  IN  A     192.31.231.165
flits.cs.vu.nl. 86400  IN  MX    1 flits.cs.vu.nl.
flits.cs.vu.nl. 86400  IN  MX    2 zephyr.cs.vu.nl.
flits.cs.vu.nl. 86400  IN  MX    3 top.cs.vu.nl.
www.cs.vu.nl.   86400  IN  CNAME star.cs.vu.nl
ftp.cs.vu.nl.   86400  IN  CNAME zephyr.cs.vu.nl

rowboat         IN  A     130.37.56.201
                IN  MX    1 rowboat
                IN  MX    2 zephyr
                IN  HINFO Sun Unix

little-sister   IN  A     130.37.62.23
                IN  HINFO Mac MacOS

laserjet        IN  A     192.31.231.216
                IN  HINFO "HP Laserjet IIISi" Proprietary
```

Fig. 7-27. A portion of a possible DNS database for *cs.vu.nl*

The first noncomment line of Fig. 7-27 gives some basic information about the domain, which will not concern us further. The next two lines give textual information about where the domain is located. Then come two entries giving the first and second places to try to deliver email sent to *person@cs.vu.nl*. The *zephyr* (a specific machine) should be tried first. If that fails, the *top* should be tried next.

After the blank line, added for readability, come lines telling that the *flits* is a Sun workstation running UNIX and giving both of its IP addresses. Then three choices are given for handling email sent to *flits.cs.vu.nl*. First choice is naturally the *flits* itself, but if it is down, the *zephyr* and *top* are the second and third choices. Next comes an alias, *www.cs.vu.nl*, so that this address can be used without designating a specific machine. Creating this alias allows *cs.vu.nl* to change its World Wide Web server without invalidating the address people use to get to it. A similar argument holds for *ftp.cs.vu.nl*.

The next four lines contain a typical entry for a workstation, in this case, *rowboat.cs.vu.nl*. The information provided contains the IP address, the primary and secondary mail drops, and information about the machine. Then comes an entry for a non-UNIX system that is not capable of receiving mail itself, followed by an entry for a laser printer.

What is not shown (and is not in this file), are the IP addresses to use to look up the top level domains. These are needed to look up distant hosts, but since they are not part of the *cs.vu.nl* domain, they are not in this file. They are supplied by the root servers, whose IP addresses are present in a system configuration file and loaded into the DNS cache when the DNS server is booted. They have very long timeouts, so once loaded, they are never purged from the cache.

7.2.3. Name Servers

In theory at least, a single name server could contain the entire DNS database and respond to all queries about it. In practice, this server would be so overloaded as to be useless. Furthermore, if it ever went down, the entire Internet would be crippled.

To avoid the problems associated with having only a single source of information, the DNS name space is divided up into nonoverlapping **zones**. One possible way to divide up the name space of Fig. 7-25 is shown in Fig. 7-28. Each zone contains some part of the tree and also contains name servers holding the authoritative information about that zone. Normally, a zone will have one primary name server, which gets its information from a file on its disk, and one or more secondary name servers, which get their information from the primary name server. To improve reliability, some servers for a zone can be located outside the zone.




Fig. 7-28. Part of the DNS name space showing the division into zones.

Where the zone boundaries are placed within a zone is up to that zone's administrator. This decision is made in large part based on how many name

servers are desired, and where. For example, in Fig. 7-28, Yale has a server for *yale.edu* that handles *eng.yale.edu* but not *cs.yale.edu*, which is a separate zone with its own name servers. Such a decision might be made when a department such as English does not wish to run its own name server, but a department such as computer science does. Consequently, *cs.yale.edu* is a separate zone but *eng.yale.edu* is not.

When a resolver has a query about a domain name, it passes the query to one of the local name servers. If the domain being sought falls under the jurisdiction of the name server, such as *ai.cs.yale.edu* falling under *cs.yale.edu*, it returns the authoritative resource records. An **authoritative record** is one that comes from the authority that manages the record, and is thus always correct. Authoritative records are in contrast to cached records, which may be out of date.

If, however, the domain is remote and no information about the requested domain is available locally, the name server sends a query message to the top-level name server for the domain requested. To make this process clearer, consider the example of Fig. 7-29. Here, a resolver on *flits.cs.vu.nl* wants to know the IP address of the host *linda.cs.yale.edu*. In step 1, it sends a query to the local name server, *cs.vu.nl*. This query contains the domain name sought, the type (A) and the class (IN).




Fig. 7-29. How a resolver looks up a remote name in eight steps.

Let us suppose the local name server has never had a query for this domain before and knows nothing about it. It may ask a few other nearby name servers, but if none of them know, it sends a UDP packet to the server for *edu* given in its database (see Fig. 7-29), *edu-server.net*. It is unlikely that this server knows the address of *linda.cs.yale.edu*, and probably does not know *cs.yale.edu* either, but it must know all of its own children, so it forwards the request to the name server for *yale.edu* (step 3). In turn, this one forwards the request to *cs.yale.edu* (step 4), which must have the authoritative resource records. Since each request is from a client to a server, the resource record requested works its way back in steps 5 through 8.

Once these records get back to the *cs.vu.nl* name server, they will be entered into a cache there, in case they are needed later. However, this information is not authoritative, since changes made at *cs.yale.edu* will not be propagated to all the caches in the world that may know about it. For this reason, cache entries should not live too long. This is the reason that the *Time_to_live* field is included in each resource record. It tells remote name servers how long to cache records. If a

certain machine has had the same IP address for years, it may be safe to cache that information for 1 day. For more volatile information, it might be safer to purge the records after a few seconds or a minute.

It is worth mentioning that the query method described here is known as a **recursive query**, since each server that does not have the requested information goes and finds it somewhere, then reports back. An alternative form is also possible. In this form, when a query cannot be satisfied locally, the query fails, but the name of the next server along the line to try is returned. This procedure gives the client more control over the search process. Some servers do not implement recursive queries and always return the name of the next server to try.

It is also worth pointing out that when a DNS client fails to get a response before its timer goes off, it normally will try another server next time. The assumption here is that the server is probably down, rather than the request or reply got lost.

7.3. SNMP—SIMPLE NETWORK MANAGEMENT PROTOCOL

In the early days of the ARPANET, if the delay to some host became unexpectedly large, the person detecting the problem would just run the Ping program to bounce a packet off the destination. By looking at the timestamps in the header of the packet returned, the location of the problem could usually be pinpointed and some appropriate action taken. In addition, the number of routers was so small, that it was feasible to ping each one to see if it was sick.

When the ARPANET turned into the worldwide Internet, with multiple backbones and multiple operators, this solution ceased to be adequate, so better tools for network management were needed. Two early attempts were defined in RFC 1028 and RFC 1067, but these were short lived. In May 1990, RFC 1157 was published, defining version 1 of **SNMP (Simple Network Management Protocol)**. Along with a companion document (RFC 1155) on management information, SNMP provided a systematic way of monitoring and managing a computer network. This framework and protocol were widely implemented in commercial products and became the de facto standards for network management.

As experience was gained, shortcomings in SNMP came to light, so an enhanced version of SNMP (SNMPv2) was defined (in RFCs 1441 to 1452) and started along the road to become an Internet standard. In the sections to follow, we will give a brief discussion of the SNMP (meaning SNMPv2) model and protocol.

Although SNMP was designed with the idea of its being simple, at least one author has managed to produce a 600-page book on it (Stallings, 1993a). For more compact descriptions (450-550 pages), see the books by Rose (1994) and Rose and McCloghrie (1995), both of whom were among the designers of SNMP. Other references are (Feit, 1995; and Hein and Griffiths, 1995).

7.3.1. The SNMP Model

The SNMP model of a managed network consists of four components:

1. Managed nodes.
2. Management stations.
3. Management information.
4. A management protocol.

These pieces are illustrated in Fig. 7-30 and discussed below.




Fig. 7-30. Components of the SNMP management model.

The managed nodes can be hosts, routers, bridges, printers, or any other devices capable of communicating status information to the outside world. To be managed directly by SNMP, a node must be capable of running an SNMP management process, called an **SNMP agent**. All computers meet this requirement, as do increasingly many bridges, routers, and peripheral devices designed for network use. Each agent maintains a local database of variables that describe its state and history and affect its operation.

Network management is done from **management stations**, which are, in fact, general-purpose computers running special management software. The management stations contain one or more processes that communicate with the agents over the network, issuing commands and getting responses. In this design, all the intelligence is in the management stations, in order to keep the agents as simple as possible and minimize their impact on the devices they are running on. Many management stations have a graphical user interface to allow the network manager to inspect the status of the network and take action when required.

Most real networks are multivendor, with hosts from one or more manufacturers, bridges and routers from other companies, and printers from still other ones.

In order to allow a management station (potentially from yet another supplier) to talk to all these diverse components, the nature of the information maintained by all the devices must be rigidly specified. Having the management station ask a router what its packet loss rate is of no use if the router does not keep track of its loss rate. Therefore, SNMP describes (in excruciating detail) the exact information each kind of agent has to maintain and the format it has to supply it in. The largest portion of the SNMP model is the definition of who has to keep track of what and how this information is communicated.

Very briefly, each device maintains one or more variables that describe its state. In the SNMP literature, these variables are called **objects**, but the term is misleading because they are not objects in the sense of an object-oriented system because they just have state and no methods (other than reading and writing their values). Nevertheless, the term is so ingrained (e.g., used in various reserved words in the specification language used) that we will use it here. The collection of all possible objects in a network is given in a data structure called the **MIB (Management Information Base)**.

The management station interacts with the agents using the SNMP protocol. This protocol allows the management station to query the state of an agent's local objects, and change them if necessary. Most of SNMP consists of this query-response type communication.

However, sometimes events happen that are not planned. Managed nodes can crash and reboot, lines can go down and come back up, congestion can occur, and so on. Each significant event is defined in a MIB module. When an agent notices that a significant event has occurred, it immediately reports the event to all management stations in its configuration list. This report is called an **SNMP trap** (for historical reasons). The report usually just states that some event has occurred. It is up to the management station to then issue queries to find out all the gory details. Because communication from managed nodes to the management station is not reliable (i.e., is not acknowledged), it is wise for the management station to poll each managed node occasionally anyway, checking for unusual events, just in case. The model of polling at long intervals with acceleration on receipt of a trap is called **trap directed polling**.

This model assumes that each managed node is capable of running an SNMP agent internally. Older devices or devices not originally intended for use on a network may not have this capability. To handle them, SNMP defines what is called a **proxy agent**, namely an agent that watches over one or more nonSNMP devices and communicates with the management station on their behalf, possibly communicating with the devices themselves using some nonstandard protocol.

Finally, security and authentication play a major role in SNMP. A management station has the capability of learning a great deal about every node under its control and also has the capability of shutting them all down. Hence it is of great importance that agents be convinced that queries allegedly coming from the management station, in fact, come from the management station. In SNMPv1, the

management station proved who it was by putting a (plaintext) password in each message. In SNMPv2, security was improved considerably using modern cryptographic techniques of the type we have already studied. However, this addition made an already bulky protocol every bulkier, and it was later thrown out.

7.3.2. ASN.1—Abstract Syntax Notation 1

The heart of the SNMP model is the set of objects managed by the agents and read and written by the management station. To make multivendor communication possible, it is essential that these objects be defined in a standard and vendor-neutral way. Furthermore, a standard way is needed to encode them for transfer over a network. While definitions in C would satisfy the first requirement, such definitions do not define a bit encoding on the wire in such a way that a 32-bit two's complement little endian management station can exchange information unambiguously with an agent on a 16-bit one's complement big endian CPU.

For this reason, a standard object definition language, along with encoding rules, is needed. The one used by SNMP is taken from OSI and called **ASN.1 (Abstract Syntax Notation One)**. Like much of OSI, it is large, complex, and not especially efficient. (The author is tempted to say that by calling it ASN.1 instead of just ASN, the designers implicitly admitted that it would soon be replaced by ASN.2, but he will politely refrain from saying this.) The one alleged strength of ASN.1 (the existence of unambiguous bit encoding rules) is now really a weakness, because the encoding rules are optimized to minimize the number of bits on the wire, at the cost of wasting CPU time at both ends encoding and decoding them. A simpler scheme, using 32-bit integers aligned on 4-byte boundaries would probably have been better. Nevertheless, for better or worse, SNMP is drenched in ASN.1, (albeit a simplified subset of it), so anyone wishing to truly understand SNMP must become fluent in ASN.1. Hence the following explanation.

Let us start with the data description language, described in International Standard 8824. After that we will discuss the encoding rules, described in International Standard 8825. The ASN.1 abstract syntax is essentially a primitive data declaration language. It allows the user to define primitive objects and then combine them into more complex ones. A series of declarations in ASN.1 is functionally similar to the declarations found in the header files associated with many C programs.

SNMP has some lexical conventions that we will follow. These are not entirely the same as pure ASN.1 uses, however. Built-in data types are written in uppercase (e.g., *INTEGER*). User-defined types begin with an uppercase letter but must contain at least one character other than an uppercase letter. Identifiers may contain upper and lowercase letters, digits, and hyphens, but must begin with a lowercase letter (e.g., *counter*). White space (tabs, carriage returns, etc.) is not

significant. Finally, comments start with -- and continue until the end of the line or the next occurrence of --.

The ASN.1 basic data types allowed in SNMP are shown in Fig. 7-31. (We will generally ignore features of ASN.1, such as *BOOLEAN* and *REAL* types, not permitted in SNMP.) The use of the codes will be described later.

Primitive type	Meaning	Code
INTEGER	Arbitrary length integer	2
BIT STRING	A string of 0 or more bits	3
OCTET STRING	A string of 0 or more unsigned bytes	4
NULL	A place holder	5
OBJECT IDENTIFIER	An officially defined data type	6

Fig. 7-31. The ASN.1 primitive data types permitted in SNMP.

A variable of type *INTEGER* may, in theory, take on any integral value, but other SNMP rules limit the range. As an example of how types are used, consider how a variable, *count*, of type *INTEGER* would be declared and (optionally) initialized to 100 in ASN.1:

```
count INTEGER ::= 100
```

Often a subtype whose variables are restricted to specific values or to a specific range is required. These can be declared as follows:

```
Status ::= INTEGER { up(1), down(2), unknown(3) }
```

```
PacketSize ::= INTEGER (0..1023)
```

Variables of type *BIT STRING* and *OCTET STRING* contain zero or more bits and bytes, respectively. A bit is either 0 or 1. A byte falls in the range 0 to 255, inclusive. For both types, a string length and an initial value may be given.

OBJECT IDENTIFIERS provide a way of identifying objects. In principle, every object defined in every official standard can be uniquely identified. The mechanism that is used is to define a standards tree, and place every object in every standard at a unique location in the tree. The portion of the tree that includes the SNMP MIB is shown in Fig. 7-32.

The top level of the tree lists all the important standards organizations in the world (in ISO's view), namely ISO and CCITT (now ITU), plus the combination of the two. From the *iso* node, four arcs are defined, one of which is for *identified-organization*, which is ISO's concession that maybe some other folks are vaguely involved with standards, too. The U.S. Dept. of Defense has been assigned a place in this subtree, and DoD has assigned the Internet number 1 in its hierarchy. Under the Internet hierarchy, the SNMP MIB has code 1.




Fig. 7-32. Part of the ASN.1 object naming tree.

Every arc in Fig. 7-32 has both a label and a number, so nodes can be identified by a list of arcs, using label(number) or numbers. Thus all SNMP MIB objects are identified by a label of the form

{iso identified-organization(3) dod(6) internet(1) mgmt(2) mib-2(1) ...}

or alternatively {1 3 6 1 2 1 ...}. Mixed forms are also permitted. For example, the above identification can also be written as

{internet(1) 2 1 ...}

In this way, every object in every standard can be represented as an *OBJECT IDENTIFIER*.

ASN.1 defines five ways to construct new types from the basic ones. *SEQUENCE* is an ordered list of types, similar to a structure in C and a record in Pascal. *SEQUENCE OF* is a one-dimensional array of a single type. *SET* and *SET OF* are analogous, but unordered. *CHOICE* creates a union from a given list of types. The two set constructors are not used in any of the SNMP documents.

Another way to create new types is to tag old ones. Tagging a type is somewhat similar to the practice in C of defining new types, say *time_t* and *size_t*, both of which are longs, but which are used in different contexts. Tags come in four

categories: universal, application-wide, context-specific and private. Each tag consists of a label and an integer identifying the tag. For example,

```
Counter32 ::= [APPLICATION 1] INTEGER (0..4294967295)
```

```
Gauge32 ::= [APPLICATION 2] INTEGER (0..4294967295)
```

define two different application-wide types, both of which are implemented by 32-bit unsigned integers, but which are conceptually different. The former might, for example, wrap around when it gets to the maximum value, whereas the latter might just continue to return the maximum value until its is decreased or reset.

A tagged type can have the keyword *IMPLICIT* after the closing square bracket when the type of what follows is obvious from the context (not true in a *CHOICE*, for example). Doing so allows a more efficient bit encoding since the tag does not have to be transmitted. In a type involving a *CHOICE* between two different types, a tag must be transmitted to tell the receiver which type is present.

ASN.1 defines a complex macro mechanism, which is heavily used in SNMP. A macro can be used as a kind of prototype to generate a set of new types and values, each with its own syntax. Each macro defines some (possibly optional) keywords, that are used in the call to identify which parameter is which (i.e., the macro parameters are identified by keyword, not by position). The details of how ASN.1 macros work is beyond the scope of this book. Suffice it to say that a macro is invoked by giving its name and then listing (some of) its keywords and their values for this invocation. Macros are expanded at compile time, not at run time. Some examples of macros will be cited below.

ASN.1 Transfer Syntax

An ASN.1 **transfer syntax** defines how values of ASN.1 types are unambiguously converted to a sequence of bytes for transmission (and unambiguously decoded at the other end). The transfer syntax used by ASN.1 is called **BER (Basic Encoding Rules)**. ASN.1 has other transfer syntaxes that SNMP does not use. The rules are recursive, so the encoding of a structured object is just the concatenation of the encodings of the component objects. In this way, all object encodings can be reduced to a well-defined sequence of encoded primitive objects. The encoding of these objects, in turn, is defined by the BER.

The guiding principle behind the basic encoding rules is that every value transmitted, both primitive and constructed ones, consists of up to four fields:

1. The identifier (type or tag).
2. The length of the data field, in bytes.
3. The data field.
4. The end-of-contents flag, if the data length is unknown.

The last one is permitted by ASN.1, but specifically forbidden by SNMP, so we will assume the data length is always known.

The first field identifies the item that follows. It, itself, has three subfields, as shown in Fig. 7-33. The high-order 2 bits identify the tag type. The next bit tells whether the value is primitive (0) or not (1). The tag bits are 00, 01, 10, and 11, for *UNIVERSAL*, *APPLICATION*, context-specific, and *PRIVATE*, respectively. The remaining 5 bits can be used to encode the value of the tag if it is in the range 0 through 30. If the tag is 31 or more, the low-order 5 bits contain 11111, with the true value in the next byte or bytes.




Fig. 7-33. The first byte of each data item sent in the ASN.1 transfer syntax.

The rule used to encode tags greater than 30 has been designed to handle arbitrarily large numbers. Each identifier byte following the first one contains 7 data bits. The high-order bit is set to 0 in all but the last one. Thus tag values up to $2^7 - 1$ can be handled in 2 bytes, and up to $2^{14} - 1$ can be handled in 3 bytes.

The encoding of the *UNIVERSAL* types is straightforward. Each primitive type has been assigned a code, as given in the third column of Fig. 7-31. *SEQUENCE* and *SEQUENCE OF* share code 16. *CHOICE* does not have a code, since any actual value sent always has a specific type. The other codes are for types not used in SNMP.

Following the 1-byte identifier field comes a field telling how many bytes the data occupy. Lengths shorter than 128 bytes are directly encoded in 1 byte whose leftmost bit is 0. Those that are longer use multiple bytes, with first byte containing a 1 in the high-order bit and the length field (up to 127 bytes) in the low-order 7 bits. For example, if the data length is 1000 bytes, the first byte contains 130 to indicate a two byte length field follows. Then come two bytes whose value is 1000, with the high-order byte first.

The encoding of the data field depends on the type of data present. Integers are encoded in two's complement. A positive integer below 128 requires 1 byte, a positive integer below 32,768 requires 2 bytes, and so forth. The most significant byte is transmitted first.

Bit strings are encoded as themselves. The only problem is how to indicate the length. The length field tells how many *bytes* the value has, not how many

bits. The solution chosen is to transmit 1 byte before the actual bit string telling how many bits (0 through 7) of the final byte are unused. Thus the encoding of the 9-bit string '010011111' would be 07, 4F, 80 (hexadecimal).

Octet strings are easy. The bytes of the string are just transmitted in standard big endian style, left to right.

The null value is indicated by setting the length field to 0. No numerical value is actually transmitted.

An *OBJECT IDENTIFIER* is encoded as the sequence of integers it represents. For example, the Internet is {1, 3, 6, 1}. However, since the first number is always 0, 1, or 2, and the second is less than 40 (by definition—ISO simply will not recognize the 41st category to show up on its doorstep), the first two numbers, *a* and *b*, are encoded as 1 byte having the value $40a + b$. For the Internet, this number is 43. As usual, numbers exceeding 127 are encoded in multiple bytes, the first of which contains the high-order bit set to 1 and a byte count in the other 7 bits.

Both sequence types are transmitted by first sending the type or tag, then the total length of the encoding for all the fields, followed by the fields themselves. The fields are sent in order.

The encoding of a *CHOICE* value is the same as the encoding of the actual data structure being transferred.

An example showing encoding of some values is given in Fig. 7-34. The values encoded are the *INTEGER* 49, the *OCTET STRING* '110', "xy", the only possible value for *NULL*, the *OBJECT IDENTIFIER* for the Internet {1, 3, 6, 1}, and a *Gauge32* value of 14.




Fig. 7-34. ASN.1 encoding of some example values.

7.3.3. SMI—Structure of Management Information

In the preceding section, we have discussed only those parts of ASN.1 that are used in SNMP. In reality, the SNMP documents are organized differently. RFC 1442 first says that ASN.1 will be used to describe SNMP data structures, then it goes on for 57 pages scratching out parts of the ASN.1 standard that it does not want and adding new definitions (in ASN.1) that are needed. In particular, RFC 1442 defines four key macros and eight new data types that are heavily used throughout SNMP. It is this sub-super-set of ASN.1, which goes by the ungainly name of **SMI (Structure of Management Information)**, that is really used to define the SNMP data structures.

Although this approach is somewhat bureaucratic, some rules and regulations are necessary if products from hundreds of vendors are expected to talk to one another and actually understand what the others are saying. A few words about SMI are therefore now in order.

At the lowest level, SNMP variables are defined as individual objects. Related objects are collected together into groups, and groups are assembled into modules. For example, groups exist for IP objects and TCP objects. A router might support the IP group, since its manager cares about how many packets it has lost. On the other hand, a low-end router might not support the TCP group, since it need not use TCP to perform its routing functions. It is the intention that vendors supporting a group support all the objects in that group. However, a vendor supporting a module need not support all of its groups, since not all may be applicable to the device.

All MIB modules start with an invocation of the *MODULE-IDENTITY* macro. Its parameters provide the name and address of the implementer, the revision history, and other administrative information. Typically, this call is followed by an invocation of the *OBJECT-IDENTITY* macro, which tells where the module fits in the naming tree of Fig. 7-32.

Later on come one or more invocations of the *OBJECT-TYPE* macro, which name the actual variables being managed and specify their properties. Grouping variables into groups is done by convention; there are no *BEGIN-GROUP* and *END-GROUP* statements in ASN.1 or SMI.

The *OBJECT-TYPE* macro has four required parameters and four (sometimes) optional ones. The first required parameter is *SYNTAX* and defines the variable's data type from among the types listed in Fig. 7-35. For the most part, these types should be self explanatory, with the following comments. The suffix 32 is used when the implementer really wants a 32-bit number, even if all the machines in sight have 64-bit CPUs. Gauges differ from counters in that they do not wrap around when they hit their limits. They stick there. If a router has lost exactly 2^{32} packets, it is better to report this as $2^{32} - 1$ than as 0. SMI also supports arrays, but we will not go into those here. For details, see (Rose, 1994).

In addition to requiring a specification of the data type used by the variable

Name	Type	Bytes	Meaning
INTEGER	Numeric	4	Integer (32 bits in current implementations)
Counter32	Numeric	4	Unsigned 32-bit counter that wraps
Gauge32	Numeric	4	Unsigned value that does not wrap
Integer32	Numeric	4	32 Bits, even on a 64-bit CPU
UInteger32	Numeric	4	Like Integer32, but unsigned
Counter64	Numeric	8	A 64-bit counter
TimeTicks	Numeric	4	In hundredths of a second since some epoch
BIT STRING	String	4	Bit map of 1 to 32 bits
OCTET STRING	String	≥ 0	Variable length byte string
Opaque	String	≥ 0	Obsolete; for backward compatibility only
OBJECT IDENTIFIER	String	>0	A list of integers from Fig. 7-32
IpAddress	String	4	A dotted decimal Internet address
NsapAddress	String	< 22	An OSI NSAP address

Fig. 7-35. Data types used for SNMP monitored variables.

being declared, the *OBJECT TYPE* macro also requires three other parameters. *MAX-ACCESS* contains information about the variable's access. The most common values are read-write and read-only. If a variable is read-write, the management station can set it. If it is read-only, the management station can read it but cannot set it.

The *STATUS* has three possible values. A current variable is conformant with the current SNMP specification. An obsolete variable is not conformant but was conformant with an older version. A deprecated variable is in between. It is really obsolete, but the committee that wrote the standard did not dare say this in public for fear of the reaction from vendors whose products use it. Nevertheless, the handwriting is on the wall.

The last required parameter is *DESCRIPTION*, which is an ASCII string telling what the variable does. If a manager buys a nice new shiny device, queries it from the management station, and discovers that it keeps track of *pktCnt*, fetching the *DESCRIPTION* field is supposed to give a clue as to what kind of packets it is counting. This field is intended exclusively for human (as opposed to computer) consumption.

A simple example of an *OBJECT TYPE* declaration is given in Fig. 7-36. The variable is called *lostPackets* and might be useful in a router or other device dealing with packets. The value after the ::= sign places it in the tree.

```

lostPackets OBJECT TYPE
SYNTAX Counter32          -- use a 32-bit counter
MAX-ACCESS read-only     -- the management station may not change it
STATUS current           -- this variable is not obsolete (yet)
DESCRIPTION
    "The number of packets lost since the last boot"
 ::= {experimental 20}

```

Fig. 7-36. An example SNMP variable.

7.3.4. The MIB—Management Information Base

The collection of objects managed by SNMP is defined in the MIB. For convenience, these objects are (currently) grouped into ten categories, which correspond to the ten nodes under *mib-2* in Fig. 7-32. (Note that *mib-2* corresponds to SNMPv2 and that object 9 is no longer present.) The ten categories are intended to provide a basis of what a management station should understand. New categories and objects will certainly be added in the future, and vendors are free to define additional objects for their products. The ten categories are summarized in Fig. 7-37.

Group	# Objects	Description
System	7	Name, location, and description of the equipment
Interfaces	23	Network interfaces and their measured traffic
AT	3	Address translation (deprecated)
IP	42	IP packet statistics
ICMP	26	Statistics about ICMP messages received
TCP	19	TCP algorithms, parameters, and statistics
UDP	6	UDP traffic statistics
EGP	20	Exterior gateway protocol traffic statistics
Transmission	0	Reserved for media-specific MIBs
SNMP	29	SNMP traffic statistics

Fig. 7-37. The object groups of the Internet MIB-II.

Although space limitations prevent us from delving into the details of all 175 objects defined in MIB-II, a few comments may be helpful. The system group allows the manager to find out what the device is called, who made it, what hardware and software it contains, where it is located, and what it is supposed to do. The time of the last boot and the name and address of the contact person are

also provided. This information means that a company can contract out system management to another company in a distant city and have the latter be able to easily figure out what the configuration being managed actually is and who should be contacted if there are problems with various devices.

The interfaces group deals with the network adapters. It keeps track of the number of packets and bytes sent and received from the network, the number of discards, the number of broadcasts, and the current output queue size.

The AT group was present in MIB-I and provided information about address mapping (e.g., Ethernet to IP addresses). This information was moved to protocol-specific MIBs in SNMPv2.

The IP group deals with IP traffic into and out of the node. It is especially rich in counters keeping track of the number of packets discarded for each of a variety of reasons (e.g., no known route to the destination or lack of resources). Statistics about datagram fragmentation and reassembly are also available. All these items are particularly important for managing routers.

The ICMP group is about IP error messages. Basically, it has a counter for each ICMP message that records how many of that type have been seen.

The TCP group monitors the current and cumulative number of connections opened, segments sent and received, and various error statistics.

The UDP group logs the number of UDP datagrams sent and received, and how many of the latter were undeliverable due to an unknown port or some other reason.

The EGP group is used for routers that support the exterior gateway protocol. It keeps track of how many packets of what kind went out, came in and were forwarded correctly, and came in and were discarded.

The transmission group is a place holder for media-specific MIBs. For example, Ethernet-specific statistics can be kept here. The purpose of including an empty group in MIB-II is to reserve the identifier {internet 2 1 9} for such purposes.

The last group is for collecting statistics about the operation of SNMP itself. How many messages are being sent, what kinds of messages are they, and so on.

MIB-II is formally defined in RFC 1213. The bulk of RFC 1213 consists of 175 macro calls similar to those of Fig. 7-36, with comments delineating the ten groups. For each of the 175 objects defined, the data type is given along with an English text description of what the variable is used for. For further information about MIB-II, the reader is referred to this RFC.

7.3.5. The SNMP Protocol

We have now seen that the model underlying SNMP is a management station that sends requests to agents in managed nodes, inquiring about the 175 variables just alluded to, and many other vendor-specific variables. Our last topic is the

actual protocol that the management station and agents speak. The protocol itself is defined in RFC 1448.

The normal way that SNMP is used is that the management station sends a request to an agent asking it for information or commanding it to update its state in a certain way. Ideally, the agent just replies with the requested information or confirms that it has updated its state as requested. Data are sent using the ASN.1 transfer syntax. However, various errors can also be reported, such as No Such Variable.

SNMP defines seven messages that can be sent. The six messages from an initiator are listed in Fig. 7-38 (the seventh message is the response message). The first three request variable values to be sent back. The first format names the variables it wants explicitly. The second one asks for the next variable, allowing a manager to step through the entire MIB alphabetically (the default is the first variable). The third is for large transfers, such as tables.

Message	Description
Get-request	Requests the value of one or more variables
Get-next-request	Requests the variable following this one
Get-bulk-request	Fetches a large table
Set-request	Updates one or more variables
Inform-request	Manager-to-manager message describing local MIB
SnmpV2-trap	Agent-to-manager trap report

Fig. 7-38. SNMP message types.

Then comes a message that allows the manager to update an agent's variables, to the extent that the object specification permits such updates, of course. Next is an informational request that allows one manager to tell another one which variables it is managing. Finally, comes the message sent from an agent to a manager when a trap has sprung.

7.4. ELECTRONIC MAIL

Having finished looking at some of the support protocols used in the application layer, we finally come to real applications. When asked: "What are you going to do now?" few people will say: "I am going to look up some names with DNS." People do say they are going to read their email or news, surf the Web, or watch a movie over the net. In the remainder of this chapter, we will explain in a fair amount of detail how these four applications work.

Electronic mail, or **email**, as it is known to its many fans, has been around for over two decades. The first email systems simply consisted of file transfer protocols, with the convention that the first line of each message (i.e., file) contained the recipient's address. As time went on, the limitations of this approach became more obvious. Some of the complaints were

1. Sending a message to a group of people was inconvenient. Managers often need this facility to send memos to all their subordinates.
2. Messages had no internal structure, making computer processing difficult. For example, if a forwarded message was included in the body of another message, extracting the forwarded part from the received message was difficult.
3. The originator (sender) never knew if a message arrived or not.
4. If someone was planning to be away on business for several weeks and wanted all incoming email to be handled by his secretary, this was not easy to arrange.
5. The user interface was poorly integrated with the transmission system requiring users first to edit a file, then leave the editor and invoke the file transfer program.
6. It was not possible to create and send messages containing a mixture of text, drawings, facsimile, and voice.

As experience was gained, more elaborate email systems were proposed. In 1982, the ARPANET email proposals were published as RFC 821 (transmission protocol) and RFC 822 (message format). These have since become the de facto Internet standards. Two years later, CCITT drafted its X.400 recommendation, which was later taken over as the basis for OSI's MOTIS. In 1988, CCITT modified X.400 to align it with MOTIS. MOTIS was to be the flagship application for OSI, a system that was to be all things to all people.

After a decade of competition, email systems based on RFC 822 are widely used, whereas those based on X.400 have disappeared under the horizon. How a system hacked together by a handful of computer science graduate students beat an official international standard strongly backed by all the PTTs worldwide, many governments, and a substantial part of the computer industry brings to mind the Biblical story of David and Goliath. The reason for RFC 822's success is not that it is so good, but that X.400 is so poorly designed and so complex that nobody could implement it well. Given a choice between a simple-minded, but working, RFC 822-based email system and a supposedly truly wonderful, but nonworking, X.400 email system, most organizations chose the former. For a long diatribe on what is wrong with X.400, see Appendix C of (Rose, 1993). Consequently, our discussion of email will focus on RFC 821 and RFC 822 as used in the Internet.

7.4.1. Architecture and Services

In this section we will provide an overview of what email systems can do and how they are organized. They normally consist of two subsystems: the **user agents**, which allow people to read and send email, and the **message transfer agents**, which move the messages from the source to the destination. The user agents are local programs that provide a command-based, menu-based, or graphical method for interacting with the email system. The message transfer agents are typically system daemons that run in the background and move email through the system.

Typically, email systems support five basic functions, as described below. **Composition** refers to the process of creating messages and answers. Although any text editor can be used for the body of the message, the system itself can provide assistance with addressing and the numerous header fields attached to each message. For example, when answering a message, the email system can extract the originator's address from the incoming email and automatically insert it into the proper place in the reply.

Transfer refers to moving messages from the originator to the recipient. In large part, this requires establishing a connection to the destination or some intermediate machine, outputting the message, and releasing the connection. The email system should do this automatically, without bothering the user.

Reporting has to do with telling the originator what happened to the message. Was it delivered? Was it rejected? Was it lost? Numerous applications exist in which confirmation of delivery is important and may even have legal significance ("Well, Your Honor, my email system is not very reliable, so I guess the electronic subpoena just got lost somewhere").

Displaying incoming messages is needed so people can read their email. Sometimes conversion is required or a special viewer must be invoked, for example, if the message is a PostScript file or digitized voice. Simple conversions and formatting are sometimes attempted as well.

Disposition is the final step and concerns what the recipient does with the message after receiving it. Possibilities include throwing it away before reading, throwing it away after reading, saving it, and so on. It should also be possible to retrieve and reread saved messages, forward them, or process them in other ways.

In addition to these basic services, most email systems provide a large variety of advanced features. Let us just briefly mention a few of these. When people move, or when they are away for some period of time, they may want their email forwarded, so the system should be able to do this automatically.

Most systems allow users to create **mailboxes** to store incoming email. Commands are needed to create and destroy mailboxes, inspect the contents of mailboxes, insert and delete messages from mailboxes, and so on.

Corporate managers often need to send a message to each of their subordinates, customers, or suppliers. This gives rise to the idea of a **mailing list**, which

is a list of email addresses. When a message is sent to the mailing list, identical copies are delivered to everyone on the list.

Registered email is another important idea, to allow the originator to know that his message has arrived. Alternatively, automatic notification of undeliverable email may be desired. In any case, the originator should have some control over reporting what happened.

Other advanced features are carbon copies, high-priority email, secret (encrypted) email, alternative recipients if the primary one is not available, and the ability for secretaries to handle their bosses' email.

Email is now widely used within industry for intracompany communication. It allows far-flung employees to cooperate on complex projects, even over many time zones. By eliminating most cues associated with rank, age, and gender, email debates tend to focus on ideas, not on corporate status. With email, a brilliant idea from a summer student can have more impact than a dumb one from an executive vice president. Some companies have estimated that email has improved their productivity by as much as 30 percent (Perry and Adam, 1992).

A key idea in all modern email systems is the distinction between the **envelope** and its contents. The envelope encapsulates the message. It contains all the information needed for transporting the message, such as the destination address, priority, and security level, all of which are distinct from the message itself. The message transport agents use the envelope for routing, just as the post office does.

The message inside the envelope contains two parts: the **header** and the **body**. The header contains control information for the user agents. The body is entirely for the human recipient. Envelopes and messages are illustrated in Fig. 7-39.

7.4.2. The User Agent

Email systems have two basic parts, as we have seen: the user agents and the message transfer agents. In this section we will look at the user agents. A user agent is normally a program (sometimes called a mail reader) that accepts a variety of commands for composing, receiving, and replying to messages, as well as for manipulating mailboxes. Some user agents have a fancy menu- or icon-driven interface that requires a mouse, while others expect 1-character commands from the keyboard. Functionally, these are the same.

Sending Email

To send an email message, a user must provide the message, the destination address, and possibly some other parameters (e.g., the priority or security level). The message can be produced with a free-standing text editor, a word processing




Fig. 7-39. Envelopes and messages. (a) Postal email. (b) Electronic mail.

program, or possibly with a text editor built into the user agent. The destination address must be in a format that the user agent can deal with. Many user agents expect DNS addresses of the form *mailbox@location*. Since we have studied these earlier in this chapter, we will not repeat that material here.

However, it is worth noting that other forms of addressing exist. In particular, X.400 addresses look radically different than DNS addresses. They are composed of *attribute = value* pairs, for example,

`/C=US/SP=MASSACHUSETTS/L=CAMBRIDGE/PA=360 MEMORIAL DR./CN=KEN SMITH/`

This address specifies a country, state, locality, personal address and a common name (Tom Smith). Many other attributes are possible, so you can send email to someone whose name you do not know, provided you know enough other attributes (e.g., company and job title). Many people feel that this form of naming is considerably less convenient than DNS names.

In all fairness, however, the X.400 designers assumed that people would use **aliases** (short user-assigned strings) to identify recipients, so that they would never even see the full addresses. However, the necessary software was never

widely available, so people sending mail to users with X.400 addresses often had to type in strings like the one above. In contrast, most email systems for the Internet have always allowed users to have alias files.

Most email systems support mailing lists, so that a user can send the same message to a list of people with a single command. If the mailing list is maintained locally, the user agent can just send a separate message to each intended recipient. However, if the list is maintained remotely, then messages will be expanded there. For example, if a group of bird watchers have a mailing list called *birders* installed on *meadowlark.arizona.edu*, then any message sent to *birders@meadowlark.arizona.edu* will be routed to the University of Arizona and expanded there into individual messages to all the mailing list members, wherever in the world they may be. Users of this mailing list cannot tell that it is a mailing list. It could just as well be the personal mailbox of Prof. Gabriel O. Birders.

Reading Email

Typically, when a user agent is started up, it will look at the user's mailbox for incoming email before displaying anything on the screen. Then it may announce the number of messages in the mailbox or display a one-line summary of each one and wait for a command.

As an example of how a user agent works, let us take a look at a typical mail scenario. After starting up the user agent, the user asks for a summary of his email. A display like that of Fig. 7-40 then appears on the screen. Each line refers to one message. In this example, the mailbox contains eight messages.

#	Flags	Bytes	Sender	Subject
1	K	1030	asw	Changes to MINIX
2	KA	6348	radia	Comments on material you sent me
3	K F	4519	Amy N. Wong	Request for information
4		1236	bal	Deadline for grant proposal
5		103610	kaashoek	Text of DCS paper
6		1223	emily E.	Pointer to WWW page
7		3110	saniya	Referee reports for the paper
8		1204	dmr	Re: My student's visit

Fig. 7-40. An example display of the contents of a mailbox.

Each display line contains several fields extracted from the envelope or header of the corresponding message. In a simple email system, the choice of fields displayed is built into the program. In a more sophisticated system, the user can specify which fields are to be displayed by providing a **user profile**, a file

describing the display format. In this example, the first field is the message number. The second field, *Flags*, can contain a *K*, meaning that the message is not new but was read previously and kept in the mailbox; an *A*, meaning that the message has already been answered; and/or an *F*, meaning that the message has been forwarded to someone else. Other flags are also possible.

The third field tells how long the message is and the fourth one tells who sent the message. Since this field is simply extracted from the message, this field may contain first names, full names, initials, login names, or whatever else the sender chooses to put there. Finally, the *Subject* field gives a brief summary of what the message is about. People who fail to include a *Subject* field often discover that responses to their email tend not to get the highest priority.

After the headers have been displayed, the user can perform any of the commands available. A typical collection is listed in Fig. 7-41. Some of the commands require a parameter. The # sign means that the number of a message (or perhaps several messages) is expected. Alternatively, the letter *a* can be used to mean all messages.

Command	Parameter	Description
h	#	Display header(s) on the screen
c		Display current header only
t	#	Type message(s) on the screen
s	address	Send a message
f	#	Forward message(s)
a	#	Answer message(s)
d	#	Delete message(s)
u	#	Undelete previously deleted message(s)
m	#	Move message(s) to another mailbox
k	#	Keep message(s) after exiting
r	mailbox	Read a new mailbox
n		Go to the next message and display it
b		Backup to the previous message and display it
g	#	Go to a specific message but do not display it
e		Exit the mail system and update the mailbox

Fig. 7-41. Typical mail handling commands.

Innumerable email programs exist. Our example email program is patterned after the one used by the UNIX Mmdf system, as it is quite straightforward. The *h* command displays one or more headers in the format of Fig. 7-40. The *c* command prints the current message's header. The *t* command types (i.e., displays on the screen) the requested message or messages. Possible commands are *t 3*, to type message 3, *t 4-6*, to type messages 4 through 6, and *t a* to type them all.

The next group of three commands deals with sending messages rather than receiving them. The *s* command sends a message by calling an appropriate editor (e.g., specified in the user's profile) to allow the user to compose the message. Spelling, grammar, and diction checkers can see if the message is syntactically correct. Unfortunately, the current generation of email programs do not have checkers to see if the sender knows what he is talking about. When the message is finished, it is prepared for transmission to the message transfer agent.

The *f* command forwards a message from the mailbox, prompting for an address to send it to. The *a* command extracts the source address from the message to be answered and calls the editor to allow the user to compose the reply.

The next group of commands is for manipulating mailboxes. Users typically have one mailbox for each person with whom they correspond, in addition to the mailbox for incoming email that we have already seen. The *d* command deletes a message from the mailbox, but the *u* command undoes the delete. (The message is not actually deleted until the email program is exited.) The *m* command moves a message to another mailbox. This is the usual way to save important email after reading it. The *k* command keeps the indicated message in the mailbox even after it is read. If a message is read but not explicitly kept, some default action is taken when the email program is exited, such as moving it to a special default mailbox. Finally, the *r* command is used to finish up with the current mailbox and go read another one.

The *n*, *b*, and *g* commands are for moving about in the current mailbox. It is common for a user to read message 1, answer, move, or delete it, and then type *n* to get the next one. The value of this command is that the user does not have to keep track of where he is. It is possible to go backward using *b* or to a given message with *g*.

Finally, the *e* command exits the email program and makes whatever changes are required, such as deleting some messages and marking others as kept. This command overwrites the mailbox, replacing its contents.

In mail systems designed for beginners, each of these commands is typically associated with an on-screen icon, so that the user does not have to remember that *a* stands for *answer*. Instead, she has to remember that the little picture of a person with his mouth open means answer and not display message.

It should be clear from this example that email has come a long way from the days when it was just file transfer. Sophisticated user agents make managing a large volume of email possible. For people such as the author who (reluctantly) receive and send thousands of messages a year, such tools are invaluable.

7.4.3. Message Formats

Let us now turn from the user interface to the format of the email messages themselves. First we will look at basic ASCII email using RFC 822. After that, we will look at multimedia extensions to RFC 822

RFC 822

Messages consist of a primitive envelope (described in RFC 821), some number of header fields, a blank line, and then the message body. Each header field (logically) consists of a single line of ASCII text containing the field name, a colon, and, for most fields, a value. RFC 822 is an old standard, and does not clearly distinguish envelope from header fields, as a new standard would do. In normal usage, the user agent builds a message and passes it to the message transfer agent, which then uses some of the header fields to construct the actual envelope, a somewhat old-fashioned mixing of message and envelope.

The principal header fields related to message transport are listed in Fig. 7-42. The *To:* field gives the DNS address of the primary recipient. Having multiple recipients is also allowed. The *Cc:* field gives the addresses of any secondary recipients. In terms of delivery, there is no distinction between the primary and secondary recipients. It is entirely a psychological difference that may be important to the people involved but is not important to the mail system. The term *Cc:* (Carbon copy) is a bit dated, since computers do not use carbon paper, but it is well established. The *Bcc:* (Blind carbon copy) field is like the *Cc:* field, except that this line is deleted from all the copies sent to the primary and secondary recipients. This feature allows people to send copies to third parties without the primary and secondary recipients knowing this.

Header	Meaning
To:	Email address(es) of primary recipient(s)
Cc:	Email address(es) of secondary recipient(s)
Bcc:	Email address(es) for blind carbon copies
From:	Person or people who created the message
Sender:	Email address of the actual sender
Received:	Line added by each transfer agent along the route
Return-Path:	Can be used to identify a path back to the sender

Fig. 7-42. RFC 822 header fields related to message transport.

The next two fields, *From:* and *Sender:* tell who wrote and sent the message, respectively. These may not be the same. For example, a business executive may write a message, but her secretary may be the one who actually transmits it. In this case, the executive would be listed in the *From:* field and the secretary in the *Sender:* field. The *From:* field is required, but the *Sender:* field may be omitted if it is the same as the *From:* field. These fields are needed in case the message is undeliverable and must be returned to the sender.

A line containing *Received:* is added by each message transfer agent along the

way. The line contains the agent's identity, the date and time the message was received, and other information that can be used for finding bugs in the routing system.

The *Return-Path:* field is added by the final message transfer agent and was intended to tell how to get back to the sender. In theory, this information can be gathered from all the *Received:* headers (except for the name of the sender's mailbox), but it is rarely filled in as such and typically just contains the sender's address.

In addition to the fields of Fig. 7-42, RFC 822 messages may also contain a variety of header fields used by the user agents or human recipients. The most common ones are listed in Fig. 7-43. Most of these are self-explanatory, so we will not go into all of them in detail.

Header	Meaning
Date:	The date and time the message was sent
Reply-To:	Email address to which replies should be sent
Message-Id:	Unique number for referencing this message later
In-Reply-To:	Message-Id of the message to which this is a reply
References:	Other relevant Message-Ids
Keywords:	User chosen keywords
Subject:	Short summary of the message for the one-line display

Fig. 7-43. Some fields used in the RFC 822 message header.

The *Reply-To:* field is sometimes used when neither the person composing the message nor the person sending the message wants to see the reply. For example, a marketing manager writes an email message telling customers about a new product. The message is sent by a secretary, but the *Reply-To:* field lists the head of the sales department, who can answer questions and take orders.

The RFC 822 document explicitly says that users are allowed to invent new headers for their own private use, provided that these headers start with the string X-. It is guaranteed that no future headers will use names starting with X-, to avoid conflicts between official and private headers. Sometimes wiseguy undergraduates include fields like *X-Fruit-of-the-Day:* or *X-Disease-of-the-Week:*, which are legal, although not always illuminating.

After the headers comes the message body. Users can put whatever they want here. Some people terminate their messages with elaborate signatures, including simple ASCII cartoons, quotations from greater and lesser authorities, political statements, and disclaimers of all kinds (e.g., The ABC Corporation is not responsible for my opinions; it cannot even comprehend them).

MIME—Multipurpose Internet Mail Extensions

In the early days of the ARPANET, email consisted exclusively of text messages written in English and expressed in ASCII. For this environment, RFC 822 did the job completely: it specified the headers but left the content entirely up to the users. Nowadays, on the worldwide Internet, this approach is no longer adequate. The problems include sending and receiving

1. Messages in languages with accents (e.g., French and German).
2. Messages in nonLatin alphabets (e.g., Hebrew and Russian).
3. Messages in languages without alphabets (e.g., Chinese and Japanese).
4. Messages not containing text at all (e.g., audio and video).

A solution was proposed in RFC 1341 and updated in RFC 1521. This solution, called **MIME (Multipurpose Internet Mail Extensions)** is now widely used. We will now describe it. For additional information about MIME, see RFC 1521 or (Rose, 1993).

The basic idea of MIME is to continue to use the RFC 822 format, but to add structure to the message body and define encoding rules for non-ASCII messages. By not deviating from 822, MIME messages can be sent using the existing mail programs and protocols. All that has to be changed are the sending and receiving programs, which users can do for themselves.

MIME defines five new message headers, as shown in Fig. 7-44. The first of these simply tells the user agent receiving the message that it is dealing with a MIME message, and which version of MIME it uses. Any message not containing a *MIME-Version*: header is assumed to be an English plaintext message, and is processed as such.

Header	Meaning
MIME-Version:	Identifies the MIME version
Content-Description:	Human-readable string telling what is in the message
Content-Id:	Unique identifier
Content-Transfer-Encoding:	How the body is wrapped for transmission
Content-Type:	Nature of the message

Fig. 7-44. RFC 822 headers added by MIME.

The *Content-Description*: header is an ASCII string telling what is in the message. This header is needed so the recipient will know whether it is worth decoding and reading the message. If the string says: "Photo of Barbara's gerbil" and the person getting the message is not a big gerbil fan, the message will probably be discarded rather than decoded into a high-resolution color photograph.

The *Content-Id:* header identifies the content. It uses the same format as the standard *Message-Id:* header.

The *Content-Transfer-Encoding:* tells how the body is wrapped for transmission through a network that may object to most characters other than letters, numbers, and punctuation marks. Five schemes (plus an escape to new schemes) are provided. The simplest scheme is just ASCII text. ASCII characters use 7 bits, and can be carried directly by the email protocol provided that no line exceeds 1000 characters.

The next simplest scheme is the same thing, but using 8-bit characters, that is, all values from 0 up to and including 255. This encoding scheme violates the (original) Internet email protocol but is used by some parts of the Internet that implement some extensions to the original protocol. While declaring the encoding does not make it legal, having it explicit may at least explain things when something goes wrong. Messages using the 8-bit encoding must still adhere to the standard maximum line length.

Even worse are messages that use binary encoding. These are arbitrary binary files that not only use all 8 bits but also do not even respect the 1000 character line limit. Executable programs fall into this category. No guarantee is given that messages in binary will arrive correctly, but many people send them anyway.

The correct way to encode binary messages is to use **base64 encoding**, sometimes called **ASCII armor**. In this scheme, groups of 24 bits are broken up into four 6-bit units, with each unit being sent as a legal ASCII character. The coding is "A" for 0, "B" for 1, and so on, followed by the 26 lowercase letters, the ten digits, and finally + and / for 62 and 63, respectively. The == and = sequences are used to indicate that the last group contained only 8 or 16 bits, respectively. Carriage returns and line feeds are ignored, so they can be inserted at will to keep the lines short enough. Arbitrary binary text can be sent safely using this scheme.

For messages that are almost entirely ASCII, but with a few non-ASCII characters, base64 encoding is somewhat inefficient. Instead, an encoding known as **quoted-printable encoding** is used. This is just 7-bit ASCII, with all the characters above 127 encoded as an equal sign followed by the character's value as two hexadecimal digits.

In summary, binary data should be sent encoded in base64 or quoted printable form. When there are valid reasons not to use one of these schemes, it is possible to specify a user-defined encoding in the *Content-Transfer-Encoding:* header.

The last header shown in Fig. 7-44 is really the most interesting one. It specifies the nature of the message body. Seven types are defined in RFC 1521, each of which has one or more subtypes. The type and subtype are separated by a slash, as in

Content-Type: video/mpeg

The subtype must be given explicitly in the header; no defaults are provided. The initial list of types and subtypes specified in RFC 1521 is given in Fig. 7-45.

Many new ones have been added since then, and additional entries are being added all the time as the need arises.

Type	Subtype	Description
Text	Plain	Unformatted text
	Richtext	Text including simple formatting commands
Image	Gif	Still picture in GIF format
	Jpeg	Still picture in JPEG format
Audio	Basic	Audible sound
Video	Mpeg	Movie in MPEG format
Application	Octet-stream	An uninterpreted byte sequence
	Postscript	A printable document in PostScript
Message	Rfc822	A MIME RFC 822 message
	Partial	Message has been split for transmission
	External-body	Message itself must be fetched over the net
Multipart	Mixed	Independent parts in the specified order
	Alternative	Same message in different formats
	Parallel	Parts must be viewed simultaneously
	Digest	Each part is a complete RFC 822 message

Fig. 7-45. The MIME types and subtypes defined in RFC 1521.

Let us now go through the list of types. The *text* type is for straight text. The *text/plain* combination is for ordinary messages that can be displayed as received, with no encoding and no further processing. This option allows ordinary messages to be transported in MIME with only a few extra headers.

The *text/richtext* subtype allows a simple markup language to be included in the text. This language provides a system-independent way to express boldface, italics, smaller and larger point sizes, indentation, justification, sub- and superscripting, and simple page layout. The markup language is based on SGML, the Standard Generalized Markup Language also used as the basis for the World Wide Web's HTML. For example, the message

The <bold> time </bold> has come the <italic> walrus </italic> said ...

would be displayed as

The **time** has come the *walrus* said ...

It is up to the receiving system to choose the appropriate rendition. If boldface and italics are available, they can be used; otherwise, colors, blinking,

underlining, reverse video, etc. can be used for emphasis. Different systems can, and do, make different choices.

The next MIME type is *image*, which is used to transmit still pictures. Many formats are widely used for storing and transmitting images nowadays, both with and without compression. Two of these, GIF and JPEG, are official subtypes, but no doubt others will be added later.

The *audio* and *video* types are for sound and moving pictures, respectively. Note that *video* includes only the visual information, not the soundtrack. If a movie with sound is to be transmitted, the video and audio portions may have to be transmitted separately, depending on the encoding system used. The only video format defined so far is the one devised by the modestly-named Moving Picture Experts Group (MPEG).

The *application* type is a catchall for formats that require external processing not covered by one of the other types. An *octet-stream* is just a sequence of uninterpreted bytes. Upon receiving such a stream, a user agent should probably display it by suggesting to the user that it be copied to a file and prompting for a file name. Subsequent processing is then up to the user.

The other defined subtype is *postscript*, which refers to the PostScript language produced by Adobe Systems and widely used for describing printed pages. Many printers have built-in PostScript interpreters. Although a user agent can just call an external PostScript interpreter to display incoming PostScript files, doing so is not without danger. PostScript is a full-blown programming language. Given enough time, a sufficiently masochistic person could write a C compiler or a database management system in PostScript. Displaying an incoming PostScript message is done by executing the PostScript program contained in it. In addition to displaying some text, this program can read, modify, or delete the user's files, and have other nasty side effects.

The *message* type allows one message to be fully encapsulated inside another. This scheme is useful for forwarding email, for example. When a complete RFC 822 message is encapsulated inside an outer message, the *rfc822* subtype should be used.

The *partial* subtype makes it possible to break an encapsulated message up into pieces and send them separately (for example, if the encapsulated message is too long). Parameters make it possible to reassemble all the parts at the destination in the correct order.

Finally, the *external-body* subtype can be used for very long messages (e.g., video films). Instead of including the MPEG file in the message, an FTP address is given and the receiver's user agent can fetch it over the network at the time it is needed. This facility is especially useful when sending a movie to a mailing list of people, only a few of whom are expected to view it (think about electronic junk mail containing advertising videos).

The final type is *multipart*, which allows a message to contain more than one part, with the beginning and end of each part being clearly delimited. The *mixed*

subtype allows each part to be different, with no additional structure imposed. In contrast, with the *alternative* subtype, each part must contain the same message but expressed in a different medium or encoding. For example, a message could be sent in plain ASCII, in richtext, and in PostScript. A properly-designed user agent getting such a message would display it in PostScript if possible. Second choice would be richtext. If neither of these were possible, the flat ASCII text would be displayed. The parts should be ordered from simplest to most complex to help recipients with pre-MIME user agents make some sense of the message (e.g., even a pre-MIME user can read flat ASCII text).

The *alternative* subtype can also be used for multiple languages. In this context, the Rosetta Stone can be thought of as an early *multipart/alternative* message.

A multimedia example is shown in Fig. 7-46. Here a birthday greeting is transmitted both as text and as a song. If the receiver has an audio capability, the user agent there will fetch the sound file, *birthday.snd*, and play it. If not, the lyrics are displayed on the screen in stony silence. The parts are delimited by two hyphens followed by the (user-defined) string specified in the *boundary* parameter.

Note that the *Content-Type* header occurs in three positions within this example. At the top level, it indicates that the message has multiple parts. Within each part, it gives the type and subtype of that part. Finally, within the body of the second part, it is required to tell the user agent what kind of an external file it is to fetch. To indicate this slight difference in usage, we have used lowercase letters here, although all headers are case insensitive. The *content-transfer-encoding* is similarly required for any external body that is not encoded as 7-bit ASCII.

Getting back to the subtypes for multipart messages, two more possibilities exist. The *parallel* subtype is used when all parts must be “viewed” simultaneously. For example, movies often have an audio channel and a video channel. Movies are more effective if these two channels are played back in parallel, instead of consecutively.

Finally, the *digest* subtype is used when many messages are packed together into a composite message. For example, some discussion groups on the Internet collect messages from subscribers and then send them out as a single *multipart/digest* message.

7.4.4. Message Transfer

The message transfer system is concerned with relaying messages from originator to the recipient. The simplest way to do this is to establish a transport connection from the source machine to the destination machine and then just transfer the message. After examining how this is normally done, we will examine some situations in which this does not work and what can be done about them.

```

From: elinor@abc.com
To: carolyn@xyz.com
MIME-Version: 1.0
Message-Id: <0704760941.AA00747@abc.com>
Content-Type: multipart/alternative; boundary=qwertyuiopasdfghjklzxcvbnm
Subject: Earth orbits sun integral number of times

```

This is the preamble. The user agent ignores it. Have a nice day.

```

--qwertyuiopasdfghjklzxcvbnm
Content-Type: text/richtext

```

```

Happy birthday to you
Happy birthday to you
Happy birthday dear <bold> Carolyn </bold>
Happy birthday to you

```

```

--qwertyuiopasdfghjklzxcvbnm
Content-Type: message/external-body;
    access-type="anon-ftp";
    site="bicycle.abc.com";
    directory="pub";
    name="birthday.snd"

```

```

content-type: audio/basic
content-transfer-encoding: base64
--qwertyuiopasdfghjklzxcvbnm--

```

Fig. 7-46. A multipart message containing richtext and audio alternatives.

SMTP—Simple Mail Transfer Protocol

Within the Internet, email is delivered by having the source machine establish a TCP connection to port 25 of the destination machine. Listening to this port is an email daemon that speaks **SMTP (Simple Mail Transfer Protocol)**. This daemon accepts incoming connections and copies messages from them into the appropriate mailboxes. If a message cannot be delivered, an error report containing the first part of the undeliverable message is returned to the sender.

SMTP is a simple ASCII protocol. After establishing the TCP connection to port 25, the sending machine, operating as the client, waits for the receiving machine, operating as the server, to talk first. The server starts by sending a line of text giving its identity and telling whether or not it is prepared to receive mail. If it is not, the client releases the connection and tries again later.

If the server is willing to accept email, the client announces whom the email is coming from and whom it is going to. If such a recipient exists at the

destination, the server gives the client the go-ahead to send the message. Then the client sends the message and the server acknowledges it. No checksums are generally needed because TCP provides a reliable byte stream. If there is more email, that is now sent. When all the email has been exchanged in both directions, the connection is released. A sample dialog for sending the message of Fig. 7-46, including the numerical codes used by SMTP, is shown in Fig. 7-47. The lines sent by the client are marked *C:*; those sent by the server are marked *S:*.

A few comments about Fig. 7-47 may be helpful. The first command from the client is indeed *HELO*. Of the two four-character abbreviations for *HELLO*, this one has numerous advantages over its competitor. Why all the commands had to be four characters has been lost in the mists of time.

In Fig. 7-47, the message is sent to only one recipient, so only one *RCPT* command is used. Multiple such commands are allowed to send a single message to multiple receivers. Each one is individually acknowledged or rejected. Even if some recipients are rejected (because they do not exist at the destination), the message can be sent to the remainder.

Finally, although the syntax of the four-character commands from the client is rigidly specified, the syntax of the replies is less rigid. Only the numerical code really counts. Each implementation can put whatever string it wants after the code.

Even though the SMTP protocol is well defined (by RFC 821), a few problems can still arise. One problem relates to message length. Some older implementations cannot handle messages exceeding 64KB. Another problem relates to timeouts. If the client and server have different timeouts, one of them may give up while the other is still busy, unexpectedly terminating the connection. Finally, in rare situations, infinite mailstorms can be triggered. For example, if host 1 holds mailing list *A* and host 2 holds mailing list *B* and each list contains an entry for the other one, then any message sent to either list will generate a never-ending amount of email traffic.

To get around some of these problems, extended SMTP (**ESMTP**) has been defined in RFC 1425. Clients wanting to use it should send an *EHLO* message instead of *HELO* initially. If this is rejected, then the server is a regular SMTP server, and the client should proceed in the usual way. If the *EHLO* is accepted, then new commands and parameters are allowed. The standardization of these commands and parameters is an ongoing process.

Email Gateways

Email using SMTP works best when both the sender and the receiver are on the Internet and can support TCP connections between sender and receiver. However, many machines that are not on the Internet still want to send and receive email from Internet sites. For example, many companies intentionally do not

```

S: 220 xyz.com SMTP service ready
C: HELO abc.com
S: 250 xyz.com says hello to abc.com
C: MAIL FROM: <elinor@abc.com>
S: 250 sender ok
C: RCPT TO: <carolyn@xyz.com>
S: 250 recipient ok
C: DATA
S: 354 Send mail; end with "." on a line by itself
C: From: elinor@abc.com
C: To: carolyn@xyz.com
C: MIME-Version: 1.0
C: Message-Id: <0704760941.AA00747@abc.com>
C: Content-Type: multipart/alternative; boundary=qwertyuiopasdfghjklzxcvbnm
C: Subject: Earth orbits sun integral number of times
C:
C: This is the preamble. The user agent ignores it. Have a nice day.
C:
C: --qwertyuiopasdfghjklzxcvbnm
C: Content-Type: text/richtext
C:
C: Happy birthday to you
C: Happy birthday to you
C: Happy birthday dear <bold> Carolyn </bold>
C: Happy birthday to you
C:
C: --qwertyuiopasdfghjklzxcvbnm
C: Content-Type: message/external-body;
C:   access-type="anon-ftp";
C:   site="bicycle.abc.com";
C:   directory="pub";
C:   name="birthday.snd"
C:
C: content-type: audio/basic
C: content-transfer-encoding: base64
C: --qwertyuiopasdfghjklzxcvbnm
C: .
S: 250 message accepted
C: QUIT
S: 221 xyz.com closing connection

```

Fig. 7-47. Transferring a message from *elinor@abc.com* to *carolyn@xyz.com*.

want to be on the Internet for security reasons. Some of them even remove themselves from the Internet by erecting firewalls between themselves and the Internet.

Another problem occurs when the sender speaks only RFC 822 and the

receiver speaks only X.400 or some proprietary vendor-specific mail protocol. Since all these worlds differ in message formats and protocols, direct communication is impossible.

Both of these problems are solved using application layer **email gateways**. In Fig. 7-48 host 1 speaks only TCP/IP and RFC 822, whereas host 2 speaks only OSI TP4 and X.400. Nevertheless, they can exchange email using an email gateway. The procedure is for host 1 to establish a TCP connection to the gateway and then use SMTP to transfer a message (1) there. The daemon on the gateway then puts the message in a buffer of messages destined for host 2. Later, a TP4 connection (the OSI equivalent to TCP) is established with host 2 and the message (2) is transferred using the OSI equivalent of SMTP. All the gateway process has to do is to extract incoming messages from one queue and deposit them in another.




Fig. 7-48. Transferring email using an application layer email gateway.

It looks easy, but it is not. The first problem is that Internet addresses and X.400 addresses are totally different. An elaborate mapping mechanism is needed between them. The second problem is that envelope or header fields that are present in one system may not be present in the other. For example, if one system requires priority classes and the other does not have this concept at all, in one direction valuable information must be dropped and in the other it must be generated out of thin air.

An even worse concept is what to do if body parts are incompatible. What should a gateway do with a message from the Internet whose body holds a reference to an audio file to be obtained by FTP if the destination system does not support this concept? What should it do when an X.400 system tells it to deliver a message to a certain address, but if that fails, to send the contents by fax? Using fax is not part of the RFC 822 model. Clearly, there are no simple solutions here. For simple unstructured text messages in ASCII, gatewaying is a reasonable solution, but for anything fancier, the idea tends to break down.

Final Delivery

Up until now, we have assumed that all users work on machines that are capable of sending and receiving email. Frequently this situation is false. For example, at many companies, users work at desktop PCs that are not on the Internet and are not capable of sending or receiving email from outside the company. Instead, the company has one or more email servers that can send and receive email. To send or receive messages, a PC must talk to an email server using some kind of delivery protocol.

A simple protocol used for fetching email from a remote mailbox is **POP3 (Post Office Protocol)**, which is defined in RFC 1225. It has commands for the user to log in, log out, fetch messages, and delete messages. The protocol itself consists of ASCII text and has something of the flavor of SMTP. The point of POP3 is to fetch email from the remote mailbox and store it on the user's local machine to be read later.

A more sophisticated delivery protocol is **IMAP (Interactive Mail Access Protocol)**, which is defined in RFC 1064. It was designed to help the user who uses multiple computers, perhaps a workstation in the office, a PC at home, and a laptop on the road. The basic idea behind IMAP is for the email server to maintain a central repository that can be accessed from any machine. Thus unlike POP3, IMAP does not copy email to the user's personal machine because the user may have several.

IMAP has many features, such as the ability to address mail not by arrival number as is done in Fig. 7-40, but by using attributes (e.g., Give me the first message from Sam). In this view, a mailbox is more like a relational database system than a linear sequence of messages.

Yet a third delivery protocol is **DMSP (Distributed Mail System Protocol)**, which is part of the PCMAIL system and described in RFC 1056. This one does not assume that all email is on one server, as do POP3 and IMAP. Instead, it allows users to download email from the server to a workstation, PC, or laptop and then disconnect. The email can be read and answered while disconnected. When reconnection occurs later, email is transferred and the system is resynchronized.

Independent of whether email is delivered directly to the user's workstation or to a remote server, many systems provide hooks for additional processing of incoming email. An especially valuable tool for many email users is the ability to set up **filters**. These are rules that are checked when email comes in or when the user agent is started. Each rule specifies a condition and an action. For example, a rule could say that any message from Andrew S. Tanenbaum should be displayed in a 24-point flashing red boldface font (or alternatively, be discarded automatically without comment).

Another delivery feature often provided is the ability to (temporarily) forward incoming email to a different address. This address can even be a computer

operated by a commercial paging service, which then pages the user by radio or satellite, displaying the *Subject:* line on his beeper.

Still another common feature of final delivery is the ability to install a **vacation daemon**. This is a program that examines each incoming message and sends the sender an insipid reply such as

Hi. I'm on vacation. I'll be back on the 24th of August. Have a nice day.

Such replies can also specify how to handle urgent matters in the interim, other people to contact for specific problems, etc. Most vacation daemons keep track of whom they have sent canned replies to and refrain from sending the same person a second reply. The good ones also check to see if the incoming message was sent to a mailing list, and if so, do not send a canned reply at all. (People who send messages to large mailing lists during the summer probably do not want to get hundreds of replies detailing everyone's vacation plans.)

The author recently ran into a most extreme form of delivery processing when he sent an email message to a person who claims to get 600 messages a day. His identity will not be disclosed here, lest half the readers of this book also send him email. Let us call him John.

John has installed an email robot that checks every incoming message to see if it is from a new correspondent. If so, it sends back a canned reply explaining that John can no longer personally read all his email. Instead he has produced a personal FAQ (Frequently Asked Questions) document that answers many questions he is commonly asked. Normally, newsgroups have FAQs, not people.

John's FAQ gives his address, fax, and telephone numbers and tells how to contact his company. It explains how to get him as a speaker and describes where to get his papers and other documents. It also provides pointers to software he has written, a conference he is running, a standard he is the editor of, and so on. Perhaps this approach is necessary, but maybe a personal FAQ is the ultimate status symbol.

7.4.5. Email Privacy

When an email message is sent between two distant sites, it will generally transit dozens of machines on the way. Any of these can read and record the message for future use. Privacy is nonexistent, despite what many people think (Weisband and Reinig, 1995). Nevertheless, many people would like to be able to send email that can be read by the intended recipient and no one else: not their boss, not hackers, not even the government. This desire has stimulated several people and groups to apply the cryptographic principles we studied earlier to email to produce secure email. In the following sections we will study two widely used secure email systems, PGP and PEM. For additional information, see (Kaufman et al., 1995; Schneier, 1995; Stallings, 1995b; and Stallings, 1995c).

PGP—Pretty Good Privacy

Our first example, **PGP (Pretty Good Privacy)** is essentially the brainchild of one person, Phil Zimmermann (Zimmermann, 1995a, 1995b). It is a complete email security package that provides privacy, authentication, digital signatures, and compression, all in easy-to-use form. Furthermore, the complete package, including all the source code, is distributed free of charge via the Internet, bulletin boards, and commercial networks. Due to its quality, price (zero), and easy availability on MS-DOS/Windows, UNIX, and Macintosh platforms, it is widely used today. A commercial version is also available for those companies requiring support.

It has also been embroiled in various controversies (Levy, 1993). Because it is freely available over the Internet, the U.S. government has claimed the ability of foreigners to obtain it constitutes a violation of the laws concerning the export of munitions. Later versions were produced outside the United States to get around this restriction. Another problem has involved an alleged infringement of the RSA patent, but that problem was settled with releases starting at 2.6. Nevertheless, not everyone likes the idea of people being able to keep secrets from them, so PGP's enemies are always lurking in the shadows, waiting to pounce. Accordingly, Zimmermann's motto is: "If privacy is outlawed, only outlaws will have privacy."

PGP intentionally uses existing cryptographic algorithms rather than inventing new ones. It is largely based on RSA, IDEA, and MD5, all algorithms that have withstood extensive peer review and were not designed or influenced by any government agency trying to weaken them. For people who tend to distrust government, this property is a big plus.

PGP supports text compression, secrecy, and digital signatures and also provides extensive key management facilities. To see how PGP works, let us consider the example of Fig. 7-49. Here, Alice wants to send a signed plaintext message, P , to Bob in a secure way. Both Alice and Bob have private (D_X) and public (E_X) RSA keys. Let us assume that each one knows the other's public key; we will cover key management later.

Alice starts out by invoking the PGP program on her computer. PGP first hashes her message, P , using MD5 and then encrypts the resulting hash using her private RSA key, D_A . When Bob eventually gets the message, he can decrypt the hash with Alice's public key and verify that the hash is correct. Even if someone else (e.g., Trudy) could acquire the hash at this stage and decrypt it with Alice's known public key, the strength of MD5 guarantees that it would be computationally infeasible to produce another message with the same MD5 hash.

The encrypted hash and the original message are now concatenated into a single message, PI , and compressed using the ZIP program, which uses the Ziv-Lempel algorithm (Ziv and Lempel, 1977). Call the output of this step $PI.Z$.

Next, PGP prompts Alice for some random input. Both the content and the




Fig. 7-49. PGP in operation for sending a message.

typing speed are used to generate a 128-bit IDEA message key, K_M (called a session key in the PGP literature, but this is really a misnomer since there is no session). K_M is now used to encrypt $P1.Z$ with IDEA in cipher feedback mode. In addition, K_M is encrypted with Bob's public key, E_B . These two components are then concatenated and converted to base64, as we discussed in the section on MIME. The resulting message then contains only letters, digits, and the symbols +, / and =, which means it can be put into an RFC 822 body and be expected to arrive unmodified.

When Bob gets the message, he reverses the base64 encoding and decrypts the IDEA key using his private RSA key. Using this key, he decrypts the message to get $P1.Z$. After decompressing it, Bob separates the plaintext from the encrypted hash and decrypts the hash using Alice's public key. If the plaintext hash agrees with his own MD5 computation, he knows that P is the correct message and that it came from Alice.

It is worth noting that RSA is only used in two places here: to encrypt the 128-bit MD5 hash and to encrypt the 128-bit IDEA key. Although RSA is slow, it has to encrypt only 256 bits, not a large volume of data. Furthermore, all 256 plaintext bits are exceedingly random, so a considerable amount of work will be required on Trudy's part just to determine if a guessed key is correct. The heavy-duty encryption is done by IDEA, which is orders of magnitude faster than RSA. Thus PGP provides security, compression, and a digital signature and does so in a much more efficient way than the scheme illustrated in Fig. 7-23.

PGP supports three RSA key lengths. It is up to the user to select the one that is most appropriate. The lengths are

1. Casual (384 bits): can be broken today by folks with large budgets.
2. Commercial (512 bits): might be breakable by three-letter organizations.
3. Military (1024): Not breakable by anyone on earth.

There has been some discussion about a fourth category: alien (2048 bits), which could not be broken by anyone or anything in the universe, but this has not yet been adopted. Since RSA is only used for two small computations, probably everyone should use military strength keys all the time, except perhaps on aged PC-XTs.

The format of a PGP message is shown in Fig. 7-50. The message has three parts, containing the IDEA key, the signature, and the message, respectively. The key part contains not only the key, but also a key identifier, since users are permitted to have multiple public keys.




Fig. 7-50. A PGP message.

The signature part contains a header, which will not concern us here. The header is followed by a timestamp, the identifier for the sender's public key that can be used to decrypt the signature hash, some type information that identifies the algorithms used (to allow MD6 and RSA2 to be used when they are invented), and the encrypted hash itself.

The message part also contains a header, the default name of the file to be used if the receiver writes the file to the disk, a message creation timestamp, and, finally, the message itself.

Key management has received a large amount of attention in PGP as it is the Achilles heel of all security systems. Each user maintains two data structures locally: a private key ring and a public key ring. The **private key ring** contains one or more personal private-public key pairs. The reason for supporting multiple pairs per user is to permit users to change their public keys periodically or when one is thought to have been compromised, without invalidating messages

currently in preparation or in transit. Each pair has an identifier associated with it, so that a message sender can tell the recipient which public key was used to encrypt it. Message identifiers consist of the low-order 64 bits of the public key. Users are responsible for avoiding conflicts in their public key identifiers. The private keys on disk are encrypted using a special (arbitrarily long) password to protect them against sneak attacks.

The **public key ring** contains public keys of the user's correspondents. These are needed to encrypt the message keys associated with each message. Each entry on the public key ring contains not only the public key, but also its 64-bit identifier and an indication of how strongly the user trusts the key.

The problem being tackled here is the following. Suppose that public keys are maintained on bulletin boards. One way for Trudy to read Bob's secret email is to attack the bulletin board and replace Bob's public key with one of her choice. When Alice later fetches the key so-called belonging to Bob, Trudy can mount a bucket brigade attack on Bob.

To prevent such attacks, or at least minimize the consequences of them, Alice needs to know how much to trust the item called "Bob's key" on her public key ring. If she knows that Bob personally handed her a floppy disk containing the key, she can set the trust value to the highest value.

However, in practice, people often receive public keys by querying a trusted key server, a number of which are already in operation on the Internet. When a key server receives a request for someone's public key, it generates a response containing the public key, a timestamp, and the expiration date of the key. It then hashes this response with MD5 and signs the response with its own private key so the requesting party can verify who sent it. It is up to the user to assign a trust level to keys maintained by the local system administrator, the phone company, ACM, the Bar Association, the government, or whoever else decides to get into the business of maintaining keys.

PEM—Privacy Enhanced Mail

In contrast to PGP, which was initially a one-man show, our second example, **PEM (Privacy Enhanced Mail)**, is an official Internet standard and described in four RFCs: RFC 1421 through RFC 1424. Very roughly, PEM covers the same territory as PGP: privacy and authentication for RFC 822-based email systems. Nevertheless, it also has some differences with PGP in approach and technology. Below we will describe PEM and then compare and contrast it to PGP. For more information about PEM, see (Kent, 1993).

Messages sent using PEM are first converted to a canonical form so they all have the same conventions about white space (e.g., tabs, trailing spaces) and the use of carriage returns and line feeds. This transformation is done to eliminate the effects of message transfer agents that modify messages not to their liking.

Without canonicalization, such modifications might affect hashes made from messages at their destinations.

Next, a message hash is computed using MD2 or MD5. It is not optional, as it is in PGP. Then the concatenation of the hash and the message is encrypted using DES. In light of the known weakness of a 56-bit key, this choice is certainly suspect. The encrypted message can then be encoded with base64 coding and transmitted to the recipient. Mailing lists are explicitly supported.

As in PGP, each message is encrypted with a one-time key that is enclosed along with the message. The key can be protected either with RSA or with triple DES using EDE. In practice, everyone uses RSA, so we will concentrate on that. In fact, we have to: PEM does not tell how to do key management with DES.

Key management is more structured than in PGP. Keys are certified by **certification authorities** in the form of certificates stating a user's name, public key, and the key's expiration date. Each certificate has a unique serial number for identifying it. Certificates include an MD5 hash signed by the certification authority's private key. These certificates conform to the ITU X.509 recommendation for public key certificates, and as such, use X.400 names like the Tom Smith example given earlier.

PGP has a similar scheme (without the use of X.509), but has a problem: Should a user believe a certification authority? PEM solves this problem by certifying the certification authorities using what are called **PCAs (Policy Certification Authorities)**. These, in turn, are certified by the **IPRA (Internet Policy Registration Authority)**, the ultimate arbiter of who's naughty and who's nice.

Each PCA must define an official policy on registration and file it with IPRA. These statements are then signed by IPRA and made public. For example, one PCA may insist on having users under its jurisdiction show up in person with a birth certificate, drivers' license, passport, two major credit cards, a live witness, and a public key on floppy disk. Another PCA may accept email registrations from strangers. By making the policy statements public, users have some basis for deciding which authorities to trust. No provision has been made for seeing if the policies are actually enforced.

Three different kinds of certification authorities are planned. An organizational one can issue certificates for its employees. Most companies will run their own. A residential one will operate on behalf of private citizens, much as current Internet service providers will provide service to anyone willing to pay for it. Finally, a scheme is planned for anonymous registration. With all these certification authorities running around, the need for the PCAs to ride herd on them should now be clear.

While rigidly hierarchical and bureaucratic, this scheme has the advantage over PGP of making certificate revocation potentially practical. Revocation is needed if a user wants to change his public key, for example, because it has been compromised or his certification authority has been burglarized (or stolen). Revocation is accomplished by a user telling his certification authority that his public

key has been compromised (or possibly vice versa). The certification authority then adds the serial number of the now-invalid certificate to a list of revoked certificates, signs it, and spreads the list far and wide.

Anyone wanting to send a PEM message to a user must therefore first check the most recent revocation list to see if the cached public key is still valid. This process is analogous to a merchant checking the list of stolen credit cards before accepting one. Critics of PEM argue that checking all the time is too much work so nobody will bother. Supporters argue that computers do not get bored; if they are programmed to check all the time, they will check all the time.

Some of the similarities and differences between PGP and PEM are listed in Fig. 7-51. Most of these points have already been covered, but a few are worth commenting on. Authentication seems more important in PEM than in PGP since it is mandatory in PEM and optional in PGP. PEM also carries the authentication information outside the encryption wrapper, which means that the network can verify the origin of every message. As a consequence, eavesdroppers can log who is sending to whom, even if they cannot read the messages.

All these technical differences aside, there is a surprising cultural difference as well. PGP, which is not an official internet standard, has the Internet culture. PEM, which is an official Internet standard, does not. PGP was based on what Dave Clark calls "rough consensus and running code." Somebody (Zimmermann) thought of a solution to a well-known problem, implemented it well, and released the source code for everyone to use. PEM began as a four-part official standard, using ASN.1 to define layouts, X.400 to define names, and X.509 to define certificates. It uses a rigid three-layer organizational hierarchy for multiple kinds of certification authorities, complete with officially certified policy statements and a requirement that everyone trust the IPRA. Implementations came later and are far behind PGP in quality, quantity, and availability on many platforms. In short, PGP looks like a typical Internet package, whereas PEM exhibits most of the characteristics of an OSI standard that Internet people hate and PTTs love. You figure.

7.5. USENET NEWS

One of the more popular applications of computer networking is the world-wide system of newsgroups called **net news**. Often net news is referred to as **USENET**, which harks back to a separate UNIX-to-UNIX physical network that once carried the traffic using a program called **uucp**. Nowadays, much of the traffic is carried on the Internet, but USENET and the Internet are not the same. Some Internet sites do not get net news, and other sites get net news without being on the Internet.

In the follow sections we will describe USENET. First we will look at it from the users' viewpoint. Then we will describe how it is implemented.

Item	PGP	PEM
Supports encryption?	Yes	Yes
Supports authentication?	Yes	Yes
Supports nonrepudiation?	Yes	Yes
Supports compression?	Yes	No
Supports canonicalization?	No	Yes
Supports mailing lists?	No	Yes
Uses base64 coding?	Yes	Yes
Current data encryption algorithm	IDEA	DES
Key length for data encryption (bits)	128	56
Current algorithm for key management	RSA	RSA or DES
Key length for key management (bits)	384/512/1024	Variable
User name space	User defined	X.400
X.509 conformant?	No	Yes
Do you have to trust anyone?	No	Yes (IPRA)
Key certification	Ad hoc	IPRA/PKA/CA hierarchy
Key revocation	Haphazard	Better
Can eavesdroppers read messages?	No	No
Can eavesdroppers read signatures?	No	Yes
Internet Standard?	No	Yes
Designed by	Small team	Standards committee

Fig. 7-51. A comparison of PGP and PEM.

7.5.1. The User View of USENET

A newsgroup is a worldwide discussion forum on some specific topic. People interested in the subject can “subscribe” to the newsgroup. Subscribers can use a special kind of user agent, a news reader, to read all the articles (messages) posted to the newsgroup. People can also post articles to the newsgroup. Each article posted to a newsgroup is automatically delivered to all the subscribers, wherever they may be in the world. Delivery typically takes between a few seconds and a few hours, depending how far off the beaten path the sender and receiver are. In effect, a newsgroup is somewhat like a mailing list, but internally it is implemented differently. It can also be thought of as a kind of high-level multicast.

The number of newsgroups is so large (probably over 10,000) that they are

arranged in a hierarchy to make them manageable. Figure 7-52 shows the top levels of the "official" hierarchies. Other hierarchies also exist, but these are typically intended for regional consumption or are in languages other than English. One of the other hierarchies *alt*, is special. *Alt* is to the official groups as a flea market is to a department store. It is a chaotic, unregulated mishmash of newsgroups on all topics, some of which are very popular, and most of which are worldwide.

Name	Topics covered
Comp	Computers, computer science, and the computer industry
Sci	The physical sciences and engineering
Humanities	Literature and the humanities
News	Discussion of USENET itself
Rec	Recreational activities, including sports and music
Misc	Everything that does not fit in somewhere else
Soc	Socializing and social issues
Talk	Diatribes, polemics, debates and arguments galore
Alt	Alternative tree covering virtually everything

Fig. 7-52. USENET hierarchies in order of decreasing signal-to-noise ratio.

The *comp* groups were the original USENET groups. These groups are populated by computer scientists, computer professionals, and computer hobbyists. Each one features technical discussions on a topic related to computer hardware or software.

The *sci* and *humanities* groups are populated by scientists, scholars, and amateurs with an interest in physics, chemistry, biology, Shakespeare, and so on. Not entirely surprisingly, the *sci* hierarchy is much larger than the *humanities* hierarchy because the very concept of instant electronic communication with colleagues all over the world is something most scientists like, and most humanists are at least skeptical about. C.P. Snow was right.

The *news* hierarchy is used to discuss and manage the news system itself. System administrators can get help here, and discussions about whether to create new newsgroups occurs here.

The hierarchies covered so far have a professional, somewhat academic tone. That changes with *rec* which is about recreational activities and hobbies. Nevertheless, many of the people who post here are fairly knowledgeable about their respective interests.

As we drift downward, we come to *soc*, which has many newsgroups concerning, politics, gender, religion, various national cultures, and genealogy. *Talk*

covers controversial topics and is populated by people who are strong on opinions, weak on facts. *Alt* is a complete alternative tree which operates under its own rules.

Each of the categories listed in Fig. 7-52 is broken into subcategories, recursively. For example, *rec.sport* is about sports, *rec.sport.basketball* is about basketball, and *rec.sport.basketball.women* is about women's basketball. A sample of some of the newsgroups in each category is given in Fig. 7-53. In many cases, the existence of additional groups can be inferred by changing the obvious parameters. For example, *comp.lang.c* is about the C programming language, but the *.c* can be replaced by just about every other programming language to generate the name of the corresponding newsgroup.

Numerous news readers exist. Like email readers, some are keyboard based; others are mouse based. In nearly all cases, when the news reader is started, it checks a file to see which newsgroups the user subscribes to. It then typically displays a one-line summary of each as-yet-unread article in the first newsgroup and waits for the user to select one or more for reading. The selected articles are then displayed one at a time. After being read, they can be discarded, saved, printed, and so on.

News readers also allow users to subscribe and unsubscribe to newsgroups. Changing a subscription simply means editing the local file listing which newsgroups the user is subscribed to. To make an analogy, subscribing to a newsgroup is like watching a television program. If you want to watch some program every week, you just do it. You do not have to register with some central authority first.

News readers also handle posting. The user composes an article and then gives a command or clicks on a icon to send the article on its way. Within a day, it will reach almost everyone in the world subscribing to the newsgroup to which it was posted. It is possible to **crosspost** an article, that is, to send it to multiple newsgroups with a single command. It is also possible to restrict the geographic distribution of a posting. An announcement of Tuesday's colloquium at Stanford will probably not be of much interest in, say, Hong Kong, so the posting can be restricted to California.

The sociology of USENET is unique, to put it mildly. Never before has it been possible for thousands of people who do not know each other to have world-wide discussions on a vast variety of topics. For example, it is now possible for someone with a problem to post it to the net. The next day, the poster may have 18 solutions, and with a little bit of luck, only 17 of them are wrong.

Unfortunately, some people use their new-found power to communicate to a large group irresponsibly. When someone posts a message saying: "People like you should be shot" tempers flare and a torrent of abusive postings, called a **flamewar**, typically follows.

This situation can be attacked in two ways, one individual and one collective. Individual users can install a **killfile**, which specifies that articles with a certain subject or from a certain person are to be discarded upon arrival, prior to being

Name	Topics covered
Comp.ai	Artificial intelligence
Comp.databases	Design and implementation of database systems
Comp.lang.c	The C programming language
Comp.os.minix	Tanenbaum's educational MINIX operating system
Comp.os.ms-windows.video	Video hardware and software for Windows
Sci.bio.entomology.lepidoptera	Research on butterflies and moths
Sci.geo.earthquakes	Geology, seismology, and earthquakes
Sci.med.orthopedics	Orthopedic surgery
Humanities.lit.authors.shakespeare	Shakespeare's plays and poetry
News.groups	Potential new newsgroups
News.lists	Lists relating to USENET
Rec.arts.poems	Free poetry
Rec.food.chocolate	Yum yum
Rec.humor.funny	Did you hear the joke about the farmer who ...
Rec.music.folk	Folks discussing folk music
Misc.jobs.offered	Announcements of positions available
Misc.health.diabetes	Day-to-day living with diabetes
Soc.culture.estonia	Life and culture in Estonia
Soc.singles	Single people and their interests
Soc.couples	Graduates of soc.singles
Talk.abortion	No signal, all noise
Talk.rumors	This is where rumors come from
Alt.alien.visitors	Place to report flying saucer rides
Alt.bermuda.triangle	If you read this, you vanish mysteriously
Alt.sex.voyeurism	Take a peek and see for yourself
Alt.tv.simpsons	Bart et al.

Fig. 7-53. A small selection of the newsgroups.

displayed. Most news readers also allow an individual discussion thread to be killed, too. This feature is useful when a discussion looks like it is starting to get into an infinite loop.

If enough subscribers to a group get annoyed with newsgroup pollution, they can propose having the newsgroup be moderated. A **moderated newsgroup** is one in which only one person, the moderator, can post articles to the newsgroup. All postings to a moderated newsgroup are automatically sent to the moderator, who posts the good ones and discards the bad ones. Some topics have both a moderated newsgroup and an unmoderated one.

Since thousands of people subscribe to USENET for the first time every day, the same beginner's questions tend to be asked over and over. To reduce this traffic, many newsgroups have constructed a **FAQ (Frequently Asked Questions)** document that tries to answer all the questions that beginners have. Some of these are highly authoritative and run to over 100 pages. The maintainer typically posts them once or twice a month.

USENET is full of jargon such as BTW (By The Way), ROFL (Rolling On the Floor Laughing), and IMHO (In My Humble Opinion). Many people also use little ASCII symbols called **smileys** or **emoticons**. A few of the more interesting ones are reproduced in Fig. 7-54. For most, rotating the book 90 degrees clockwise will make them clearer. For a minibook giving over 650 smileys, see (Sanderson and Dougherty, 1993).

Smiley	Meaning	Smiley	Meaning	Smiley	Meaning
: -)	I'm happy	= : -)	Abe Lincoln	: +)	Big nose
: -(I'm sad/angry	=) : -)	Uncle Sam	: -))	Double chin
: -	I'm apathetic	* < : -)	Santa Claus	: -{}	Mustache
; -)	I'm winking	< : -(Dunce	# : -)	Matted hair
: -(O)	I'm yelling	(- :	Australian	8 -)	Wears glasses
: -(*)	I'm vomiting	: -)X	Man with bowtie	C : -)	Large brain

Fig. 7-54. Some smileys.

Although most people use their real names in postings, some people wish to remain totally anonymous, especially when posting to controversial newsgroups or when posting personal ads to newsgroups dealing with finding partners. This desire has led to the creation of **anonymous remailers**, which are servers that accept email messages (including postings) and change the *From:*, *Sender:*, and *Reply-To:* fields to make them point to the remailer instead of the sender. Some of the remailers assign a number to each user and forward email addressed to these numbers, so people can send email replies to anonymous postings like "SWF 25 seeks SWM/DWM 20-30" Whether these remailers can keep their secrets when the local police become curious about the identity of some user is doubtful (Barlow, 1995).

As more and more people subscribe to USENET, there is a constant demand for new and more specialized newsgroups. Consequently, a procedure has been established for creating new ones. Suppose that somebody likes cockroaches and wants to talk to other cockroach fans. He posts a message to *news.groups* naming the proposed group, say *rec.animals.wildlife.cockroaches*, and describing why it is so important (cockroaches are fascinating; there are 3500 species of them; they come in red, yellow, green, brown, and black; they appeared on earth long before

the first dinosaurs; they were probably the first flying animals, and so on). He also specifies whether or not it should be moderated.

Discussion then ensues. When it settles down, an email vote is taken. The votes are posted, identifying who voted which way (to prevent fraud). If the yeas outnumber the nays by more than 2:1 and there were at least 100 more yeas than nays, the moderator of *news.groups* posts a message accepting the new newsgroup. This message is the signal to system administrators worldwide that the new newsgroup has been blessed by the powers that be and is now official.

New group creation is less formal in the *alt* hierarchy and this is, in fact, the reason *alt* exists. Some of the newsgroups there are so close to the legal and moral edge of what is tolerable that they would never have been accepted in a public vote. In effect, the people who supported them, just bypassed the normal procedure and created their own hierarchy. Nevertheless, much of the *alt* hierarchy is fairly conventional.

7.5.2. How USENET Is Implemented

Some of the smaller news groups are implemented as mailing lists. To post an article to such a mailing list, one sends it to the mailing list address, which causes copies to be sent to each address on the mailing list.

However, if half the undergraduates at a large university subscribed to *alt.sex*, the servers there would collapse under the weight of the incoming email. Consequently, USENET is not generally implemented using mailing lists. Instead each site (campus, company, or Internet service provider) stores incoming mail in a single directory, say, *news*, with subdirectories for *comp*, *sci*, etc. These, in turn have subdirectories such as *news/comp/os/minix*. All incoming news is deposited in the appropriate directory. News readers just fetch the articles from there as they need them. This arrangement means that each site needs only one copy of each news article, no matter how many people subscribe to its newsgroup. After a few days, articles time out and are removed from the disk.

To get on USENET, a site must have a **newsfeed** from another site on USENET. One can think of the set of all sites that get net news as the nodes of a directed graph. The transmission lines connecting pairs of nodes form the arcs of the graph. This graph is USENET. Note that being on the Internet is neither necessary nor sufficient for being on USENET.

Periodically, each site that wants news can poll its newsfeed(s), asking if any new news has arrived since the previous contact. If so, that news is collected and stored in the appropriate subdirectory of *news*. In this manner, news diffuses around the network. It is equally possible for the newsfeed, rather than the receiver, to take the initiative and make contact when there is enough new news. Initially, most sites polled their newsfeeds, but now it is mostly the other way.

Not every site gets all newsgroups. There are several reasons here. First, the total newsfeed exceeds 500 MB per day and is growing rapidly. Storing it all

would require a very large amount of disk space. Second, transmission time and cost are issues. At 28.8 kbps, it takes 39 hours and a dedicated telephone line to transmit 24 hours worth of news. Even at 56 kbps, getting everything requires having a dedicated line for almost 20 hours a day. In fact, the total volume has now gotten so large that newsfeeds via satellite have been created.

Third, not every site is interested in every topic. For example, it is unlikely that many people at companies in Finland want to read *rec.arts.manga* (about Japanese comic books). Finally, some newsgroups are a bit too funky for the tastes of many system administrators, who then ban them, despite considerable local interest. In Dec. 1995, the worldwide CompuServe network (temporarily) stopped carrying all newsgroups with “sex” in the name because some minor German official thought this would be a good way to combat pornography. The ensuing uproar was predictable, instantaneous, worldwide, and very loud.

News articles have the same format as RFC 822 email messages, but with the addition of a few extra headers. This property makes them easy to transport and compatible with most of the existing email software. The news headers are defined in RFC 1036. An example article is shown in Fig. 7-55.

```
From: Vogel@nyu.edu
Message-Id: <54731@nyu.edu>
Subject: Bird Sighting
Path: cs.vu.nl!sun4n!EU.net!news.sprintlink.net!in2.uu.net!pc144.nyu.edu!news
Newsgroups: rec.birds
Followup-To: rec.birds
Distribution: world
Nntp-Posting-host: nuthatch.bio.nyu.edu
References:
Organization: New York University
Lines: 4
Summary: Guess what I saw
```

I just saw an ostrich on 52nd St. and Fifth Ave. in New York. Is this their migration season? Did anybody else see it?

Jay Vogel

Fig. 7-55. A sample news article.

A few words about the news headers are perhaps in order. The *Path:* header is the list of nodes the message traversed to get from the poster to the recipient. At each hop, the forwarding machine puts its name at the front of the list. This list gives a path back to the poster. The use of exclamation marks (pronounced: bang) go back to USENET addresses, which predate DNS.

The *Newsgroups:* header tells which newsgroups the message belongs to. It may contain more than one newsgroup name. Any message crossposted to

multiple newsgroups will contain all of their names. Because multiple names are allowed here, the *Followup-To:* header is needed to tell people where to post comments and reactions to put all of the subsequent discussion in one newsgroup.

The *Distribution:* header tells how far to spread the posting. It may contain one or more state or country codes, the name of a specific site or network, or "world."

The *Nntp-Posting-Host:* header is analogous to the RFC 822 *Sender:* header. It tells which machine actually posted the article, even if it was composed on a different machine (NNTP is the news protocol, described below).

The *References:* header indicates that this article is a response to an earlier article and gives the ID of that article. It is required on all follow-up articles and prohibited when starting a new discussion.

The *Organization:* header can be used to tell what company, university, or agency the poster is affiliated with. Articles that fill in this header often have a disclaimer at the end saying that if the article is goofy, it is not the organization's fault.

The *Lines:* header gives the length of the body. The header lines and the blank line separating the header from the body do not count.

The *Subject:* lines tie discussion threads together. Many news readers have a command to allow the user to see the next article on the current subject, rather than the next article that came in. Also, killfiles and kill commands use this header to know what to reject.

Finally, the *Summary:* is normally used to summarize the follow-up article. On follow-up articles, the *Subject:* header contains "Re: " followed by the original subject.

NNTP—Network News Transfer Protocol

Now let us look at how articles diffuse around the network. The initial algorithm just flooded articles onto every line within USENET. While this worked for a while, eventually the volume of traffic made this scheme impractical, so something better had to be worked out.

Its replacement was a protocol called **NNTP (Network News Transfer Protocol)**, which is defined in RFC 977. NNTP has something of the same flavor as SMTP, with a client issuing commands in ASCII and a server issuing responses as decimal numbers coded in ASCII. Most USENET machines now use NNTP.

NNTP was designed for two purposes. The first goal was to allow news articles to propagate from one machine to another over a reliable connection (e.g., TCP). The second goal was to allow users whose desktop computers cannot receive news to read news remotely. Both are widely used, but we will concentrate on how news articles spread out over the network using NNTP.

As mentioned above, two general approaches are possible. In the first one, news pull, the client calls one of its newsfeeds and asks for new news. In the

second one, news push, the newsfeed calls the client and announces that it has news. The NNTP commands support both of these approaches, as well as having people read news remotely.

To acquire recent articles, a client must first establish a TCP connection with port 119 on one of its newsfeeds. Behind this port is the NNTP daemon, which is either there all the time waiting for clients or is created on the fly as needed. After the connection has been established, the client and server communicate using a sequence of commands and responses. These commands and responses are used to ensure that the client gets all the articles it needs, but no duplicates, no matter how many newsfeeds it uses. The main ones used for moving articles between news daemons are listed in Fig. 7-56.

Command	Meaning
LIST	Give me a list of all newsgroups and articles you have
NEWGROUPS date time	Give me a list of newsgroups created after date/time
GROUP grp	Give me a list of all articles in grp
NEWNEWS grps date time	Give me a list of new articles in specified groups
ARTICLE id	Give me a specific article
POST	I have an article for you that was posted here
IHAVE id	I have article id. Do you want it?
QUIT	Terminate the session

Fig. 7-56. The principal NNTP commands for news diffusion.

The *LIST* and *NEWGROUPS* commands allow the client to find out which groups the server has. The former gives the complete list. The latter gives only those groups created after the date and time specified. If the client knows the list is long, it is more efficient for the client to keep track of what each of its newsfeeds has and just ask for updates. The responses to each of these commands is a list, in ASCII, one newsgroup per line, giving the name of the newsgroup, the number of the last article the server has, the number of the first article the server has, and a flag telling whether posting to this newsgroup is allowed.

Once the client knows which newsgroups the server has, it can begin asking about what articles the server has (e.g., for old newsgroups when *NEWGROUPS* is used). The *GROUP* and *NEWNEWS* commands are used for this purpose. Again, the former gives the full list and the latter gives only updates subsequent to the indicated date and time, normally the time of the last connection to this newsfeed. The first parameter may contain asterisks, meaning all of them. For example, *comp.os.** means all the newsgroups that start with the string *comp.os*.

After the client has assembled a complete list of which articles exist in which groups (or even before it has the full list), it can begin to ask for the articles it

needs using the *ARTICLE* command. Once all the required articles are in, the client can offer articles it has acquired from other newsfeeds using the *IHAVE* command and articles that were posted locally using the *POST* command. The server can accept or decline these, as it wishes. When the client is done, it can terminate the session using *QUIT*. In this way, each machine has complete control over which articles it gets from which newsfeeds, eliminating all duplicate articles.

As an example of how NNTP works, consider an information provider, *wholesome.net* that wants to avoid controversy at all costs, so the only newsgroups it offers are *soc.couples* and *misc.kids*. Nevertheless, management is open minded and willing to carry other newsgroups, provided they contain no material potentially offensive to anyone. Therefore, it wants to be informed of all newly created groups so it can make an informed decision for its customers. A possible scenario between *wholesome.com* acting as the client and its newsfeed, *feeder.com*, acting as the server, is shown in Fig. 7-57. This scenario uses the news pull approach (the client initiates the connection to ask for news). The remarks in parentheses are comments and not part of the NNTP protocol.

In this session, *wholesome.com* first asks if there is any news for *soc.couples*. When it is told there are two articles, it fetches both of them and stores them in *news/soc/couples* as separate files. Each file is named by its article number. Then *wholesome.com* asks about *misc.kids* and is told there is one article. It fetches that one and puts it in *news/misc/kids*.

Having gotten all the news about the groups it carries, it now checks for new groups and is told that two new groups have appeared since the last session. One of them looks promising, so its articles are fetched. The other looks scary, so it is not taken. (*Wholesome.com* has made a big investment in AI software to be able to figure out what to carry just by looking at the names.)

After having acquired all the articles it wants, *wholesome.com* offers *feeder.com* a new article posted by someone at its site. The offer is accepted and the article is transferred. Now *wholesome.com* offers another article, one that came from its other newsfeed. Since *feeder.com* already has this one, it declines. Finally, *wholesome.com* ends the session and releases the TCP connection.

The news push approach is similar. It begins with the newsfeed calling the machine that is to receive the news. The newsfeed normally keeps track of which newsgroups its customers subscribe to and begins by announcing its first article in the first of these newsgroups using the *IHAVE* command. The potential recipient then checks its tables to see whether it already has the article, and can accept or reject it. If the article is accepted, it is transmitted, followed by a line containing a period. Then the newsfeed advertises the second article, and so forth, until all the news has been transferred.

A problem with both news pull and news push is that they use stop and wait. Typically 100 msec are lost waiting for an answer to a question. With 100,000 or more news articles per day, this lost time adds up to a substantial overhead.

```

S: 200 feeder.com NNTP server at your service(response to new connection)
C: NEWNEWS soc.couples 960901 030000 (any new news in soc.couples?)
S: 230 List of 2 articles follows
S: <13281@psyc.berkeley.edu> (article 1 of 2 in soc.couples is from Berkeley)
S: <162721@aol.com> (article 2 of 2 in soc.couples is from AOL)
S: . (end of list)
C: ARTICLE <13281@psyc.berkeley.edu> (please give me the Berkeley article)
S: 220 <13281@psyc.berkeley.edu> follows
S: (entire article <13281@psyc.berkeley.edu> is sent here)
S: . (end of article)
C: ARTICLE <162721@aol.com> (please give me the AOL article)
S: 220 <162721@aol.com> follows
S: (entire article <162721@aol.com> is sent here)
S: . (end of article)
C: NEWNEWS misc.kids 960901 030000 (any new news in misc.kids?)
S: 230 List of 1 article follows
S: <43222@bio.rice.edu> (1 article from Rice)
S: . (end of list)
C: ARTICLE <43222@bio.rice.edu> (please give me the Rice article)
S: 220 <43222@bio.rice.edu> follows
S: (entire article <43222@bio.rice.edu> is sent here)
S: . (end of article)
C: NEWGROUPS 960901 030000
S: 231 2 new groups follow
S: rec.pets
S: rec.nude
S: .
C: NEWNEWS rec.pets 0 0 (list everything you have)
S: 230 List of 1 article follows
S: <124@fido.net> (1 article from fido.net)
S: . (end of list)
C: ARTICLE <124@fido.net> (please give me the fido.net article)
S: 220 <124@fido.net> follows
S: (entire article is sent here)
S: .
C: POST
S: 340 (please send your posting)
C: (article posted on wholesome.com sent here)
S: 240 (article received)
C: IHAVE <5321@foo.com>
S: 435 (I already have it, please do not send it)
C: QUIT
S: 205 (Have a nice day)

```

Fig. 7-57. How *wholesome.com* might acquire news articles from its newsfeed.

7.6. THE WORLD WIDE WEB

The World Wide Web is an architectural framework for accessing linked documents spread out over thousands of machines all over the Internet. In 5 years, it went from being a way to distribute high-energy physics data to the application that millions of people think of as being "The Internet." Its enormous popularity stems from the fact that it has a colorful graphical interface that is easy for beginners to use, and it provides an enormous wealth of information on almost every conceivable subject, from aboriginals to zoology.

The Web (also known as **WWW**) began in 1989 at CERN, the European center for nuclear research. CERN has several accelerators at which large teams of scientists from the participating European countries carry out research in particle physics. These teams often have members from half a dozen or more countries. Most experiments are highly complex, and require years of advance planning and equipment construction. The Web grew out of the need to have these large teams of internationally dispersed researchers collaborate using a constantly changing collection of reports, blueprints, drawings, photos, and other documents.

The initial proposal for a web of linked documents came from CERN physicist Tim Berners-Lee in March 1989. The first (text-based) prototype was operational 18 months later. In December 1991, a public demonstration was given at the Hypertext '91 conference in San Antonio, Texas. Development continued during the next year, culminating in the release of the first graphical interface, Mosaic, in February 1993 (Vetter et al., 1994).

Mosaic was so popular that a year later, its author, Marc Andreessen left the National Center for Supercomputing Applications, where Mosaic was developed, to form a company, Netscape Communications Corp., whose goal was to develop clients, servers, and other Web software. When Netscape went public in 1995, investors, apparently thinking this was the next Microsoft, paid 1.5 billion dollars for the stock. This record was all the more surprising because the company had only one product, was operating deeply in the red, and had announced in its prospectus that it did not expect to make a profit for the foreseeable future.

In 1994, CERN and M.I.T. signed an agreement setting up the World Wide Web Consortium, an organization devoted to further developing the Web, standardizing protocols, and encouraging interoperability between sites. Berners-Lee became the director. Since then, hundreds of universities and companies have joined the consortium. M.I.T. runs the U.S. part of the consortium and the French research center, INRIA, runs the European part. Although there are more books about the Web than you can shake a stick at, the best place to get up-to-date information about the Web is (naturally) on the Web itself. The consortium's home page can be found at <http://www.w3.org>. Interested readers are referred there for links to pages covering all of the consortium's documents and activities.

In the following sections we will describe how the Web appears to the user, and, especially, how it works inside. Since the Web is basically a client-server

system, we will discuss both the client (i.e., user) side and the server side. Then we will examine the language in which Web pages are written (HTML and Java). Finally, comes an examination of how to find information on the Web.

7.6.1. The Client Side

From the users' point of view, the Web consists of a vast, worldwide collection of documents, usually just called **pages** for short. Each page may contain links (pointers) to other, related pages, anywhere in the world. Users can follow a link (e.g., by clicking on it), which then takes them to the page pointed to. This process can be repeated indefinitely, possibly traversing hundreds of linked pages while doing so. Pages that point to other pages are said to use **hypertext**.

Pages are viewed with a program called a **browser**, of which Mosaic and Netscape are two popular ones. The browser fetches the page requested, interprets the text and formatting commands that it contains, and displays the page, properly formatted, on the screen. An example is given in Fig. 7-58(a). Like many Web pages, this one starts with a title, contains some information, and ends with the email address of the page's maintainer. Strings of text that are links to other pages, called **hyperlinks**, are highlighted, either by underlining, displaying them in a special color, or both. To follow a link, the user places the cursor on the highlighted area (using the mouse or the arrow keys) and selects it (by clicking a mouse button or hitting ENTER). Although nongraphical browsers, such as Lynx, exist, they are not as popular as graphical browsers, so we will concentrate on the latter. Voice-based browsers are also being developed.

Users who are curious about the Department of Animal Psychology can learn more about it by clicking on its (underlined) name. The browser then fetches the page to which the name is linked and displays it, as shown in Fig. 7-58(b). The underlined items here can also be clicked on to fetch other pages, and so on. The new page can be on the same machine as the first one, or on a machine halfway around the globe. The user cannot tell. Page fetching is done by the browser, without any help from the user. If the user ever returns to the main page, the links that have already been followed may be shown with a dotted underline (and possibly a different color) to distinguish them from links that have not been followed. Note that clicking on the *Campus Information* line in the main page does nothing. It is not underlined, which means that it is just text and is not linked to another page.

Most browsers have numerous buttons and features to make it easier to navigate the Web. Many have a button for going back to the previous page, a button for going forward to the next page (only operative after the user has gone back from it), and a button for going straight to the user's own home page. Most browsers have a button or menu item to set a bookmark on a given page and another one to display the list of bookmarks, making it possible to revisit any of

WELCOME TO THE UNIVERSITY OF EAST PODUNK'S WWW HOME PAGE

- Campus Information
 - [Admissions information](#)
 - [Campus map](#)
 - [Directions to campus](#)
 - [The UEP student body](#)
- Academic Departments
 - [Department of Animal Psychology](#)
 - [Department of Alternative Studies](#)
 - [Department of Microbiotic Cooking](#)
 - [Department of Nontraditional Studies](#)
 - [Department of Traditional Studies](#)

Webmaster@eastpodunk.edu

(a)

THE DEPARTMENT OF ANIMAL PSYCHOLOGY

- [Information for prospective majors](#)
- Personnel
 - [Faculty members](#)
 - [Graduate students](#)
 - [Nonacademic staff](#)
- [Research Projects](#)
- [Positions available](#)
- Our most popular courses
 - [Dealing with herbivores](#)
 - [Horse management](#)
 - [Negotiating with your pet](#)
 - [User-friendly doghouse construction](#)
- [Full list of courses](#)

Webmaster@animalpsyc.eastpodunk.edu

(b)

Fig. 7-58. (a) A Web page. (b) The page reached by clicking on Department of Animal Psychology

them with a single mouse click. Pages can also be saved to disk or printed. Numerous options are generally available for controlling the screen layout and setting various user preferences. A comparison of nine browsers is given in (Berghel, 1996).

In addition to having ordinary text (not underlined) and hypertext (underlined), Web pages can also contain icons, line drawings, maps, and photographs. Each of these can (optionally) be linked to another page. Clicking on one of these elements causes the browser to fetch the linked page and display it, the same as clicking on text. With images such as photos and maps, which page is fetched next may depend on what part of the image was clicked on.

Not all pages are viewable in the conventional way. For example, some pages consist of audio tracks, video clips, or both. When hypertext pages are mixed with other media, the result is called **hypermedia**. Some browsers can display all kinds of hypermedia, but others cannot. Instead they check a configuration file to see how to handle the received data. Normally, the configuration file gives the name of a program, called an **external viewer**, or a **helper application**, to be run with the incoming page as input. If no viewer is configured, the browser usually asks the user to choose one. If no viewer exists, the user can tell the browser to save the incoming page to a disk file, or to discard it. Helper applications for producing speech are making it possible for even blind users to access the Web. Other helper applications contain interpreters for special Web languages, making it possible to download and run programs from Web pages. This mechanism makes it possible to extend the functionality of the Web itself.

Many Web pages contain large images, which take a long time to load. For example, fetching an uncompressed 640×480 (VGA) image with 24 bits per pixel (922 KB) takes about 4 minutes over a 28.8-kbps modem line. Some browsers deal with the slow loading of images by first fetching and displaying the text, then getting the images. This strategy gives the user something to read while the images are coming in and also allows the user to kill the load if the page is not sufficiently interesting to warrant waiting. An alternative strategy is to provide an option to disable the automatic fetching and display of images.

Some page writers attempt to placate potentially bored users by displaying images in a special way. First the image quickly appears in a coarse resolution. Then the details are gradually filled in. For the user, seeing the whole image after a few seconds, albeit at low resolution, is often preferable to seeing it built up slowly from the top, scan line by scan line.

Some Web pages contain forms that request the user to enter information. Typical applications of these forms are searching a database for a user-supplied item, ordering a product, or participating in a public opinion survey. Other Web pages contain maps that allow users to click on them to zoom in or get information about some geographical area. Handling forms and active (clickable) maps requires more sophisticated processing than just fetching a known page. We will describe later how these features are implemented.

Some browsers use the local disk to cache pages that they have fetched. Before a page is fetched, a check is made to see if it is in the local cache. If so, it is only necessary to check if the page is still up to date. If so, the page need not be loaded again. As a result, clicking on the BACK button to see the previous page is normally very fast.

To host a Web browser, a machine must be directly on the Internet, or at least have a SLIP or PPP connection to a router or other machine that is directly on the Internet. This requirement exists because the way a browser fetches a page is to establish a TCP connection to the machine where the page is, and then send a message over the connection asking for the page. If it cannot establish a TCP connection to an arbitrary machine on the Internet, a browser will not work.

Sometimes the lengths that people will go to get Web access are amazing. At least one company is offering Web-by-Fax service. A client without Internet access calls up the Web-by-Fax server and logs in using the telephone keypad. He then types in a code identifying the Web page desired and it is faxed to the caller's fax machine.

7.6.2. The Server Side

Every Web site has a server process listening to TCP port 80 for incoming connections from clients (normally browsers). After a connection has been established, the client sends one request and the server sends one reply. Then the connection is released. The protocol that defines the legal requests and replies is called HTTP. We will study it in some detail below, but a simple example using it may provide a reasonable idea of how Web servers work. Figure 7-59 shows how the various parts of the Web model fit together.

For this example, we can imagine that the user has just clicked on some piece of text or perhaps on an icon that points to the page whose name (URL—Uniform Resource Locator) is *http://www.w3.org/hypertext/WWW/TheProject.html*. We will also explain URLs later on in this chapter. For the moment, it is sufficient to know that a URL has three parts: the name of the protocol (*http*), the name of the machine where the page is located (*www.w3.org*), and the name of the file containing the page (*hypertext/WWW/TheProject.html*). The steps that occur between the user's click and the page being displayed are as follows:

1. The browser determines the URL (by seeing what was selected).
2. The browser asks DNS for the IP address of *www.w3.org*.
3. DNS replies with 18.23.0.23.
4. The browser makes a TCP connection to port 80 on 18.23.0.23.
5. It then sends a *GET /hypertext/WWW/TheProject.html* command.
6. The *www.w3.org* server sends the file *TheProject.html*.




Fig. 7-59. The parts of the Web model.

7. The TCP connection is released.
8. The browser displays all the text in *TheProject.html*.
9. The browser fetches and displays all images in *TheProject.html*.

Many browsers display which step they are currently executing in a status line at the bottom of the screen. In this way, when the performance is poor, the user can see if it is due to DNS not responding, the server not responding, or simply network congestion during page transmission.

It is worth noting that for each in-line image (icon, drawing, photo, etc.) on a page, the browser establishes a new TCP connection to the relevant server to fetch the image. Needless to say, if a page contains many icons, all on the same server, establishing, using, and releasing a new connection for each one is not wildly efficient, but it keeps the implementation simple. Future revisions of the protocol will address the efficiency issue. One proposal is given in (Mogul, 1995).

Because HTTP is an ASCII protocol like SMTP, it is quite easy for a person at a terminal (as opposed to a browser) to directly talk to Web servers. All that is needed is a TCP connection to port 80 on the server. The simplest way to get such a connection is to use the Telnet program. Figure 7-60 shows a scenario of how this can be done. In this example, the lines marked *C:* are typed in by the user (client), the lines marked *T:* are produced by the Telnet program, and the lines marked *S:* are produced by the server at M.I.T.


```
C: telnet www.w3.org 80
T: Trying 18.23.0.23 ...
T: Connected to www.w3.org.
T: Escape character is '^]'.
C: GET /hypertext/WWW/TheProject.html HTTP/1.0
C:
S: HTTP/1.0 200 Document follows
S: MIME-Version: 1.0
S: Server: CERN/3.0
S: Content-Type: text/html
S: Content-Length: 8247
S:
S: <HEAD> <TITLE> The World Wide Web Consortium (W3C) </TITLE> </HEAD>
S: <BODY>
S: <H1> <IMG ALIGN=MIDDLE ALT="W3C" SRC="icons/WWW/w3c_96x67.gif">
S: The World Wide Web Consortium </H1> <P>
S:
S: The World Wide Web is the universe of network-accessible information.
S: The <A HREF="Consortium/"> World Wide Web Consortium </A>
S: exists to realize the full potential of the Web. <P>
S:
S: W3C works with the global community to produce
S: <A HREF="#Specifications"> specifications </A> and
S: <A HREF="#Reference"> reference software </A> .
S: W3C is funded by industrial
S: <A HREF="Consortium/Member/List.html"> members </A>
S: but its products are freely available to all. <P>
S:
S: In this document:
S: <menu>
S: <LI> <A HREF="#Specifications"> Web Specifications and Development Areas </A>
S: <LI> <A HREF="#Reference"> Web Software </A>
S: <LI> <A HREF="#Community"> The World Wide Web and the Web Community </A>
S: <LI> <A HREF="#Joining"> Getting involved with the W3C </A>
S: </menu>
S: <P> <HR>
S: <P> W3C is hosted by the
S: <A HREF="http://www.lcs.mit.edu/"> Laboratory for Computer Science </A> at
S: <A HREF="http://web.mit.edu/"> MIT </A> , and
S: in Europe by <A HREF="http://www.inria.fr/"> INRIA </A> .
S: </BODY>
```

Fig. 7-60. A sample scenario for obtaining a Web page.

Readers are encouraged to try this scenario personally (preferably from a UNIX system, because some other systems do not return the connection status). Be sure to note the spaces and the protocol version on the *GET* line, and the blank line following the *GET* line. As an aside, the actual text that will be received will differ from what is shown in Fig. 7-60 for three reasons. First, the example output here has been abridged and edited to make it fit on one page. Second, it has been cleaned up somewhat to avoid embarrassing the author, who no doubt expected thousands of people to examine the formatted page, but zero people to scrutinize the HTML that produced it. Third, the contents of the page are constantly being revised. Nevertheless, this example should give a reasonable idea of how HTTP works.

What the example shows is the following. The client, in this case a person, but normally a browser, first connects to a particular host and then sends a command asking for a particular page and specifying a particular protocol and version to use (HTTP/1.0). On line 7, the server responds with a status line telling the protocol it is using (the same as the client) and the code 200, meaning OK. This line is followed by an RFC 822 MIME message, of which five of the header lines are shown in the figure (several others have been omitted to save space). Then comes a blank line, followed by the message body. For sending a picture, the *Content-Type* field might be

```
Content-Type: Image/GIF
```

In this way, the MIME types allow arbitrary objects to be sent in a standard way. As an aside, the MIME *Content-Transfer-Encoding* header is not needed because TCP allows arbitrary byte streams, even pictures, to be sent without modification. The meaning of the commands within angle brackets used in the sample page will be discussed later in this chapter.

Not all servers speak HTTP. In particular, many older servers use the FTP, Gopher, or other protocols. Since a great deal of useful information is available on FTP and Gopher servers, one of the design goals of the Web was to make this information available to Web users. One solution is to have the browser use these protocols when speaking to an FTP or Gopher server. Some of them, in fact, use this solution, but making browsers understand every possible protocol makes them unnecessarily large.

Instead, a different solution is often used: proxy servers (Luotonen and Altis, 1994). A **proxy server** is a kind of gateway that speaks HTTP to the browser but FTP, Gopher, or some other protocol to the server. It accepts HTTP requests and translates them into, say, FTP requests, so the browser does not have to understand any protocol except HTTP. The proxy server can be a program running on the same machine as the browser, but it can also be on a free-standing machine somewhere in the network serving many browsers. Figure 7-61 shows the difference between a browser that can speak FTP and one that uses a proxy.




Fig. 7-61. (a) A browser that speaks FTP. (b) A browser that does not.

Often users can configure their browsers with proxies for protocols that the browsers do not speak. In this way, the range of information sources to which the browser has access is increased.

In addition to acting as a go-between for unknown protocols, proxy servers have a number of other important functions, such as caching. A caching proxy server collects and keeps all the pages that pass through it. When a user asks for a page, the proxy server checks to see if it has the page. If so, it can check to see if the page is still current. In the event that the page is still current, it is passed to the user. Otherwise, a new copy is fetched.

Finally, an organization can put a proxy server inside its firewall to allow users to access the Web, but without giving them full Internet access. In this configuration, users can talk to the proxy server, but it is the proxy server that contacts remote sites and fetches pages on behalf of its clients. This mechanism can be used, for example, by high schools, to block access to Web sites the principal feels are inappropriate for tender young minds.

For information about one of the more popular Web servers (NCSA's HTTP daemon) and its performance, see (Katz et al., 1994; and Kwan et al., 1995).

HTTP—HyperText Transfer Protocol

The standard Web transfer protocol is **HTTP (HyperText Transfer Protocol)**. Each interaction consists of one ASCII request, followed by one RFC 822 MIME-like response. Although the use of TCP for the transport connection is very common, it is not formally required by the standard. If ATM networks become reliable enough, the HTTP requests and replies could be carried in AAL 5 messages just as well.

HTTP is constantly evolving. Several versions are in use and others are under development. The material presented below is relatively basic and is unlikely to change in concept, but some details may be a little different in future versions.

The HTTP protocol consists of two fairly distinct items: the set of requests from browsers to servers and the set of responses going back the other way. We will now treat each of these in turn.

All the newer versions of HTTP support two kinds of requests: simple requests and full requests. A simple request is just a single *GET* line naming the page desired, without the protocol version. The response is just the raw page, with no headers, no MIME, and no encoding. To see how this works, try making a Telnet connection to port 80 of *www.w3.org* (as shown in the first line of Fig. 7-60) and then type

```
GET /hypertext/WWW/TheProject.html
```

but without the HTTP/1.0 this time. The page will be returned with no indication of its content type. This mechanism is needed for backward compatibility. Its use will decline as browsers and servers based on full requests become standard.

Full requests are indicated by the presence of the protocol version on the *GET* request line, as in Fig. 7-60. Requests may consist of multiple lines, followed by a blank line to indicate the end of the request, which is why the blank line was needed in Fig. 7-60. The first line of a full request contains the command (of which *GET* is but one of the possibilities), the page desired, and the protocol/version. Subsequent lines contain RFC 822 headers.

Although HTTP was designed for use in the Web, it has been intentionally made more general than necessary with an eye to future object-oriented applications. For this reason, the first word on the full request line is simply the name of the **method** (command) to be executed on the Web page (or general object). The built-in methods are listed in Fig. 7-62. When accessing general objects, additional object-specific methods may also be available. The names are case sensitive, so, *GET* is a legal method but *get* is not.

Method	Description
GET	Request to read a Web page
HEAD	Request to read a Web page's header
PUT	Request to store a Web page
POST	Append to a named resource (e.g., a Web page)
DELETE	Remove the Web page
LINK	Connects two existing resources
UNLINK	Breaks an existing connection between two resources

Fig. 7-62. The built-in HTTP request methods.

The *GET* method requests the server to send the page (by which we mean object, in the most general case), suitably encoded in MIME. However, if the

GET request is followed by an *If-Modified-Since* header, the server only sends the data if it has been modified since the date supplied. Using this mechanism, a browser that is asked to display a cached page can conditionally ask for it from the server, giving the modification time associated with the page. If the cache page is still valid, the server just sends back a status line announcing that fact, thus eliminating the overhead of transferring the page again.

The *HEAD* method just asks for the message header, without the actual page. This method can be used to get a page's time of last modification, to collect information for indexing purposes, or just to test a URL for validity. Conditional *HEAD* requests do not exist.

The *PUT* method is the reverse of *GET*: instead of reading the page, it writes the page. This method makes it possible to build a collection of Web pages on a remote server. The body of the request contains the page. It may be encoded using MIME, in which case the lines following the *PUT* might include *Content-Type* and authentication headers, to prove that the caller indeed has permission to perform the requested operation.

Somewhat similar to *PUT* is the *POST* method. It too bears a URL, but instead of replacing the existing data, the new data is "appended" to it in some generalized sense. Posting a message to a news group or adding a file to a bulletin board system are examples of appending in this context. It is clearly the intention here to have the Web take over the functionality of the USENET news system.

DELETE does what you might expect: it removes the page. As with *PUT*, authentication and permission play a major role here. There is no guarantee that *DELETE* succeeds, since even if the remote HTTP server is willing to delete the page, the underlying file may have a mode that forbids the HTTP server from modifying or removing it.

The *LINK* and *UNLINK* methods allow connections to be established between existing pages or other resources.

Every request gets a response consisting of a status line, and possibly additional information (e.g., all or part of a Web page). The status line can bear the code 200 (OK), or any one of a variety of error codes, for example 304 (not modified), 400 (bad request), or 403 (forbidden).

The HTTP standards describe message headers and bodies in considerable detail. Suffice it to say that these are very close to RFC 822 MIME messages, so we will not look at them here.

7.6.3. Writing a Web Page in HTML

Web pages are written in a language called **HTML (HyperText Markup Language)**. HTML allows users to produce Web pages that include text, graphics, and pointers to other Web pages. We will begin our study of HTML with these pointers, since they are the glue that holds the Web together.

URLs—Uniform Resource Locators

We have repeatedly said that Web pages may contain pointers to other Web pages. Now it is time to see how these pointers are implemented. When the Web was first created, it was immediately apparent that having one page point to another Web page required mechanisms for naming and locating pages. In particular, there were three questions that had to be answered before a selected page could be displayed:

1. What is the page called?
2. Where is the page located?
3. How can the page be accessed?

If every page were somehow assigned a unique name, there would not be any ambiguity in identifying pages. Nevertheless, the problem would not be solved. Consider a parallel between people and pages. In the United States, almost everyone has a social security number, which is a unique identifier, as no two people have the same one. Nevertheless, armed only with a social security number, there is no way to find the owner's address, and certainly no way to tell whether you should write to the person in English, Spanish, or Chinese. The Web has basically the same problems.

The solution chosen identifies pages in a way that solves all three problems at once. Each page is assigned a **URL (Uniform Resource Locator)** that effectively serves as the page's worldwide name. URLs have three parts: the protocol (also called a scheme), the DNS name of the machine on which the page is located, and a local name uniquely indicating the specific page (usually just a file name on the machine where it resides). For example, the URL for the author's department is

`http://www.cs.vu.nl/welcome.html`

This URL consists of three parts: the protocol (*http*), the DNS name of the host (*www.cs.vu.nl*), and the file name (*welcome.html*), with certain punctuation separating the pieces.

Many sites have certain shortcuts for file names built in. For example, *~user/* might be mapped onto *user's* WWW directory, with the convention that a reference to the directory itself implies a certain file, say, *index.html*. Thus the author's home page can be reached at

`http://www.cs.vu.nl/~ast/`

even though the actual file name is different. At many sites, a null file name defaults to the organization's home page.

Now it should be clear how hypertext works. To make a piece of text clickable, the page writer must provide two items of information: the clickable text to

be displayed and the URL of the page to go to if the text is selected. When the text is selected, the browser looks up the host name using DNS. Now armed with the host's IP address, the browser then establishes a TCP connection to the host. Over that connection, it sends the file name using the specified protocol. Bingo. Back comes the page. This is precisely what we saw in Fig. 7-60.

This URL scheme is open-ended in the sense that it is straightforward to have protocols other than HTTP. In fact, URLs for various other common protocols have been defined, and many browsers understand them. Slightly simplified forms of the more common ones are listed in Fig. 7-63.

Name	Used for	Example
http	Hypertext (HTML)	http://www.cs.vu.nl/~ast/
ftp	FTP	ftp://ftp.cs.vu.nl/pub/minix/README
file	Local file	/usr/suzanne/prog.c
news	News group	news:comp.os.minix
news	News article	news:AA0134223112@cs.utah.edu
gopher	Gopher	gopher://gopher.tc.umn.edu/11/Libraries
mailto	Sending email	mailto:kim@acm.org
telnet	Remote login	telnet://www.w3.org:80

Fig. 7-63. Some common URLs.

Let us briefly go over the list. The *http* protocol is the Web's native language, the one spoken by HTTP servers. It supports all the methods of Fig. 7-62, as well as whatever object-specific methods are needed.

The *ftp* protocol is used to access files by FTP, the Internet's file transfer protocol. FTP has been around more than two decades and is well entrenched. Numerous FTP servers all over the world allow people anywhere on the Internet to log in and download whatever files have been placed on the FTP server. The Web does not change this; it just makes obtaining files by FTP easier, as FTP has a somewhat arcane interface. In due course, FTP will probably vanish, as there is no particular advantage for a site to run an FTP server instead of an HTTP server, which can do everything that the FTP server can do, and more (although there are some arguments about efficiency).

It is possible to access a local file as a Web page, either by using the *file* protocol, or more simply, by just naming it. This approach is similar to using FTP but does not require having a server. Of course, it only works for local files.

The *news* protocol allows a Web user to call up a news article as though it were a Web page. This means that a Web browser is simultaneously a news reader. In fact, many browsers have buttons or menu items to make reading USENET news even easier than using standard news readers.

Two formats are supported for the *news* protocol. The first format specifies a newsgroup and can be used to get a list of articles from a preconfigured news site. The second one requires the identifier of a specific news article to be given, in this case *AA0134223112@cs.utah.edu*. The browser then fetches the given article from its preconfigured news site using the NNTP protocol.

The *gopher* protocol is used by the Gopher system, which was designed at the University of Minnesota and named after the school's athletic teams, the Golden Gophers (as well as being a slang expression meaning "go for", i.e., go fetch). Gopher predates the Web by several years. It is an information retrieval scheme, conceptually similar to the Web itself, but supporting only text and no images. When a user logs into a Gopher server, he is presented with a menu of files and directories, any of which can be linked to another Gopher menu anywhere in the world.

Gopher's big advantage over the Web is that it works very well with 25×80 ASCII terminals, of which there are still quite a few around, and because it is text based, it is very fast. Consequently, there are thousands of Gopher servers all over the world. Using the *gopher* protocol, Web users can access Gopher and have each Gopher menu presented as a clickable Web page. If you are not familiar with Gopher, try the example given in Fig. 7-63 or have your favorite Web search engine look for "gopher."

Although the example given does not illustrate it, it is also possible to send a complete query to a Gopher server using the *gopher+* protocol. What is displayed is the result of querying the remote Gopher server.

The last two protocols do not really have the flavor of fetching Web pages, and are not supported by all browsers, but are useful anyway. The *mailto* protocol allows users to send email from a Web browser. The way to do this is to click on the OPEN button and specify a URL consisting of *mailto:* followed by the recipient's email address. Most browsers will respond by popping up a form containing slots for the subject and other header lines and space for typing the message.

The *telnet* protocol is used to establish an on-line connection to a remote machine. It is used the same way as the Telnet program, which is not surprising, since most browsers just call the Telnet program as a helper application. As an exercise, try the scenario of Fig. 7-60 again, but now using a Web browser.

In short, the URLs have been designed to not only allow users to navigate the Web, but to deal with FTP, news, Gopher, email, and telnet as well, making all the specialized user interface programs for those other services unnecessary, and thus integrating nearly all Internet access into a single program, the Web browser. If it were not for the fact that this scheme was designed by a physics researcher, it could easily pass for the output of some software company's advertising department.

Despite all these nice properties, the growing use of the Web has turned up an inherent weakness in the URL scheme. A URL points to one specific host. For

pages that are heavily referenced, it is desirable to have multiple copies far apart, to reduce the network traffic. The trouble is that URLs do not provide any way to reference a page without simultaneously telling where it is. There is no way to say: "I want page xyz, but I do not care where you get it." To solve this problem and make it possible to replicate pages, the IETF is working on a system of **URIs (Universal Resource Identifiers)**. A URI can be thought of as a generalized URL. This topic is the subject of much current research.

Although we have discussed only absolute URLs here, relative URLs also exist. The difference is analogous to the difference between the absolute file name */usr/ast/foobar* and just *foobar* when the context is unambiguously defined.

HTML—HyperText Markup Language

Now that we have a good idea of how URLs work, it is time to look at HTML itself. HTML is an application of ISO standard 8879, **SGML (Standard Generalized Markup Language)**, but specialized to hypertext and adapted to the Web.

As mentioned earlier, HTML is a markup language, a language for describing how documents are to be formatted. The term "markup" comes from the old days when copyeditors actually marked up documents to tell the printer—in those days, a human being—which fonts to use, and so on. Markup languages thus contain explicit commands for formatting. For example, in HTML, ** means start boldface mode, and means leave boldface mode.** The advantage of a markup language over one with no explicit markup is that writing a browser for it is straightforward: the browser simply has to understand the markup commands. TeX and troff are other well-known examples of markup languages.

Documents written in a markup language can be contrasted to documents produced with a WYSIWYG (What You See Is What You Get) word processor, such as MS-Word[®] or WordPerfect[®]. These systems may store their files with hidden embedded markup so they can reproduce them later, but not all of them work this way. Word processors for the Macintosh, for example, keep the formatting information in separate data structures, not as commands embedded in the user files.

By embedding the markup commands within each HTML file and standardizing them, it becomes possible for any Web browser to read and reformat any Web page. Being able to reformat Web pages after receiving them is crucial because a page may have been produced full screen on a 1024 × 768 display with 24-bit color but may have to be displayed in a small window on a 640 × 480 screen with 8-bit color. Proprietary WYSIWYG word processors cannot be used on the Web because their internal markup languages (if any) are not standardized across vendors, machines and operating systems. Also, they do not handle reformatting for different-sized windows and different resolution displays. However, word processing program can offer the option of saving documents in HTML instead of in the vendor's proprietary format, and some of them already do.

Like HTTP, HTML is in a constant state of flux. When Mosaic was the only browser, the language it interpreted, HTML 1.0, was the de facto standard. When new browsers came along, there was a need for a formal Internet standard, so the HTML 2.0 standard was produced. HTML 3.0 was initially created as a research effort to add many new features to HTML 2.0, including tables, toolbars, mathematical formulas, advanced style sheets (for defining page layout and the meaning of symbols), and more.

The official standardization of HTML is being managed by the WWW Consortium, but various browser vendors have added their own ad hoc extensions. These vendors hope to get people to write Web pages using their extensions, so readers of these pages will need the vendor's browser to properly interpret the pages. This tendency does not make HTML standardization any easier.

Below we will give a brief introduction to HTML, just to give an idea of what it is like. While it is certainly possible to write HTML documents with any standard editor, and many people do, it is also possible to use special HTML editors that do most of the work (but correspondingly give the user less control over all the details of the final result).

A proper Web page consists of a head and a body enclosed by `<HTML>` and `</HTML>` **tags** (formatting commands), although most browsers do not complain if these tags are missing. As can be seen from Fig. 7-64(a), the head is bracketed by the `<HEAD>` and `</HEAD>` tags and the body is bracketed by the `<BODY>` and `</BODY>` tags. The commands inside the tags are called **directives**. Most HTML tags have this format, that is, `<SOMETHING>` to mark the beginning of something and `</SOMETHING>` to mark its end. Numerous other examples of HTML are easily available. Most browsers have a menu item VIEW SOURCE or something like that. Selecting this item displays the current page's HTML source, instead of its formatted output.

Tags can be in either lowercase or uppercase. Thus `<HEAD>` and `<head>` mean the same thing, but the former stands out better for human readers. Actual layout of the HTML document is irrelevant. HTML parsers ignore extra spaces and carriage returns since they have to reformat the text to make it fit the current display area. Consequently, white space can be added at will to make HTML documents more readable, something most of them are badly in need of. As another consequence, blank lines cannot be used to separate paragraphs, as they are simply ignored. An explicit tag is required.

Some tags have (named) parameters. For example

```
<IMG SRC="abc" ALT="foobar">
```


is a tag, ``, with parameter `SRC` set equal to `abc` and parameter `ALT` set equal to `foobar`. For each tag, the HTML standard gives a list of what the permitted parameters, if any, are, and what they mean. Because each parameter is named, the order in which the parameters are given is not significant.

```

<HTML> <HEAD> <TITLE> AMALGAMATED WIDGET, INC. </TITLE> </HEAD>
<BODY> <H1> Welcome to AWI's Home Page </H1>
<IMG SRC="http://www.widget.com/images/logo.gif" ALT="AWI Logo"> <BR>
We are so happy that you have chosen to visit <B> Amalgamated Widget's</B>
home page. We hope <l> you </l> will find all the information you need here.
<P>Below we have links to information about our many fine products.
You can order electronically (by WWW), by telephone, or by fax. <HR>
<H2> Product information </H2>
<UL> <LI> <A HREF="http://widget.com/products/big"> Big widgets </A>
      <LI> <A HREF="http://widget.com/products/little"> Little widgets </A>
</UL>
<H2> Telephone numbers </H2>
<UL> <LI> By telephone: 1-800-WIDGETS
      <LI> By fax: 1-415-765-4321
</UL> </BODY> </HTML>

```

(a)



(b)

Fig. 7-64. (a) The HTML for a sample Web page. (b) The formatted page.

Technically, HTML documents are written in the ISO 8859-1 Latin-1 character set, but for users whose keyboards only support ASCII, escape sequences are present for the special characters, such as è. The list of special characters is given in the standard. All of them begin with an ampersand and end with a semicolon. For example, ` produces è and ´ produces é. Since <, >, and & have special meanings, they can be expressed only with their escape sequences, < > and & respectively.

The main item in the head is the title, delimited by <TITLE> and </TITLE>, but certain kinds of meta-information may also be present. The title itself is not displayed on the page. Some browsers use it to label the page's window.

Let us now take a look at some of the other features illustrated in Fig. 7-64. All of the tags used in Fig. 7-64 and some others are shown in Fig. 7-65. Headings are generated by an <Hn> tag, where *n* is a digit in the range 1 to 6. <H1> is the most important heading; <H6> is the least important one. It is up to the browser to render these appropriately on the screen. Typically the lower numbered headings will be displayed in a larger and heavier font. The browser may also choose to use different colors for each level of heading. Typically <H1> headings are large and boldface with at least one blank line above and below. In contrast, <H2> headings are in a smaller font, and with less space above and below.

The tags and <I> are used to enter boldface and italics mode, respectively. If the browser is not capable of displaying boldface and italics, it must use some other method of rendering them, for example, using a different color for each or perhaps reverse video. Instead of specifying physical styles such as boldface and italics, authors can also use logical styles such as <DN> (define), (weak emphasis), (strong emphasis), and <VAR> (program variables). The logical styles are defined in a document called a **style sheet**. The advantage of the logical styles is that by changing one definition, all the variables can be changed, for example, from italics to a constant width font.

HTML provides various mechanisms for making lists, including nested lists. The tag starts an unordered list. The individual items, which are marked with the tag in the source, appear with bullets (•) in front of them. A variant of this mechanism is , which is for ordered lists. When this tag is used, the items are numbered by the browser. A third option is <MENU>, which typically produces a more compact list on the screen, with no bullets and no numbers. Other than the use of different starting and ending tags, , , and <MENU> have the same syntax and similar results.

In addition to the list mechanisms shown in Fig. 7-65, there are two others that are worth mentioning briefly. <DIR> can be used for making short tables. Also, <DL> and </DL> can make definition lists (glossaries) with two-part entries, whose parts are defined by <DT> and <DD> respectively. The first is for the name, the second for its meaning. These features are largely superseded by the (more general and complex) table mechanism, described below.

Tag	Description
<HTML> ... </HTML>	Declares the Web page to be written in HTML
<HEAD> ... </HEAD>	Delimits the page's head
<TITLE> ... </TITLE>	Defines the title (not displayed on the page)
<BODY> ... </BODY>	Delimits the page's body
<H <i>n</i> > ... </H <i>n</i> >	Delimits a level <i>n</i> heading
 ... 	Set ... in boldface
<I> ... </I>	Set ... in italics
 ... 	Brackets an unordered (bulleted) list
 ... 	Brackets a numbered list
<MENU> ... </MENU>	Brackets a menu of items
	Start of a list item (there is no)
 	Force a break here
<P>	Start of paragraph
<HR>	Horizontal rule
<PRE> ... </PRE>	Preformatted text; do not reformat
	Load an image here
 ... 	Defines a hyperlink

Fig. 7-65. A selection of common HTML tags. Some have additional parameters.

The
, <P>, and <HR> tags all indicate a boundary between sections of text. The precise format can be determined by the style sheet associated with the page. The
 tag just forces a line break. Typically, browsers do not insert a blank line after
. In contrast, <P> starts a paragraph, which might, for example, insert a blank line and possibly some indentation. (Theoretically, </P> exists to mark the end of a paragraph, but it is rarely used; most HTML authors do not even know it exists.) Finally, <HR> forces a break and draws a horizontal line across the screen.

HTML 1.0 had no ability to display tables or other formatted information. Worse yet, if the HTML writer carefully formatted a table by judicious use of spaces and carriage returns, browsers would ignore all the layout and display the page as if all the formatted material were unformatted. To prevent browsers from messing up carefully laid out text, the <PRE> and </PRE> tags were provided. They are instructions to the browser to just display everything in between literally, character for character, without changing anything. As the table and other fancy layout features become more widely implemented, the need for <PRE> will

diminish, except for program listings, for which most programmers will tolerate no formatting other than their own.

HTML allows images to be included in-line on a Web page. The `` tag specifies that an image is to be loaded at the current position in the page. It can have several parameters. The `SRC` parameter gives the URL (or URI) of the image. The HTML standard does not specify which graphic formats are permitted. In practice, all browsers support GIF files and many support JPEG files as well. Browsers are free to support other formats, but this extension is a two-edged sword. If a user is accustomed to a browser that supports, say, BMP files, he may include these in his Web pages and later be surprised when other browsers just ignore all of his wonderful art.

Other parameters of `` are `ALIGN`, which controls the alignment of the image with respect to the text baseline (`TOP`, `MIDDLE`, `BOTTOM`), `ALT`, which provides text to use instead of the image when the user has disabled images, and `ISMAP`, a flag indicating that the image is an active map.

Finally, we come to hyperlinks, which use the `<A>` (anchor) and `` tags. Like ``, `<A>` has various parameters, including `HREF` (the URL), `NAME` (the hyperlink's name), and `METHODS` (access methods), among others. The text between the `<A>` and `` is displayed. If it is selected, the hyperlink is followed to a new page. It is also permitted to put an `` image there, in which case clicking on the image also activates the hyperlink.

As an example, consider the following HTML fragment:

```
<A HREF="http://www.nasa.gov"> NASA's home page </A>
```

When a page with this fragment is displayed, what appears on the screen is

NASA's home page

If the user subsequently clicks on this text, the browser immediately fetches the page whose URL is `http://www.nasa.gov` and displays it.

As a second example, now consider

```
<A HREF="http://www.nasa.gov"> <IMG SRC="shuttle.gif" ALT="NASA"> </A>
```

When displayed, this page shows a picture (e.g., of the space shuttle). Clicking on the picture switches to NASA's home page, just as clicking on the underlined text did in the previous example. If the user has disabled automatic image display, the text NASA will be displayed where the picture belongs.

The `<A>` tag can take a parameter `NAME` to plant a hyperlink, so it can be referred to from within the page. For example, some Web pages start out with a clickable table of contents. By clicking on an item in the table of contents, the user jumps to the corresponding section of the page.

One feature that HTML 2.0 did not include and which many page authors missed, was the ability to create tables whose entries could be clicked on to active hyperlinks. As a consequence, a large amount of work was done to add tables to

HTML 3.0. Below we give a very brief introduction to tables, just to capture the essential flavor.

An HTML table consists of one or more rows, each consisting of one or more **cells**. Cells can contain a wide range of material, including text, figures, and even other tables. Cells can be merged, so, for example, a heading can span multiple columns. Page authors have limited control over the layout, including alignment, border styles, and cell margins, but the browsers have the final say in rendering tables.

An HTML table definition is listed in Fig. 7-66(a) and a possible rendition is shown in Fig. 7-66(b). This example just shows a few of the basic features of HTML tables. Tables are started by the `<TABLE>` tag. Additional information can be provided to describe general properties of the table.

The `<CAPTION>` tag can be used to provide a figure caption. Each row is started with a `<TR>` (Table Row) tag. The individual cells are marked as `<TH>` (Table Header) or `<TD>` (Table Data). The distinction is made to allow browsers to use different renditions for them, as we have done in the example.

Numerous other tags are also allowed in tables. They include ways to specify horizontal and vertical cell alignments, justification within a cell, borders, grouping of cells, units, and more.

Forms

HTML 1.0 was basically one way. Users could call up pages from information providers, but it was difficult to send information back the other way. As more and more commercial organizations began using the Web, there was a large demand for two-way traffic. For example, many companies wanted to be able to take orders for products via their Web pages, software vendors wanted to distribute software via the Web and have customers fill out their registration cards electronically, and companies offering Web searching wanted to have their customers be able to type in search keywords.

These demands led to the inclusion of **forms** starting in HTML 2.0. Forms contain boxes or buttons that allow users to fill in information or make choices and then send the information back to the page's owner. They use the `<INPUT>` tag for this purpose. It has a variety of parameters for determining the size, nature, and usage of the box displayed. The most common forms are blank fields for accepting user text, boxes that can be checked, active maps, and SUBMIT buttons. The example of Fig. 7-67 illustrates some of these choices.

Let us start our discussion of forms by going over this example. Like all forms, this one is enclosed between the `<FORM>` and `</FORM>` tags. Text not enclosed in a tag is just displayed. All the usual tags (e.g., ``) are allowed in a form. Three kinds of input boxes are used in this form.

The first kind of input box follows the text "Name". The box is 46 characters wide and expects the user to type in a string, which is then stored in the variable

```

<HTML> <HEAD> <TITLE> A sample page with a table </TITLE> </HEAD>
<BODY>
<TABLE BORDER=ALL RULES=ALL>
<CAPTION> Some Differences between HTML Versions </CAPTION>
<COL ALIGN=LEFT>
<COL ALIGN=CENTER>
<COL ALIGN=CENTER>
<COL ALIGN=CENTER>
<TR> <TH>Item <TH>HTML 1.0 <TH>HTML 2.0 <TH>HTML 3.0
<TR> <TH> Active Maps and Images <TD> <TD> x <TD> x
<TR> <TH> Equations <TD> <TD> x <TD> x
<TR> <TH> Forms <TD> <TD> x <TD> x
<TR> <TH> Hyperlinks x <TD> <TD> x <TD> x
<TR> <TH> Images <TD> x <TD> x <TD> x
<TR> <TH> Lists <TD> x <TD> x <TD> x
<TR> <TH> Toolbars <TD> <TD> <TD> x
<TR> <TH> Tables <TD> <TD> <TD> x
</TABLE> </BODY> </HTML>

```

(a)

Some Differences between HTML Versions

Item	HTML 1.0	HTML 2.0	HTML 3.0
Active Maps and Images		x	x
Equations			x
Forms		x	x
Hyperlinks	x	x	x
Images	x	x	x
Lists	x	x	x
Toolbars			x
Tables			x

Fig. 7-66. (a) An HTML table. (b) A possible rendition of this table.

customer for later processing. The <P> tag instructs the browser to display subsequent text and boxes on the next line, even if there is room on the current line. By using <P> and other layout tags, the author of the page can control the look of the form on the screen.

The next line of the form asks for the user's street address, 40 columns wide, also on a line by itself. Then comes a line asking for the city, state, and country. No <P> tags are used between the fields here, so the browser displays them all on one line if they will fit. As far as the browser is concerned, this paragraph just contains six items: three strings alternating with three boxes. It displays them linearly from left to right, going over to a new line whenever the current line


```

<HTML> <HEAD> <TITLE> AWI CUSTOMER ORDERING FORM </TITLE> </HEAD>
<BODY>
<H1> Widget Order Form </H1>
<FORM ACTION="http://widget.com/cgi-bin/widgetorder" METHOD=POST>
Name <INPUT NAME="customer" SIZE=46> <P>
Street Address <INPUT NAME="address" SIZE=40> <P>
City <INPUT NAME="city" SIZE=20> State <INPUT NAME="state" SIZE =4>
Country <INPUT NAME="country" SIZE=10> <P>
Credit card # <INPUT NAME="cardno" SIZE=10>
Expires <INPUT NAME="expires" SIZE=4>
M/C <INPUT NAME="cc" TYPE=RADIO VALUE="mastercard">
VISA <INPUT NAME="cc" TYPE=RADIO VALUE="visacard"> <P>
Widget size Big <INPUT NAME="product" TYPE=RADIO VALUE="expensive">
Little <INPUT NAME="product" TYPE=RADIO VALUE="cheap">
Ship by express courier <INPUT NAME="express" TYPE=CHECKBOX> <P>
<INPUT TYPE=SUBMIT VALUE="Submit order"> <P>
Thank you for ordering an AWI widget, the best widget money can buy!
</FORM> </BODY> </HTML>

```

(a)

Widget Order Form

Name

Street address

City State Country

Credit card # Expires M/C Visa

Widget size Big Little Ship by express courier

Thank you for ordering an AWI widget, the best widget money can buy!

(b)

Fig. 7-67. (a) The HTML for an order form. (b) The formatted page.

cannot hold the next item. Thus it is conceivable that on a 1024×768 screen all three strings and their corresponding boxes will appear on the same line, but on a 640×480 screen they might be split over two lines. In the worst scenario, the word "Country" is at the end of one line and its box is at the beginning of the next line. There is no way to tell the browser to force the box adjacent to the text.

The next line asks for the credit card number and expiration date. Transmitting credit card numbers over the Internet should only be done when adequate security measures have been taken. For example, some, but not all, browsers encrypt information sent by users. Even then, secure communication and key management are complicated matters and are subject to many kinds of attacks, as we saw earlier.

Following the expiration date we encounter a new feature: radio buttons. These are used when a choice must be made among two or more alternatives. The intellectual model here is a car radio with half a dozen buttons for choosing stations. The browser displays these boxes in a form that allows the user to select and deselect them by clicking on them (or using the keyboard). Clicking on one of them turns off all the other ones in the same group. The visual presentation depends on the graphical interface being used. It is up to the browser to choose a form that is consistent with Windows, Motif, OS/2 Warp, or whatever windowing system is being used. The widget size also uses two radio buttons. The two groups are distinguished by their *NAME* field, not by static scoping using something like `<RADIOBUTTON> ... </RADIOBUTTON>`.

The *VALUE* parameters are used to indicate which radio button was pushed. Depending on which of the credit card options the user has chosen, the variable *cc* will be set to either the string "mastercard" or the string "visacard".

After the two sets of radio buttons, we come to the shipping option, represented by a box of type *CHECKBOX*. It can be either on or off. Unlike radio buttons, where exactly one out of the set must be chosen, each box of type *CHECKBOX* can be on or off, independently of all the others. For example, when ordering a pizza via Electropizza's Web page, the user can choose sardines *and* onions *and* pineapple (if she can stand it), but she cannot choose small *and* medium *and* large for the same pizza. The pizza toppings would be represented by three separate boxes of type *CHECKBOX*, whereas the pizza size would be a set of radio buttons.

As an aside, for very long lists from which a choice must be made, radio buttons are somewhat inconvenient. Therefore, the `<SELECT>` and `</SELECT>` tags are provided to bracket a list of alternatives, but with the semantics of radio buttons (unless the *MULTIPLE* parameter is given, in which case the semantics are those of checkable boxes). Some browsers render the items between `<SELECT>` and `</SELECT>` as a pop-up menu.

We have now seen two of the built-in types for the `<INPUT>` tag: *RADIO* and *CHECKBOX*. In fact, we have already seen a third one as well: *TEXT*. Because this type is the default, we did not bother to include the parameter *TYPE = TEXT*, but we could have. Two other types are *PASSWORD* and *TEXTAREA*. A *PASSWORD* box is the same as a *TEXT* box, except that the characters are not displayed as they are typed. A *TEXTAREA* box is also the same as a *TEXT* box, except that it can contain multiple lines.

Getting back to the example of Fig. 7-67, we now come across an example of

a *SUBMIT* button. When this is clicked, the user information on the form is sent back to the machine that provided the form. Like all the other types, *SUBMIT* is a reserved word that the browser understands. The *VALUE* string here is the label on the button and is displayed. All boxes can have values; we only needed that feature here. For *TEXT* boxes, the contents of the *VALUE* field are displayed along with the form, but the user can edit or erase it. *CHECKBOX* and *RADIO* boxes can also be initialized, but with a field called *CHECKED* (because *VALUE* just gives the text, but does not indicate a preferred choice).

The browser also understands the *RESET* button. When clicked, it resets the form to its initial state.

Two more types are worth noting. The first is the *HIDDEN* type. This is output only; it cannot be clicked or modified. For example, when working through a series of pages throughout which choices have to be made, previously made choices might be of *HIDDEN* type, to prevent them from being changed.

Our last type is *IMAGE*, which is for active maps (and other clickable images). When the user clicks on the map, the (x, y) coordinates of the pixel selected (i.e., the current mouse position) are stored in variables and the form is automatically returned to the owner for further processing.

Forms can be submitted in three ways: using the submit button, clicking on an active map, or typing ENTER on a one-item *TEXT* form. When a form is submitted, some action must be taken. The action is specified by the parameters of the <FORM> tag. The *ACTION* parameter specifies the URL (or URI) to tell about the submission, and the *METHOD* parameter tells which method to use. The order of these (and all other) parameters is not significant.

The way the form's variables are sent back to the page's owner depends on the value of the *METHOD* parameter. For *GET*, the only way to return values is to cheat: they are appended to the URL, separated by a question mark. This approach can result in URLs that are thousands of characters long. Nevertheless, this method is frequently used because it is simple.

If the *POST* method (see Fig. 7-62) is used, the body of the message contains the form's variables and their values. The & is used to separate fields; the + represents the space character. For example, the response to the widget form might be

```
customer=John+Doe&address=100+Main+St.&city=White+Plains&
state=NY&country=USA&cardno=1234567890&expires=6/98&cc=mastercard&
product=cheap&express=on
```

The string would be sent back to the server as one line, not three. If a *CHECKBOX* is not selected, it is omitted from the string. It is up to the server to make sense of this string.

Fortunately, a standard for handling forms' data is already available. It is called **CGI (Common Gateway Interface)**. Let us consider a common way of

using it. Suppose that someone has an interesting database (e.g., an index of Web pages by keyword and topic) and wants to make it available to Web users. The CGI way to make the database available is to write a script (or program) that interfaces (i.e., gateways) between the database and the Web. This script is given a URL, by convention in the directory *cgi-bin*. HTTP servers know (or can be told) that when they have to invoke a method on a page located in *cgi-bin*, they are to interpret the file name as being an executable script or program and start it up.

Eventually, some user opens the form associated with our widget script and has it displayed. After the form has been filled out, the user clicks on the SUBMIT button. This action causes the browser to establish a TCP connection to the URL listed in the form's *ACTION* parameter—the script in the *cgi-bin* directory. Then the browser invokes the operation specified by the form's *METHOD*, usually *POST*. The result of this operation is that the script is started and presented (via the TCP connection, on standard input) with the long string given above. In addition, several environment variables are set. For example, the environment variable *CONTENT_LENGTH* tells how long the input string is.

At this point, most scripts need to parse their input to put it in a more convenient form. This goal can be accomplished by calling one of the many libraries or script procedures available. The script can then interact with its database in any way it wishes. For example, active maps normally use CGI scripts to take different actions depending on where the user pointed.

CGI scripts can also produce output and do many other things as well as accepting input from forms. If a hyperlink points to a CGI script, when that link is invoked, the script is started up, with several environment variables set to provide some information about the user. The script then writes a file (e.g. an HTML page) on standard output, which is shipped to the browser and interpreted there. This mechanism makes it possible for the script to generate custom Web pages on the spot.

For better or worse, some Web sites that answer queries have a database of advertisements that can be selectively included in the Web page being constructed, depending on what the user is looking for. If the user is searching for “car” a General Motors ad might be displayed, whereas a search for “vacation” might produce an ad from United Airlines. These ads usually include clickable text and pictures.

7.6.4. Java

HTML makes it possible to describe how static Web pages should appear, including tables and pictures. With the *cgi-bin* hack, it is also possible to have a limited amount of two-way interaction (forms, etc.). However, rapid interaction with Web pages written in HTML is not possible. To make it possible to have

highly interactive Web pages, a different mechanism is needed. In this section we will describe one such mechanism, the JavaTM language and interpreter.

Java originated when some people at Sun Microsystems were trying to develop a new language that was suitable for programming information-oriented consumer appliances. Later it was reoriented toward the World Wide Web. Although Java borrows many ideas and some syntax from C and C++, it is a new object-oriented language, compatible with neither. It is sometimes said that in the large, Java is like Smalltalk, but that in the small it is like C or C++.

The main idea of using Java for interactive Web pages is that a Web page can point to a small Java program, called an **applet** (SAT I verbal analogy question: Pig is to piglet as application is to ?). When the browser reaches it, the applet is downloaded to the client machine and executed there in a secure way. It must be structurally impossible for the applet to read or write any files that it is not authorized to access. It must also be impossible for the applet to introduce viruses or cause any other damage. For these reasons, and to achieve portability across machines, applets are compiled to a bytecode after being written and debugged. It is these bytecode programs that are pointed to by Web pages, similar to the way images are pointed to. When an applet arrives, it is executed interpretively in a secure environment.

Before getting into the details of the Java language, it is worth saying a few words about what the whole Java system is good for and why people want to include Java applets in their Web pages. For one thing, applets allow Web pages to become interactive. For example, a web page can contain a board for tic tac toe, othello, or chess, and play a game with the user. The game-playing program (written in Java) is just downloaded along with its Web page. As a second example, complex forms (e.g., spreadsheets) can be displayed, with the user filling in items and seeing calculations made instantly.

It is entirely possible that in the long run, the model of people buying programs, installing them, and running them locally will be replaced by a model in which people click on Web pages, get applets downloaded to do work for them, possibly in conjunction with a remote server or data base. Instead of filling out the income tax form by hand or using a special program, people may be able to click on the IRS home page to get a tax applet downloaded. This applet might ask some questions, then contact the person's employer, bank, and stockbroker to collect the required salary, interest, and dividend information, fill the tax form in, and then display it for verification and submission.

Another reason for running applets on the client machine is they make it possible to add animation and sound to Web pages without having to spawn external viewers. The sound can be played when the page is loaded, as background music, or when some specific event happens (e.g., clicking on the cat makes it meow). The same is true for animation. Because the applet is running locally, even if it is being interpreted, it can write all over (its portion) of the screen any way it wants to, and at very high speed (compared to a remote cgi-bin shell script).

The Java system has three parts:

1. A Java-to-bytecode compiler.
2. A browser that understands applets.
3. A bytecode interpreter.

The developer writes the applet in Java, then compiles it to bytecode. To include this compiled applet on a Web page, a new HTML tag, `<APPLET>`, has been invented. A typical use is

```
<APPLET CODE=game.class WIDTH=100 HEIGHT=200> </APPLET>
```

When the browser sees the `<APPLET>` tag, it fetches the compiled applet *game.class* from the current Web page's site (or if another parameter, *CODE-BASE*, is present, from the URL it specifies). The browser then passes the applet to the local bytecode interpreter for execution (or interprets the applet itself, if it has an internal interpreter). The *WIDTH* and *HEIGHT* parameters give the size of the applet's default window, in pixels.

In a sense, the `<APPLET>` tag is analogous to the `` tag. In both cases, the browser goes and gets a file and then hands it off to a (possibly internal) interpreter for display within a bounded area of the screen. Then it continues processing the Web page.

For applications that need very high performance, some Java interpreters have the ability to compile bytecode programs to actual machine language on-the-fly, as needed.

As a consequence of this model, Java-based browsers are extensible in a way that first-generation browsers are not. First generation browsers are basically HTML interpreters that have built-in modules for speaking the various protocols needed, such as HTTP 1.0, FTP, etc., as well as decoders for various image formats. An example is shown in Fig. 7-68(a). If someone invents or popularizes a new format, such as audio or MPEG-2, these old browsers are not able to read pages containing them. At best, the user has to find, download, and install an appropriate external viewer.

With a Java-based browser, the situation is different. At startup, the browser is effectively an empty Java virtual machine, as shown in Fig. 7-68(b). By loading HTML and HTTP applets, it becomes able to read standard Web pages. However, as new protocols and decoders are required, their classes are loaded dynamically, possibly over the network from sites specified in Web pages. After a while, the browser might look like Fig. 7-68(c).

Thus if someone invents a new format, all that person has to do is include the URL of an applet for handling it in a Web page, and the browser will automatically fetch and load the applet. No first-generation browser is capable of automatically downloading and installing new external viewers on-the-fly. The ability to




Fig. 7-68. (a) A first generation browser. (b) A Java-based browser at startup. (c) The browser of (b) after running for a while.

load applets dynamically means that people can easily experiment with new formats without first having to have endless standardization meetings to reach a consensus.

This extensibility also applies to protocols. For some applications, special protocols are needed, for example, secure protocols for banking and commerce. With Java, these protocols can be loaded dynamically as needed, and there is no need to achieve universal standardization. To communicate with company X, you just download its protocol applet. To talk to company Y, you get its protocol applet. There is no need for X and Y to agree on a standard protocol.

Introduction to the Java Language

The objectives listed above have led to a type-safe, object-oriented language with built-in multithreading and no undefined or system dependent features. What follows is a highly simplified description of Java, just to give a feel for it. Many features, details, options, and special cases have been omitted for the sake of brevity. The complete language specification, and much more about Java, is available on the Web itself (naturally) at <http://java.sun.com>. For tutorials on Java, see (Campione and Walrath, 1996; and Van der Linden, 1996). For the full story, see (Arnold and Gosling, 1996; and Gosling et al., 1996). For a brief comparison between Java and Microsoft's answer to it (Blackbird), see (Singleton, 1996).

As we mentioned above, in the small, Java is similar to C and C++. The lexical rules, for example, are pretty much the same (e.g., tokens are delimited by white space, and new lines can be inserted between any two tokens). Comments can be inserted using either the C syntax (*/* ... */*) or the C++ syntax (*// ...*).

Java has eight primitive data types, as listed in Fig. 7-69. Each type has a specific size, independent of the local implementation. Thus unlike C, where an integer may be 16, 32, or 64 bits, depending on the underlying machine

architecture, a Java `int` is always 32 bits, no more and no less, no matter what kind of machine the interpreter is running on. This consistency is essential since the same applet must run on 16-bit, 32-bit, and 64-bit machines, and give the same result on all of them.

Type	Size	Description
Byte	1 Byte	A signed integer between -128 and +127
Short	2 Bytes	A signed 2-byte integer
Int	4 Bytes	A signed 4-byte integer
Long	8 Bytes	A signed 8-byte integer
Float	4 Bytes	A 4-byte IEEE floating-point number
Double	8 Bytes	An 8-byte IEEE floating-point number
Boolean	1 Bit	The only values are true and false
Char	2 Bytes	A character in Unicode

Fig. 7-69. The basic Java data types.

Arithmetic variables (the first 6 types) can be combined using the usual arithmetic operators (including `++` and `--`) and compared using the usual relational operators (e.g., `<`, `<=`, `==`, `!=`). Conversions between types are permitted where they make sense.

Java uses the 16-bit Unicode instead of ASCII for characters, so character variables are 2 bytes long. The first 127 Unicode characters are the same as ASCII for backward compatibility. Above these are some graphic symbols, and then the characters needed for Russian, Arabic, Hebrew, Japanese (kanji, katakana, and hiragana), and virtually every other language. Characters not present in ASCII can be represented with `\u` followed by four hexadecimal digits. For example, `\u0a66` is the Gujarati zero.

Java allows one dimensional arrays to be declared. For example,

```
int[] table;
```

declares an array, *table*, but does not allocate any space for it. That can be done later on, as in C++, for example, by

```
table = new int [1024];
```

to allocate an array with 1024 entries. It is not necessary (or even possible) to return arrays that are no longer needed; the garbage collector reclaims them. Thus the highly error-prone *malloc* and *free* library routines are not needed for storage management. Arrays can be initialized, and arrays of arrays can be used to get higher dimensionality, as in C. Strings are available, but they are defined in a class, rather than being simply character arrays ending with a null byte.

The Java control statements are shown in Fig. 7-70. The first nine have essentially the same syntax and semantics as in C, except that where a Boolean expression is required, the language actually insists upon a Boolean expression. Also, the `break` and `continue` statements now can take labels indicating which of the labeled loops to exit or repeat.

Statement	Description	Example
Assignment	Assign a value	<code>n = i + j;</code>
If	Boolean choice	<code>if (k < 0) k = 0; else k = 2*k;</code>
Switch	Select a case	<code>switch (b) {case 1: n++; case 2: n--;}</code>
For	Iteration	<code>for (i = 0; i < n; i++) a[i] = b[i];</code>
While	Repetition	<code>while (n < k) n += i;</code>
Do	Repetition	<code>do {n = n + n} while (n < m);</code>
Break	Exit statement	<code>break label;</code>
Return	Return	<code>return n;</code>
Continue	Next iteration	<code>continue label;</code>
Throw	Raise exception	<code>throw new IllegalArgumentException();</code>
Try	Exception scoping	<code>try { ... } catch (Exception e) {return -1};</code>
Synchronized	Mutual exclusion	<code>synchronized void update(int s) { ... }</code>

Fig. 7-70. The Java statements. The notation { ... } indicates a block of code.

The next two statements are in C++ but not in C. The `throw` and `try` statements deal with exception handling. Java defines a variety of standard exceptions, such as attempting to divide by zero, and allows programmers to define and raise their own exceptions. Programmers can write handlers to catch exceptions, making it unnecessary to constantly test if something has gone wrong (e.g., when reading from a file). The `throw` statement raises an exception, and the `try` statement defines a scope to associate exception handlers with a block of code in which an exception might occur.

The `synchronized` statement is new to Java and has to do with the fact that Java programs can have multiple threads of control. To avoid race conditions, this statement is used to delimit a block of code (or a whole procedure) that must not have more than one thread active in it at once. Such blocks of code are usually called **critical regions**. When the `synchronized` statement is executed, the thread executing it must acquire the lock associated with the critical region, execute the code, and then release the lock. If the lock is not available, the thread waits until it is free. By guarding entire procedures this way and using condition variables, programmers have the full power of monitors (Hoare, 1974).

Java programs can be called with arguments. Command-line processing is similar to C, except that the argument array is called *args* instead of *argv* and *args[0]* is the first parameter, not the program name. Figure 7-71 illustrates a small Java program that computes a table of factorials, just to give an idea of what a small Java program looks like.

```
class Factorial { /* This program consists of a single class with two methods. */

    public static void main (int argc, String args[]) { // main program
        long i, f, lower = 1, upper = 20;           // declarations of four longs

        for (i = lower; i <= upper; i++) {         // loop from lower to upper
            f = factorial(i);                       // f = i!
            System.out.println(i + " " + f);        // print i and f
        }
    }

    static long factorial (long k) {               // recursive factorial function
        if (k == 0)
            return 1;                               // 0! = 1
        else
            return k * factorial(k-1);              // k! = k * (k-1)!
    }
}
```

Fig. 7-71. A Java program for computing and printing 0! to 20!.

Despite both being object-oriented languages based on C, Java and C++ differ in some ways. Some features were removed from Java to make it typesafe or easier to read. These include `#define`, `typedef`, `enums`, `unions`, `structs`, `operator overloading`, `explicit pointers`, `global variables`, `standalone functions`, and `friend functions`. It almost goes without saying that the `goto` statement has been sent to that special place reserved for obsolete programming language features. Other features were added to give the language more power. The features added include `garbage collection`, `multithreading`, `object interfaces`, and `packages`.

Object Orientation in Java

In traditional procedural languages such as Pascal or C, a program consists of a collection of variables and procedures, without any general organizing principle. In contrast in **object-oriented languages**, (almost) everything is an object. An **object** normally contains some internal (i.e., hidden) state variables along with some public procedures, called **methods**, for accessing them. Programs that use the object are expected (and can be forced) to invoke the methods to manipulate the object's state. In this way, the object writer can control how programs use the

information inside the object. This principle is called **encapsulation**, and is the basis of all object-oriented programming.

Java tries to capture the best of both worlds. It can be used as a traditional procedural language or as an object-oriented language. The Java example of Fig. 7-71, for example, could equally well have been written in C, and in essentially the same way. In this view, a subset of Java can be regarded as a cleaned-up version of C. However, for writing Web pages, Java is better regarded as an object-oriented language, so we will study its object orientation in this section.

A Java program consists of one or more **packages**, each of which contains some class definitions. Packages can be accessed remotely over a network, so those intended for use by a wide audience must have unique names. Normally, hierarchical names are used, starting with the reverse of their machine's DNS name, for example

```
EDU.univ.cs.catie.games.chess
```

A **class** definition is a template for stamping out object instances, each of which contains the same state variables and same methods as all the other object instances of its class. The values of the state variables within different objects are independent, however. Classes are thus like cookie cutters: they are not cookies themselves, but are used to stamp out structurally identical cookies, with each cookie cutter producing a different shape of cookie. Once produced, different cookies (objects) are independent of one another.

Java objects can be produced dynamically during execution, for example by

```
object = new ClassName()
```

These objects are stored on the heap and removed by the garbage collector when no longer needed. In this way, storage management in Java is handled by the system, with no need for the dreaded *malloc* and *free* procedures, or even for explicit pointers, for that matter.

Each class is based on another class. A newly defined class is said to be a **subclass** of the class on which it is based, the **superclass**. A (sub)class always **inherits** the methods of its superclass. It may or may not have direct access to the superclass' internal variables, depending on whether or not the superclass wants that. For example, if a superclass, *A*, has methods *M1*, *M2*, and *M3*, and a subclass, *B*, defines a new method, *M4*, then objects created from *B*, will have methods *M1*, *M2*, *M3*, and *M4*. The property of a class automatically acquiring all the methods of its superclass is called **inheritance**, and is an important property of Java. Adding new methods to the superclass' methods is called **extending** the superclass. As an aside, some object-oriented languages allow classes to inherit methods from two or more superclasses (multiple inheritance), but the Java designers thought this property to be too messy and intentionally left it out.

Since every class has exactly one immediate superclass, the set of all classes in a Java program form a tree. The class at the top of the tree is called **Object**.

All other classes inherit its methods. Any class whose superclass is not explicitly mentioned in its definition defaults to being a subclass of the *Object* class. The *Factorial* class of Fig. 7-71, for example, is thus a subclass of *Object*.

Let us now take a look at an example of the object-oriented concepts presented so far. In Fig. 7-72 we have a package defining two classes, *ComplexNumber*, for defining and using complex numbers (i.e., numbers with a real part and an imaginary part), and *test*, for showing how *ComplexNumber* can be used.

```
class ComplexNumber {           // Define a subclass of Object called ComplexNumber
    // Hidden data.
    protected double re, im;    // real and imaginary parts

    // Five methods that manage the hidden data.
    public void Complex(double x, double y) {re = x; im = y;}
    public double Real() {return re;}
    public double Imaginary() {return im;}
    public double Magnitude() {return Math.sqrt(re*re + im*im);}
    public double Angle() {return Math.atan(im/re);}
}

class test {                   // A second class, for testing ComplexNumber
    public static void main (String args[]) {
        ComplexNumber c;      // declare an object of class ComplexNumber

        c = new ComplexNumber(); // actually allocate storage for c
        c.Complex(3.0, 4.0);    // invoke the Complex method to initialize c
        System.out.println("The magnitude of c is " + c.Magnitude() );
    }
}
```

Fig. 7-72. A package defining two classes.

Like *Factorial*, the class *ComplexNumber* is based on *Object*, because no other superclass is named in its definition. Each object of class *ComplexNumber* represents one complex number. Each object of this class contains two hidden variables, *re*, and *im*, both 64-bit floating-point numbers, for representing the real and imaginary parts, respectively. They cannot be accessed outside the class definition (and its subclasses), because they have been declared *protected*. Had they been declared *private*, then they would have been visible only to *ComplexNumber* and not to any subclasses. For the moment, *private* would have been fine, but we will soon define a subclass. Had they been declared *public*, they would have been visible everywhere the package was visible, thus destroying much of the value of object-oriented programming. Nevertheless, situations do exist in which having the internal state of an object be public is sometimes needed.

Five methods are defined on objects belonging to class *ComplexNumber*. Users of the class are thus restricted to the operations provided by these five methods, and cannot get at the state directly. An example of how objects of class *ComplexNumber* are created, initialized, and used is given in *test*.

When this package is compiled, the Java compiler produces two binary (bytecode) files, one containing each of the classes and named after its class. Typing the command

```
java test
```

results in invoking the Java interpreter with class *test* as parameter. The interpreter then looks for a method called *main*, and upon finding it, executes it. The result of execution is that the line

```
The magnitude of c is 5
```

is printed out.

Now let us define a subclass of *ComplexNumber*, just to see how that works. It starts out by importing the original class, to learn what it does and what its methods are. Then it defines an extension of *ComplexNumber*, which we will call *HairyNumber*. The new class automatically inherits the five methods present in the superclass. To make things interesting, we will define a sixth method, *AddTo*, in the subclass, which adds a complex number to the object, increasing its real and imaginary parts.

The subclass definition is shown in Fig. 7-73, along with another test program showing how an object belonging to class *HairyNumber* can be used. When the new test program is run, it will print out

```
h = (-0.5,6)
```

Remember that the six methods are usable on the objects *a* and *h*, without regard to which method was defined where. If we now define yet another subclass based on *HairyNumber* and give it, say, three new methods, objects produced from it will have nine valid methods.

In addition to adding new methods to its superclass, a subclass can override (replace) existing methods by simply redefining them. Thus it is possible for a subclass to redefine all the methods inherited from its superclass, so objects belonging to the two classes have nothing in common. Doing so, however, is in poor taste, and should be avoided.

Finally, a Java class may define multiple methods with the same name but different parameters and different definitions. When the compiler sees a method invocation using this name, it has to use the parameter types to determine which method to use. This property is called **overloading** or **polymorphism**. Unlike C++, where operators can also be overloaded, in Java, only methods, not operators, can be overloaded, to make programs easier to understand.

```

import ComplexNumber;           // import the ComplexNumber package

class HairyNumber extends ComplexNumber { // define a new class
    public void AddTo(ComplexNumber z) { // with one method
        re = re + z.Real();
        im = im + z.Imaginary();
    }
}

class test2 { // test program for HairyNumber
    public static void main(String args[]) {
        HairyNumber a, h; // declare two HairyNumbers

        a = new HairyNumber(); // allocate storage for a
        h = new HairyNumber(); // allocate storage for h
        a.Complex(1.0, 2.0); // assign a value to a
        h.Complex(-1.5, 4.0); // assign a value to h
        h.AddTo(a); // invoke the AddTo method on h
        System.out.println("h = (" + h.Real() + ", " + h.Imaginary() + ")");
    }
}

```

Fig. 7-73. A subclass of *ComplexNumber* defining a new method.

The Application Programmers Interface

In addition to the bare language, the Java designers have defined and implemented about 200 classes with the initial release. The methods contained in these classes form a kind of standard environment for Java program developers. The classes are written in Java, so they are portable to all platforms and operating systems.

While a detailed discussion of all these classes and methods is clearly beyond the scope of this book, a brief description may be of some interest. The 200 classes are grouped into seven packages of uneven size, each of which is focused on some central theme. Applets that need a particular package can include it using the Java import statement. The methods contained within can just be used as needed. This mechanism replaces the need for including header files in C. It also replaces the need for libraries, since the packages are dynamically loaded during execution when they are invoked.

The seven packages are summarized in Fig. 7-74. The *java.lang* package contains classes that can be viewed as part of the language, but are technically not. These include classes for managing the classes themselves, threads, and exception handling. The standard mathematical and string libraries are also here.

Package	Example functionality
Java.lang	Classes, threads, exceptions, math, strings
Java.io	I/O on streams and random access files, printing
Java.net	Sockets, IP addresses, URLs, datagrams
Java.util	Stacks, hash tables, vectors, time, date
Java.applet	Getting and displaying Web pages, audio, Object class
Java.awt	Events, dialog, menus, fonts, graphics, window management
Java.awt.image	Colors, image cropping, filtering, and conversion
Java.awt.peer	Access to the underlying window system

Fig. 7-74. The packages included in the standard API.

Like C, the Java language contains no I/O primitives. I/O is done by loading and using the *java.io* package. It is analogous to the standard I/O library in C. Methods are provided for reading and writing streams, random access files, and doing the formatting needed for printing. In Fig. 7-71 we saw one of these methods, *println*, which does formatted printing.

Closely related to I/O is network transport. Methods that look up and manage IP addresses are located here. Access to sockets is also part of this package. So is datagram preparation. The actual transmission is handled in *java.io*.

The next class is *java.util*. It contains classes and methods for common data structures, such as stacks and hash tables, so programmers do not constantly have to reinvent the wheel. Time and date management is also here.

The *java.applet* package contains some of the basic machinery for applets, including methods for getting Web pages starting from their URLs. It also has methods for displaying Web pages and playing audio clips (e.g., background music). The *java.applet* package also contains the *Object* class. All objects inherit its methods, unless they are overridden. These methods include cloning an object, comparing two objects for equality, converting an object to a string, and various others.

Finally, we come to *java.awt* and its two subpackages. AWT stands for **Abstract Window Toolkit**, and is designed to make applets portable across window systems. For example, how should an applet draw a rectangle on the screen in such a way that the same compiled (bytecode) version of the applet can run on UNIX, Windows, and the Macintosh, even though each one has its own window system? Part of the package deals with drawing on the screen, so there are methods for placing lines, geometric figures, text, menus, buttons, scroll bars, and many other items on the screen. Java programmers call these methods to write on the screen. It is up to the *java.awt* package to make the appropriate calls to the local operating system to get the job done. This strategy means that *java.awt* has

to be rewritten for each new platform, but that applets are then platform independent, which is far more important.

Another important task of this class is event management. Most window systems are fundamentally event driven. What this means is that the operating system detects keystrokes, mouse motion, button pushes and releases, and other events, and converts these into calls to user procedures. In the case of Java, a large library of methods for dealing with these events is provided in *java.awt*. Using them makes it easier to write programs that interact with the local window system and still be 100 percent portable to machines with different operating systems and different window systems.

Some of the work of this package is done in *java.awt.image*, such as image management, and in *java.awt.peer*, which allows access to the underlying window system.

Security

One of the most important aspects of Java is its security properties. When a Web page containing an applet is fetched, the applet is automatically executed on the client's machine. Ideally, it should not crash or otherwise bring down the client's machine.

Furthermore, it does not take much imagination to envision some enterprising undergraduate producing a Web page containing some nifty new game, then publicizing its URL widely (e.g., crossposting it to every newsgroup). Not mentioned in the posting is the small detail that the page also contains an applet that upon arrival immediately encrypts all the files on the user's hard disk. When it is finished, the applet announces what it has done and politely mentions that users wishing to purchase the decryption key can do so by sending 1000 dollars in small unmarked bills to a certain post office box in Panama.

In addition to the above get-rich-quick scheme, there are other dangers inherent in letting foreign code run on your machine. An applet could hunt around for interesting information (saved email, the password file, the local environment strings, etc.) and send or email them back over the network. It could also consume resources (e.g., filling up the disk), display naughty pictures or political slogans on the screen, or make an earsplitting racket using the sound card.

The Java designers were well aware of these problems, of course, and erected a series of barriers against them. The first line of defense is a typesafe language. Java has strong typing, true arrays with bounds checking and no pointers. These restrictions make it impossible for a Java program to construct a pointer to read and write arbitrary memory locations.

However, Trudy, who has given up on trying to break cryptographic protocols and gotten into the much more interesting business of writing malicious Java applets, can just write or modify a C compiler to produce Java bytecode, thus bypassing all the safeguards provided by the Java language and compiler.

The second line of defense is that before an incoming applet is executed, it is run through a bytecode verifier. The bytecode verifier looks for attempts to manufacture pointers, execute instructions or call methods with invalid parameters, use variables before they are initialized, and so on. These checks are supposed to guarantee that only legal applets get executed, but Trudy will certainly work hard on finding tricks the verifier does not check for.

The third line of defense is the class loader. Since classes can be loaded on the fly, there is a danger that an applet could load one of its own classes to replace a critical system class, thus bypassing that class' security checks. This Trojan horse attack has been rendered impossible by virtue of giving each class its own name space (like a kind of abstract directory), and carefully searching for system classes before looking for user classes. In other words, if the user loads a malicious version of *println*, it will never be used because the official *println* will always be found first.

The fourth line of defense is that some standard classes have their own security measures built in. For example, the file access class maintains a list of files that may be accessed by applets, and pops up a dialog box any time an applet tries to do something that violates the protection rules.

Despite all these measures, security problems are to be expected. First, there can be bugs in the Java software that clever programmers can exploit to bypass the security. The infamous Internet worm of 1988 used a bug in the UNIX Finger daemon to bring thousands of machines all over the Internet to a grinding halt (Hafner and Markoff, 1991; and Spafford, 1989).

Second, while it may be possible to prevent an applet from doing anything except writing to the screen, many applets will need more power, so when they ask for additional privileges, users may grudgingly (or naively) grant them. For example, applets may need to write temporary files, so users may give them access to the */tmp* directory, thinking that nothing important is there. Unfortunately, most editors keep the temporary versions of documents and email being edited there, so malicious applets can copy them and try to send them over the network. Of course, it may be possible to block applets' access to the network, but many may not work then, so they will have to be granted this power too.

But even in the unlikely event that applets are allowed no network access at all, they may be able to transmit information using **covert channels** (Lampson, 1973). For example, after acquiring some information, an applet can form a bit stream by using the local system's real time clock. To send a 1, it computes very hard for ΔT ; to send a 0, it just waits for ΔT .

To acquire this information, the applet's owner can establish a connection to the client's machine to read some of its public Web pages or FTP some of its public files. By carefully monitoring the incoming data rate, the applet owner's can see whether the applet is computing (and thus slowing down the observed output stream) or resting. Of course, this channel is noisy, but that can be handled by standard techniques. The stream can be divided into frames delimited by flag

bytes, individual frames can use a strong error-correcting code, and all frames can be sent two or three times. Many other covert channels exist, and it is extremely difficult to discover and block them all. For more information about the security problems in Java see (Dean and Wallach, 1995).

In short, Java introduces many new possibilities and opportunities into the World Wide Web. It allows Web pages to be interactive, and to contain animation and sound. It also permits browsers to be infinitely extensible. However, the Java model of downloading applets also introduces some serious new security problems that have not been entirely solved yet.

7.6.5. Locating Information on the Web

Although the Web contains a vast amount of information, finding the right item is not always easy. To make it easier for people to find pages that are useful to them, several researchers have written programs to index the Web in various ways. Some of these have become so popular that they have gone commercial. Programs that search the Web are sometimes called **search engines**, **spiders**, **crawlers**, **worms**, or **knowbots** (knowledge robots). In this section we will give a brief introduction to this subject. For more information, see (Pinkerton, 1994; and McBryan, 1994).

Although the Web is huge, reduced to its barest essentials, the Web is a big graph, with the pages being the nodes and the hyperlinks being the arcs. Algorithms for visiting all the nodes in a graph are well known. What makes Web indexing difficult is the enormous amount of data that must be managed and the fact that it is constantly changing.

Let us start our discussion with a simple goal: indexing all the keywords in Web pages' titles. For our algorithm, we will need three data structures. First, we need a large, linear array, *url_table*, that contains millions of entries, ultimately one per Web page. It should be kept in virtual memory, so parts not heavily used will automatically be paged to disk. Each entry contains two pointers, one to the page's URL and one to the page's title. Both of these items are variable length strings and can be kept on a heap (a large unstructured chunk of virtual memory to which strings can be appended). The heap is our second data structure.

The third data structure is a hash table of size n entries. It is used as follows. Any URL can be run through a hash function to produce a nonnegative integer less than n . All URLs that hash to the value k are hooked together on a linked list starting at entry k of the hash table. Whenever a URL is entered into *url_table*, it is also entered into the hash table. The main use of the hash table is to start with a URL and be able to quickly determine whether it is already present in *url_table*. These three data structures are illustrated in Fig. 7-75.

Building the index requires two phases: searching and indexing. Let us start with a simple engine for doing the searching. The heart of the search engine is a recursive procedure *process_url*, which takes a URL string as input. It operates as




Fig. 7-75. Data structures used in a simple search engine.

follows. First, it hashes the URL to see if it is already present in *url_table*. If so, it is done and returns immediately. Each URL is processed only once.

If the URL is not already known, its page is fetched. The URL and title are then copied to the heap and pointers to these two strings are entered in *url_table*. The URL is also entered into the hash table.

Finally, *process_url* extracts all the hyperlinks from the page and calls *process_url* once per hyperlink, passing the hyperlink's URL as the input parameter.

To run the search engine, *process_url* is called with some starting URL. When it returns, all pages reachable from that URL have been entered into *url_table* and the search phase has been completed.

Although this design is simple and theoretically correct, it has a serious problem in a system as large as the Web. The problem is that this algorithm does a depth-first search, and will ultimately go into recursion as many times as the longest noncyclic path on the Web. No one knows how long this path is, but it is probably thousands of hyperlinks long. As a consequence, any search engine that uses this depth-first search will probably hit stack overflow before finishing the job.

In practice, actual search engines first collect all the hyperlinks on each page they read, remove all the ones that have already been processed, and save the rest. The Web is then searched breadth-first; that is, each link on a page is followed and all the hyperlinks on all the pages pointed to are collected, but they are not traced in the order obtained.

The second phase does the keyword indexing. The indexing procedure goes down *url_table* linearly, processing each entry in turn. For each entry, it examines the title and selects out all words not on the stop list. (The stop list prevents indexing of prepositions, conjunctions, articles, and other words with many hits and little value.) For each word selected, it writes a line consisting of the word

followed by the current *url_table* entry number to a file. When the whole table has been scanned, the file is sorted by word.

The index will have to be stored on disk and can be used as follows. The user fills in a form listing one or more keywords and clicks on the SUBMIT button. This action causes a *POST* request to be done to a CGI script on the machine where the index is located. This script (or, more likely, program) then looks up the keywords in the index to find the set of *url_table* indices for each one. If the user wants the BOOLEAN AND of the keywords, the set intersection is computed. If the BOOLEAN OR is desired, the set union is computed.

The script now indexes into *url_table* to find all the titles and URLs. These are then combined to form a Web page and are sent back to the user as the response to the *POST*. The browser now displays the page, allowing the user to click on any items that appear interesting.

Sounds easy? It's not. The following problems have to be solved in any practical system:

1. Some URLs are obsolete (i.e., point to pages that no longer exist).
2. Some machines will be temporarily unreachable.
3. Not all pages may be reachable from the starting URL.
4. Some pages may be reachable only from active maps.
5. Some documents cannot be indexed (e.g., audio clips).
6. Not all documents have (useful) titles.
7. The search engine could run out of memory or disk space.
8. The entire process might take too long.

Obsolete URLs waste time but are mostly a nuisance because the server on which they are supposed to be located replies immediately with an error code. In contrast, when the server is down, all the search engine observes is a long delay in establishing the TCP connection. To prevent it from hanging indefinitely, it must have a timeout. If the timeout is too short, valid URLs will be missed. If it is too long, searching will be slowed down appreciably.

Choosing the starting URL is clearly a major issue. If the search engine starts with the home page of some astrophysicist, it may eventually find everything on astronomy, physics, chemistry and space science, but it may miss pages about veterinary medicine, Middle English, and rock 'n roll completely. These sets may simply be disjoint. One solution is to gather as large a set of URLs as possible, and use each of them as a starting page. Starting URLs can be gathered from USENET news articles and last week's version of the *url_table*, since some of these pages may have changed recently (e.g., one of the astrophysicists married a veterinarian and they solemnly updated their home pages to point to each other).

Indexing works well on text, but increasingly, many pages contain items other than text, including pictures, audio, and video. One approach here is to probe each new-found URL with the *HEAD* method, just to get back its MIME header. Anything not of type *text* is not searched.

About 20 percent of all Web pages have no title, and many of those that do have a title have a quasi-useless one ("Joe's page"). A big improvement to the basic index is to not only include titles, but also all the hypertext. In this approach, when a page is scanned, all the hyperlinks are also recorded, along with the page they came from and the page they point to. After the search phase has been completed, all the hyperwords can be indexed too.

Even more ambitious is to index all the important words in each page. To determine the important words, the occurrence frequency of all words not on the stop list can be computed (per Web page). The top 10 or 20 words are probably worth indexing. After all, if the word "liver" is the most common word on a page, there is a chance that the page will be of interest to biliary surgeons (or to cooks). Some search engines (e.g., Lycos) use this strategy.

Finally, the search engine can run out of memory or time. One attack is to redesign the algorithms more carefully. A completely different approach is to do what Harvest does and offload the work (Bowman et al., 1994, 1996). In particular, Harvest provides a program to run on cooperating servers. This program does all the searching locally and transmits back the finished local index. At the central site, all the local indices are merged into the master index. This approach reduces by orders of magnitude the amount of memory, CPU time, and network bandwidth required but has the major disadvantage of requiring all Web servers to cooperate by running foreign software. Given the potential problems with viruses and worms, when a system administrator is approached with the request: "Will you please run this program on your machine for me?" it should not be surprising if many of them decline.

One small request is in order. Although writing a search engine sounds easy, a buggy one can wreak havoc with the network by generating vast numbers of spurious requests, not only wasting bandwidth but bringing many servers to their knees due to the load. If you cannot resist the temptation to write your own search engine, proper netiquette requires restricting it to your own local DNS domain until it is totally debugged.

7.7. MULTIMEDIA

Multimedia is the holy grail of networking. When the word is mentioned, both the propeller heads and the suits begin salivating as if on cue. The former see immense technical challenges in providing (interactive) video on demand to every home. The latter see equally immense profits in it. No book on networking would be complete without at least an introduction to the subject. Given the

length of this one so far, our introduction will of necessity be brief. For additional information about this fascinating and potentially profitable subject, see (Buford, 1994; Deloddere et al., 1994; Dixit and Skelly, 1995; Fluckiger, 1995; Minoli, 1995; and Steinmetz and Nahrstedt, 1995).

Literally, multimedia is just two or more media. If the publisher of this book wanted to join the current hype about multimedia, it could advertise the book as using multimedia technology. After all, it contains two media: text and graphics (the figures). Nevertheless, when most people refer to multimedia, they generally mean the combination of two or more **continuous media**, that is, media that have to be played during some well-defined time interval, usually with some user interaction. In practice, the two media are normally audio and video, that is, sound plus moving pictures. For this reason, we will begin our study with an introduction to audio and video technology. Then we will combine them and move on to true multimedia systems, including video on demand and the Internet's multimedia system, MBone.

7.7.1. Audio

An audio (sound) wave is a one-dimensional acoustic (pressure) wave. When an acoustic wave enters the ear, the eardrum vibrates, causing the tiny bones of the inner ear to vibrate along with it, sending nerve pulses to the brain. These pulses are perceived as sound by the listener. In a similar way, when an acoustic wave strikes a microphone, the microphone generates an electrical signal, representing the sound amplitude as a function of time. The representation, processing, storage, and transmission of such audio signals are a major part of the study of multimedia systems.

The frequency range of the human ear runs from 20 Hz to 20,000 Hz, although some animals, notably dogs, can hear higher frequencies. The ear hears logarithmically, so the ratio of two sounds with amplitudes A and B is conventionally expressed in **dB (decibels)** according to the formula

$$\text{dB} = 20 \log_{10}(A/B)$$

If we define the lower limit of audibility (a pressure of about 0.0003 dyne/cm^2) for a 1-kHz sine wave as 0 dB, an ordinary conversation is about 50 dB and the pain threshold is about 120 dB, a dynamic range of a factor of 1 million. To avoid any confusion, A and B above are *amplitudes*. If we were to use the power level, which is proportional to the square of the amplitude, the coefficient of the logarithm would be 10, not 20.

The ear is surprisingly sensitive to sound variations lasting only a few milliseconds. The eye, in contrast, does not notice changes in light level that last only a few milliseconds. The result of this observation is that jitter of only a few milliseconds during a multimedia transmission affects the perceived sound quality more than it affects the perceived image quality.

Audio waves can be converted to digital form by an **ADC (Analog Digital Converter)**. An ADC takes an electrical voltage as input and generates a binary number as output. In Fig. 7-76(a) we see an example of a sine wave. To represent this signal digitally, we can sample it every ΔT seconds, as shown by the bar heights in Fig. 7-76(b). If a sound wave is not a pure sine wave, but a linear superposition of sine waves where the highest frequency component present is f , then the Nyquist theorem (see Chap. 2) states that it is sufficient to make samples at a frequency $2f$. Sampling more often is of no value since the higher frequencies that such sampling could detect are not present.




Fig. 7-76. (a) A sine wave. (b) Sampling the sine wave. (c) Quantizing the samples to 3 bits.

Digital samples are never exact. The 3-bit samples of Fig. 7-76(c) allow only eight values, from -1.00 to $+1.00$ in steps of 0.25 . An 8-bit sample would allow 256 distinct values. A 16-bit sample would allow 65,536 distinct values. The error introduced by the finite number of bits per sample is called the **quantization noise**. If it is too large, the ear detects it.

Two well-known examples of sampled sound are the telephone and audio compact discs. Pulse code modulation, as used within the telephone system, uses 7-bit (North America and Japan) or 8-bit (Europe) samples 8000 times per second. This system gives a data rate of 56,000 bps or 64,000 bps. With only 8000 samples/sec, frequencies above 4 kHz are lost.

Audio CDs are digital with a sampling rate of 44,100 samples/sec, enough to capture frequencies up to 22,050 Hz, which is good for people, bad for dogs. The samples are 16 bits each, and are linear over the range of amplitudes. Note that 16-bit samples allow only 65,536 distinct values, even though the dynamic range of the ear is about 1 million when measured in steps of the smallest audible sound. Thus using only 16 bits per sample introduces some quantization noise (although the full dynamic range is not covered—CDs are not supposed to hurt). With 44,100 samples/sec of 16 bits each, an audio CD needs a bandwidth of 705.6 kbps

for monaural and 1.411 Mbps for stereo. While this is lower than what video needs (see below), it still takes almost a full T1 channel to transmit uncompressed CD quality stereo sound.

Digitized sound can be easily processed by computers in software. Dozens of programs exist for personal computers to allow users to record, display, edit, mix, and store sound waves from multiple sources. Virtually all professional sound recording and editing are digital nowadays.

Many musical instruments even have a digital interface now. When digital instruments first came out, each one had its own interface, but after a while, a standard, **MIDI (Music Instrument Digital Interface)**, was developed and adopted by virtually the entire music industry. This standard specifies the connector, the cable, and the message format. Each MIDI message consists of a status byte followed by zero or more data bytes. A MIDI message conveys one musically significant event. Typical events are a key being pressed, a slider being moved, or a foot pedal being released. The status byte indicates the event, and the data bytes give parameters, such as which key was depressed and with what velocity it was moved.

Every instrument has a MIDI code assigned to it. For example, a grand piano is 0, a marimba is 12, and a violin is 40. This is needed to avoid having a flute concerto be played back as a tuba concerto. The number of "instruments" defined is 127. However, some of these are not instruments, but special effects such as chirping birds, helicopters, and the canned applause that accompanies many television programs.

The heart of every MIDI system is a synthesizer (often a computer) that accepts messages and generates music from them. The synthesizer understands all 127 instruments, so it generates a different power spectrum for middle C on a trumpet than for a xylophone. The advantage of transmitting music using MIDI compared to sending a digitized waveform is the enormous reduction in bandwidth, often by a factor of 1000. The disadvantage of MIDI is that the receiver needs a MIDI synthesizer to reconstruct the music again, and different ones may give slightly different renditions.

Music, of course, is just a special case of general audio, but an important one. Another important special case is speech. Human speech tends to be in the 600-Hz to 6000-Hz range. Speech is made up of vowels and consonants, which have different properties. Vowels are produced when the vocal tract is unobstructed, producing resonances whose fundamental frequency depends on the size and shape of the vocal system and the position of the speaker's tongue and jaw. These sounds are almost periodic for intervals of about 30 msec. Consonants are produced when the vocal tract is partially blocked. These sounds are less regular than vowels.

Some speech generation and transmission systems make use of models of the vocal system to reduce speech to a few parameters (e.g., the sizes and shapes of various cavities), rather than just sampling the speech waveform.

7.7.2. Video

The human eye has the property that when an image is flashed on the retina, it is retained for some number of milliseconds before decaying. If a sequence of images is flashed at 50 or more images/sec, the eye does not notice that it is looking at discrete images. All video (i.e., television) systems exploit this principle to produce moving pictures.

Analog Systems

To understand video systems, it is best to start with simple, old-fashioned black-and-white television. To represent the two-dimensional image in front of it as a one-dimensional voltage as a function of time, the camera scans an electron beam rapidly across the image and slowly down it, recording the light intensity as it goes. At the end of the scan, called a **frame**, the beam retraces. This intensity as a function of time is broadcast, and receivers repeat the scanning process to reconstruct the image. The scanning pattern used by both the camera and the receiver is shown in Fig. 7-77. (As an aside, CCD cameras integrate rather than scan, but some cameras and all monitors do scan.)




Fig. 7-77. The scanning pattern used for NTSC video and television.

The exact scanning parameters vary from country to country. The system used in North and South America and Japan has 525 scan lines, a horizontal to

vertical aspect ratio of 4:3, and 30 frames/sec. The European system has 625 scan lines, the same aspect ratio of 4:3, and 25 frames/sec. In both systems, the top few and bottom few lines are not displayed (to approximate a rectangular image on the original round CRTs). Only 483 of the 525 NTSC scan lines (and 576 of the 625 PAL/SECAM scan lines) are displayed. The beam is turned off during the vertical retrace, so many stations (especially in Europe) use this interval to broadcast TeleText (text pages containing news, weather, sports, stock prices, etc.).

While 25 frames/sec is enough to capture smooth motion, at that frame rate many people, especially older ones, will perceive the image to flicker (because the old image has faded off the retina before the new one appears). Rather than increase the frame rate, which would require using more scarce bandwidth, a different approach is taken. Instead of displaying the scan lines in order, first all the odd scan lines are displayed, then the even ones are displayed. Each of these half frames is called a **field**. Experiments have shown that although people notice flicker at 25 frames/sec, they do not notice it at 50 fields/sec. This technique is called **interlacing**. Noninterlaced television or video is said to be **progressive**.

Color video uses the same scanning pattern as monochrome (black and white), except that instead of displaying the image with one moving beam, three beams moving in unison are used. One beam is used for each of the three additive primary colors: red, green, and blue (RGB). This technique works because any color can be constructed from a linear superposition of red, green, and blue with the appropriate intensities. However, for transmission on a single channel, the three color signals must be combined into a single **composite** signal.

When color television was invented, various methods for displaying color were technically possible, and different countries made different choices, leading to systems that are still incompatible. (Note that these choices have nothing to do with VHS versus Betamax versus P2000, which are recording methods.) In all countries, a political requirement was that programs transmitted in color had to be receivable on existing black-and-white television sets. Consequently, the simplest scheme, just encoding the RGB signals separately, was not acceptable. RGB is also not the most efficient scheme.

The first color system was standardized in the United States by the National Television Standards Committee, which lent its acronym to the standard: **NTSC**. Color television was introduced in Europe several years later, by which time the technology had improved substantially, leading to systems with greater noise immunity and better colors. These are called **SECAM (SEquentiel Couleur Avec Memoire)**, which is used in France and Eastern Europe, and **PAL (Phase Alternating Line)** used in the rest of Europe. The difference in color quality between the NTSC and PAL/SECAM has led to an industry joke that NTSC really stands for Never Twice the Same Color.

To allow color transmissions to be viewed on black-and-white receivers, all three systems linearly combine the RGB signals into a **luminance** (brightness)

signal, and two **chrominance** (color) signals, although they all use different coefficients for constructing these signals from the RGB signals. Interestingly enough, the eye is much more sensitive to the luminance signal than to the chrominance signals, so the latter need not be transmitted as accurately. Consequently, the luminance signal can be broadcast at the same frequency as the old black-and-white signal, so it can be received on black-and-white television sets. The two chrominance signals are broadcast in narrow bands at higher frequencies. Some television sets have controls labeled brightness, hue, and saturation (or brightness, tint and color) for controlling these three signals separately. Understanding luminance and chrominance is necessary for understanding how video compression works.

In the past few years, there has been considerable interest in **HDTV (High Definition TeleVision)**, which produces sharper images by roughly doubling the number of scan lines. The United States, Europe, and Japan have all developed HDTV systems, all different and all mutually incompatible. The basic principles of HDTV in terms of scanning, luminance, chrominance, and so on, are similar to the existing systems. However, all three formats have a common aspect ratio of 16:9 instead of 4:3 to match them better to the format used for movies (which are recorded on 35 mm film).

For an introduction to television technology, see (Buford, 1994).

Digital Systems

The simplest representation of digital video is a sequence of frames, each consisting of a rectangular grid of picture elements, or **pixels**. Each pixel can be a single bit, to represent either black or white. The quality of such a system is similar to what you get by sending a color photograph by fax—awful. (Try it if you can; otherwise photocopy a color photograph on a copying machine that does not rasterize.)

The next step up is to use 8 bits per pixel to represent 256 gray levels. This scheme gives high-quality black-and-white video. For color video, good systems use 8 bits for each of the RGB colors, although nearly all systems mix these into composite video for transmission. While using 24 bits per pixel limits the number of colors to about 16 million, the human eye cannot even distinguish this many colors, let alone more. Digital color images are produced using three scanning beams, one per color. The geometry is the same as for the analog system of Fig. 7-77 except that the continuous scan lines are now replaced by neat rows of discrete pixels.

To produce smooth motion, digital video, like analog video, must display at least 25 frames/sec. However, since good quality computer monitors often rescan the screen from images stored in memory at 75 times per second or more, interlacing is not needed and consequently is not normally used. Just repainting (i.e., redrawing) the same frame three times in a row is enough to eliminate flicker.

In other words, smoothness of motion is determined by the number of *different* images per second, whereas flicker is determined by the number of times the screen is painted per second. These two parameters are different. A still image painted at 20 frames/sec will not show jerky motion but it will flicker because one frame will decay from the retina before the next one appears. A movie with 20 different frames per second, each of which is painted four times in a row, will not flicker, but the motion will appear jerky.

The significance of these two parameters becomes clear when we consider the bandwidth required for transmitting digital video over a network. Current computer monitors all use the 4:3 aspect ratio so they can use inexpensive, mass-produced picture tubes designed for the consumer television market. Common configurations are 640×480 (VGA), 800×600 (SVGA), and 1024×768 (XGA). An XGA display with 24 bits per pixel and 25 frames/sec needs to be fed at 472 Mbps. Even OC-9 is not quite good enough, and running an OC-9 SONET carrier into everyone's house is not exactly on the agenda. Doubling this rate to avoid flicker is even less attractive. A better solution is to transmit 25 frames/sec and have the computer store each one and paint it twice. Broadcast television does not use this strategy because television sets do not have memory, and in any event, analog signals cannot be stored in RAM without first converting them to digital form, which requires extra hardware. As a consequence, interlacing is needed for broadcast television but not for digital video.

7.7.3. Data Compression

It should be obvious by now that transmitting multimedia material in uncompressed form is completely out of the question. The only hope is that massive compression is possible. Fortunately, a large body of research over the past few decades has led to many compression techniques and algorithms that make multimedia transmission feasible. In this section we will study some methods for compressing multimedia data, especially images. For more detail, see (Fluckiger, 1995; and Steinmetz and Nahrstedt, 1995).

All compression systems require two algorithms: one for compressing the data at the source, and another for decompressing it at the destination. In the literature, these algorithms are referred to as the **encoding** and **decoding** algorithms, respectively. We will use this terminology here, too.

These algorithms have certain asymmetries that are important to understand. First, for many applications, a multimedia document, say, a movie will only be encoded once (when it is stored on the multimedia server) but will be decoded thousands of times (when it is viewed by customers). This asymmetry means that it is acceptable for the encoding algorithm to be slow and require expensive hardware provided that the decoding algorithm is fast and does not require expensive hardware. After all, the operator of a multimedia server might be quite willing to rent a parallel supercomputer for a few weeks to encode its entire video

library, but requiring consumers to rent a supercomputer for 2 hours to view a video is not likely to be a big success. Many practical compression systems go to great lengths to make decoding fast and simple, even at the price of making encoding slow and complicated.

On the other hand, for real-time multimedia, such as video conferencing, slow encoding is unacceptable. Encoding must happen on-the-fly, in real time. Consequently, real-time multimedia uses different algorithms or parameters than storing videos on disk, often with appreciably less compression.

A second asymmetry is that the encode/decode process need not be invertible. That is, when compressing a file, transmitting it, and then decompressing it, the user expects to get the original back, accurate down to the last bit. With multimedia, this requirement does not exist. It is usually acceptable to have the video signal after encoding and then decoding be slightly different than the original. When the decoded output is not exactly equal to the original input, the system is said to be **lossy**. If the input and output are identical, the system is **lossless**. Lossy systems are important because accepting a small amount of information loss can give a huge payoff in terms of the compression ratio possible.

Entropy Encoding

Compression schemes can be divided into two general categories: entropy encoding and source encoding. We will now discuss each in turn.

Entropy encoding just manipulates bit streams without regard to what the bits mean. It is a general, lossless, fully reversible technique, applicable to all data. We will illustrate it by means of three examples.

Our first example of entropy encoding is **run-length encoding**. In many kinds of data, strings of repeated symbols (bits, numbers, etc.) are common. These can be replaced by a special marker not otherwise allowed in the data, followed by the symbol comprising the run, followed by how many times it occurred. If the special marker itself occurs in the data, it is duplicated (as in character stuffing). For example, consider the following string of decimal digits:

```
3150000000000084587111111111111163546740000000000000000000065
```

If we now introduce A as the marker and use two-digit numbers for the repetition count, we can encode the above digit string as

```
315A01284587A11316354674A02265
```

Here run-length encoding has cut the data string in half.

Runs are common in multimedia. In audio, silence is often represented by runs of zeros. In video, runs of the same color occur in shots of the sky, walls, and many flat surfaces. All of these runs can be greatly compressed.

Our second example of entropy encoding is **statistical encoding**. By this we mean using a short code to represent common symbols and long ones to represent

infrequent ones. Morse code uses this principle, with E being • and Q being --•-- and so on. Huffman coding and the Ziv-Lempel algorithm used by the UNIX Compress program also use statistical encoding.

Our third example of entropy encoding is **CLUT (Color Look Up Table)** encoding. Consider an image using RGB encoding with 3 bytes/pixel. In theory, the image might contain as many as 2^{24} different color values. In practice, it will normally contain many fewer values, especially if the image is a cartoon or computer-generated drawing, rather than a photograph. Suppose that only 256 color values are actually used. A factor of almost three compression can be achieved by building a 768-byte table listing the RGB values of the 256 colors actually used, and then representing each pixel by the index of its RGB value in the table. Here we see a clear example where encoding is slower than decoding because encoding requires searching the table whereas decoding can be done with a single indexing operation.

Source Encoding

Now we come to **source encoding**, which takes advantage of properties of the data to produce more (usually lossy) compression. Here, too, we will illustrate the idea with three examples. Our first example is **differential encoding**, in which a sequence of values (e.g., audio samples) are encoded by representing each one as the difference from the previous value. Differential pulse code modulation, which we saw in Chap. 2, is an example of this technique. It is lossy because the signal might jump so much between two consecutive values that the difference does not fit in the field provided for expressing differences, so at least one incorrect value will be recorded and some information lost.

Differential encoding is a kind of source encoding because it takes advantage of the property that large jumps between consecutive data points are unlikely. Not all sequences of numbers have this property. An example lacking this property is a computer-generated list of random telephone numbers to be used by telemarketers for bothering people during dinner. The differences between consecutive telephone numbers in the list will take as many bits to represent as the numbers themselves.

Our second example of source encoding consists of **transformations**. By transforming signals from one domain to another, compression may become much easier. Consider, for example, the Fourier transformation of Fig. 2-1(e). Here a function of time is represented as a list of amplitudes. Given the exact values of all the amplitudes, the original function can be reconstructed perfectly. However, given only the values of the first, say, eight amplitudes rounded off to two decimal places, it may still be possible to reconstruct the signal so well that the listener cannot tell that some information has been lost. The gain is that transmitting eight amplitudes requires many fewer bits than transmitting the sampled waveform.

Transformations are also applicable to two-dimensional image data. Suppose that the 4×4 matrix of Fig. 7-78(a) represents the gray-scale values of a monochrome image. We can transform these data by subtracting the value in the upper left-hand corner from all elements except itself, as shown in Fig. 7-78(b). This transformation might be useful if variable-length encoding is used. For example, values between -7 and $+7$ could be encoded with 4-bit numbers and values between 0 and 255 could be encoded as a special 4-bit code (-8) followed by an 8-bit number.

Pixel value		4 pixels			
160	160	161	160		
161	165	166	158		
160	167	165	161		
159	160	160	160		
(a)		160	0	1	0
		1	5	6	-2
		0	7	5	1
		-1	0	1	0
		(b)			

Fig. 7-78. (a) Pixel values for part of an image. (b) A transformation in which the upper left-hand element is subtracted from all elements except itself.

Although this simple transformation is lossless, other, more useful ones are not. An especially important two-dimensional spatial transformation is the **DCT (Discrete Cosine Transformation)** (Feig and Winograd, 1992). This transformation has the property that for images without sharp discontinuities, most of the spectral power is in the first few terms, allowing the later ones to be ignored without much information loss. We will come back to DCT shortly.

Our third example of source encoding is **vector quantization**, which is also directly applicable to image data. Here, the image is divided up into fixed-size rectangles. In addition to the image itself, we also need a table of rectangles of the same size as the image rectangles (possibly constructed from the image). This table is called the **code book**. Each rectangle is transmitted by looking it up in the code book and just sending the index instead of the rectangle. If the code book is created dynamically (i.e., per image), it must be transmitted, too. Clearly, if a small number of rectangles dominate the image, large savings in bandwidth are possible here.

An example of vector quantization is shown in Fig. 7-79. In Fig. 7-79(a) we have a grid of rectangles of unspecified size. In Fig. 7-79(b) we have the code book. The output stream is just the list of integers 001022032200400 shown in Fig. 7-79(c). Each one represents an entry from the code book.




Fig. 7-79. An example of vector quantization. (a) An image divided into squares. (b) A code book for the image. (c) The encoded image.

In a sense, vector quantization is just a two-dimensional generalization of CLUT. The real difference, however, is what happens if no match can be found. Three strategies are possible. The first one is just to use the best match. The second one is to use the best match, and append some information about how to improve the match (e.g., append the true mean value). The third one is use the best match and append whatever is necessary to allow the decoder to reconstruct the data exactly. The first two strategies are lossy but exhibit high compression. The third is lossless but less effective as a compression algorithm. Again, we see that encoding (pattern matching) is far more time consuming than decoding (indexing into a table).

The JPEG Standard

The **JPEG (Joint Photographic Experts Group)** standard for compressing continuous-tone still pictures (e.g., photographs) was developed by photographic experts working under the joint auspices of ITU, ISO, and IEC, another standards body. It is important for multimedia because, to a first approximation, the multimedia standard for moving pictures, MPEG, is just the JPEG encoding of each frame separately, plus some extra features for interframe compression and motion detection. JPEG is defined in International Standard 10918.

JPEG has four modes and many options. It is more like a shopping list than a single algorithm. For our purposes, though, only the lossy sequential mode is relevant, and that one is illustrated in Fig. 7-80. Furthermore, we will concentrate on the way JPEG is normally used to encode 24-bit RGB video images and will leave out some of the minor details for the sake of simplicity.

Step 1 of encoding an image with JPEG is block preparation. For the sake of specificity, let us assume that the JPEG input is a 640×480 RGB image with 24 bits/pixel, as shown in Fig. 7-81(a). Since using luminance and chrominance




Fig. 7-80. The operation of JPEG in lossy sequential mode.

gives better compression, we first compute the luminance, Y , and the two chrominances, I and Q (for NTSC), according to the following formulas:

$$Y = 0.30R + 0.59G + 0.11B$$

$$I = 0.60R - 0.28G - 0.32B$$

$$Q = 0.21R - 0.52G + 0.31B$$

For PAL, the chrominances are called U and V and the coefficients are different, but the idea is the same. SECAM is different from both NTSC and PAL.




Fig. 7-81. (a) RGB input data. (b) After block preparation.

Separate matrices are constructed for Y , I , and Q , each with elements in the range 0 to 255. Next, square blocks of four pixels are averaged in the I and Q matrices to reduce them to 320×240 . This reduction is lossy, but the eye barely notices it since the eye responds to luminance more than to chrominance. Nevertheless, it compresses the data by a factor of two. Now 128 is subtracted from each element of all three matrices to put 0 in the middle of the range. Finally, each matrix is divided up into 8×8 blocks. The Y matrix has 4800 blocks; the other two have 1200 blocks each, as shown in Fig. 7-81(b).

Step 2 of JPEG is to apply a discrete cosine transformation to each of the 7200 blocks separately. The output of each DCT is an 8×8 matrix of DCT coefficients. DCT element (0, 0) is the average value of the block. The other elements tell how much spectral power is present at each spatial frequency. In theory, a DCT is lossless, but in practice using floating-point numbers and

transcendental functions always introduces some roundoff error that results in a little information loss. Normally, these elements decay rapidly with distance from the origin, (0, 0), as suggested by Fig. 7-82.




Fig. 7-82. (a) One block of the Y matrix. (b) The DCT coefficients.

Once the DCT is complete, JPEG moves on to step 3, called **quantization**, in which the less important DCT coefficients are wiped out. This (lossy) transformation is done by dividing each of the coefficients in the 8×8 DCT matrix by a weight taken from a table. If all the weights are 1, the transformation does nothing. However, if the weights increase sharply from the origin, higher spatial frequencies are dropped quickly.

An example of this step is given in Fig. 7-83. Here we see the initial DCT matrix, the quantization table, and the result obtained by dividing each DCT element by the corresponding quantization table element. The values in the quantization table are not part of the JPEG standard. Each application must supply its own, allowing it to control the loss-compression trade-off.

Step 4 reduces the (0, 0) value of each block (the one in the upper left-hand corner) by replacing it with the amount it differs from the corresponding element in the previous block. Since these elements are the averages of their respective blocks, they should change slowly, so taking the differential values should reduce most of them to small values. No differentials are computed from the other values. The (0, 0) values are referred to as the DC components; the other values are the AC components.

Step 5 linearizes the 64 elements and applies run-length encoding to the list. Scanning the block from left to right and then top to bottom will not concentrate the zeros together, so a zig zag scanning pattern is used, as shown in Fig. 7-84. In this example, the zig zag pattern ultimately produces 38 consecutive 0s at the end of the matrix. This string can be reduced to a single count saying there are 38 zeros.

Now we have a list of numbers that represent the image (in transform space). Step 6 Huffman encodes the numbers for storage or transmission.

DCT Coefficients								Quantized coefficients							
150	80	40	14	4	2	1	0	150	80	20	4	1	0	0	0
92	75	36	10	6	1	0	0	92	75	18	3	1	0	0	0
52	38	26	8	7	4	0	0	26	19	13	2	1	0	0	0
12	8	6	4	2	1	0	0	3	2	2	1	0	0	0	0
4	3	2	0	0	0	0	0	1	0	0	0	0	0	0	0
2	2	1	1	0	0	0	0	0	0	0	0	0	0	0	0
1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Quantization table

1	1	2	4	8	16	32	64
1	1	2	4	8	16	32	64
2	2	2	4	8	16	32	64
4	4	4	4	8	16	32	64
8	8	8	8	8	16	32	64
16	16	16	16	16	16	32	64
32	32	32	32	32	32	32	64
64	64	64	64	64	64	64	64

Fig. 7-83. Computation of the quantized DCT coefficients.

150	80	20	4	1	0	0	0
92	75	18	3	1	0	0	0
26	19	13	2	1	0	0	0
3	2	2	1	0	0	0	0
1	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

Fig. 7-84. The order in which the quantized values are transmitted.

JPEG may seem complicated, but that is because it *is* complicated. Still, since it often produces a 20:1 compression or better, it is widely used. Decoding a JPEG image requires running the algorithm backward. Unlike some of the other compression algorithms we have seen, JPEG is roughly symmetric: decoding takes as long as encoding.

Interestingly enough, due to the mathematical properties of the DCT, it is possible to perform certain geometric transformations (e.g. image rotation) directly on the transformed matrix, without regenerating the original image. These transformations are discussed in (Shen and Sethi, 1995). Similar properties also apply to MPEG compressed audio (Broadhead and Owen, 1995).

The MPEG Standard

Finally, we come to the heart of the matter: the **MPEG (Motion Picture Experts Group)** standards. These are the main algorithms used to compress videos and have been international standards since 1993. Because movies contain both images and sound, MPEG can compress both audio and video, but since video takes up more bandwidth and also contains more redundancy than audio, we will primarily focus on MPEG video compression below.

The first standard to be finalized was MPEG-1 (International Standard 11172). Its goal was to produce video recorder-quality output (352×240 for NTSC) using a bit rate of 1.2 Mbps. Since we saw earlier that uncompressed video alone can run to 472 Mbps, getting it down to 1.2 Mbps is not entirely trivial, even at this lower resolution. MPEG-1 can be transmitted over twisted pair transmission lines for modest distances. MPEG-1 is also used for storing movies on CD-ROM in CD-I and CD-Video format.

The next standard in the MPEG family was MPEG-2 (International Standard 13818), which was originally designed for compressing broadcast quality video into 4 to 6 Mbps, so it could fit in a NTSC or PAL broadcast channel. Later, MPEG-2 was expanded to support higher resolutions, including HDTV. MPEG-4 is for medium-resolution videoconferencing with low frame rates (10 frames/sec) and at low bandwidths (64 kbps). This will permit videoconferences to be held over a single N-ISDN B channel. Given this numbering, one might think that the next standard will be MPEG-8. Actually, ISO is numbering them linearly, not exponentially. Originally MPEG-3 existed. It was intended for HDTV, but that project was later canceled, and HDTV was added to MPEG-2.

The basic principles of MPEG-1 and MPEG-2 are similar, but the details are different. To a first approximation, MPEG-2 is a superset of MPEG-1, with additional features, frame formats and encoding options. It is likely that in the long run MPEG-1 will dominate for CD-ROM movies and MPEG-2 will dominate for long-haul video transmission. We will discuss MPEG-1 first and then MPEG-2.

MPEG-1 has three parts: audio, video, and system, which integrates the other two, as shown in Fig. 7-85. The audio and video encoders work independently,

which raises the issue of how the two streams get synchronized at the receiver. This problem is solved by having a 90-kHz system clock that outputs the current time value to both encoders. These values are 33 bits, to allow films to run for 24 hours without wrapping around. These timestamps are included in the encoded output and propagated all the way to the receiver, which can use them to synchronize the audio and video streams.




Fig. 7-85. Synchronization of the audio and video streams in MPEG-1.

MPEG audio compression is done by sampling the waveform at 32 kHz, 44.1 kHz, or 48 kHz. It can handle monaural, disjoint stereo (each channel compressed separately), or joint stereo (interchannel redundancy exploited). It is organized as three layers, each one applying additional optimizations to get more compression (and at greater cost). Layer 1 is the basic scheme. This layer is used, for example, in the DCC digital tape system. Layer 2 adds advanced bit allocation to the basic scheme. It is used for CD-ROM audio and movie soundtracks. Layer 3 adds hybrid filters, nonuniform quantization, Huffman coding, and other advanced techniques.

MPEG audio can compress a rock 'n roll CD down to 96 kbps with no perceptible loss in quality, even for rock 'n roll fans with no hearing loss. For a piano concert, at least 128 kbps are needed. These differ because the signal-to-noise ratio for rock 'n roll is much higher than for a piano concert (in an engineering sense, anyway).

Audio compression is carried out by performing a fast Fourier transformation on the audio signal to transform it from the time domain to the frequency domain. The resulting spectrum is then divided up into 32 frequency bands, each of which is processed separately. When two stereo channels are present, the redundancy inherent in having two highly overlapping audio sources is also exploited. The resulting MPEG-1 audio stream is adjustable from 32 kbps to 448 kbps. An introduction to the process is given in (Pan, 1995).

Now let us consider MPEG-1 video compression. Two kinds of redundancies exist in movies: spatial and temporal. MPEG-1 uses both. Spatial redundancy can be utilized by simply coding each frame separately with JPEG. This approach is sometimes used, especially when random access to each frame is needed, as in

editing video productions. In this mode, a compressed bandwidth in the 8- to 10-Mbps range is achievable.

Additional compression can be achieved by taking advantage of the fact that consecutive frames are often almost identical. This effect is smaller than it might first appear since many movie makers cut between scenes every 3 or 4 seconds (time a movie and count the scenes). Nevertheless, even a run of 75 highly similar frames offers the potential of a major reduction over simply encoding each frame separately with JPEG.

For scenes where the camera and background are stationary and one or two actors are moving around slowly, nearly all the pixels will be identical from frame to frame. Here, just subtracting each frame from the previous one and running JPEG on the difference would do fine. However, for scenes where the camera is panning or zooming, this technique fails badly. What is needed is some way to compensate for this motion. This is precisely what MPEG does; it is the main difference between MPEG and JPEG.

MPEG-1 output consists of four kinds of frames:

1. I (Intracoded) frames: Self-contained JPEG-encoded still pictures.
2. P (Predictive) frames: Block-by-block difference with the last frame.
3. B (Bidirectional) frames: Differences with the last and next frame.
4. D (DC-coded) frames: Block averages used for fast forward.

I-frames are just still pictures coded using JPEG, also using full-resolution luminance and half-resolution chrominance along each axis. It is necessary to have I-frames appear in the output stream periodically for three reasons. First, MPEG-1 can be used for a multicast transmission, with viewers tuning it at will. If all frames depended on their predecessors going back to the first frame, anybody who missed the first frame could never decode any subsequent frames. Second, if any frame were received in error, no further decoding would be possible. Third, without I-frames, while doing a fast forward or rewind, the decoder would have to calculate every frame passed over so it would know the full value of the one it stopped on. For these reasons, I-frames are inserted into the output once or twice per second.

P-frames, in contrast, code interframe differences. They are based on the idea of **macroblocks**, which cover 16×16 pixels in luminance space and 8×8 pixels in chrominance space. A macroblock is encoded by searching the previous frame for it or something only slightly different from it.

An example of where P-frames would be useful is given in Fig. 7-86. Here we see three consecutive frames that have the same background, but differ in the position of one person. The macroblocks containing the background scene will match exactly, but the macroblocks containing the person will be offset in position by some unknown amount and will have to be tracked down.




Fig. 7-86. Three consecutive frames.

The MPEG-1 standard does not specify how to search, how far to search, or how good a match has to be to count. This is up to each implementation. For example, an implementation might search for a macroblock at the current position in the previous frame, and all other positions offset $\pm\Delta x$ in the x direction and $\pm\Delta y$ in the y direction. For each position, the number of matches in the luminance matrix could be computed. The position with the highest score would be declared the winner, provided it was above some predefined threshold. Otherwise, the macroblock would be said to be missing. Much more sophisticated algorithms are also possible, of course.

If a macroblock is found, it is encoded by taking the difference with its value in the previous frame (for luminance and both chrominances). These difference matrices are then subject to the discrete cosine transformation, quantization, run-length encoding, and Huffman encoding, just as with JPEG. The value for the macroblock in the output stream is then the motion vector (how far the macroblock moved from its previous position in each direction), followed by the Huffman encoded list of numbers. If the macroblock is not located in the previous frame, the current value is encoded with JPEG, just as in an I-frame.

Clearly, this algorithm is highly asymmetric. An implementation is free to try every plausible position in the previous frame if it wants to, in a desperate attempt to locate every last macroblock. This approach will minimize the encoded MPEG-1 stream at the expense of very slow encoding. This approach might be fine for a one-time encoding of a film library but would be terrible for real-time videoconferencing.

Similarly, each implementation is free to decide what constitutes a “found” macroblock. This freedom allows implementers to compete on the quality and speed of their algorithms, but always produce compliant MPEG-1. No matter what search algorithm is used, the final output is either the JPEG encoding of the current macroblock, or the JPEG encoding of the difference between the current macroblock and one in the previous frame at a specified offset from the current one.

So far, decoding MPEG-1 is straightforward. Decoding I-frames is the same as decoding JPEG images. Decoding P-frames requires the decoder to buffer the previous frame and then build up the new one in a second buffer based on fully

encoded macroblocks and macroblocks containing differences with the previous frame. The new frame is assembled macroblock by macroblock.

B-frames are similar to P-frames, except that they allow the reference macroblock to be in either a previous frame or in a succeeding frame. This additional freedom allows improved motion compensation, and is also useful when objects pass in front of, or behind, other objects. To do B-frame encoding, the encoder needs to hold three decoded frames in memory at once: the past one, the current one, and the future one. Although B-frames give the best compression, not all implementations support them.

D-frames are only used to make it possible to display a low-resolution image when doing a rewind or fast forward. Doing the normal MPEG-1 decoding in real time is difficult enough. Expecting the decoder to do it when slewing through the video at ten times normal speed is asking a bit much. Instead, the D-frames are used to produce low-resolution images. Each D-frame entry is just the average value of one block, with no further encoding, making it easy to display in real time. This facility is important to allow people to scan through a video at high speed in search of a particular scene.

Having finished our treatment of MPEG-1, let us move on to MPEG-2. MPEG-2 encoding is fundamentally similar to MPEG-1 encoding, with I-frames, P-frames, and B-frames. D-frames are not supported, however. Also, the discrete cosine transformation is 10×10 instead of 8×8 , to give 50 percent more coefficients, hence better quality. Since MPEG-2 is targeted at broadcast television as well as CD-ROM applications, it supports both progressive and interlaced images, whereas MPEG-1 supports only progressive images. Other minor details also differ between the two standards.

Instead of supporting only one resolution level, MPEG-2 supports four: low (352×240), main (720×480), high-1440 (1440×1152), and high (1920×1080). Low resolution is for VCRs and backward compatibility with MPEG-1. Main is the normal one for NTSC broadcasting. The other two are for HDTV.

In addition to having four resolution levels, MPEG-2 also supports five **profiles**. Each profile targets some application area. The main profile is for general-purpose use, and probably most chips will be optimized for the main profile and the main resolution level. The simple profile is similar to the main one, except that it excludes the use of B-frames, to make software encoding and decoding easier. The other profiles deal with scalability and HDTV. The profiles differ in terms of the presence or absence of B-frames, chrominance resolution, and scalability of the encoded bit stream to other formats.

The compressed data rate for each combination of resolution and profile is different. These range from about 3 Mbps up to 100 Mbps for HDTV. The normal case is about 3 to 4 Mbps. Some performance data for MPEG are given in (Pancho and El Zarki, 1994).

MPEG-2 has a more general way of multiplexing audio and video than the MPEG-1 model of Fig. 7-85. It defines an unlimited number of elementary

streams, including video and audio, but also including data streams that must be synchronized with the audio and video, for example, subtitles in multiple languages. Each of the streams is first packetized with timestamps. A simple two-stream example is shown in Fig. 7-87.




Fig. 7-87. Multiplexing of two streams in MPEG-2.

The output of each packetizer is a **PES (Packetized Elementary Stream)**. Each PES packet has about 30 header fields and flags, including lengths, stream identifiers, encryption control, copyright status, timestamps, and a CRC.

The PES streams for audio, video, and possibly data are then multiplexed together on a single output stream for transmission. Two types of streams are defined. The MPEG-2 **program stream** is similar to the MPEG-1 systems stream of Fig. 7-85. It is used for multiplexing together elementary streams that have a common time base and have to be displayed in a synchronized way. The program stream uses long variable-length packets.

The other MPEG-2 stream is the **transport stream**. It is used for multiplexing together streams (including program streams) that do not use a common time base. The transport stream packets are fixed length (188 bytes), to make it easier to limit the effect of packets damaged or lost during transmission.

It is worth noting that all the encoding schemes we have discussed are based on the model of lossy encoding followed by lossless transmission. Neither JPEG nor MPEG, for example, can recover from lost or damaged packets in a graceful way. A different approach to image transmission is to transform the images in a way that separates the important information from the less important information (as the DCT does, for example). Then add a considerable amount of redundancy (even duplicate packets) to the important information and none to the less

important information. If some packets are lost or garbled, it may still be possible to display reasonable images without retransmission. These ideas are described further in (Danskin et al., 1995). They are especially applicable to multicast transmissions, where feedback from each receiver is impossible anyway.

7.7.4. Video on Demand

Video on demand is sometimes compared to an electronic video rental store. The user (customer) selects any one of a large number of available videos and takes it home to view. Only with video on demand, the selection is made at home using the television set's remote control, and the video starts immediately. No trip to the store is needed. Needless to say, implementing video on demand is a wee bit more complicated than describing it. In this section, we will give an overview of the basic ideas and their implementation. A description of one real implementation can be found in (Nelson and Linton, 1995). A more general treatment of interactive television is in (Hodge, 1995). Other relevant references are (Chang et al., 1994; Hodge et al., 1993; and Little and Venkatesh, 1994).

Is video on demand really like renting a video, or is it more like picking a movie to watch from a 500- or 5000-channel cable system? The answer has important technical implications. In particular, video rental users are used to the idea of being able to stop a video, make a quick trip to the kitchen or bathroom, and then resume from where the video stopped. Television viewers do not expect to put programs on pause.

If video on demand is going to compete successfully with video rental stores, it may be necessary to allow users to stop, start, and rewind videos at will. Giving users this ability virtually forces the video provider to transmit a separate copy to each one.

On the other hand, if video on demand is seen more as advanced television, then it may be sufficient to have the video provider start each popular video, say, every 10 minutes, and run these nonstop. A user wanting to see a popular video may have to wait up to 10 minutes for it to start. Although pause/resume is not possible here, a viewer returning to the living room after a short break can switch to another channel showing the same video but 10 minutes behind. Some material will be repeated, but nothing will be missed. This scheme is called **near video on demand**. It offers the potential for much lower cost, because the same feed from the video server can go to many users at once. The difference between video on demand and near video on demand is similar to the difference between driving your own car and taking the bus.

Watching movies on (near) demand is but one of a vast array of potential new services possible once wideband networking is available. The general model that many people use is illustrated in Fig. 7-88. Here we see a high-bandwidth, (national or international) wide area backbone network at the center of the system. Connected to it are thousands of local distribution networks, such as cable TV or

telephone company distribution systems. The local distribution systems reach into people's houses, where they terminate in **set-top boxes**, which are, in fact, powerful, specialized personal computers.




Fig. 7-88. Overview of a video-on-demand system.

Attached to the backbone by high-bandwidth optical fibers are thousands of information providers. Some of these will offer pay-per-view video or pay-per-hear audio CDs. Others will offer specialized services, such as home shopping (with the ability to rotate a can of soup and zoom in on the list of ingredients or view a video clip on how to drive a gasoline-powered lawn mower). Sports, news, reruns of "I Love Lucy," WWW access, and innumerable other possibilities will no doubt quickly become available.

Also included in the system are local spooling servers that allow videos to be prepositioned closer to the users, to save bandwidth during peak hours. How these pieces will fit together and who will own what are matters of vigorous debate within the industry. Below we will examine the design of the main pieces of the system: the video servers, the distribution network, and the set-top boxes.

Video Servers

To have (near) video on demand, we need **video servers** capable of storing and outputting a large number of movies simultaneously. The total number of movies ever made is estimated at 65,000 (Minoli, 1995). When compressed in

MPEG-2, a normal movie occupies roughly 4 GB of storage, so 65,000 of them would require something like 260 terabytes. Add to this all the old television programs ever made, sports films, newsreels, talking shopping catalogs, etc., and it is clear that we have an industrial-strength storage problem on our hands.

The cheapest way to store large volumes of information is on magnetic tape. This has always been the case and probably always will be. A DAT tape can store 8 GB (two movies) at a cost of about 5 dollars/gigabyte. Large mechanical tape servers that hold thousands of tapes and have a robot arm for fetching any tape and inserting it into a tape drive are commercially available now. The problem with these systems is the access time (especially for the second movie on a tape), the transfer rate, and the limited number of tape drives (to serve n movies at once, the unit would need n drives).

Fortunately, experience with video rental stores, public libraries, and other such organizations shows that not all items are equally popular. Experimentally, when there are N movies available, the fraction of all requests being for the k th most popular one is approximately C/k (Chervenak, 1994). Here C is computed to normalize the sum to 1, namely

$$C = 1/(1 + 1/2 + 1/3 + 1/4 + 1/5 + \cdots + 1/N)$$

Thus the most popular movie is seven times as popular as the number seven movie. This result is known as **Zipf's law** (Zipf, 1949).

The fact that some movies are much more popular than others suggests a possible solution in the form of a storage hierarchy, as shown in Fig. 7-89. Here, the performance increases as one moves up the hierarchy.




Fig. 7-89. A video server storage hierarchy.

An alternative to tape is optical storage. Current CD-ROMs hold only 650 MB, but the next generation will hold about 4 GB, to make them suitable for distributing MPEG-2 movies. Although seek times are slow compared to magnetic disks (100 msec versus 10 msec), their low cost and high reliability make optical juke boxes containing thousands of CD-ROMs a good alternative to tape for the more heavily used movies.

Next come magnetic disks. These have short access times (10 msec), high transfer rates (10 MB/sec), and substantial capacities (10 GB), which makes them well suited to holding movies that are actually being transmitted (as opposed to

just being stored in case somebody ever wants them). Their main drawback is the high cost for storing movies that are rarely accessed.

At the top of the pyramid of Fig. 7-89 is RAM. RAM is the fastest storage medium, but also the most expensive. It is best suited to movies for which different parts are being sent to different destinations at the same time (e.g., true video on demand to 100 users who all started at different times). When RAM prices drop to 10 dollars/megabyte, a 4-GB movie will occupy 40,000 dollars worth of RAM, so having 100 movies in RAM will cost 4 million dollars for the 400 GB of memory. Still, for a 10 million dollar video server, this expense might well be worthwhile if each movie has enough simultaneous paying customers.

Since a video server is really just a massive real-time I/O device, it needs a different hardware and software architecture than a PC or a UNIX workstation. The hardware architecture of a typical video server is illustrated in Fig. 7-90. The server has one or more high-performance RISC CPUs, each with some local memory, a shared main memory, a massive RAM cache for popular movies, a variety of storage devices for holding the movies, and some networking hardware, normally an optical interface to an ATM (or SONET) network at OC-3 or higher. These subsystems are connected by an extremely high-speed bus (at least 1 GB/sec).




Fig. 7-90. The hardware architecture of a typical video server.

Now let us take a brief look at video server software. The CPUs are used for accepting user requests, locating movies, moving data between devices, customer billing, and many other functions. Some of these are not time critical, but many others are, so some, if not all, the CPUs will have to run a real-time operating system, such as a real-time microkernel. These systems normally break work up into

small tasks, each with a known deadline. The scheduler can then run an algorithm such as nearest deadline next or the rate monotonic algorithm (Liu and Layland, 1973).

The CPU software also defines the nature of the interface that the server presents to the clients (spooling servers and set-top boxes). Two designs are popular. The first one is a traditional file system, in which the clients can open, read, write, and close files. Other than the complications introduced by the storage hierarchy and real-time considerations, such a server can have a file system modeled after that of UNIX.

The second kind of interface is based on the video recorder model. The commands to the server request it to open, play, pause, fast forward, and rewind files. The difference with the UNIX model is that once a PLAY command is given, the server just keeps pumping out data at a constant rate, with no new commands required.

The heart of the video server software is the disk management software. It has two main jobs: placing movies on the magnetic disk when they have to be pulled up from optical or tape storage, and handling disk requests for the many output streams. Movie placement is important because it can greatly affect performance.

Two possible ways of organizing disk storage are the disk farm and the disk array. With the **disk farm**, each drive holds a few entire movies. For performance and reliability reasons, each movie should be present on at least two drives, maybe more. The other storage organization is the **disk array** or **RAID (Redundant Array of Inexpensive Disks)**, in which each movie is spread out over multiple drives, for example, block 0 on drive 0, block 1 on drive 1, and so on, with block $n - 1$ on drive $n - 1$. After that, the cycle repeats, with block n on drive 0, and so forth. This organizing is called **striping**.

A striped disk array has several advantages over a disk farm. First, all n drives can be running in parallel, increasing the performance by a factor of n . Second, it can be made redundant by adding an extra drive to each group of n , where the redundant drive contains the block-by-block EXCLUSIVE OR of the other drives, to allow full data recover in the event one drive fails. Finally, the problem of load balancing is solved (manual placement is not needed to avoid having all the popular movies on the same drive). On the other hand, the disk array organization is more complicated than the disk farm and highly sensitive to multiple failures. It is also ill-suited to video recorder operations such as rewinding or fast forwarding a movie. A simulation study comparing the two organizations is given in (Chervenak et al., 1995).

Closely related to block placement is finding disk blocks. The UNIX scheme of having an unbalanced tree of disk blocks pointed to by the i-node is usually unacceptable because video files are huge, so most blocks can only be located by going through a triple indirect block, which means many extra disk accesses (Tanenbaum, 1992). Instead, it is common to link the blocks together on a singly-

or doubly-linked list. Sometimes a UNIX-style index (i-node) is also used to allow random access.

The other job of the disk software is to service all the real-time output streams and meet their timing constraints. An MPEG-2 video stream at 25 frames/sec needs to fetch and transmit about 14 KB every 40 msec, but the actual amount varies considerably because I-, P-, and B-frames have different compression ratios. Consequently, to maintain a uniform output rate, buffering is needed at both ends of the stream.

In Fig. 7-91 we see a staircase showing the total amount of data fetched from the disk for a given video stream (assuming that the movie is on disk). It moves up in discrete jumps, one jump for each block read. Nevertheless, transmission must occur at a more uniform rate, so the disk reading process must keep ahead of the transmission process. The shaded area in the figure shows data that have been fetched from disk but not yet transmitted.




Fig. 7-91. Disk buffering at the server.

Normally, disks are scheduled using the elevator algorithm, which starts the arm moving inward and keeps going until it hits the innermost cylinder, processing all requests it hits in cylinder order. When it gets as far in as it can, the arm reverses and starts moving outward, again processing all pending requests along the way in order. While this algorithm minimizes seek time, it makes no guarantees about real-time performance, so is not useful for a video server.

A better algorithm is to keep track of all video streams and make a list of the next block needed by each one. These block numbers are then sorted and the blocks read in cylinder order. When the last block is read, the next round begins by collecting the number of the block now at the head of each stream. These are also sorted and read in cylinder order, and so on. This algorithm maintains real-time performance for all streams but also minimizes seek time compared to a pure first-come, first-served algorithm.

Another software issue is admission control. If a request for a new stream comes in, can it be accepted without ruining the real-time performance of the

existing streams? One algorithm that can be used for making a decision examines the worst case to see if going from k streams to $k + 1$ streams is guaranteed to be possible, based on the known properties of the CPU, RAM, and disk. Another algorithm just looks at the statistical properties.

Another server software issue is how to manage the display during a fast forward or fast backward (so people can search visually). The D-frames provide the necessary information for MPEG-1, but unless they are marked and stored in some special way, the server will not be able to find them without decoding the entire stream, and normally servers do not perform MPEG decoding during transmission. For MPEG-2, some other mechanism will be needed, at the very least to make it easy to find and decode I-frames.

Finally, encryption is an issue. When movies are multicast (e.g., if the local distribution network is a cable TV system), encryption is needed to ensure that only paying customers can watch movies. Two approaches are possible: pre-encryption and encryption on the fly. If movies are stored encrypted, then anyone learning a movie's key may be able to watch it for free because the same key is used every time. Separate encryption for each stream is more secure, but also more costly of computing resources.

Key management is also an issue. The usual approach is to encrypt on the fly with a simple algorithm, but change the key often, so even if an intruder can break the key in 10 minutes, it will be obsolete by then.

The Distribution Network

The distribution network is the set of switches and lines between the source and destination. As we saw in Fig. 7-88, it consists of a SONET or ATM (or ATM over SONET) backbone, connected to a local distribution network. Usually, the backbone is switched and the local distribution network is not.

The main requirements imposed on the backbone are high bandwidth and low jitter. For a pure SONET backbone, these are trivial to achieve—the bandwidth is guaranteed and the jitter is zero because the network is synchronous. For an ATM backbone, or ATM over SONET, the quality of service is very important. It is managed by the leaky bucket algorithm and all the other techniques we studied in great detail in Chap. 5, so we will not repeat that discussion here. For additional information about real-time MPEG over ATM backbones, see (Dixit and Skelly, 1995; and Morales et al., 1995). Below we will focus on the local distribution network, a topic we have barely touched upon so far.

Local distribution is highly chaotic, with different companies trying out different networks in different regions. Telephone companies, cable TV companies, and new entrants are all convinced that whoever gets there first will be the big winner, so we are now seeing a proliferation of technologies being installed. The four main local distribution schemes for video on demand go by the acronyms ADSL, FTTC, FTTH, and HFC. We will now explain each of these in turn.

ADSL (Asymmetric Digital Subscriber Line) was the telephone industry's first entrant in the local distribution sweepstakes (Chen and Waring, 1994). The idea is that virtually every house in the United States, Europe, and Japan already has a copper twisted pair going into it (for analog telephone service). If these wires could be used for video on demand, the telephone companies could clean up.

The problem, of course, is that these wires cannot support even MPEG-1 over their typical 10-km length, let alone MPEG-2. The ADSL solution takes advantage of advances in digital signal processing to eliminate echoes and other line noise electronically. As shown in Fig. 7-92, each ADSL subscriber is given an in-house ADSL subscriber unit containing a digital signal processing chip. The telephone and set-top box plug into the ADSL unit. At the other end of the local loop, another ADSL unit is attached. This one may either be in the telephone company end office, or, if the local loop is too long, at the end of an optical fiber in the neighborhood of the house.




Fig. 7-92. ADSL as the local distribution network.

ADSL-1 offers a 1.536-Mbps downlink channel (T1 minus the 193rd bit), but only a 16-kbps uplink channel. In addition, the old 4-kHz analog telephone channel (or in some cases, two N-ISDN digital channels) is also on there. The idea is that the uplink has enough bandwidth for the user to order movies, and the downlink has enough bandwidth to send them encoded in MPEG-1. ADSL should be regarded more as a quick-and-dirty hack than a long-term solution, but it is being installed in various cities. Improved versions, called ADSL-2 and ADSL-3 are also being worked on. The latter allows MPEG-2 over local loops of up to about 2 km.

The second telephone company design is **FTTC (Fiber To The Curb)**. We saw this design in Fig. 2-23(a). In FTTC, the telephone company runs optical fiber from the end office into each residential neighborhood, terminating in a device called an **ONU (Optical Network Unit)**. The ONU is labeled "junction

box" in Fig. 2-23(a). On the order of 16 copper local loops can terminate in an ONU. These loops are now so short that it is possible to run full-duplex T1 or T2 over them, allowing MPEG-1 and MPEG-2 movies, respectively. In addition, videoconferencing for home workers and small businesses is now possible because FTTC is symmetric.

The third telephone company solution is to run fiber into everyone's house. It is called **FTTH (Fiber To The Home)**. In this scheme, everyone can have an OC-1, OC-3, or even higher carrier if that is required. FTTH is very expensive and will not happen for years but clearly will open a vast range of new possibilities when it finally happens.

ADSL, FTTC, and FTTH are all point-to-point local distribution networks, which is not surprising given how the current telephone system is organized. A completely different approach is **HFC (Hybrid Fiber Coax)**, which is the preferred solution currently being installed by cable TV providers. It is illustrated in Fig. 2-23(b). The story goes something like this. The current 300- to 450-MHz coax cables will be replaced by 750-MHz coax cables, upgrading the capacity from 50 to 75 6-MHz channels to 125 6-MHz channels. Seventy-five of the 125 channels will be used for transmitting analog television.

The 50 new channels will each be modulated using QAM-256, which provides about 40 Mbps per channel, giving a total of 2 Gbps of new bandwidth. The head-ends will be moved deeper into the neighborhoods, so each cable runs past only 500 houses. Simple division show that each house can then be allocated a dedicated 4 Mbps channel, which can be used for some combination of MPEG-1 programs, MPEG-2 programs, upstream data, analog and digital telephony, and so on.

While this sounds wonderful, it does require the cable providers to replace all the existing cables with 750 MHz coax, install new head-ends, and remove all the one-way amplifiers—in short, replace the entire cable TV system. Consequently, the amount of new infrastructure here is comparable to what the telephone companies need for FTTC. In both cases the local network provider has to run fiber into residential neighborhoods. Again, in both cases, the fiber terminates at an optoelectrical converter. In FTTC, the final segment is a point-to-point local loop using twisted pairs. In HFC, the final segment is a shared coaxial cable. Technically, these two systems are not really as different as their respective proponents often make out.

Nevertheless, there is one real difference that is worth pointing out. HFC uses a shared medium without switching and routing. Any information put onto the cable can be removed by any subscriber without further ado. FTTC, which is fully switched, does not have this property. As a result, HFC operators want video servers to send out encrypted streams, so customers who have not paid for a movie cannot see it. FTTC operators do not especially want encryption because it adds complexity, lowers performance, and provides no additional security in their system. From the point of view of the company running a video server, is it a

good idea to encrypt or not? A server operated by a telephone company or one of its subsidiaries or partners might intentionally decide not to encrypt its videos, claiming efficiency as the reason but really to cause economic losses to its HFC competitors.

For all these local distribution networks, it is likely that each neighborhood will be outfitted with one or more spooling servers. These are, in fact, just smaller versions of the video servers we discussed above. The big advantage of these local servers is that since the local distribution networks are short and generally not switched, they do not introduce jitter as an ATM backbone network would.

They can be preloaded with movies either dynamically or by reservation. For example, when a user selects a movie, the first minute could be transmitted to the local server in under 2 seconds at OC-3. After 55 seconds, the next minute could be shipped to the local server in 2 seconds, and so on. In this way, the traffic over the ATM backbone no longer has to be jitter free, making it possible to use ABR service instead of the more expensive CBR service.

If people tell the provider which movies they want well in advance, they can be downloaded to the local server during off-peak hours, giving even bigger savings. This observation is likely to lead the network operators to lure away airline executives to do their pricing. One can envision tariffs in which movies ordered 24 to 72 hours in advance for viewing on a Tuesday or Thursday evening before 6 P.M. or after 11 P.M. get a 27 percent discount. Movies ordered on the first Sunday of the month before 8 A.M. for viewing on a Wednesday afternoon on a day whose date is a prime number get a 43 percent discount, and so on.

The choice of the protocol stack to use for video on demand is still up in the air. ATM is clearly the technology of choice, but which adaptation protocol should be used? AAL 1 was designed for video, so it is a strong candidate, but it corresponds to the CBR service category. Dedicating the maximum possible bandwidth needed is expensive, especially since MPEG is inherently VBR traffic so the virtual circuit will have to be overdimensioned.

AAL 2 is not finished (and probably never will be) and AAL 3/4 is too clumsy, so AAL 5 is the only remaining contender. It is not tied to CBR service, and sending a large block of MPEG in each message would be extremely efficient, getting nearly 100 percent of the user bandwidth for the video stream. On the downside, AAL 5 does error detection. Having an entire block discarded due to a 1-bit error is highly unattractive, especially since most errors are single bit errors in the middle of the data. As a consequence, there is some movement toward changing AAL 5 to allow applications to ask for all blocks, along with a bit telling whether or not the checksum was correct.

The video on demand protocol stack we have sketched above is illustrated in Fig. 7-93. Above the AAL layer, we see the MPEG program and transport stream layer. Then come the encoding and decoding of MPEG audio and video, respectively. Finally, we have the application on top.




Fig. 7-93. A video-on-demand protocol stack.

Set-Top Boxes

All of the above local distribution methods ultimately bring one or more MPEG streams into the home. To decode and view them, a network interface, MPEG decoder, and other electronic components are needed. Two approaches are possible here.

In approach one, people use their personal computers for decoding and viewing movies. Doing this requires buying a special plug-in board containing a few special chips and a connector for interfacing to the local distribution network. The movies then appear on the computer's monitor, possibly even in a window. One might call this the set-bottom box since with computers, the box is usually under the monitor instead of on top of it. This approach is cheaper (all that is needed is one plug-in board and the software), uses a high-resolution noninterlaced monitor, has a sophisticated mouse-oriented user interface, and can easily be integrated with the WWW and other computer-oriented information and entertainment sources. On the other hand, PCs usually have small screens, are located in studies or dens rather than in living rooms, and are traditionally used by one person at a time. They also emit significantly less light than television sets.

In approach two, the local network operator rents or sells the user a **set-top box** to which the network and television set are connected. This approach has the advantage that everyone has a television but not everyone has a PC, and many of the PCs that people do have are old, peculiar, or otherwise unsuited to MPEG decoding. Furthermore, the television is often located in a room intended for group viewing.

On the down side, the monitor has a low-resolution interlaced display (making it unsuited for text-oriented material, such as the WWW). In addition, it has a dreadful user interface (the remote control), making it virtually impossible for the user to do anything except select items from simple menus. Even typing in the name of a movie is painful, let alone engaging in a dialog asking the server to

search for all the films made by a certain actor, director, or production company during a certain time period. Finally, set-top boxes with the required performance are not easy to produce for an acceptable price (thought to be a few hundred dollars).

All these factors considered, most video-on-demand systems have opted for the set-top box model, primarily because mass marketers hate to exclude any potential customers (people without a PC). Also, there may be money to be made renting or selling set-top boxes. Nevertheless, the PC plug-in board market is large enough so no doubt these boards will be produced, too.

The primary functions of the set-top box are interfacing with the local distribution network, decoding the MPEG signal, synchronizing the audio and video streams, producing a composite NTSC, PAL, or SECAM signal for the television set, listening to the remote control, and handling the user interface. Additional functions might include interfacing with stereos, telephones, and other devices. A major battle is raging within the industry about how much functionality should be put in the set-top box and how much should be in the network. How that turns out remains to be seen.

A possible architecture for a simple set-top box is shown in Fig. 7-94. The device consists of a CPU, ROM, RAM, I/O controller, MPEG decoder, and network interface. Optionally, a security chip can also be added for decryption of incoming movies and encryption of outgoing messages (credit card numbers for home shopping, etc.).




Fig. 7-94. The hardware architecture of a simple set-top box.

An important issue for video on demand is audio/video synchronization and jitter management. Adding an additional 500 KB of RAM allows for 1 second of MPEG-2 buffering, but at additional expense in a device that the manufacturers are hoping to sell for a few hundred dollars, at most.

Since the set-top box is just a computer, it will need software, probably a microkernel-based real-time operating system kept in the ROM. To provide flexibility and adaptability, it is probably a good idea to make it possible to download other software from the network. This possibility then raises the problem of what happens when the owner of a MIPS-based set-top box wants to play a game written for a SPARC-based set-top box? Using an interpreted language such as Java solves the compatibility problem but severely lowers performance in a real-time environment in which high performance is crucial.

Standards

The economics of video on demand cannot be ignored. A large video server can easily cost more than a mainframe, certainly 10 million dollars. Suppose that it serves 100,000 homes, each of which has rented a 300-dollar set-top box. Now throw in 10 million dollars worth of networking equipment and a 4-year depreciation period, and the system has to generate 10 dollars per home per month. At 5 dollars/movie, everyone has to buy two movies a month for the operator to break even (excluding salaries, marketing, and all other costs). Whether this will actually happen is far from obvious.

The numbers given above can be rearranged in many ways (e.g., charging 6 dollars per month rental for the set-top box and 2 dollars per movie), and the costs are changing all the time, but it should be clear that without a mass market, there is no way that video on demand makes economic sense. For a mass market to develop, it is essential that all parts of the system be standardized. If each video provider, network operator, and set-top box manufacturer designs its own interface, nothing will interwork with the rest of the system. So far, the only standard that everyone agrees on is the use of MPEG-2 for video encoding. Everything else is up for grabs. A few of the many questions that have to be answered before a national system can be built are listed in Fig. 7-95.

If all these areas can be standardized, we can easily imagine many vendors producing products consisting of a box with a telephone jack, monitor, keyboard, and mouse that can be used for watching videos, computing, or maybe doing both at once. The much-discussed convergence of the computing, communication, and entertainment industries will then be a reality.

7.7.5. MBone—Multicast Backbone

While all these industries are making great—and highly publicized—plans for future (inter)national digital video on demand, the Internet community has been quietly implementing its own digital multimedia system, **MBone (Multicast Backbone)**. In this section we will give a brief overview of what it is and how it works. For an entire book on MBone, see (Kumar, 1996). For articles on MBone, see (Eriksson, 1994; and Macedonia and Brutzman, 1994).

What technology will the backbone use (SONET, ATM, SONET + ATM)?
What speed will the backbone run at (OC-3, OC-12)?
How will local distribution be done (HFC, FTTC)?
How much upstream bandwidth will there be (16 kbps, 1.5 Mbps)?
Will movies be encrypted, and if so, how?
Will error correction be present (mandatory, optional, absent)?
Who will own the set-top box (user, network operator)?
Will telephony be part of the system (analog, N-ISDN)?
Will high-resolution hypertext applications be supported (e.g., WWW)?

Fig. 7-95. A few areas in which standards are needed.

MBone can be thought of as Internet radio and television. Unlike video on demand, where the emphasis is on calling up and viewing precompressed movies stored on a server, MBone is used for broadcasting live audio and video in digital form all over the world via the Internet. It has been operational since early 1992. Many scientific conferences, especially IETF meetings, have been broadcast, as well as newsworthy scientific events, such as space shuttle launches. A Rolling Stones concert was once broadcast over MBone. Whether this qualifies as a newsworthy scientific event is arguable. For people who want to digitally record an MBone broadcast, software for accomplishing that is also available (Holfelder, 1995).

Most of the research concerning MBone has been about how to do multicasting efficiently over the (datagram-oriented) Internet. Little has been done on audio or video encoding. MBone sources are free to experiment with MPEG or any other encoding technology they wish. There are no Internet standards on content or encoding.

Technically, MBone is a virtual overlay network on top of the Internet. It consists of multicast-capable islands connected by tunnels, as shown in Fig. 7-96. In this figure, MBone consists of six islands, *A* through *F*, connected by seven tunnels. Each island (typically a LAN or group of interconnected LANs) supports hardware multicast to its hosts. The tunnels propagate MBone packets between the islands. Some day in the future, when all the routers are capable of handling multicast traffic directly, this superstructure will no longer be needed, but for the moment, it does the job.

Each island contains one or more special routers called **mrouter**s (**multicast routers**). Some of these are actually normal routers, but most are just UNIX workstations running special user-level software (but as the root). The mrouter are logically connected by tunnels. In the past, MBone packets were tunneled from mrouter to mrouter (usually through one or more routers that did not know




Fig. 7-96. MBone consists of multicast islands connected by tunnels.

about MBone) using loose source routing. Nowadays, MBone packets are encapsulated within IP packets and sent as regular unicast packets to the destination mrouter's IP address. If all the intervening routers support multicast, however, tunneling is not needed.

Tunnels are configured manually. Usually, a tunnel runs above a path for which a physical connection exists, but this is not a requirement. If, by accident, the physical path underlying a tunnel goes down, the mrouter using the tunnel will not even notice it, since the Internet will automatically reroute all the IP traffic between them via other lines.

When a new island appears and wishes to join MBone, such as *G* in Fig. 7-96, its administrator sends a message announcing its existence to the MBone mailing list. The administrators of nearby sites then contact him to arrange to set up tunnels. Sometimes existing tunnels are reshuffled to take advantage of the new island to optimize the topology. After all, tunnels have no physical existence. They are defined by tables in the mrouter and can be added, deleted, or moved simply by changing these tables. Typically, each country on MBone has a backbone, with regional islands attached to it. Normally, MBone is configured with one or two tunnels crossing the Atlantic and Pacific oceans, making MBone global in scale.

Thus at any instant, MBone consists of a specific topology consisting of islands and tunnels, independent of the number of multicast addresses currently in use and who is listening to them or watching them. This situation is very similar to a normal (physical) subnet, so the normal routing algorithms apply to it. Consequently, MBone initially used a routing algorithm, **DVMRP (Distance Vector Multicast Routing Protocol)** based on the Bellman-Ford distance vector

algorithm. For example, in Fig. 7-96, island *C* can route to *A* either via *B* or via *E* (or conceivably via *D*). It makes its choice by taking the values those nodes give it about their respective distances to *A* and then adding its distance to them. In this way, every island determines the best route to every other island. The routes are not actually used in this way, however, as we will see shortly.

Now let us consider how multicasting actually happens. To multicast an audio or video program, a source must first acquire a class D multicast address, which acts like a station frequency or channel number. Class D addresses are reserved by using a program that looks in a database for free multicast addresses. Many multicasts may be going on at once, and a host can “tune” to the one it is interested in by listening to the appropriate multicast address.

Periodically, each mrouter sends out an IGMP broadcast packet limited to its island asking who is interested in which channel. Hosts wishing to (continue to) receive one or more channels send another IGMP packet back in response. These responses are staggered in time, to avoid overloading the local LAN. Each mrouter keeps a table of which channels it must put out onto its LAN, to avoid wasting bandwidth by multicasting channels that nobody wants.

Multicasts propagate through Mbone as follows. When an audio or video source generates a new packet, it multicasts it to its local island using the hardware multicast facility. This packet is picked up by the local mrouter, which then copies it into all the tunnels to which it is connected.

Each mrouter getting such a packet via a tunnel then checks to see if the packet came along the best route, that is, the route that its table says to use to reach the source (as if it were a destination). If the packet came along the best route, the mrouter copies the packet to all its other tunnels. If the packet arrived via a suboptimal route, it is discarded. Thus, for example, in Fig. 7-96, if *C*'s tables tell it to use *B* to get to *A*, then when a multicast packet from *A* reaches *C* via *B*, the packet is copied to *D* and *E*. However, when a multicast packet from *A* reaches *C* via *E* (not the best path), it is simply discarded. This algorithm is just the reverse path forwarding algorithm that we saw in Chap. 5. While not perfect, it is fairly good and very simple to implement.

In addition to using reverse path forwarding to prevent flooding the Internet, the IP *Time to live* field is also used to limit the scope of multicasting. Each packet starts out with some value (determined by the source). Each tunnel is assigned a weight. A packet is only passed through a tunnel if it has enough weight. Otherwise it is discarded. For example, transoceanic tunnels are normally configured with a weight of 128, so packets can be limited to the continent of origin by giving them an initial *Time to live* of 127 or less. After passing through a tunnel, the *Time to live* field is decremented by the tunnel's weight.

While the Mbone routing algorithm works, much research has been devoted to improving it. One proposal keeps the idea of distance vector routing, but makes the algorithm hierarchical by grouping Mbone sites into regions and first routing to them (Thyagarajan and Deering, 1995).

Another proposal is to use a modified form of link state routing instead of distance vector routing. In particular, an IETF working group is busy modifying OSPF to make it suitable for multicasting within a single autonomous system. The resulting multicast OSPF is called **MOSPF** (Moy, 1994). What the modifications do is have the full map built by MOSPF keep track of multicast islands and tunnels, in addition to the usual routing information. Armed with the complete topology, it is straightforward to compute the best path from every island to every other island using the tunnels. Dijkstra's algorithm can be used, for example.

A second area of research is inter-AS routing. Here an algorithm called **PIM (Protocol Independent Multicast)** is being developed by another IETF working group (Huitema, 1995). PIM comes in two versions, depending on whether the islands are dense (almost everyone wants to watch) or sparse (almost nobody wants to watch). Both versions use the standard unicast routing tables, instead of creating an overlay topology as DVMRP and MOSPF do.

In PIM-dense, the idea is to prune useless paths. Pruning works as follows. When a multicast packet arrives via the "wrong" tunnel, a prune packet is sent back through the tunnel telling the sender to stop sending it packets from the source in question. When a packet arrives via the "right" tunnel, it is copied to all the other tunnels that have not previously pruned themselves. If all the other tunnels have pruned themselves and there is no interest in the channel within the local island, the mrouter sends a prune message back through the "right" channel. In this way, the multicast adapts automatically and only goes where it is wanted.

PIM-sparse works differently. The idea here is to prevent saturating the Internet because three people in Berkeley want to hold a conference call over a class D address. PIM-sparse works by setting up rendezvous points. Each of the sources in a PIM-sparse multicast group send their packets to the rendezvous points. Any site interested in joining up asks one of the rendezvous points to set up a tunnel to it. In this way, all PIM-sparse traffic is transported by unicast instead of by multicast.

All in all, multimedia is an exciting and rapidly moving field. New technologies and applications are announced daily, but the area as a whole is likely to remain important for decades to come.

7.8. SUMMARY

Computer networks are inherently insecure. To keep information secret, it must be encrypted. Encryption protocols fall into two general classes: secret key (e.g., DES, IDEA), and public key (e.g., RSA). Using these protocols is straightforward; the hard part is key management.

In addition to providing secrecy, cryptographic protocols can also provide authentication, so that when Alice thinks she is communicating with Bob, she really is communicating with Bob, and not with Trudy. Finally, cryptography can

also be used to allow messages to be signed in such a way that the sender cannot repudiate them after they have been sent.

Naming in the Internet uses a distributed database system, DNS. DNS holds records with IP addresses, mail exchanges, and other information. By querying a DNS server, a process can map an Internet domain name onto the IP address used to communicate with that domain.

As networks grow larger, they become harder to manage. For this reason, special network management systems and protocols have been devised, the most popular of which is SNMP. This protocol allows managers to communicate with agents inside devices to read out their status and issue commands to them.

Four major network applications are electronic mail, USENET news, the World Wide Web, and multimedia (video on demand and Mbone). Most email systems use the mail system defined in RFCs 821 and 822. Messages sent in this system use system ASCII headers to define message properties. These messages are sent using SMTP. Two systems for securing email exist, PGP and PEM.

USENET news consists of thousands of newsgroups on all manner of topics. People can join newsgroups locally, and can then post messages all over the world using the NNTP protocol, which has some resemblance to SMTP.

The World Wide Web is a system for linking up hypertext documents. Each document is a page written in HTML, possible with hyperlinks to other documents. A browser can display a document by establishing a TCP connection to its server, asking for the document, and then closing the connection. When a hyperlink is selected by the user, that document can also be fetched in the same way. In this manner, documents all over the world are linked together in a giant web.

Multimedia is the rising star in the networking firmament. It allows audio and video to be digitized and transported electronically for display. Most multimedia projects use the MPEG standards and transmit the data over ATM connections. The Mbone is an experimental worldwide digital radio and television service on the Internet.

PROBLEMS

1. Break the following monoalphabetic cipher. The plaintext, consisting of letters only, is a well-known excerpt from a poem by Lewis Carroll.

```
kfd ktbd fzm eubd kfd pzyiom mztz ku kzyg ur bzha kfthem
ur mfudm zhx mftnm zhx mdzythe pzq ur ezsscdm zhx gthem
zhx pfa kfd mdz tm sutythe fuk zhx pfdkfdi ntem fzld pthem
sok pztz z stk kfd uamkdim eitdx sdruidd pd fzld uoi efzk
rui mubd ur om zid uok ur sidzfk zhx zyy ur om zid rzk
hu foia mztz kfd ezindhkdi kfda kfzhdgx ftb boef rui kfzk
```

2. Break the following columnar transposition cipher. The plaintext is taken from a popular computer textbook, so “computer” is a probable word. The plaintext consists entirely of letters (no spaces). The ciphertext is broken up into blocks of five characters for readability.
 aauan cvlre rurnn dltnm aeepb ytust iceat nrmey iicgo gorch srsoc
 nntii imiha oofpa gsvit tpsit lbolr otoex
3. In Fig. 7-4, the P-boxes and S-boxes alternate. Although this arrangement is esthetically pleasing, is it any more secure than first having all the P-boxes and then all the S-boxes?
4. Suppose that a message has been encrypted using DES in ciphertext block chaining mode. One bit of ciphertext in block C_i is accidentally transformed from a 0 to a 1 during transmission. How much plaintext will be garbled as a result?
5. Now consider ciphertext block chaining again. Instead of a single 0 bit being transformed into a 1 bit, an extra 0 bit is inserted into the ciphertext stream after block C_i . How much plaintext will be garbled as a result?
6. Design an attack on DES based on the knowledge that the plaintext consists exclusively of uppercase ASCII letters, plus space, comma, period, semicolon, carriage return, and line feed. Nothing is known about the plaintext parity bits.
7. Compare cipher block chaining with cipher feedback mode in terms of the number of encryption operations needed to transmit a large file. Which one is more efficient and by how much?
8. Using the RSA public key cryptosystem, with $a = 1$, $b = 2$, etc.,
 - (a) If $p = 7$ and $q = 11$, list five legal values for d .
 - (b) If $p = 13$, $q = 31$ and $d = 7$, find e .
 - (c) Using $p = 5$, $q = 11$, and $d = 27$, find e and encrypt “abcdefghij”
9. The Diffie-Hellman key exchange is being used to establish a secret key between Alice and Bob. Alice sends Bob (719, 3, 191). Bob responds with (543). Alice’s secret number, x , is 16. What is the secret key?
10. Change one message in protocol of Fig. 7-14 in a minor way to make it resistant to the reflection attack. Explain why your change works.
11. In the protocol of Fig. 7-17, why is A sent in plaintext along with the encrypted session key?
12. In the protocol of Fig. 7-17, we pointed out that starting each plaintext message with 32 zero bits is a security risk. Suppose that each message begins with a per-user random number, effectively a second secret key known only to its user and the KDC. Does this eliminate the known plaintext attack?
13. In the Needham-Schroeder protocol, Alice generates two challenges, R_A and R_{A2} . This seems like overkill. Would one not have done the job?
14. In the public-key authentication protocol of Fig. 7-21, in message 3, R_B is encrypted with K_S . Is this encryption necessary, or would it have been adequate to send it back in plaintext?

15. The signature protocol of Fig. 7-22 has the following weakness. If Bob crashes, he may lose the contents of his RAM. What problems does this cause and what can he do to prevent them?
16. After Ellen confessed to Marilyn about tricking her in the matter of Tom's tenure, Marilyn resolved to avoid this problem by dictating the contents of future messages into a dictating machine and having her new secretary just type them in. Marilyn then planned to examine the messages on her terminal after they have been typed in to make sure they contain her exact words. Can the new secretary still use the birthday attack to falsify a message, and if so, how? *Hint:* She can.
17. Point-of-sale terminals that use magnetic-stripe cards and PIN codes have a fatal flaw: a malicious merchant can modify his card reader to capture and store all the information on the card as well as the PIN code in order to post additional (fake) transactions in the future. The next generation of point-of-sale terminals will use cards with a complete CPU, keyboard, and tiny display on the card. Devise a protocol for this system that malicious merchants cannot break.
18. According to the information given in Fig. 7-27, is *little-sister.cs.vu.nl* on a class A, B, or C network?
19. In Fig. 7-27, there is no period after *rowboat*? Why not?
20. What is the *OBJECT IDENTIFIER* for the tcp object?
21. An SNMP integer whose value is 200 has to be transmitted. Show the binary representation of the bits sent in the ASN.1 transfer syntax.
22. What is the representation of the 11-bit binary bit string '11100001111' in the ASN.1 transfer syntax?
23. Suppose that you are hired by a bridge vendor to write SNMP-conformant code for one of their bridges. You read all the RFCs and still have questions. You suggest to IAB that a complete, formal grammar of the language used to describe SNMP variables be given in one place. IAB's reaction is to agree and appoint you to do the job. Should the grammar be added to RFC 1442 or RFC 1213? Why? *Hint:* You do not need to fetch the RFCs; enough information is given in the text.
24. Some email systems support a header field *Content Return:*. It specifies whether the body of a message is to be returned in the event of nondelivery. Does this field belong to the envelope or to the header?
25. Electronic mail systems need directories so people's email addresses can be looked up. To build such directories, names should be broken up into standard components (e.g., first name, last name) to make searching possible. Discuss some problems that must be solved for a worldwide standard to be acceptable.
26. A binary file is 3072 bytes long. How long will it be if encoded using base64 encoding, with a CR+LF pair inserted after every 80 bytes sent and at the end?
27. Consider the quoted-printable MIME encoding scheme. Mention a problem not discussed in the text and propose a solution.
28. Give *two* reasons why PGP compresses messages.

29. Suppose that someone sets up a vacation daemon and then sends a message just before logging out. Unfortunately, the recipient has been on vacation for a week and also has a vacation daemon in place. What happens next? Will canned replies go back and forth until somebody returns?
30. Assuming that everyone on the Internet used PGP, could a PGP message be sent to an arbitrary Internet address and be decoded correctly by all concerned? Discuss your answer.
31. PGP does not support canonicalization as does PEM. Why not?
32. Make a guess about what the smiley :-X (sometimes written as :-#) might mean.
33. How long does it take to distribute a days' worth of news over a 50-Mbps satellite channel?
34. Which of the commands listed in Fig. 7-56 are theoretically redundant?
35. A large network consists of an $n \times n$ grid of machines. All the interior nodes have four neighbors; the ones on the edges (corners) have three (two) neighbors. If an m -byte article is posted on some machine using NNTP, how many bytes of bandwidth are consumed getting it to all other machines (ignoring the NNTP overhead and just counting the message bytes)?
36. Repeat the previous problem, but now compute the approximate bandwidth that would be needed to distribute the message using a mailing list. How much more is it than in the previous problem?
37. When Web pages are sent out, they are prefixed by MIME headers. Why?
38. When are external viewers needed? How does a browser know which one to use?
39. Imagine that someone in the CS Department at Stanford has just written a new program that he wants to distribute by FTP. He puts the program in the FTP directory *ftp/pub/freebies/newprog.c*. What is the URL for this program likely to be?
40. In Fig. 7-60, the *ALT* parameter is set in the tag. Under what conditions does the browser use it, and how?
41. How do you make an image clickable in HTML? Given an example.
42. Show the <A> tag that is needed to make the string "ACM" be a hyperlink to *http://www.acm.org*.
43. Design a form for a new company, Interburger, that allows hamburgers to be ordered via the Internet. The form should include the customer's name, address, and city, as well as a choice of size (either gigantic or immense) and a cheese option. The burgers are to be paid for in cash upon delivery, so no credit card information is needed.
44. Java does not have structures as in C or records as in Pascal. Is there some other way to achieve the same effect of bundling a group of dissimilar variables together to form a single data type? If so, what is it?
45. Using the data structures of Fig. 7-75, list the exact steps needed to check a new URL to see if it is already in *url_table*.

46. Suppose that in its effort to become more market oriented, the KGB goes commercial and hires an advertising agency that designs a Web page for it. Your company has been hired as an outside consultant to implement it. Write the HTML to produce the Web page below.

WELCOME TO THE KGB'S WWW HOME PAGE

As a consequence of its recent privatization, the KGB is pleased to announce the commercial availability of many fine products and services previously available only to major governments.

Competitive prices! Discreet service ensured!

- Products
 - Nuclear weapons (small, medium, large, jumbo)
 - Spy satellites (keep tabs on your neighbors)
 - Low-radar-profile supersonic aircraft (buzz your friends' houses unseen)
- Services
 - Mole placement in the organization of your choice
 - Coups (corporate as well as governmental)
 - Assistance in setting up your very own germ-warfare laboratory
- Bargain basement specials
 - The collected works of Felix Dzerzhinsky (limited edition)
 - Aerial photographs of Afghanistan (ca. 1984)
 - Quality Bulgarian-made tanks (95 percent discount)

Webmaster@kgb.ru

47. In C and C++, the size of an integer is not specified by the language. In Java it is. Give an argument for the C way and one for the Java way.
48. Suppose that the Web contains 10 million pages, each with an average of 10 hyperlinks. Fetching a page averages 100 msec. What is the minimum time to index the entire Web?
49. A compact disc holds 650 MB of data. Is compression used for audio CDs? Explain your reasoning.
50. What is the bit rate for transmitting uncompressed VGA color with 8 bits/pixel at 40 frames/sec?
51. In Fig. 7-76(c) quantization noise occurs due to the use of 3-bit samples. The first sample, at 0, is exact, but the next few are not. What is the percent error for the samples at $1/32$, $2/32$, and $3/32$ of the period?
52. Can a 1-bit error in an MPEG frame affect more than the frame in which the error occurs? Explain your answer.
53. Consider the 100,000 customer video server example given in the text. Suppose that half of all movies are served from 8 P.M. to 10 P.M. How many movies does the server have to transmit at once during this time period? If each movie requires 4 Mbps, how many OC-12 connections does the server need to the network?

54. Suppose that Zipf's law holds for accesses to a 10,000-movie video server. If the server holds the most popular 1000 movies on magnetic disk and the remaining 9000 on optical disk, give an expression for the fraction of all references that will be to magnetic disk. Write a little program to evaluate this expression numerically.
55. MPEG PES packets contain a field giving the copyright status of the current transmission. Of what conceivable use is this field?

8

READING LIST AND BIBLIOGRAPHY

We have now finished our study of computer networks, but this is only the beginning. Many interesting topics have not been treated in as much detail as they deserve, and others have been omitted altogether for lack of space. In this chapter we provide some suggestions for further reading and a bibliography, for the benefit of readers who wish to continue their study of computer networks.

8.1. SUGGESTIONS FOR FURTHER READING

There is an extensive literature on all aspects of computer networks and distributed systems. Four journals that frequently publish papers in this area are *IEEE Transactions on Communications*, *IEEE Journal on Selected Areas in Communications*, *Computer Communication Review*, and *Computer Networks and ISDN Systems*. Many other journals also publish occasional papers on the subject.

IEEE also publishes two magazines, *IEEE Network Magazine* and *IEEE Communications Magazine*, that contain surveys, tutorials, and case studies on networking. The former emphasizes architecture, standards, and software, and the latter tends toward communications technology (fiber optics, satellites, and so on).

In addition, there are several annual or biannual conferences that often attract many papers on networks and distributed systems, in particular, *SIGCOMM '9x*, *The International Conference on Distributed Computer Systems*, *The Symposium on Operating Systems Principles* and *The N-th Data Communications Symposium*.

Furthermore, IEEE has published several volumes of network paper reprints in convenient paperback form.

Below we list some suggestions for supplementary reading, keyed to the chapters of this book.

8.1.1. Introduction and General Works

Bell, "Communications"

For an excellent overview of trends in communication, including telephone, ATM, ISDN, wireless LANs, the Internet, and pagers, this article is a must.

Comer, *The Internet Book*

Anyone looking for an easy-going introduction to the Internet should look here. Comer describes the history, growth, technology, protocols, and services of the Internet in terms that novices can understand, but so much material is covered that the book is also of interest to more technical readers as well.

Jabbari et al., "Network Issues for Wireless Communication"

This introduction to cellular radio systems covers call control, routing, signaling, and other aspects of modern mobile communication systems.

Kwok, "A Vision for Residential Broadband Service"

If you want to know how Microsoft thinks video on demand should be organized, this article is for you. In it, Microsoft's chief ATM architect explains his company's vision. Briefly summarized, Microsoft's idea is: ATM to the home is the way to go. Forget all the "realistic" (i.e., ad hoc) solutions, like ADSL and do it right.

Le Boudec, "The Asynchronous Transfer Mode: A tutorial"

ATM is an up-and-coming technology, and this paper gives a thorough introduction to it. The physical layer, ATM layer, and AAL layer are all covered. In addition, the final section discusses the debate about ATM.

Pahlavan et al., "Trends in Local Wireless Networks"

Wireless LANs will no doubt become increasingly important in the future. In this paper, the authors discuss the state of the art and trends in spectrum use and technologies for wireless LANs.

Siu and Jain, "A Brief Overview of ATM"

Many features of ATM systems are covered in this introductory paper, but the focus is on LAN emulation and traffic management. It also serves as the introduction to a special issue of *Computer Communication Review* devoted to ATM technology.

8.1.2. The Physical Layer

Awdeh and Mouftah, "Survey of ATM Switch Architectures"

Anyone interested in learning more about ATM switch design should look here. After introducing switches in general and buffering strategies, the authors discuss many kinds of crossbar, disjoint-path, and banyan switches. The paper also provides over 200 references to other papers.

Bellamy, *Digital Telephony*

Everything you ever wanted to know about the telephone system and more is contained in this authoritative book. Particularly interesting are the chapters on transmission and multiplexing, digital switching, fiber optics, and ISDN.

De Prycker, *Asynchronous Transfer Mode*, 2nd ed.

Chapter 4 contains a wealth of information on ATM switches. The principles are illustrated by numerous example switches, including the knockout, Roxanne, Coprin, and Athena switches.

Held, *The Complete Modem Reference*, 2nd ed.

Everything you might conceivably want to know about modems is here, from the U.S. and Canadian governments' compliance rules, through modulation techniques and standards, to how to troubleshoot a sick modem.

IEEE Communications Mag., Jan. 1995, "Wireless Personal Communications"

This special issue contains seven papers on different aspects of wireless personal communication. Collectively they cover propagation, access methods, receiver principles, system aspects, and network issues.

Metcalf, "Computer/Network Interface Design: Lessons from Arpanet & Ethernet"

Although engineers have been building network interfaces for decades now, one often wonders if they have learned anything from all this experience. In this paper, the designer of the Ethernet tells how to build a network interface, and what to do with it once you have built it. He pulls no punches, telling what he did wrong as well as what he did right.

Padgett et al., "Overview of Wireless Personal Communications"

An introduction to cellular and cordless communication systems and a comparison of the two. Both the American and European standards are covered.

Palais, *Fiber Optic Communication*, 3rd ed.

Books on fiber optic technology tend to be aimed at the specialist, but this one is more accessible than most. It covers waveguides, light sources, light detectors, couplers, modulation, noise, and many other topics.

Pandya, "Emerging Mobile and Personal Communications Systems"

For a short and sweet introduction to hand-held personal communication systems, this article is worth looking at. One of the nine pages contains a list of 70 acronyms used on the other eight pages.

Partridge, *Gigabit Networking*

In addition to describing several kinds of ATM switches, Chap. 5 also compares input buffering and output buffering and derives formulas for the performance of each.

Spragins et al., *Telecommunications Protocols and Design*

Chapter 2 contains a good introduction to transmission technology, including copper wires, fiber optics, cellular radio, and satellites. It also has extended discussions of the Nyquist and Shannon limits and their implications.

8.1.3. The Data Link Layer

Black, *Data Link Protocols*

Here is an entire book on the data link layer. It has a practical emphasis, with a large amount of material on HDLC, LLC, PPP, and other commercially important protocols.

Holzmann, *Design and Validation of Computer Protocols*

Readers interested in the more formal aspects of data link (and similar) protocols should look here. The specification, modeling, correctness, and testing of such protocols are all covered in this book.

Spragins et al., *Telecommunications Protocols and Design*

Readers interested in learning more about error-detecting and error-correcting codes should look at Chap. 6 of this book. It also covers the principles of elementary data link protocols at about the same level as this book does. Chapter 7 continues the discussion and discusses various data link protocols in detail.

Walrand, *Communication Networks: A First Course*

Chapter 4 covers data link protocols, with an emphasis on performance analysis. The finite state machine and Petri net approaches to protocol correctness are also treated.

8.1.4. The Medium Access Control Sublayer

Abeyundara and Kamal, "High-Speed Local Area Networks and Their Performance"

Since high-speed LANs are of interest due to their high speed, a paper

discussing and analyzing the performance is welcome. In this one, the focus is on different kinds of bus, ring, tree, and star LANs, and their delay and utilization characteristics.

Jain, *FDDI Handbook—High-Speed Networking Using Fiber and other Media*

For a thorough treatment of FDDI (including nice tutorials on fiber optics and SONET), this book is a good choice. In addition to long sections on FDDI hardware and software, it has a section on performance and even advice on shopping for fiber optic cables.

Perlman, *Interconnections: Bridges and Routers*

For an authoritative, but entertaining, treatment of bridges (and routers), Perlman's book is the place to look. The author designed the algorithms used in the IEEE 802 spanning tree bridge as well as the DECnet routing algorithms and is clearly an expert on the subject.

Stallings, *Local and Metropolitan Area Networks*, 4th ed.

The three IEEE 802 LANs form the core of this book, but material on other LANs and MANs is also present.

Walrand, *Communication Networks: A First Course*

Like Stallings book above, Chap. 5 of this one covers the basic 802 material, plus FDDI and DQDB. The emphasis is on analyzing protocol performance.

8.1.5. The Network Layer

Comer, *Internetworking with TCP/IP*, Vol. 1, 3rd ed.

Comer has written the definitive work on the TCP/IP protocol suite. Chapters 4 through 11 deal with IP and related protocols in the network layer. The other chapters deal primarily with the higher layers, and are also worth reading.

Huitema, *Routing in the Internet*

If you want to know everything there is to know about routing in the Internet, this is the book for you. Both pronounceable algorithms (e.g., RIP, CIDR, and MBONE) and unpronounceable algorithms (e.g., OSPF, IGRP, EGP, and BGP) are treated in great detail. New features, such as multicast, mobile IP, and resource reservation, are also here.

Perlman, *Interconnections: Bridges and Routers*

In Chap. 9, Perlman describes many of the issues involved in unicast and multicast routing algorithm design, both for WANs and networks of LANs, and their implementation in various devices. The author clearly cares about the subject, having entitled Sec. 9.13.10 "My Opinion on IP-Style Network Layer Multicast."

Sterbenz et al., "Report on the IEEE ComSoc Gigabit Networking Workshop"

Before gigabit networking is usable, a number of basic questions have to be resolved. A key one is whether these networks will use ATM, TCP/IP, or both. To better understand these issues, IEEE organized a workshop in April 1995, a summary of which is presented here. The critique of ATM by Schulzrinne is worth reading by anyone who believes that ATM is the solution to the world's telecommunication problems.

Stevens, *TCP/IP Illustrated*, Vol. 1

Chapters 3-10 provide a comprehensive treatment of IP and related protocols (ARP, RARP, and ICMP) illustrated by examples.

Yang and Reddy, "A Taxonomy for Congestion Control Algorithms in Packet Switching Networks"

The authors have devised a taxonomy for congestion control algorithms. The main categories are open loop with source control, open loop with destination control, closed loop with explicit feedback, and closed loop with implicit feedback. They use this taxonomy to describe and classify 23 existing algorithms.

8.1.6. The Transport Layer

Comer, *Internetworking with TCP/IP*, Vol. 1, 3rd ed.

As mentioned above, Comer has written the definitive work on the TCP/IP protocol suite. Chap. 12 is about UDP; Chap. 13 is about TCP.

Mogul, "IP Network Performance"

Despite the title of this article, it is at least, if not more, about TCP and network performance in general, than about IP performance in particular. It is full of useful guidelines and rules of thumb.

Stallings, *Data and Computer Communications*, 4th ed.

Chapter 12 is about transport protocols and covers services and mechanisms in the abstract, as well as the OSI and TCP transport protocols in detail.

Stevens, *TCP/IP Illustrated*, Vol. 1

Chapters 17-24 provide a comprehensive treatment of TCP illustrated by examples.

8.1.7. The Application Layer

Anderson, R., "Why Cryptosystems Fail"

According to Anderson, security in banking systems is poor, but not due to clever intruders breaking DES on their PCs. The real problems range from

dishonest employees (a bank clerk's changing a customer's mailing address to his own to intercept the bank card and PIN number) to programming errors (giving all customers the same PIN code). What is especially interesting is the response banks give when confronted with an error: our systems are perfect and therefore all errors must be due to customer errors or fraud.

Berghel, "The Client Side of the Web"

An easygoing introduction to Web browsers and the features they (can) support. The main topics are HTML/HTTP compliance, performance, reconfigurability, integration with the desktop, and navigational aids. Nine popular browsers are compared on these issues.

Berners-Lee et al., "The World Wide Web"

A perspective on the Web and where it is going by the person who invented it. The article focuses on the Web architecture, HTTP, and HTML, as well as future directions.

Carl-Mitchell and Quarterman, *Practical Internetworking with TCP/IP and UNIX*

Chapter 5 presents a nice introduction to naming and DNS, including naming authorities, the operational architecture, and the DNS database.

Choudbury et al., "Copyright Protection for Electronic Publishing on Computer Networks"

Although numerous books and articles describe cryptographic algorithms, few describe how they could be used to prevent users from further distributing documents which they are allowed to decrypt. This paper describes a variety of mechanisms that might help protect authors' copyrights in the electronic era.

Furht et al., "Design Issues for Interactive Television Systems"

Video on demand raises many complex technical issues related to the system architecture, network topology, server design, and set-top box design. In this paper, the authors present a tutorial on some of the key problems and some solutions that are being investigated.

Handley and Crowcroft, *The World Wide Web—Beneath the Surf*

While 99 percent of WWW books just tell you how to use some browser or list interesting URLs, this one explains how the Web works inside. The client side, the server side, and HTML are all explained in nice bite-sized chunks.

Kaufman et al., *Network Security*

This authoritative and frequently witty book is the first place to look for more information on network security. Secret and public key algorithms and protocols, message hashes, authentication, Kerberos, and email are all explained at length.

The best parts are the interauthor (and even intra-author) discussions, labeled by subscripts, as in: “ I_2 could not get me₁ to be very specific ...”

Kumar, *MBone: Interactive Multimedia on the Internet*

The cover of this book says: “Discover how you can broadcast, advertise, and display your products on the Internet.” Fortunately, this subject is not mentioned anywhere else in the book. What is covered is the architecture and implementation of the MBone, including a lot of material about how it works and how to use it.

Nemeth et al., *UNIX System Administration Handbook*

Chapter 16 is a long introduction to DNS. It gets into all the nitty-gritty details, illustrating the various files and resource records with numerous examples. Programs and other tools used for managing a DNS server are also covered in some detail.

Rose, *The Internet Message*

If you like your email served with a dash of iconoclasm, this book is a good bet. The author is not above getting up on a soapbox from time to time to announce what is wrong with the world. When you come right down to it, his taste is not bad.

Schneier, *Applied Cryptography*, 2nd ed.

This monumental compendium is NSA’s worst nightmare: a single book that describes every known cryptographic algorithm. To make it worse (or better, depending on your point of view), the book contains most of the algorithms as runnable programs (in C). Furthermore, over 1600 references to the cryptographic literature are provided. If you *really* want to keep your files secret, read this book.

Steinmetz and Nahrstedt, *Multimedia: Computing, Communications and Applications*

Although somewhat chaotic, this book does cover a lot of ground in multimedia. Topics treated at length include audio, still pictures, moving pictures, compression, optical storage, multimedia operating systems, networking, hypertext, synchronization of streams, and multimedia applications.

Van der Linden, *Just Java*

When Chap. 1 of a book is entitled “Come into my parlor, said the spider to the fly,” it is a safe bet that it is either a children’s fairy tale or about the World Wide Web. This one is about the Web, specifically about the Java language and its environment. For people who want to play with Java, the book comes complete with the full Java system on CD-ROM.

8.2. ALPHABETICAL BIBLIOGRAPHY

- ABEYSUNDARA, B.W., and KAMAL, A.E.: "High-Speed Local Area Networks and Their Performance" *Computing Surveys*, vol. 23, pp. 221-264, June 1991.
- ABRAMSON, N.: "Development of the ALOHANET," *IEEE Trans. on Information Theory*, vol. IT-31, pp. 119-123, March 1985.
- ADAM, J.A.: "Privacy and Computers," *IEEE Spectrum*, vol. 32, pp. 46-52, Dec. 1995.
- ADAMS, N., GOLD, R., SCHILIT, B.N., TSO, M.M., and WANT, R.: "An Infrared Network for Mobile Computers," *Proc. USENIX Mobile and Location-Independent Computing Symposium*, USENIX, pp. 41-51, 1993.
- ANDERSON, R.J.: "Why Cryptosystems Fail," *Commun. of the ACM*, vol. 37, pp. 32-40, Nov. 1994.
- ARMBRUSTER, H.: "The Flexibility of ATM: Supporting Future Multimedia and Mobile Communications," *IEEE Commun. Magazine*, vol. 33, pp. 76-84, Feb. 1995.
- ARMITAGE, G.J., and ADAMS, K.M.: "How Efficient is IP over ATM Anyway?" *IEEE Network Magazine*, vol. 9, pp. 18-26, Jan./Feb. 1995.
- ARNOLD, K., and GOSLING, J.: *The Java Programming Language*, Reading, MA: Addison-Wesley, 1996.
- AT&T and BELLCORE: "Observations of Error Characteristics of Fiber Optic Transmission Systems," CCITT SG XVIII, San Diego, Jan. 1989.
- AWDEH, R.Y., and MOUFTAH, H.T.: "Survey of ATM Switch Architectures," *Computer Networks and ISDN Systems*, vol. 27, pp. 1567-1613, Nov. 1995.
- BAKNE, A., and BADRINATH, B.R.: "I-TCP: Indirect TCP for Mobile Hosts," *Proc. Fifteenth Int'l. Conf. on Distr. Computer Systems*, IEEE, pp. 136-143, 1995.
- BALAKRISHNAN, H., SESHAN, S., and KATZ, R.H.: "Improving Reliable Transport and Hand-off Performance in Cellular Wireless Networks," *Proc. ACM Mobile Computing and Networking Conf.*, ACM, pp. 2-11, 1995.
- BALLARDIE, T., FRANCIS, P., and CROWCROFT, J.: "Core Based Trees (CBT)," *Proc. SIGCOMM '93 Conf.*, ACM, pp. 85-95, 1993.
- BANTZ, D.F., and BAUCHOT, F.J.: "Wireless LAN Design Alternatives," *IEEE Network Magazine*, vol. 8, pp. 43-53, March/April, 1994.
- BARANSEL, C., DOBOSIEWICZ, W., and GBURZYNSKI, P.: "Routing in Multihop Packet Switching Networks: Gb/s Challenge," *IEEE Network Magazine*, vol. 9, pp. 38-61, May/June, 1995.
- BARLOW, J.P.: "Property and Speech: Who Owns What You Say in Cyberspace," *Commun. of the ACM*, vol. 38, pp. 19-22, Dec. 1995.
- BATCHER, K.E.: "Sorting Networks and Their Applications," *Proc. AFIPS Spring Joint Computer Conf.*, vol. 32, pp. 307-315, 1968.

- BATES, R.J.: *Wireless Networked Communications*, New York: McGraw-Hill, 1994.
- BERGHEL, H.L.: "The Client Side of the Web," *Commun. of the ACM*, vol. 39, pp. 33-40, Jan. 1996.
- BELL, T.E. "Communications," *IEEE Spectrum*, vol. 33, pp. 30-41, Jan 1996.
- BELLAMY, J.: *Digital Telephony*, New York: John Wiley, 1991.
- BELLMAN, R.E.: *Dynamic Programming*, Princeton, NJ: Princeton University Press, 1957.
- BELSNES, D.: "Flow Control in the Packet Switching Networks," *Communications Networks*, Uxbridge, England: Online, pp. 349-361, 1975.
- BERNERS-LEE, T., CAILLAU, A., LOUITONEN, A., NIELSEN, H.F., and SECRET, A.: "The World Wide Web," *Commun. of the ACM*, vol. 37, pp. 76-82, Aug. 1994.
- BERTSEKAS, D., and GALLAGER, R.: *Data Networks*, 2nd ed., Englewood Cliffs, NJ: Prentice Hall, 1992.
- BHARGHAVAN, V., DEMERS, A., SHENKER, S., and ZHANG, L.: "MACAW: A Media Access Protocol for Wireless LANs," *Proc. SIGCOMM '94 Conf.*, ACM, pp. 212-225, 1994.
- BIHAM, E., and SHAMIR, A.: *Differential Cryptanalysis of the Data Encryption Standard*, New York: Springer-Verlag, 1993.
- BINDER, R.: "A Dynamic Packet Switching System for Satellite Broadcast Channels," *Proc. Int'l. Conf. on Commun.*, pp. 41-1 to 41-5a, 1975.
- BLACK, U.D.: *TCP/IP and Related Protocols*, New York: McGraw-Hill, 1995.
- BLACK, U.D.: *Emerging Commun. Technol.*, Englewood Cliffs, NJ: Prentice Hall, 1994.
- BLACK, U.D.: *Data Link Protocols*, Englewood Cliffs, NJ: Prentice Hall, 1993.
- BLAZE, M.: "Protocol Failure in the Escrowed Encryption Standard," *Proc. Second ACM Conf. on Computer and Commun. Security*, ACM, pp. 59-67, 1994.
- BOGINENI, K., SIVALINGAM, K.M., and DOWD, P.W.: "Low-Complexity Multiple Access Protocols for Wavelength-Division Multiplexed Photonic Networks," *IEEE Journal on Selected Areas in Commun.*, vol. 11, pp. 590-604, May 1993.
- BONOMI, F., and FENDICK, K.W.: "The Rate-Based Flow Control Framework for the Available Bit-rate ATM Service," *IEEE Network Magazine*, vol. 9, pp. 25-39, March/April 1995.
- BOWMAN, C.M., DANZIG, P.B., HARDY, D.R., MANBER, U., and SCHWARTZ, M.F.: "The Harvest Information Discovery and Access System," *Computer Networks and ISDN Systems*, vol. 28, pp. 119-125, Dec. 1995.
- BOWMAN, C.M., DANZIG, P.B., MANBER, U., and SCHWARTZ, M.F.: "Scalable Internet Resource Discovery: Research Problems and Approaches," *Commun. of the ACM*, vol. 37, pp. 98-107, Aug. 1994.
- BRAKMO, L.S., O'MALLEY, S.W., and PETERSON, L.L.: "TCP Vegas: New Techn. for Congestion Detection and Avoidance," *Proc. SIGCOMM '94 Conf.*, ACM, pp. 24-35, 1994.

- BROADHEAD, M.A. and OWEN, C.B.:** "Direct Manipulation of MPEG Compressed Digital Audio," *Proc. of ACM Multimedia '95*, ACM, pp. 499-507, 1995.
- BROWN, L., KWAN, M., PIEPRZYK, J., and SEBERRY, J.:** "Improving Resistance to Differential Cryptanalysis and the Redesign of LOKI," *ASIACRYPT '91 Abstracts*, pp. 25-30, 1991.
- BUFORD, J.F.K. (Ed.):** *Multimedia Systems*, Reading, MA: Addison-Wesley, 1994. DEC System Research Center Report, Feb. 1989.
- CAMPBELL, A., COULSON, G., and HUTCHISON, D.:** "A Quality of Service Architecture," *Computer Commun. Rev.*, vol. 24, pp. 6-27, April 1994.
- CAMPIONE, M., and WALRATH, K.:** *The Java Language Tutorial: Object-Oriented Programming for the Internet*, Reading, MA: Addison-Wesley, 1996.
- CAPETANAKIS, J.I.:** "Tree Algorithms for Packet Broadcast Channels," *IEEE Trans. on Information Theory*, vol. IT-25, pp. 505-515, Sept. 1979.
- CARL-MITCHELL, S., and QUARTERMAN, J.S.:** *Practical Internetworking with TCP/IP and UNIX*, Reading, MA: Addison-Wesley, 1993.
- CATLETT, C.E.:** "In Search of Gigabit Applications," *IEEE Commun. Magazine*, vol. 30, pp. 42-51, April 1992.
- CERF, V., and KAHN, R.:** "A Protocol for Packet Network Interconnection," *IEEE Trans. on Commun.*, vol. COM-22, pp. 637-648, May 1974.
- CHANDRANMENON, G.P., and VARGHESE, G.:** "Trading Packet Headers for Packet Processing," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 162-173, 1995.
- CHANG, Y.-H., COGGINS, D., PITT, D., SKELLERN, D., THAPAR, M., and VENKATRAMAN, C.:** "An Open-System Approach to Video on Demand," *IEEE Commun. Magazine*, vol. 32, pp. 68-80, May 1994.
- CHAO, J.J., GHOSAL, D., SAHA, D., and TRIPATHI, S.K.:** "IP on ATM Local Area Networks," *IEEE Commun. Magazine*, vol. 32, pp. 52-59, Aug. 1994.
- CHAPMAN, D.E., and ZWICKY, E.D.:** *Building Internet Firewalls*, Sebastopol, CA: O'Reilly, 1995.
- CHEN, K.-C.:** "Medium Access Control of Wireless LANs for Mobile Computing," *IEEE Network Magazine*, vol. 8, pp. 50-63, Sept./Oct. 1994.
- CHEN, M., and YUM, T.-S.:** "A Conflict-Free Protocol for Optical WDMA Networks," *Proc. Globecom '91*, pp. 1276-1281, 1991.
- CHEN, W.Y., and WARING, D.L.:** "Applicability of ADSL to Support Video Dial Tone in the Copper Loop," *IEEE Commun. Magazine*, vol. 32, pp. 102-106, May 1994.
- CHERITON, D., and WILLIAMSON, C.:** "VMTP as the Transport Layer for High-Performance Distributed Systems," *IEEE Commun. Magazine*, vol. 27, pp. 37-44, June 1989.

- CHERVENAK, A.L.: *Tertiary Storage: An Evaluation of New Applications*, Ph.D. thesis, CSD, Univ. of California at Berkeley, 1994.
- CHERVENAK, A.L., PATTERSON, D.A., and KATZ, R.H.: "Choosing the Best Storage System for Video Service," *Proc. of ACM Multimedia '95*, ACM, pp. 109-119, 1995.
- CHESSON, G.L.: "XTP/PE Design Considerations," *IFIP Workshop on Protocols for High-Speed Networks*, IFIP, pp. 27-33, 1989.
- CHESWICK, W.R. and BELLOVIN, S.M.: *Firewalls and Interwalls—Repelling the Wily Hacker*, Reading, MA: Addison-Wesley, 1994.
- CHOUDBURY, A.K., MAXEMCHUK, N.F., PAUL, S., and SCHULZRINNE, H.G.: "Copyright Protection for Electronic Publishing on Computer Networks," *IEEE Network Magazine*, vol. 9, pp. 12-20, May/June, 1995.
- CLARK, D.D.: "The Design Philosophy of the DARPA Internet Protocols," *Proc. SIGCOMM '88 Conf.*, ACM, pp. 106-114, 1988.
- CLARK, D.D.: "NETBLT: A Bulk Data Transfer Protocol," RFC 998, 1987.
- CLARK, D.D.: "Window and Acknowledgement Strategy in TCP," RFC 813, July 1982.
- CLARK, D.D., DAVIE, B.S., FARBER, D.J., GOPAL, I.S., KADABA, B.K., SINCOSKIE, W.D., SMITH, J.M., and TENNENHOUSE, D.L.: "The Aurora Gigabit Testbed," *Computer Networks and ISDN Systems*, vol. 25, pp. 599-621, Jan. 1993.
- CLARK, D.D., JACOBSON, V., ROMKEY, J., and SALWEN, H.: "An Analysis of TCP Processing Overhead," *IEEE Commun. Magazine*, vol. 27, pp. 23-29, June 1989.
- CLARK, D.D., LAMBERT, M., and ZHANG, L.: "NETBLT: A High Throughput Transport Protocol," *Proc. SIGCOMM '87 Conf.*, ACM, pp. 353-359, 1987.
- CLOS, C.: "A Study of Non-Blocking Switching Networks," *Bell System Tech. J.*, vol. 32, pp. 406-424, March 1953.
- COMER, D.E.: *The Internet Book*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- COMER, D.E.: *Internetworking with TCP/IP*, vol. 1, 3rd ed., Englewood Cliffs, NJ: Prentice Hall, 1995.
- COOK, A., and STERN, J.: "Optical Fiber Access—Perspectives Toward the 21st Century," *IEEE Commun. Magazine*, vol. 32, pp. 78-86, Feb. 1994.
- COOPER, E.: *Broadband Network Technology*, Englewood Cliffs, NJ: Prentice Hall, 1986.
- COULOURIS, G.F., DOLLIMORE, J., and KINDBERG, T.: *Distributed Systems Concepts and Design*, 2nd ed. Reading, MA: Addison-Wesley, 1994.
- CRESPO, P.M., HONIG, M.L., and SALEHI, J.A.: "Spread-Time Code-Division Multiple Access," *IEEE Trans. on Commun.*, vol. 43, pp. 2139-2148, June 1995.
- CRONIN, W.J., HUTCHINSON, J.D., RAMAKRISHNAN, K.K., and YANG, H.: "A Comparison of High Speed LANs," *Proc. Nineteenth Conf. on Local Computer Networks*, IEEE, pp. 40-49, 1994.

- CROWCROFT, J., WANG, Z., SMITH, A., and ADAMS, J.: "A Rough Comparison of the IETF and ATM Service Models," *IEEE Network Magazine*, vol. 9, pp. 12-16, Nov./Dec. 1995.
- CROWTHER, W., RETTBERG, R., WALDEN, D., ORNSTEIN, S., and HEART, F.: "A System for Broadcast Communication: Reservation-Aloha," *Proc. Sixth Hawaii Int. Conf. System Sci.*, pp. 371-374, 1973.
- CUSICK, T.W., and WOOD, M.C.: "The REDOC-II Cryptosystem," *Advances in Cryptology—CRYPTO '90 Proceedings*, NY: Springer-Verlag, pp. 545-563, 1991.
- DAGDEVIREN, N., NEWELL, J.A., SPINDEL, L.A., and STEFANICK, M.J.: "Global Networking with ISDN," *IEEE Commun. Magazine*, vol. 32, pp. 26-32, June 1994.
- DANSKIN, J.M., DAVIS, G.M., and SONG, X.: "Fast Lossy Internet Image Transmission," *Proc. of ACM Multimedia '95*, ACM, pp. 321-332, 1995.
- DANTHINE, A.A.S.: "Protocol Representation with Finite-State Models," *IEEE Trans. on Commun.*, vol. COM-28, pp. 632-643, April 1980.
- DAVIS, P.T., and MCGUFFIN, C.R.: *Wireless Local Area Networks*, New York: McGraw-Hill, 1995.
- DAY, J.D.: "The (Un)Revised OSI Reference Model," *Computer Commun. Rev.*, vol. 25, pp. 39-55, Oct. 1995.
- DAY, J.D., and ZIMMERMANN, H.: "The OSI Reference Model," *Proc. of the IEEE*, vol. 71, pp. 1334-1340, Dec. 1983.
- DE JONGE, W., and CHAUM, D.: "Some Variations on RSA Signatures and Their Security," in *Advances in Cryptology—CRYPTO '86 Proceedings*, Odlyzko, A.M. (Ed.), New York: Springer Verlag, 1987.
- DE PRYCKER, M.: *Asynchronous Transfer Mode*, 2nd. ed., New York: Ellis Horwood, 1993.
- DEAN, D., and WALLACH, D.S.: "Security Flaws in the HotJava Web Browser," Technical Report 502, Dept. of Computer Science, Princeton Univ., 1995.
- DEERING, S.E.: "SIP: Simple Internet Protocol," *IEEE Network Magazine*, vol. 7, pp. 16-28, May/June 1993.
- DEERING, S.E., and CHERITON, D.R.: "Multicast Routing in Datagram Internetworks and Extended LANs," *ACM Trans. on Computer Systems*, vol. 8, pp. 85-110, May 1990.
- DEERING, S.E., ESTRIN, D., FARINACCI, D., JACOBSON, V., LIU, C.-G., and WEI, L.: "An Architecture for Wide-Area Multicast Routing," *Proc. SIGCOMM '94 Conf.*, ACM, pp. 126-135, 1994.
- DELODDERE, D., VERBIEST, W., and VERHILLE, H.: "Interactive Video on Demand," *IEEE Commun. Magazine*, vol. 32, pp. 82-88, May 1994.
- DEMERS, A., KESHAV, S., and SHENKER, S.: "Analysis and Simulation of a Fair Queuing Algorithm," *Internetwork: Research and Experience*, vol. 1, pp. 3-26, Sept. 1990.

- DENNING, D.E., and SACCO, G.M.: "Timestamps in Key Distribution Protocols," *Commun. of the ACM*, vol. 24, pp. 533-536, Aug. 1981.
- DIFFIE, W., and HELLMAN, M.E.: "Exhaustive Cryptanalysis of the NBS Data Encryption Standard," *IEEE Computer Magazine*, vol. 10, pp. 74-84, June 1977.
- DIFFIE, W., and HELLMAN, M.E.: "New Directions in Cryptography," *IEEE Trans. on Information Theory*, vol. IT-22, pp. 644-654, Nov. 1976.
- DIJKSTRA, E.W.: "A Note on Two Problems in Connexion with Graphs," *Numer. Math.*, vol. 1, pp. 269-271, Oct. 1959.
- DIRVIN, R.A., and MILLER, A.R.: "The MC68824 Token Bus Controller: VLSI for the Factory LAN," *IEEE Micro Magazine*, vol. 6, pp. 15-25, June 1986.
- DIXIT, S., and SKELLY, P.: "MPEG-2 over ATM for Video Dial Tone Network," *IEEE Network Magazine*, vol. 9, pp. 30-40, Sept./Oct. 1995.
- DIXON, R.C.: "Lore of the Token Ring," *IEEE Network Magazine*, vol. 1, pp. 11-18, Jan./Feb. 1987.
- DOERINGER, W.A., DYKEMAN, D., KAISERSWERTH, M., MEISTER, B.W., RUDIN, H., and WILLIAMSON, R.: "A Survey of Light-Weight Transport Protocols for High-Speed Networks," *IEEE Trans. on Commun.*, vol. 38, pp. 2025-2039, Nov. 1990.
- DORFMAN, R.: "Detection of Defective Members of a Large Population," *Annals Math. Statistics*, vol. 14, pp. 436-440, 1943.
- ECKBERG, A.E.: "B-ISDN/ATM Traffic and Congestion Control," *IEEE Network Magazine*, vol. 6, pp. 28-37, Sept./Oct. 1992.
- ECKBERG, A.E., DOSHI, B.T., and ZOCCOLILLO, R.: "Controlling Congestion in B-ISDN/ATM: Issues and Strategies," *IEEE Commun. Magazine*, vol. 29, pp. 64-70, Sept. 1991.
- EDWARDS, A., and MUIR, S.: "Experience Implementing a High-Performance TCP in User-Space," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 197-205, 1995.
- EL GAMAL, T.: "A Public-Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms," *IEEE Trans. on Information Theory*, vol. IT-31, pp. 469-472, July 1985.
- ERIKSSON, H.: "MBone: The Multicast Backbone," *Commun. of the ACM*, vol. 37, pp. 54-60, Aug. 1994.
- ESTRIN, D., REKHTER, Y., and HOTZ, S.: "Scalable Inter-Domain Routing Architecture," *Proc. SIGCOMM '92 Conf.*, ACM, pp. 40-52, 1992.
- FEIG, E., and WINOGRAD, S.: "Fast Algorithms for Discrete Cosine Transformations," *IEEE Trans. on Signal Processing*, vol. 40, Sept. 1992.
- FEIT, S.: *SNMP—A Guide to Network Management*, New York: McGraw-Hill, 1995.
- FIORINI, D., CHIARI, M., TRALLI, V., and SALATI, C.: "Problems with HDLC," *Computer Commun. Rev.*, vol. 25, pp. 61-80, Oct. 1995.

- FISCHER, W., WALLMEIER, E., WORSTER, T., DAVIS, S.P., HAYTER, A.: "Data Communications Using ATM: Architectures, Protocols, and Resource Management," *IEEE Commun. Magazine*, vol. 32, pp. 24-33, Aug. 1994.
- FLOYD, S., and JACOBSON, V.: "Random Early Detection for Congestion Avoidance," *IEEE/ACM Trans. on Networking*, vol. 1, pp. 397-413, Aug. 1993.
- FLUCKIGER, F.: *Understanding Networked Multimedia*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- FORD, L.R., Jr., and FULKERSON, D.R.: *Flows in Networks*, Princeton, NJ: Princeton University Press, 1962.
- FORD, P.S., REKHTER, Y., and BRAUN, H.-W.: "Improving the Routing and Addressing of IP," *IEEE Network Magazine*, vol. 7, pp. 10-15, May/June 1993.
- FORMAN, G.H., and ZAHORJAN, J.: "The Challenges of Mobile Computing," *IEEE Computer Magazine*, vol. 27, pp. 38-47, April 1994.
- FRANCIS, P.: "A Near-Term Architecture for Deploying Pip," *IEEE Network Magazine*, vol. 7, pp. 30-37, May/June 1993.
- FRASER, A.G.: "Early Experiments with Asynchronous Time Division Networks," *IEEE Network Magazine*, vol. 7, pp. 12-27, Jan./Feb. 1993.
- FRASER, A.G.: "Towards a Universal Data Transport System," in *Advances in Local Area Networks*, Kummerle, K., Tobagi, F., and Limb, J.O. (Eds.), New York: IEEE Press, 1987.
- FURHT, B., KALRA, D., KITSON, F.L., RODRIGUEZ, and WALL, W.E.: "Design Issues for Interactive Televisions Systems," *IEEE Computer Magazine*, vol. 28, pp. 25-39, May 1995.
- GARCIA-HARO, J., and JAJSZCZYK, A.: "ATM Shared-Memory Switching Architectures," *IEEE Network Magazine*, vol. 8., pp. 18-26, July/Aug. 1994.
- GARG, V., and WILKES, J.E.: *Wireless and Personal Communication Systems*, Englewood Cliffs, NJ: Prentice Hall, 1996.
- GASMAN, L.: *Broadband Networking*, New York: Van Nostrand Reinhold, 1994.
- GIACOPELLI, J.N., HICKEY, J.J., MARCUS, W.S., SINOSKIE, W.D., and LITTLEWOOD, M.: "Sunshine: A High-Performance Self-Routing Broadband Packet Switch Architecture," *IEEE Journal on Selected Areas in Commun.*, vol. 9, pp. 1289-1298, Oct. 1991.
- GOODMAN, D.J.: "Trends in Cellular and Cordless Communications," *IEEE Commun. Magazine*, vol. 29, pp. 31-40, June 1991.
- GORALSKI, W.J.: *Introduction to ATM Networking*, New York: McGraw-Hill, 1995.
- GOSLING, J., JOY, B., and STEELE, G.: *The Java Language Specification*, Reading, MA: Addison-Wesley, 1996.
- GREEN, P.E., Jr.: *Fiber Optic Networks*, Englewood Cliffs, NJ: Prentice Hall, 1993.

- HAC, ANNA: "Wireless and Cellular Architecture and Services," *IEEE Commun. Magazine*, vol. 33, pp. 98-104, Nov. 1995.
- HAFNER, K., and MARKOFF, J.: *Cyberpunk*, New York: Simon and Schuster, 1991.
- HAMMING, R.W.: "Error Detecting and Error Correcting Codes," *Bell System Tech. J.*, vol. 29, pp. 147-160, April 1950.
- HANDEL, R., HUBER, M.N., and SCHRODER, S.: *ATM Concepts, Protocols, and Applications*, 2nd ed., Reading, MA: Addison-Wesley, 1994.
- HANDLEY, M., and CROWCROFT, J.: *The World Wide Web—Beneath the Surf*, London: UCL Press, 1994.
- HAWLEY, G.T.: "Historical Perspectives on the U.S. Telephone System," *IEEE Commun. Magazine*, vol. 29, pp. 24-28, March 1991.
- HEIN, M., and GRIFFITHS, D.: *SNMP*, London: Thompson, 1995.
- HELD, G.: *The Complete Modem Reference*, 2nd ed., New York: John Wiley, 1994.
- HELLMAN, M.E.: "A Cryptanalytic Time-Memory Tradeoff," *IEEE Trans. on Information Theory*, vol. IT-26, pp. 401-406, July 1980.
- HENDERSON, T.R.: "Design Principles and Performance Analysis of SSCOP: A New ATM Adaptation Layer Protocol," *Computer Commun. Review*, vol. 25, pp. 47-59, April 1995.
- HOARE, C.A.R.: "Monitors, An Operating System Structuring Concept," *Commun. of the ACM*, vol. 17, pp. 549-557, Oct. 1974; Erratum in *Commun. of the ACM*, vol. 18, p. 95, Feb. 1975.
- HODGE, W.W.: *Interactive Television*, New York: McGraw-Hill, 1995.
- HODGE, W.W., Martin, S., POWERS, J.T., Jr.: "Video on Demand: Architectures, Systems, and Applications," *Society of Motion Picture and Television Engineers Journal*, vol. 102, pp. 791-803, Sept. 1993.
- HOFFMAN, L.J. (ed.): *Building in Big Brother: The Cryptographic Policy Debate*, New York: Springer-Verlag, 1995.
- HOLFELDER, W.: "MBone VCR—Video Conference Recording on the MBone," *Proc. of ACM Multimedia '95*, ACM, pp. 237-238, 1995.
- HOLZMANN, G.J.: *Design and Validation of Computer Protocols*, Englewood Cliffs, NJ: Prentice Hall, 1991.
- HONG, D., and SUDA, T.: "Congestion Control and Prevention in ATM Networks," *IEEE Network Magazine*, vol. 5, pp. 10-16, July/Aug. 1991.
- HUANG, A., and KNAUER, S.: "Starlite: A Wideband Digital Switch," *Proc. Globecom '84*, pp. 121-125, 1984.
- HUGHES, J.P., and FRANTA, W.R.: "Geographic Extension of HIPPI Channels," *IEEE Network Magazine*, vol. 8, pp. 42-53, May/June 1994.

- HUI, J.: "A Broadband Packet Switch for Multi-rate Services," *Proc. Int'l. Conf. on Communications*, IEEE, pp. 782-788, 1987.
- HUITEMA, C.: *IPv6: The New Internet Protocol*, Englewood Cliffs, NJ: Prentice Hall, 1996.
- HUITEMA, C.: *Routing in the Internet*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- HUMBLET, P.A., RAMASWAMI, R., and SIVARAJAN, K.N.: "An Efficient Communication Protocol for High-Speed Packet-Switched Multichannel Networks," *Proc. SIGCOMM '92 Conf.*, ACM, pp. 2-13, 1992.
- IEEE: *Communications Magazine*, vol. 33, Jan. 1995.
- IEEE: 802.3: *Carrier Sense Multiple Access with Collision Detection*, New York: IEEE, 1985a.
- IEEE: 802.4: *Token-Passing Bus Access Method*, New York: IEEE, 1985b.
- IEEE: 802.5: *Token Ring Access Method*, New York: IEEE, 1985c.
- IOANNIDIS, J., and MAQUIRE, G.Q., Jr.: "The Design and Implementation of a Mobile Internetworking Architecture," *Proc. Winter USENIX Conf.*, USENIX, pp. 491-502, Jan. 1993.
- IRMER, T.: "Shaping Future Telecommunications: The Challenge of Global Standardization," *IEEE Commun. Magazine*, vol. 32, pp. 20-28, Jan. 1994.
- IVANCIC, W.D., SHALKHAUSER, M.J., and QUINTANA, J.A.: "A Network Architecture for a Geostationary Communication Satellite," *IEEE Commun. Magazine*, vol. 32, pp. 72-84, July 1994.
- JABBARI, B., COLOMBO, G., NAKAJIMA, A., and KULKARNI, J.: "Network Issues for Wireless Communications," *IEEE Commun. Magazine*, vol. 33, pp. 88-98, Jan. 1995.
- JACOBSON, V.: "Congestion Avoidance and Control," *Proc. SIGCOMM '88 Conf.*, ACM, pp. 314-329, 1988.
- JAIN, R.: "Congestion Control and Traffic Management in ATM Networks: Recent Advances and a Survey," *Computer Networks and ISDN Systems*, vol. 27, Nov. 1995.
- JAIN, R.: *FDDI Handbook—High-Speed Networking Using Fiber and other Media*, Reading, MA: Addison-Wesley, 1994.
- JAIN, R.: *The Art of Computer Systems Performance Analysis*, New York: John Wiley, 1991.
- JAIN, R.: "Congestion Control in Computer Networks: Issues and Trends," *IEEE Network Magazine*, vol. 4, pp. 24-30, May/June 1990.
- JIA, F., and MUKHERJEE, B.: "The Receiver Collision Avoidance (RCA) Protocol for a Single-Hop WDM Lightwave Network," *Journal of Lightwave Technology*, vol. 11, pp. 1053-1065, May/June 1993.
- JOHNSON, D.B.: "Scalable Support for Transparent Mobile Host Internetworking," *Wireless Networks*, vol. 1, pp. 311-321, Oct. 1995.

- JOHNSON, H.W.: *Fast Ethernet—Dawn of a New Network*, Englewood Cliffs, NJ: Prentice Hall, 1996.
- KAHN, D.: "Cryptology Goes Public," *IEEE Commun. Magazine*, vol. 18, pp. 19-28, March 1980.
- KAHN, D.: *The Codebreakers*, New York: Macmillan, 1967.
- KALISKI, B.S., and ROBshaw, M.J.B.: "Fast Block Cipher Proposal," *Proc. Cambridge Security Workshop*, Springer-Verlag, pp. 26-39, 1994.
- KAMOUN, F., and KLEINROCK, L.: "Stochastic Performance Evaluation of Hierarchical Routing for Large Networks," *Computer Networks*, vol. 3, pp. 337-353, Nov. 1979.
- KARN, P.: "MACA—A New Channel Access Protocol for Packet Radio," *ARRL/CRRL Amateur Radio Ninth Computer Networking Conf.*, pp. 134-140, 1990.
- KAROL, M.J., HLUCHYJ, M.G., and MORGAN, S.P.: "Input Versus Output Queueing on a Space-Division Packet Switch," *IEEE Trans. on Commun.*, vol. 35, pp. 1347-1356, Dec. 1987.
- KARSHMER, A.I., and THOMAS, J.N.: "Computer Networking on Cable TV Plants," *IEEE Commun. Magazine*, vol. 30, pp. 32-40, Nov. 1992.
- KATZ, D., and FORD, P.S.: "TUBA: Replacing IP with CLNP," *IEEE Network Magazine*, vol. 7, pp. 38-47, May/June 1993.
- KATZ, E.D., BUTLER, M., and McGRATH, R.: "A Scalable HTTP Server: The NCSA Prototype," *Computer Networks and ISDN Systems*, vol. 27, pp. 155-164, Nov. 1994.
- KAUFMAN, C., PERLMAN, R., and SPECINER, M.: *Network Security*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- KAVAK, N.: "Data Communication in ATM Networks," *IEEE Network Magazine*, vol. 9, pp. 28-37, May/June 1995.
- KENT, C.A., and MOGUL, J.C.: "Fragmentation Considered Harmful," *Proc. SIGCOMM '87 Conf.*, ACM, pp. 390-401, 1987.
- KENT, S.T.: "Internet Privacy Enhanced Mail," *Commun. of the ACM*, vol. 36, pp. 48-60, Aug. 1993.
- KESSLER, G.C.: *ISDN*, 2nd ed., New York: McGraw-Hill, 1993.
- KESSLER, G.C., and TRAIN, D.: *Metropolitan Area Networks: Concepts, Standards, and Services*, New York: McGraw-Hill, 1992.
- KIM, J.B., SUDA, T., and YOSHIMURA, M.: "International Standardization of B-ISDN," *Computer Networks and ISDN Systems*, vol. 27, pp. 5-27, Oct. 1994.
- KLEINROCK, L., and TOBAGI, F.: "Random Access Techniques for Data Transmission over Packet-Switched Radio Channels," *Proc. Nat. Computer Conf.*, pp. 187-201, 1975.
- KOHO, R., MEIDAN, R., and MILSTEIN, L.B.: "Spread Spectrum Access Methods for Wireless Communication," *IEEE Commun. Magazine*, vol. 33, pp. 58-67, Jan. 1995.

- KUMAR, V.: *MBone: Interactive Multimedia on the Internet*, Indianapolis, IN: New Riders, 1996.
- KUNG, H.T., and MORRIS, R.: "Credit-Based Flow Control for ATM Networks," *IEEE Network Magazine*, vol. 9, pp. 40-48, March/April 1995.
- KWAN, T.T., McGRATH, R.E., and REED, D.A.: "NCSA's WWW Server: Design and Performance," *IEEE Computer Magazine*, vol. 28, pp. 68-74, Nov. 1995.
- KWOK, T.: "A Vision for Residential Broadband Service: ATM to the Home," *IEEE Network Magazine*, vol. 9, pp. 14-28, Sept./Oct. 1995.
- KYAS, O.: *ATM Networks*, London: International Thomson Publishing, 1995.
- LAI, X.: *On the Design and Security of Block Ciphers*, Konstanz, Germany: Hartung-Gorre, 1992.
- LAI, X., and MASSEY, J.: "A Proposal for a New Block Encryption Standard," *Advances in Cryptology—Eurocrypt '90 Proceedings*, New York: Springer-Verlag, pp. 389-404, 1990.
- LAMPSON, B.W.: "A Note on the Confinement Problem," *Commun. of the ACM*, vol. 10, pp. 613-615, Oct. 1973.
- LANDAU, S.: "Zero-Knowledge and the Department of Defense," *Notices of the American Mathematical Society*, vol. 35, pp. 5-12, Jan. 1988.
- LANGSFORD, A.: "The Open System User's Programming Interfaces," *Computer Networks*, vol. 8, pp. 3-12, 1984.
- LA PORTA, T.F., VEERARAGHAVAN, M., AYANOGLU, E., KAROL, M., and GITLIN, R.D.: "B-ISDN: A Technological Discontinuity," *IEEE Commun. Magazine*, vol. 32, pp. 84-97, Oct. 1994.
- LATIF, A., ROWLANCE, E.J., and ADAMS, R.H.: "The IBM 8209 LAN Bridge," *IEEE Network Magazine*, vol. 6, pp. 28-37, May/June 1992.
- LAUDON, K.C.: "Ethical Concepts and Information Technology," *Commun. of the ACM*, vol. 38, pp. 33-39, Dec. 1995.
- LE BOUDEC, J.-Y.: "The Asynchronous Transfer Mode: A Tutorial," *Computer Networks and ISDN Systems*, vol. 24, pp. 279-309, May 1992.
- LEINER, B.M., COLE, R., POSTEL, J., and MILLS, D.: "The DARPA Internet Protocol Suite," *IEEE Commun. Magazine*, vol. 23, pp. 29-34, March 1985.
- LEVINE, D.A., and AKYILDIZ, I.A.: "PROTON: A Media Access Control Protocol for Optical Networks with Star Topology," *IEEE/ACM Trans. on Networking*, vol. 3, pp. 158-168, April 1995.
- LEVY, S.: "Crypto Rebels," *Wired*, pp. 54-61, May/June 1993.
- LIN, F., CHU, P., and LIU, M.: "Protocol Verification Using Reachability Analysis: The State Space Explosion Problem and Relief Strategies," *Proc. SIGCOMM '87 Conf.*, ACM, pp. 126-135, 1987.

- LIPPER, E.H., and RUMSEWICZ, M.P.: "Teletraffic Considerations for Widespread Deployment of PCS," *IEEE Network Magazine*, vol. 8, pp. 40-49, Sept./Oct. 1994.
- LITTLE, T.D.C., and VENKATESH, D.: "Prospects for Interactive Video on Demand," *IEEE Multimedia Magazine*, vol. 1, pp. 14-24, Fall 1994.
- LIU, C.L., and LAYLAND, J.W.: "Scheduling Algorithms for Multiprogramming in a Hard Real-Time Environment," *Journal of the ACM*, vol. 20, pp. 46-61, Jan. 1973.
- LUOTONEN, A., and ALTIIS, K.: "World Wide Web Proxies," *Computer Networks and ISDN Systems*, vol. 27, pp. 147-154, Nov. 1994.
- MACARIO, R.C.V.: *Cellular Radio—Principles and Design*, New York: McGraw-Hill, 1993.
- MACEDONIA, M.R., and BRUTZMAN, D.P.: "MBone Provides Audio and Video Across the Internet," *IEEE Computer Magazine*, vol. 27, pp. 30-36, April 1994.
- MASSEY, J.L.: "SAFER K-64: A Byte-Oriented Block Ciphering Algorithm," *Proc. Cambridge Security Workshop*, Springer-Verlag, pp. 1-17, 1994.
- MATSUI, M.: "Linear Cryptanalysis Method for DES Cipher," *Advances in Cryptology—Eurocrypt '93 Proceedings*, New York: Springer-Verlag, pp. 386-397, 1994.
- McBRYAN, O.: "GENVL and WWW: Tools for Taming the Web," *Proc. First Int'l WWW Conference*, pp. 79-90, 1994.
- McDYSAN, D.E., and SPOHN, D.L.: *ATM—Theory and Application*, NY: McGraw-Hill, 1995.
- McKENNEY, P.E., and DOVE, K.F.: "Efficient Demultiplexing of Incoming TCP Packets," *Proc. SIGCOMM '92 Conf.*, ACM, pp. 269-279, 1992.
- MENEZES, A.J., and VANSTONE, S.A.: "Elliptic Curve Cryptosystems and Their Implementation," *Journal of Cryptology*, vol. 6, pp. 209-224, 1993.
- MERKLE, R.C.: "Fast Software Encryption Functions," *Advances in Cryptology—CRYPTO '90 Proceedings*, New York: Springer-Verlag, pp. 476-501, 1991.
- MERKLE, R.C., and HELLMAN, M.: "On the Security of Multiple Encryption," *Commun. of the ACM*, vol. 24, pp. 465-467, July 1981.
- MERKLE, R.C., and HELLMAN, M.: "Hiding and Signatures in Trapdoor Knapsacks," *IEEE Trans. on Information Theory*, vol. IT-24, pp. 525-530, Sept. 1978.
- METCALFE, R.M.: "On Mobile Computing," *Byte*, vol. 20, p. 110, Sept. 1995.
- METCALFE, R.M.: "Computer/Network Interface Design: Lessons from Arpanet and Ethernet," *IEEE Journal on Selected Areas in Commun.*, vol. 11, pp. 173-179, Feb. 1993.
- METCALFE, R.M., and BOGGS, D.R.: "Ethernet: Distributed Packet Switching for Local Computer Networks," *Commun. of the ACM*, vol. 19, pp. 395-404, July 1976.
- MIKI, T.: "The Potential of Photonic Networks," *IEEE Commun. Magazine*, vol. 32, pp. 23-27, Dec. 1994a.
- MIKI, T.: "Toward the Service-Rich Era," *IEEE Commun. Magazine*, vol. 32, pp. 34-39, Feb. 1994b.

- MINOLI, D.: *Video Dialtone Technology*, New York: McGraw-Hill, 1995
- MINOLI, D., and VITELLA, M.: *ATM & Cell Relay for Corporate Environments*, New York: McGraw-Hill, 1994.
- MIRCHANDANI, S., and KHANNA, R. (eds): *FDDI Technologies and Applications*, New York: John Wiley, 1993.
- MISHRA, P.P. and KANAKIA, H.: "A Hop by Hop Rate-Based Congestion Control Scheme," *Proc. SIGCOMM '92 Conf.*, ACM, pp. 112-123, 1992.
- MOCHIDA, Y.: "Technologies for Local-Access Fiberling," *IEEE Commun. Magazine*, vol. 32, pp. 64-73, Feb. 1994.
- MOGUL, J.C.: "The Case for Persistent-Connection HTTP," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 299-314, 1995.
- MOGUL, J.C.: "IP Network Performance," in *Internet System Handbook*, Lynch, D.C. and Rose, M.T. (eds.), Reading, MA: Addison-Wesley, pp. 575-675, 1993.
- MOK, A.K., and WARD, S.A.: "Distributed Broadcast Channel Access," *Computer Networks*, vol. 3, pp. 327-335, Nov. 1979.
- MORALES, J., PATKA, A., CHO, P., and KUI, J.: "Video Dial Tone Sessions," *IEEE Network Magazine*, vol. 9, pp. 42-47, Sept./Oct. 1995.
- MOY, J.: "Multicast Routing Extensions," *Commun. of the ACM*, vol. 37, pp. 61-66, Aug. 1994.
- MULLENDER, S.J. (ed.): *Distributed Systems*, 2nd ed., New York: ACM Press, 1993.
- MYLES, A., and SKELLERN, D.: "Comparison of Mobile Host Protocols for IP," *Computer Networks and ISDN Systems*, vol. 26, pp. 349-355, Dec. 1993.
- NAGLE, J.: "On Packet Switches with Infinite Storage," *IEEE Trans. on Commun.*, vol. COM-35, pp. 435-438, April 1987.
- NAGLE, J.: "Congestion Control in TCP/IP Internetworks," *Computer Commun. Rev.*, vol. 14, pp. 11-17, Oct. 1984.
- NEEDHAM, R.M., and SCHROEDER, M.D.: "Authentication Revisited," *Operating Systems Rev.*, vol. 21, p. 7, Jan. 1987.
- NEEDHAM, R.M., and SCHROEDER, M.D.: "Using Encryption for Authentication in Large Networks of Computers," *Commun. of the ACM*, vol. 21, pp. 993-999, Dec. 1978.
- NELSON, M.N., and LINTON, M.: "A Highly Available, Scalable ITV System," *Proc. Fifteenth Symp. on Operating Systems Prin.*, ACM, pp. 54-67, 1995.
- NEMETH, E., SNYDER, G., SEEBASS, S., and HEIN, T.R.: *UNIX System Administration Handbook*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- NEMZOW, M.: *Implementing Wireless Networks*, New York: McGraw-Hill, 1995.
- NEUMAN, B.C., and TS'O, T.: "Kerberos: An Authentication Service for Computer Networks," *IEEE Commun. Magazine*, vol. 32, pp. 33-38, Sept. 1994.

- NEWMAN, P.: "Traffic Management for ATM Local Area Networks," *IEEE Commun. Magazine*, vol. 32, pp. 44-50, Aug. 1994.
- NEWMAN, P.: "ATM Local Area Networks," *IEEE Commun. Magazine*, vol. 32, pp. 86-98, March 1994.
- NIST: "Secure Hash Algorithm," U.S. Government Federal Information Processing Standard 180, 1993.
- OMIDYAR, C.G., and ALDRIDGE, A.: "Introduction to SDH/SONET," *IEEE Commun. Magazine*, vol. 31, pp. 30-33, Sept. 1993.
- OTWAY, D., and REES, O.: "Efficient and Timely Mutual Authentication," *Operating Systems Rev.*, pp. 8-10, Jan. 1987.
- PADGETT, J.E., GUNTHER, C.G., and HATTORI, T.: "Overview of Wireless Personal Communications," *IEEE Commun. Magazine*, vol. 33, pp. 28-41, Jan. 1995.
- PAFF, A.: "Hybrid Fiber/Coax in the Public Telecommunications Infrastructure," *IEEE Commun. Magazine*, vol. 33, pp. 40-45, April 1995.
- PAHLAVAN, K., PROBERT, T.H., and CHASE, M.E.: "Trends in Local Wireless Networks," *IEEE Commun. Magazine*, vol. 33, pp. 88-95, March 1995.
- PALAIS, J.C.: *Fiber Optic Commun.*, 3rd ed., Englewood Cliffs, NJ: Prentice Hall, 1992.
- PALMER, L.C., and WHITE, L.W.: "Demand Assignment in the ACTS LBR System," *IEEE Trans. on Commun.*, vol. 38, pp. 684-692, May 1990.
- PAN, D.: "A Tutorial on MPEG/Audio Compression," *IEEE Multimedia Magazine*, vol. 2, pp. 60-74, Summer 1995.
- PANCHA, P., and EL ZARKI, M.: "MPEG Coding for Variable Bit Rate Video Transmission," *IEEE Commun. Magazine*, vol. 32, pp. 54-66, May 1994.
- PANDYA, R.: "Emerging Mobile and Personal Communication Systems," *IEEE Commun. Magazine*, vol. 33, pp. 44-52, June 1995.
- PARTRIDGE, C.: *Gigabit Networking*, Reading, MA: Addison-Wesley, 1994.
- PARTRIDGE, C.: "A Proposed Flow Specification," Internet RFC 1363, Sept. 1992.
- PARTRIDGE, C., HUGHES, J., and STONE, J.: "Performance of Checksums and CRCs over Real Data," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 68-76, 1995.
- PARULKAR, G., SCHMIDT, D.C., and TURNER, J.S.: "AITPM: A Strategy for Integrating IP with ATM," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 49-58, 1995.
- PAXSON, V.: "Growth Trends in Wide-Area TCP Connections," *IEEE Network Magazine*, vol. 8, pp. 8-17, July/Aug. 1994.
- PAXSON, V., and FLOYD, S.: "Wide-Area Traffic: The Failure of Poisson Modeling," *Proc. SIGCOMM '94 Conf.*, ACM, pp. 257-268, 1995.
- PERKINS, C.: "Providing Continuous Network Access to Mobile Hosts Using TCP/IP," *Computer Networks and ISDN Systems*, vol. 26, pp. 357-370, Nov. 1993.

- PERLMAN, R.: *Interconnections: Bridges and Routers*, Reading, MA: Addison-Wesley, 1992.
- PERLMAN, R.: *Network Layer Protocols with Byzantine Robustness*, Ph.D. thesis, M.I.T., 1988.
- PERRY, T.S., and ADAM, J.A.: "E-Mail: Pervasive and Persuasive," *IEEE Spectrum*, vol. 29, pp. 22-28, Oct. 1992.
- PETERSON, W.W., and BROWN, D.T.: "Cyclic Codes for Error Detection," *Proc. IRE*, vol. 49, pp. 228-235, Jan. 1961.
- PICKHOLTZ, R.L., SCHILLING, D.L., and MILSTEIN, L.B.: "Theory of Spread Spectrum Communication—A Tutorial," *IEEE Trans. on Commun.*, vol. COM-30, pp. 855-884, May 1982.
- PIERCE, J.: "How Far Can Data Loops Go?" *IEEE Trans. on Commun.*, vol. COM-20, pp. 527-530, June 1972.
- PINKERTON, B.: "Finding What People Want: Experiences with the WebCrawler," *Proc. First Int'l. WorldWide Web Conference*, 1994.
- PISCITELLO, D.M., and CHAPIN, A.L.: *Open Systems Networking: TCP/IP and OSI*, Reading, MA: Addison-Wesley, 1993.
- PITT, D.A.: "Bridging—The Double Standard," *IEEE Network Magazine*, vol. 2, pp. 94-95, Jan. 1988.
- QUICK, R. F., Jr., and BALACHANDRAN, K.: "An Overview of the Cellular Digital Packet Data (CDPD) System," *Fourth Int'l. Symp. on Personal, Indoor, and Mobile Radio Commun.*, pp. 338-343, 1993.
- QUISQUATER, J.-J., and GIRAULT, M.: "Chinese Lotto as an Exhaustive Code-Breaking Machine," *IEEE Computer Magazine*, vol. 24, pp. 14-22, Nov. 1991.
- RABIN, M.O.: "Digital Signatures and Public-Key Functions as Intractable as Factorization," Technical Report LCS-TR-212, M.I.T., Jan 1979.
- RAHNEMA, M.: "Overview of the GSM System and Protocol Architecture," *IEEE Commun. Magazine*, vol. 31, pp. 92-100, April 1993.
- RAJAGOPALAN, B.: "Reliability and Scaling Issues in Multicast Communication," *Proc. SIGCOMM '92 Conf.*, ACM, pp. 188-198, 1992.
- RANSOM, M.N.: "The VISTAnet Gigabit Network Testbed," *Journal of High Speed Networks*, vol. 1, pp. 49-60, 1992.
- RAO, S.K., and HATAMIAN, M.: "The ATM Physical Layer," *Computer Commun. Rev.*, vol. 25, pp. 73-81, April 1995.
- RIVEST, R.L.: "The MD5 Message-Digest Algorithm," RFC 1320, April 1992.
- RIVEST, R.L., and SHAMIR, A.: "How to Expose an Eavesdropper," *Commun. of the ACM*, vol. 27, pp. 393-395, April 1984.

- RIVEST, R.L., SHAMIR, A., and ADLEMAN, L.: "On a Method for Obtaining Digital Signatures and Public Key Cryptosystems," *Commun. of the ACM*, vol. 21, pp. 120-126, Feb. 1978.
- ROBERTS, L.: "Dynamic Allocation of Satellite Capacity through Packet Reservation," *Proc. NCC, AFIPS*, pp. 711-716, 1973.
- ROBERTS, L.: "Extensions of Packet Communication Technology to a Hand Held Personal Terminal," *Proc. Spring Joint Computer Conference, AFIPS*, pp. 295-298, 1972.
- ROMANOW, A., and FLOYD, S.: "Dynamics of TCP Traffic over ATM Networks," *Proc. SIGCOMM '84 Conf., ACM*, pp. 79-88, 1994.
- ROSE, M.T.: *The Simple Book*, Englewood Cliffs, NJ: Prentice Hall, 1994.
- ROSE, M.T.: *The Internet Message*, Englewood Cliffs, NJ: Prentice Hall, 1993.
- ROSE, M.T., and McCLOGHRIE, K.: *How to Manage Your Network Using SNMP*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- ROSS, F.E., and HAMSTRA, J.R.: "Forging FDDI," *IEEE Journal on Selected Areas in Commun.*, vol. 11, pp. 181-190, Feb. 1993.
- SADIKU, M.N.O., and ARVIND, A.S.: "Annotated Bibliography on Distributed Queue Dual Bus (DQDB)," *Computer Commun. Rev.*, vol. 24, pp. 21-36, Jan. 1994.
- SALTZER, J.H., POGGRAN, K.T., and CLARK, D.D.: "Why a Ring?" *Computer Networks*, vol. 7, pp. 223-230, Aug. 1983.
- SALTZER, J.H., REED, D.P., and CLARK, D.D.: "End-to-End Arguments in System Design," *ACM Trans. on Computer Systems*, vol. 2, pp. 277-288, Nov. 1984.
- SANDERSON, D.W., and DOUGHERTY, D.: *Smileys*, Sebastopol, CA: O'Reilly, 1993.
- SANTIFALLER, M.: "TCP/IP and ONC/NFS," Reading, MA: Addison-Wesley, 1994.
- SCHNEIER, B.: *Applied Cryptography*, 2nd ed., New York: John Wiley, 1996.
- SCHNEIER, B.: *E-Mail Security*, New York: John Wiley, 1995.
- SCHNEIER, B.: "Description of a New Variable-Length Key, 64-Bit Block Cipher [Blowfish]," *Proc. of the Cambridge Security Workshop*, Springer-Verlag, pp. 191-204, 1994.
- SCHNORR, C.P.: "Efficient Signature Generation for Smart Cards," *Journal of Cryptology*, vol. 4, pp. 161-174, 1991.
- SCHOLTZ, R.A.: "The Origins of Spread-Spectrum Communications," *IEEE Trans. on Commun.*, vol. COM-30, pp. 822-854, May 1982.
- SCOTT, R.: "Wide Open Encryption Design Offers Flexible Implementations," *Cryptologia*, vol. 9, pp. 75-90, Jan. 1985.
- SELFRIDGE, O.G., and SCHWARTZ, R.T.: "Telephone Technology and Privacy," *Technology Rev.*, vol. 82, pp. 56-65, May 1980.

- SEYBOLD, A.M.: *Using Wireless Communications in Business*, New York: Van Nostrand Reinhold, 1994.
- SHACHAM, N., and MCKENNEY, P.: "Packet Recovery in High-Speed Networks Using Coding and Buffer Management," *Proc. INFOCOM '90*, IEEE, pp. 124-130, 1990.
- SHAH, A., and RAMAKRISHNAN, G.: *FDDI—A High Speed Network*, Englewood Cliffs, NJ: Prentice Hall, 1994.
- SHANNON, C.: "A Mathematical Theory of Communication," *Bell System Tech. J.*, vol. 27, pp. 379-423, July 1948; and pp. 623-656, Oct. 1948.
- SHEN, B., and SETHI, I.K.: "Inner-Block Operations on Compressed Images," *Proc. of ACM Multimedia '95*, ACM, pp. 489-498, 1995.
- SHIMIZU, A., and MIYAGUCHI, S.: "Fast Data Encipherment Algorithm FEAL," *Advances in Cryptology—Eurocrypt '87 Proceedings*, NY: Springer-Verlag, pp. 267-278, 1988.
- SHREEDHAR, M., and VARGHESE, G.: "Efficient Fair Queueing Using Deficit Round Robin," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 231-243, 1995.
- SINGLETON, A.: "Wired on the Web," *Byte*, vol. 21, pp. 77-80, Jan. 1996.
- SIPIOR, J.C., and WARD, B.T.: "The Ethical and Legal Quandary of Email Privacy," *Commun. of the ACM*, vol. 38, pp. 48-54, Dec. 1995.
- SIU, K.-Y., and JAIN, R.: "A Brief Overview of ATM: Protocol Layers, LAN Emulation, and Traffic Management," *Computer Commun. Rev.*, vol. 25, pp. 6-20, April 1995.
- SMITH, P.: *Frame Relay*, Reading, MA: Addison-Wesley, 1993.
- SOHA, M., and PERLMAN, R.: "Comparison of Two LAN Bridge Approaches," *IEEE Network Magazine*, vol. 2, pp. 37-43, Jan./Feb. 1988.
- SPAFFORD, E.H.: "The Internet Worm: Crisis and Aftermath," *Commun. of the ACM*, vol. 32, pp. 678-687, June 1989.
- SPRAGINS, J.D., with HAMMOND, J.L., and PAWLIKOWSKI, K.: *Telecommunications Protocols and Design*, Reading, MA: Addison-Wesley, 1991.
- STALLINGS, W.: *ISDN and Broadband ISDN with Frame Relay and ATM*, Englewood Cliffs, NJ: Prentice Hall, 1995a.
- STALLINGS, W.: *Network and Internetwork Security*, Englewood Cliffs, NJ: Prentice Hall, 1995b.
- STALLINGS, W.: *Protect Your Privacy: The PGP User's Guide*, Englewood Cliffs, NJ: Prentice Hall, 1995c.
- STALLINGS, W.: *Data and Computer Communications*, 4th ed., New York: Macmillan, 1994.
- STALLINGS, W.: *SNMP, SNMPv2, and CMIP*, Reading, MA: Addison-Wesley, 1993a.
- STALLINGS, W.: *Local and Metropolitan Area Networks*, 4th ed., New York: Macmillan, 1993b.

- STEELE, R., WHITEHEAD, J., and WONG, W.C.: "System Aspects of Cellular Radio," *IEEE Commun. Magazine*, vol. 33, pp. 80-86, Jan. 1995a.
- STEELE, R., WILLIAMS, J., CHANDLER, D., DEGHAN, S., and COLLARD, A.: "Teletraffic Performance of GSM900/DCS1800 in Street Microcells," *IEEE Commun. Magazine*, vol. 33, pp. 102-108, March 1995b.
- STEINER, J.G., NEUMAN, B.C., and SCHILLER, J.I.: "Kerberos: An Authentication Service for Open Network Systems," *Proc. Winter USENIX Conf.*, USENIX, pp. 191-201, 1988.
- STEINMETZ, R., and NAHRSTEDT, K.: *Multimedia: Computing, Communications and Applications*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- STEPHENS, W.E., and BANWELL, T.C.: "155.52 Mb/s Data Transmission on Category 5 Cable Plant," *IEEE Commun. Magazine*, vol. 33, pp. 62-69, April 1995.
- STERBENZ, J.P.G., SCHULZRINNE, H.G., and TOUCH, J.D.: "Report and Discussion of the IEEE ComSoc TCGN Gigabit Networking Workshop 1995," *IEEE Network Magazine*, vol. 9, pp. 9-29, July/Aug. 1995.
- STEVENS, W.R.: *TCP/IP Illustrated*, Vol. 1, Reading, MA: Addison-Wesley, 1994.
- STILLER, B.: "A Survey of UNI Signaling Systems and Protocols," *Computer Commun. Rev.*, vol. 25, pp. 21-33, April 1995.
- STINSON, D.R.: *Cryptography Theory and Practice*, Boca Raton, FL: CRC Press, 1995.
- SUNSHINE, C.A., and DALAL, Y.K.: "Connection Management in Transport Protocols," *Computer Networks*, vol. 2, pp. 454-473, 1978.
- SUZUKI, T.: "ATM Adaptation Layer Protocol," *IEEE Commun. Magazine*, vol. 32., pp. 80-83, April 1994.
- TANENBAUM, A.S.: *Distributed Operating Systems*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- TANENBAUM, A.S.: *Modern Operating Systems*, Englewood Cliffs, NJ: Prentice Hall, 1992.
- TERAOKA, F., YOKTE, Y., and TOKORO, M.: "Host Migration Transparency in IP Networks," *Computer Commun. Rev.*, vol. 23, pp. 45-65, Jan. 1993.
- THYAGARAJAN, A.S., and DEERING, S.E.: "Hierarchical Distance-Vector Multicast Routing for the Mbone," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 60-66, 1995.
- TOKORO, M., and TAMARU, K.: "Acknowledging Ethernet," *Compton*, IEEE, pp. 320-325, Fall 1977.
- TOLMIE, D.E.: "Gigabit LAN Issues—HIPPI, Fibre Channel, and ATM," in *Proc. High-Performance Computing and Networking*, Hertzberger, B., and Serazzi, G. (Eds.), Berlin: Springer Verlag, pp. 45-53, 1995.
- TOLMIE, D.E.: "Gigabit Networking," *IEEE LTS*, vol. 3, pp. 28-36, May 1992.
- TOLMIE, D.E., and RENWICK, J.: "HIPPI: Simplicity Yields Success," *IEEE Network Magazine*, vol. 7, pp. 28-32, Jan./Feb. 1993.

- TOMLINSON, R.S.: "Selecting Sequence Numbers," *Proc. SIGCOMM/SIGOPS Interprocess Commun. Workshop*, ACM, pp. 11-23, 1975.
- TOUCH, J.D.: "Performance Analysis of MD5," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 77-86, 1995.
- TRUONG, H.L., ELLINGTON, W.W. Jr., LE BOUDEC, J.-Y., MEIER, A.X., and PACE, J.W.: "LAN Emulation on an ATM Network," *IEEE Commun. Magazine*, vol. 33, pp. 70-85, May 1995.
- TUCHMAN, W.: "Hellman Presents No Shortcut Solutions to DES," *IEEE Spectrum*, vol. 16, pp. 40-41, July 1979.
- TURNER, J.S.: "New Directions in Communications (or Which Way to the Information Age)," *IEEE Commun. Magazine*, vol. 24, pp. 8-15, Oct. 1986.
- VAN DER LINDEN, P.: *Just Java*, Englewood Cliffs, NJ: Prentice Hall, 1996.
- VAN OORSCHOT, P.C., and WIENER, M.J.: "A Known-Plaintext Attack on Two-Key Triple Encryption," *Advances in Cryptology—CRYPTO '88 Proceedings*, New York: Springer-Verlag, pp. 119-131, 1988.
- VAN RENESSE, R., VAN STAVEREN, H., and TANENBAUM, A.S.: "Performance of the World's Fastest Distributed Operating System," *Operating Systems Rev.*, vol. 22, pp. 25-34, Oct. 1988.
- VARGHESE, G., and LAUCK, T.: "Hashed and Hierarchical Timing Wheels: Data Structures for the Efficient Implementation of a Timer Facility," *Proc. Eleventh Symp. on Operating Systems Prin.*, ACM, pp. 25-38, 1987.
- VENKATRAMANI, C., and CHIUEH, T.: "Design, Implementation, and Evaluation of a Software-Based Real-Time Ethernet Protocol," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 27-37, 1995.
- VETTER, R.J., SPELL, C., and WARD, C.: "Mosaic and the World-Wide Web," *IEEE Computer Magazine*, vol. 27, pp. 49-57, Oct. 1994.
- VILLAMIZAN, C., and SONG, C.: "High Performance TCP in ANSNET," *Computer Commun. Rev.*, vol. 25, pp. 45-60, Oct. 1995.
- VITERBI, A.J.: *CDMA Principles of Spread Spectrum Communication*, Reading, MA: Addison-Wesley, 1995.
- WADA, H., YOZAWA, T., OHNISHI, T., and TANAKA, Y.: "Mobile Computing Environment Based on Internet Packet Forwarding," *Proc. Winter USENIX Conf.*, USENIX, pp. 503-517, Jan. 1993.
- WALRAND, J.: *Communication Networks: A First Course*, Homewood, IL: Irwin, 1991.
- WATSON, R.W.: "Timer-Based Mechanisms in Reliable Transport Protocol Connection Management," *Computer Networks*, vol. 5, pp. 47-56, Feb. 1981.
- WAYNER, P.: "Picking the Crypto Lock," *Byte*, pp. 77,80, Oct. 1995.

- WEISBAND, S.P., and REINIG, B.A.: "Managing User Perceptions of Email Privacy," *Commun. of the ACM*, vol. 38, pp. 40-47, Dec. 1995.
- WIENER, M.J.: "Efficient DES Key Search," Technical Report TR-244, School of Computer Science, Carleton Univ., Ottawa, 1994.
- WILLIAMS, K.A., DAM, T.Q., and DU, D.H.-C.: "A Media Access Protocol for Time and Wavelength-Division Multiplexed Passive Star Networks," *IEEE Journal on Selected Areas in Commun.*, vol. 11, pp. 560-567, May 1993.
- WILLINGER, W., TAQQU, M.S., SHERMAN, R., and WILSON, D.V.: "Self-Similarity through High Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," *Proc. SIGCOMM '95 Conf.*, ACM, pp. 100-113, 1995.
- WOLTER, M.S.: "Fiber Distributed Data Interface—A Tutorial," *ConneXions*, pp. 16-26, Oct. 1990.
- YANG, C.-Q., and REDDY, A.V.S.: "A Taxonomy for Congestion Control Algorithms in Packet Switching Networks," *IEEE Network Magazine*, vol. 9, pp. 34-45, July/Aug. 1995.
- YEH, Y.-S., HLUCHYJ, M.G., and ACAMPORA, A.S.: "The Knockout Switch: A Simple, Modular Architecture for High-Performance Packet Switching," *IEEE Journal on Selected Areas in Commun.*, vol. 5, pp. 1274-1283, Oct. 1987.
- YOUSSEF, A.M., KALMAN, E., BENZONI, L.: "Technico-Economic Methods of Radio Spectrum Assignment," *IEEE Commun. Magazine*, vol. 33, pp. 88-94, June 1995.
- YUVAL, G.: "How to Swindle Rabin," *Cryptologia*, vol. 3, pp. 187-190, July 1979.
- ZHANG, L.: "Comparison of Two Bridge Routing Approaches," *IEEE Network Magazine*, vol. 2, pp. 44-48, Jan./Feb. 1988.
- ZHANG, L.: "RSVP A New Resource ReSerVation Protocol," *IEEE Network Magazine*, vol. 7, pp. 8-18, Sept./Oct. 1993.
- ZIMMERMANN, P.R.: *The Official PGP User's Guide*, Cambridge, MA: M.I.T. Press, 1995a.
- ZIMMERMANN, P.R.: *PGP: Source Code and Internals*, Cambridge, MA: M.I.T. Press, 1995b.
- ZIPF, G.K.: *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*, Cambridge, MA: Addison-Wesley, 1949.
- ZIV, J., and LEMPEL, Z.: "A Universal Algorithm for Sequential Data Compression," *IEEE Trans. on Information Theory*, vol. IT-23, pp. 337-343, May 1977.

INDEX

A

- A-side carrier, 160
- AAL (*see* ATM Adaptation Layer)
- AAL 1, 547-549, 753
- AAL 2, 549-550, 753
- AAL 3/4, 550-552, 753
- AAL 5, 552-554, 753
- AAL Layer, 64, 545-555
- AAL protocols, comparison, 554-555
- ABR (*see* Available Bit Rate service)
- Abstract Syntax Notation-1, 633-636
- Abstract window toolkit, 717-718
- Acknowledgement frame, 30
- ACR (*see* Actual Cell Rate)
- Active map, 684
- Active repeater, 91
- ACTS (*see* Advanced Communication Technology Satellite)
- Actual cell rate, 471
- Adaptive routing, 347
- ADC (*see* Analog Digital Converter)
- ADCCP (*see* Advanced Data Communications Control Procedure)
- Addenda, 42, 322
- Address resolution protocol, 420-423
 - gratuitous ARP, 433
- Address, 492
 - transport, 489-492
- Admission control, 386, 468
- ADSL (*see* Asymmetric Digital Subscriber Line)
- Advanced Communication Technology Satellite, 331
- Advanced Data Communications Control Procedure, 226
- Advanced mobile phone system, 158-161
- Advertisements, mobile IP, 433
- Agent, SNMP, 631
- Alias, email, 647-648
- ALOHA, 246-250
 - pure, 247-249,
 - satellite, 329
 - slotted, 249-250

- American National Standards Institute, 70
 - Amplitude modulation, 110
 - AMPS (*see* Advanced Mobile Phone System)
 - Analog cellular telephone, 157-161
 - Analog digital converter, 725
 - Anonymous remailer, 674
 - ANSI (*see* American National Standards Institute)
 - ANSNET, 51
 - Anycasting, 442
 - Apocalypse of the two elephants, 40-41
 - Applet, 707-709
 - Application gateway, 398, 411
 - Application layer, 33-34, 37, 577-766
 - domain name system, 622-630
 - email, 643-669
 - multimedia, 723-760
 - net news, 669-680
 - network management, 630-643
 - network security, 577-622
 - World Wide Web, 681-723
 - Architecture, network, 18
 - Area, OSPF, 425
 - ARP (*see* Address Resolution Protocol)
 - ARPANET, 35, 47-50, 71, 355, 569, 622
 - ARQ (*see* Automatic Repeat reQuest)
 - ASCII armor, 654
 - ASN.1 (*see* Abstract Syntax Notation-1)
 - ASN.1 transfer syntax, 637-638
 - Asymmetric digital subscriber line, 751
 - Asynchronous transfer mode, 61-65
 - congestion control, 467-471
 - control plane, 64
 - CS sublayer, 65
 - cell format, 450-452
 - connection setup, 452-455
 - data link layer, 235-239
 - leaky bucket, 466
 - NNI, 450-451
 - perspective, 65
 - PMD sublayer, 64
 - quality of service, 460-463
 - routing and switching, 455-458
 - SAR sublayer, 65
 - service categories, 458-460
 - TC sublayer, 64-65
 - traffic shaping, 463-468
 - UNI, 450-451
 - Asynchronous transfer mode (*cont.*)
 - user plane, 64
 - virtual channel, 450
 - virtual path, 450
 - ATM (*see* Asynchronous Transfer Mode)
 - ATM adaptation layer, 545-555
 - ATM Forum, 65
 - ATM LAN, 471-473
 - ATM layer, 449-473, 63
 - ATM network, 144-155
 - ATM switch, 147-155
 - ATMARP server, 473
 - Attenuation, 109
 - in fiber, 89
 - Audio CD, 724-725
 - Audio, digital, 724-726
 - Aurora, 55
 - Authentication protocol, 601-613
 - Kerberos, 610-612
 - using KDC, 607-620
 - public-key, 612-613
 - Authoritative record, 629
 - Automatic repeat request, 200-202
 - Autonomous system, 406, 412
 - Available bit rate service, 459-460
- B**
- B-frame, MPEG, 742
 - B-side carrier, 160
 - Backbone, OSPF, 425
 - Backward learning algorithm, 311-312
 - Balanced transmission, 115
 - Bandwidth-delay product, 557
 - Base64, 654
 - Baseband cable, 85, 86, 276, 277, 280
 - Basic rate, ISDN, 142-143
 - Batcher-banyan switch, 151-155
 - Baud, 79
 - Bell Operating Company, 106
 - Bell System, 103
 - Bellman-Ford routing, 355
 - BGP (*see* Border Gateway Protocol)
 - Big endian computer, 413
 - Binary countdown, 255-256
 - Binary exponential backoff, 282-283
 - Birthday attack, 618-620

Bit pipe, 140
 Bit stuffing, 181
 Bit-map protocol, 254-255
 BITNET, 53
 Blanca, 56
 Block cipher, 585, 590, 595-596
 BOC (*see* Bell Operating Company)
 Body, email, 646
 BOOTP, 424
 Border gateway protocol, 429-431
 Bridge, 304-318, 398

- between IEEE 802 LANs, 307-310
- remote, 317-318
- source routing, 314-316
- spanning tree, 310-313
- transparent, 310-313

 Broadband cable, 85-86
 Broadband ISDN, 61-65, 144-155
 Broadcast address, 280
 Broadcast network, 7-8
 Broadcast routing, 370-372
 Broadcast storm, 557
 Broadcast/unknown server, 472
 Broadcasting, 8
 Browser, World Wide Web, 682
 Bucket brigade attack, 606
 BUS (*see* Broadcast/Unknown Server)

C

Cable TV, 85-86, 107, 144, 172
 Care-of address, 433
 Carnegie-Mellon University, 7
 Carrier

- common, 67, 119
- modem, 110-111, 114

 Carrier sense multiple access protocols, 250-254
 CASA, 56
 Caesar cipher, 582-583
 CBR (*see* Constant Bit Rate service)
 CCITT, 68, 119, 121, 122, 124, 142, 644
 CDMA (*see* Code Division Multiple Access)
 CDPD (*see* Cellular Digital Packet Data)
 CDV (*see* Cell Delay Variation)

Cell

- ATM, 62
- cellular radio, 158
- HTML, 701

 Cell delay variation, 462
 Cell error rate, 463
 Cell loss ratio, 462
 Cell misinsertion rate, 463
 Cell relay, 62
 Cell transfer delay, 462
 Cell variation delay tolerance, 462
 Cellular digital packet data, 15, 269-271
 Cellular radio, 155-163

- digital, 266-275

 Cellular telephone

- AMPS, 158-161
- analog, 157-161
- call management, 160-161
- digital, 162
- security, 161

 Censorship

- by CMU, 7
- by CompuServe, 676

 Central office, 104
 CER (*see* Cell Error Rate)
 Certification authority, 668
 CGI (*see* Common Gateway Interface)
 Challenge-response protocol, 602-604
 Channel, 11
 Channel allocation in LANs, 244-246
 Channel associated signaling, 122
 Character stuffing, 180-181
 Checksum, 179, 182, 187-191, 235
 Chinese lottery, 593
 Chip, 272
 Chip sequence, 272
 Choke packet, 387-391
 Chosen plaintext attack, 582
 Chrominance, 728-729
 CIDR (*see* Classless InterDomain Routing)
 Cipher

- block, 588-597
- Caesar, 582-583
- substitution, 582-583
- transposition, 583-585

 Cipher block chaining, 590-591
 Cipher feedback mode, 591-592
 Ciphertext, 580

- Ciphertext only attack, 582
 - Circuit, 11
 - Circuit switching, 130-134
 - Classless interdomain routing, 434-437
 - Client-server model, 3-4
 - CLR (*see* Cell Loss Ratio)
 - CLUT (*see* Color Look Up Table)
 - CMR (*see* Cell Misinsertion rate)
 - Coaxial cable, 84-86
 - baseband, 84-85
 - broadband, 85-86
 - Code division multiple access, 271-275, 333
 - Codec, 121
 - Collision-free protocol, 254-256
 - Color look up table, 732
 - Common carrier, 67
 - Common gateway interface, 705-706
 - Common-channel signaling, 122
 - Communication satellite, 163-170
 - Communication subnet, 11
 - Composite video, 728
 - CompuServe, 676
 - Computer network, 2
 - use, 3-7
 - Concatenated virtual circuits, 401-402
 - Confirm, 25-27
 - Confirmed service, 26-27
 - Congestion control algorithm, 374-395
 - choke packets, 387-391
 - general principles, 376-378
 - in virtual circuit networks, 386-387
 - leaky bucket, 380-381
 - multicasting, 393-395
 - token bucket, 381-384
 - weighted fair queuing, 388-389
 - ATM, 467-471
 - rate-based, 469-471
 - TCP, 536-539
 - Congestion prevention, 378-379
 - Congestion threshold, 538
 - Congestion window, 537-538
 - Connection establishment, 493-498
 - TCP, 529-530
 - Connection release, 498-502, 530-533
 - Connection-oriented service, 23-25
 - Connectionless service, 23-25
 - Constant bit rate service, 458-459
 - Constellation pattern, 111-112
 - Contention system, 246-247, 252-258
 - Continuous media, 724
 - Control plane, ATM, 64
 - Convergence sublayer, AAL, 546
 - Copper wire, compared to fiber, 92-94
 - Cordless telephone, 157
 - Core-based tree, 374
 - Count-to-infinity problem, 357-358
 - Covert channel, 719-720
 - Crash recovery, 508-510
 - Crawler, 720
 - Credit message, 519
 - Crossbar switch, 135-138
 - Crosspoint, 136
 - Crosspoint switch, 135-138
 - Crossposting news, 672
 - Cryptanalysis, 581
 - Cryptography, 577-622
 - public-key, 597-601
 - secret-key, 587-597
 - traditional, 580-585
 - Cryptology, 581
 - CS sublayer, ATM, 65
 - CSMA (*see* Carrier Sense Multiple Access protocols)
 - CSMA/CD, 252-254
 - CSNET, 50
 - CTD (*see* Cell Transfer Delay)
 - CVDT (*see* Cell Variation Delay Tolerance)
 - Cyclic redundancy code, 187
- D**
- Data circuit-terminating equipment, 114
 - Data compression, 730-744
 - CLUT, 732
 - differential encoding, 732
 - discrete cosine transformation, 733
 - entropy encoding, 731-732
 - lossless, 731-732
 - lossy, 732-734
 - run-length encoding, 731
 - source encoding, 732-734
 - statistical encoding, 731-732
 - transformation encoding, 732
 - vector quantization, 733

- Data encryption standard, 588-595
 - attacking, 592-595
 - chaining, 589-592
 - controversy, 593
 - Data flow machine, 8
 - Data link layer, 175-242
 - ATM, 235-239
 - bit stuffing, 181
 - character stuffing, 180-181
 - design issues, 176-183
 - elementary protocols, 190-202
 - example protocols, 225-239
 - error control, 182-183
 - flow control, 183
 - framing, 179-182
 - HDLC protocol, 225-228
 - LLC, 275, 302-304
 - OSI, 30
 - services provided, 176-179
 - sliding window protocols, 202-219
 - Data switching exchange, 12
 - Data terminating equipment, 114
 - Datagram, 342
 - compared to virtual circuit, 344-345
 - Datagram service, 24-25
 - DCE (*see* Data Circuit-Terminating Equipment)
 - DCS 1800, 266
 - De facto standard, 67
 - De jure standard, 67
 - Deadlock, protocol, 222
 - Decibel, 81, 724
 - Decoding, 730
 - Delay distortion, 109
 - Delta modulation, 123-124
 - DES (*see* Data Encryption Standard)
 - Designated router, 428
 - Differential cryptanalysis, 595
 - Differential encoding, 732
 - Differential Manchester encoding, 279-280
 - Differential PCM, 123
 - Diffie-Hellman key exchange, 605-606
 - Digital cellular radio, 266-275
 - Digital cellular telephone, 162
 - Digital sense multiple access, 270-271
 - Digital signature, 613-620
 - public-key, 615-616
 - secret-key, 614-615
 - Digital signature standard, 616
 - Digram, 583
 - Directive, HTML, 696-699
 - Directory server, 491
 - Discrete cosine transformation, 733
 - Disk farm, 748
 - Disk striping, 748
 - Dispersion, in fiber, 89
 - Distance vector multicast routing protocol, 758-759
 - Distance vector routing, 355-359
 - Distributed mail system protocol, 662
 - Distributed queue dual bus, 11, 301-303
 - Distributed system, 2
 - Distribution network, 750-754
 - DMSP (*see* Distributed Mail System Protocol)
 - DNS (*see* Domain Name System)
 - Domain, 623
 - Domain name system, 421, 622-630
 - Dotted decimal notation, 417
 - DQDB (*see* Distributed Queue Dual Bus)
 - DS1, 121,
 - DSMA (*see* Digital Sense Multiple Access)
 - DSS (*see* Digital Signature standard)
 - DTE (*see* Data Terminating Equipment)
 - DVMRP (*see* Distance Vector Multicast Routing Protocol)
- E**
- EARN, 53
 - Establishing a connection, TCP, 529-530
 - Echo canceller, 113
 - Echo suppressor, 112-113
 - Electromagnetic spectrum, 94-97
 - Electronic code book mode, 590
 - Electronic mail (*see* Email)
 - Elephants, apocalypse, 40
 - Email, 5, 53, 643-670
 - architecture and services, 645-646
 - body, 646
 - early systems, 644
 - envelope, 646
 - filter, 662
 - final delivery, 662-663
 - functions, 645
 - gateway, 659-661

- Email (*cont.*)
 - header, 646
 - message format, 650-658
 - message transfer agent, 645
 - message transfer, 657-663
 - MIME format, 653-657
 - privacy, 663-669
 - reading, 648-650
 - RFC 822 format, 651-653
 - sending, 646-648
 - user agent, 645, 646-648
 - user commands, 648-650
- Emoticon, 674
- Encapsulation, Java, 713
- Encoding, 730
 - entropy, 731-732
 - source, 732-734
- Encryption (*see* cryptography)
- End office, 104
- End system, 11
- Entity, 22
- Entropy encoding, 731-732
- Envelope, email, 646
- ER (*see* Explicit Rate)
- Error control, 182-190
- Error correction, 184-190
- Error detection, 183-184
- Error-correcting code, 184
- Error-detecting code, 184, 186-190
- Establishing a connection, 493-498
 - TCP, 529-531
- Ethernet, 10, 276 (*see also* IEEE 802.3)
- Explicit rate, 471
- Exposed station problem, 264
- Extended SMTP, 659
- Exterior gateway protocol, 405-406, 424, 429-431
- External viewer, 684

- F**
- Fabry-Perot interferometer, 91, 261
- FAQ (*see* Frequently Asked Questions)
- Fast Ethernet, 322-324
- Fast TPDU processing, 565-568
- FCC (*see* Federal Communications Commission)
- FDDI (*see* Fiber Distributed Data Interface)
- FDM (*see* Frequency Division Multiplexing)
- Federal Communications Commission, 96, 100, 167
- Fiber cable, 90-91
- Fiber channel (*see* Fibre channel)
- Fiber distributed data interface, 319-322
- Fiber optic network, 91-94
- Fiber optics, 87-94
 - attenuation, 88-89
 - basic principles, 87-88
 - compared to copper, 92-94
 - compared to satellite, 168-170
 - dispersion, 89
 - FDDI, 319-322
 - multimode, 88
 - single-mode, 88
 - SONET, 125-130
 - system components, 86
 - WDM, 119-121
- Fiber to the curb, 116-118, 120, 751-752
- Fiber to the home, 116-118, 752-753
- Fibre channel, 326-327
- Field, video, 728
- File server, 3
- File transfer, 53
- File transfer protocol, 693
- Finite state machine, 219-223, 519-521
- Firewall, 410-412
- Flamewar, 672
- Flat address, 492
- Flooding, 351
- Flow control, 183, 502-506
- Flow specification, 384-386
- Flow-based routing, 353-355
- Flying LAN, 15
- Ford-Fulkerson routing, 355
- Foreign agent, 368
- Form, HTML, 701-706
- Fourier analysis, 78
- Fragmentation, internetwork, 406-409
- Frame
 - acknowledgement, 30
 - data, 30
 - video, 727
- Frame header, 192
- Frame relay, 60-61
- Framing, 179-182

Frequency, 94
 Frequency band, 95
 Frequency division multiplexing, 118-121, 330
 Frequency modulation, 110
 Frequently asked questions, 674
 FTP (*see* File Transfer Protocol)
 FTTC (*see* Fiber To The Curb)
 FTTH (*see* Fiber To The Home)
 Full-duplex communication, 21, 113
 Fuzzball, 50

G

Gateway, 16
 GCRA (*see* Generic Cell Rate Algorithm)
 Generator polynomial, 187
 Generic cell rate algorithm, 463-466
 Geosynchronous satellite, 164-167
 Gigabit network, 54-56, 568-572
 Global system for mobile communications, 266-275
 Go back n protocol, 207-213
 Gopher, 693
 Gratuitous ARP, 433
 Group, 119
 GSM (*see* Global System for Mobile Communications)

H

Half-duplex communication, 21, 113
 Half-gateway, 398
 Hamming distance, 184
 Handoff, cellular, 169
 HDLC (*see* High-Level Data Link Control)
 HDTV (*see* High Definition TeleVision)
 Head-end, cable, 85-86
 Head-of-line blocking, 149
 Header, 19
 email, 646
 frame, 192
 Header error control, ATM, 235-238
 Header prediction, 567
 HEC (*see* Header Error Control, ATM)
 Helper application, 684
 HEPNET, 53

HFC (*see* Hybrid Fiber Coax)
 Hidden station problem, 264
 Hierarchical address, 492
 Hierarchical routing, 365-367
 High definition television, 729
 High-level data link control, 225-228
 High-performance parallel interface, 325-326
 High-speed LAN, 318-327
 HIPPI (*see* High-Performance Parallel Interface)
 Home agent, 368
 Host, 11
 Host-to-network layer, 38
 HTML (*see* HyperText Markup Language)
 HTTP (*see* HyperText Transfer Protocol)
 Hub, satellite, 165
 Hybrid fiber coax, 752
 Hyperlink, 682
 Hypermedia, 684
 Hypertext, 682
 Hypertext markup language, 691-706
 forms, 701-706
 versions, 699-701
 Hypertext transfer protocol, 689-691

I

I Love Lucy, 745
 I-frame, MPEG, 740
 IAB (*see* Internet Architecture Board)
 IBM, 41, 226, 307, 588, 593-594
 ICMP (*see* Internet Control Message Protocol)
 IDEA (*see* International Data Encryption Algorithm)
 IDU (*see* Interface Data Unit)
 IEEE, 70
 IEEE 802, 275-301
 comparison of LANs, 299-301
 IEEE 802.2, 302-304
 IEEE 802.3, 276-287
 cabling, 276-279
 fast Ethernet, 322-324
 frame format, 281
 performance, 283-285
 protocol, 280-283
 signal encoding
 switched, 285-287

- IEEE 802.3u, 322-324
- IEEE 802.4, 287-292
 - protocol, 288-290
 - ring maintenance, 290-292
- IEEE 802.5, 292-299
 - protocol, 296-298
 - ring maintenance, 298-299
- IEEE 802.6, 301-303
- IETF (*see* Internet Engineering Task Force)
- IGMP (*see* Internet Group Management Protocol)
- IMAP (*see* Interactive Mail Access Protocol)
- IMP (*see* Interface Message Processor)
- Improved Mobile Telephone Service, 157
- IMTS (*see* Improved Mobile Telephone Service)
- In-band signaling, 113
- Indication, 25-27
- Indirect TCP, 543-544
- Industrial/Scientific/Medical band, 99
- Information frame, 226-227
- Infrared transmission, 100
- Inheritance, Java, 713
- Initial connection protocol, 490
- Initialization vector, 590
- Integrated Services Digital Network, 61, 139-155
- Interactive mail access protocol, 662
- Interexchange carrier, 106-107
- Interface, between layers, 18
- Interface data unit, 22
- Interface message processor, 47
- Interferometer
 - Fabry-Perot, 91, 261
 - Mach-Zehnder, 91, 261
- Interior gateway protocol, 405-406, 424-429
- Interlaced video, 728
- Intermediate system, 12
- International data encryption algorithm, 596-597
- International standard, 70
- International Standards Organization, 69
- International Telecommunication Union, 68
 - 453, 471, 545, 634, 668, 734
- Internet, 16
 - CIDR, 434-437
 - connection management, 529-533
 - data link layer, 229-235
- Internet (*cont.*)
 - history, 52-54
 - internet layer, 35-36, 412-449
 - IP, 36, 412-419
 - IPv6, 437-449
 - mobile IP, 432-434
 - multicasting, 431-432
 - routing protocols, 424-431
 - TCP, 36-37, 521-542
- Internet applications
 - email, 643-670
 - MBone, 756-760
 - net news, 669-680
 - World Wide Web, 681-723
- Internet Architecture Board, 71
- Internet control message protocol, 419-420
- Internet Engineering Task Force, 71
- Internet group management protocol, 431-432, 759
- Internet layer, 35-36
- Internet policy registration authority, 668
- Internet protocols
 - ARP, 420-423, 433
 - BGP, 429-431
 - DVMRP, 758-759
 - HTTP, 689-691
 - ICMP, 419-420
 - IGMP, 431-432, 759
 - IP, 36, 412-419
 - NNTP 677-680
 - OSPF, 424-429
 - PIM, 760
 - PPP, 231-235, 685
 - RARP, 423-424
 - RSVP, 394-395
 - SLIP, 229-230, 685
 - SMTP 658-660
 - TCP 36-37, 521-542, 658, 678, 685
 - UDP 37, 542-544
- Internet service provider, 229
- Internet Society, 53, 71
- Internet transport protocol, 521-545
- Internetwork routing, 405-406
- Internetwork, 16
- Internetworking, 396-412
 - connection-oriented, 401-402
 - connectionless, 401-402
 - why needed, 399-400

Interoffice trunk, 104
 Intertoll trunk, 104
 Intruder, 580
 IP (*see* Internet Protocol)
 IP address, 416-419
 IPRA (*see* Internet Policy Registration Authority)
 IPv4, 413-419
 IPv5, 438
 IPv6, 437-449
 addresses, 441
 controversies, 447-449
 extension header, 443-446
 jumbogram, 445
 main header, 439-443
 IPX, 46
 IS-IS routing, 365
 ISDN (*see* Integrated Services Digital Network)
 ISM band (*see* Industrial/Scientific/Medical band)
 ISO (*see* International Standards Organization)
 ISO standards
 ISO 3166, 623
 ISO 8802, 70, 275
 ISO 8859-1, 696
 ITU (*see* International Telecommunication Union)
 ITU-R, 68
 ITU-T, 68
 IXC (*see* IntereXchange Carrier)

J

Jacobson's slow start algorithm, 538-539
 Java, 706-720
 abstract window toolkit, 717-718
 API, 716-718
 class, 713-716
 language description, 709-718
 object orientation, 712-716
 polymorphism, 715
 security, 718-720
 Jitter, 385, 724
 Jitter control, 392-393
 JPEG standard, 734-738
 Jumbogram, 445

K

Karn's algorithm, 541
 KDC (*see* Key Distribution Center)
 Keepalive timer, TCP, 542
 Kerberos, 610-612
 Key, cryptographic, 580
 Key distribution center, 607-610
 Killfile, 672
 Knockout switch, 150-151
 Knowbot, 720
 Known plaintext attack, 582

L

LAN (*see* Local Area Network)
 LAN Emulation Server, 472
 LANs, comparison, 299-301
 LAP (*see* Link Access Procedure)
 LATA (*see* Local Access and Transport Area)
 Layer, 17
 application, 33-34, 37, 577-766
 data link, 175-242
 network, 31, 35-36, 339-478
 physical, 29-30, 77-174
 presentation, 33
 session, 32-33
 transport, 31-32, 36-37, 479-576
 LCP (*see* Link Control Protocol)
 Leaky bucket algorithm, 380-381
 LEC (*see* Local Exchange Carrier)
 LES (*see* LAN Emulation Server)
 Lightwave transmission, 100-102
 Limited contention protocol, 256-259
 Line, SONET, 126
 Line sublayer, SONET, 129-130
 Linear cryptanalysis, 595
 Link access procedure, 226
 Link control protocol, 231
 Link encryption, 579
 Link state routing, 359-365
 LIS (*see* Logical IP Subnet)
 Little endian computer, 413
 LLC (*see* Logical Link Control)
 Load shedding, 390-392
 Local Access and Transport Area, 106

- Local area network, 9-10
 - ATM, 471-473
 - channel allocation, 244-246
 - Ethernet, 10, 276-287
 - fast Ethernet, 322-324
 - IEEE 802, 275-304
 - high-speed, 318-327
 - token bus, 287-292
 - token ring, 292-299
 - Local central office, 104
 - Local exchange carrier, 106-107
 - Local loop, 104, 108-118
 - fiber, 115-118
 - Logical IP subnet, 473
 - Logical Link Control, 275, 302-304
 - Low-orbit satellite, 167-170
 - Luminance, 728-729
 - Luminiferous ether, 276
- M**
- MAC sublayer (*see* Medium Access Control sublayer)
 - MACA (*see* Multiple Access with Collision Avoidance)
 - MACAW, 265
 - Mach-Zehnder interferometer, 91, 261
 - Macroblock, 740-741
 - Mailbox, 645
 - Mailing list, 645-646
 - Mailto, 693
 - MAN (*see* Metropolitan Area Network)
 - Man-in-the-middle attack, 606
 - Management information base, 632, 641-642
 - Management station, 631
 - Manchester encoding, 279-280
 - Markup language, 695
 - Mastergroup, 119
 - Maximum transfer unit, 525
 - MBone (*see* Multicast Backbone)
 - MCR (*see* Minimum Cell Rate)
 - MD5, 618, 665
 - Medium access sublayer, 243-335
 - Meet-in-the-middle attack, 594
 - Message digest, 617-618
 - Message switching, 131-133
 - Message transfer agent, 645
 - Method, HTTP, 690
 - Metropolitan area network, 10-11
 - MIB (*see* Management Information Base)
 - Microwave transmission, 98-99
 - MIDI (*see* Music Instrument Digital Interface)
 - Midsplit cable, 86
 - Milk policy, 390
 - Millimeter wave, 100
 - MILNET, 50
 - MIME (*see* Multipurpose Internet Mail Extensions)
 - Minimum cell rate, 461
 - MNP 5, 112
 - Mobile host, routing algorithm, 367-370
 - Mobile IP, 432-434
 - Mobile switching center, 159
 - Mobile telephone switching office, 159
 - Modem, 109-113
 - Moderated newsgroup, 673
 - Modified final judgment, 104
 - Modulation, 110-112
 - amplitude, 110
 - frequency, 110
 - phase, 110
 - Monoalphabetic substitution cipher, 582-583
 - Mosaic, 696
 - MOSPF (*see* Multicast OSPF)
 - MOTIS, 644
 - MPEG standard, 738-744, 753-754
 - B-frame, 742
 - I-frame, 740
 - macroblock, 740-741
 - MPEG-1, 738-742
 - MPEG-2, 742-744
 - P-frame, 740-741
 - profiles, 742
 - streams, 743
 - Mrouter (*see* Multicast router)
 - MSC (*see* Mobile Switching Center)
 - MTSO (*see* Mobile Telephone Switching Office)
 - MTU (*see* Maximum Transfer Unit)
 - Multiaccess channel, 243
 - Multiaccess network, 425
 - Multicast addresses, 280
 - Multicast backbone, 756-760
 - Multicast OSPF, 760
 - Multicast router, 757-758

Multicast routing, 372-374
 Multicasting, 8, 372, 393-395
 Internet, 431-432
 Multicomputer, 8
 Multidestination routing, 370
 Multimedia, 723-760
 audio, 724-726
 data compression, 730-744
 MBone, 756-760
 video, 727-730
 video on demand, 744-756
 Multimode fiber, 88
 Multipath fading, 99
 Multiple access protocols, 246-275
 Multiple access with collision avoidance,
 264-265
 Multiplexing, 118-130, 506-508
 downward, 507
 upward, 506
 Multiprotocol router, 398
 Multipurpose internet mail extensions,
 653-657
 Music instrument digital interface, 726

N

N-ISDN (*see* Narrowband ISDN)
 Nagle's algorithm, 534-535
 NAK (*see* Negative AcKnowledgeMENT)
 Name server, 491, 628-630
 NAP (*see* Network Access Point)
 Narrowband ISDN, 139-144
 National Institute of Standards and
 Technology, 70
 National Security Agency, 593
 National Television Standards
 Committee, 728-729
 NCP (*see* Network Control Protocol)
 NCP (*see* Network Core Protocol)
 Near video on demand, 744
 Nectar, 56
 Needham-Schroeder protocol, 608-609
 Negative acknowledgement, 215
 Negotiation, 26
 NETBLT, 572
 NetWare, Novell, 45-47
 Network access point, 52
 Network architecture, 18
 Network control protocol, 231
 Network core protocol, 46
 Network information center, 417
 Network layer, 339-478
 ATM networks, 449-473
 congestion control, 374-395
 design issues, 339-345
 internal organization, 342-345
 Internet, 412-449
 internetworking, 396-412
 OSI, 31
 routing algorithms, 345-374
 services provided, 340-342
 Network news (*see* USENET)
 Network news transfer protocol, 677-680
 Network performance, 555-572
 IEEE 802.3, 283-285
 Network security, 577-622
 Network service access point, 489
 Network standardization, 66-72
 Network virtual terminal, 33
 Network, fiber optic, 91-94
 News, 53, 669-680, 693, 694
 News article, example, 676
 News headers, 676-677
 Newsfeed, 675
 Newsgroup
 creation, 674-675
 example, 673
 NIC (*see* Network Information Center)
 NIST (*see* National Institute of Standards
 and Technology)
 NNTP (*see* Network News Transfer Protocol)
 Noise, 109
 Nonadaptive routing, 347
 Nonce, 608
 Novell NetWare, 45-47
 NREN, 51-52
 NSA (*see* National Security Agency)
 NSAP (*see* Network Service Access Point)
 NSFNET, 50-52
 NT1, 140-142
 NT2 141, 142
 NTSC (*see* National Television Standards
 Committee)
 Null modem, 114
 Nyquist limit, 81

O

OAM cell (*see* Operation And Maintenance cell)
 Object
 Java, 713
 SNMP, 632, 641-642
 OC-*n* (*see* Optical Carrier)
 One-bit sliding window protocol, 205-207
 One-time pad, 585
 ONU (*see* Optical Network Unit)
 Open shortest path first, 424-429
 Operation and maintenance cell, 236
 Optical carrier, 128-129
 Optical fiber (*see also* Fiber optics)
 multimode, 88
 single-mode, 88
 Optical network unit, 751-752
 Optimality principle, 347-348
 Option negotiation, 483
 Oryctolagus cuniculus, 18
 OSI reference mode, 28-35
 compared to TCP/IP, 38-39
 critique, 40-43
 OSPF (*see* Open Shortest Path First)
 Otway-Rees protocol, 609-610
 Output feedback mode, 593
 Overloading, Java, 715

P

P-box, 587-588
 P-frame, MPEG, 740-741
 Package, Java, 713
 Packet, 7
 Packet assembler disassembler, 60
 Packet filter, 411
 Packet switching, 133-134
 Packet switching node, 12
 Packet-switched subnet, 12
 Packetized elementary stream, 743
 PAD (*see* Packet Assembler Disassembler)
 Paging system, 155-156
 PAL (*see* Phase Alternating Line)
 Parity bit, 185
 Passive star, 92
 Path, SONET, 126

Path sublayer, SONET, 129-130
 PBX (*see* Private Branch eXchange)
 PCA (*see* Policy Certification Authority)
 PCM (*see* Pulse Code Modulation)
 PCN (*see* Personal Communications Network)
 PCR (*see* Peak Cell Rate)
 PCS (*see* Personal Communications Services)
 PDU (*see* Protocol Data Unit)
 Peak cell rate, 461
 Peer, 17
 Peer entity, 22
 PEM (*see* Privacy Enhanced Mail)
 Performance issues, 555-572
 Permanent virtual circuit, 60, 145-146
 Persistence timer, TCP, 542
 Personal communications network, 162-163
 Personal communications services, 162-163
 PES (*see* Packetized Elementary Stream)
 Petri net model, 223-224
 PGP (*see* Pretty Good Privacy)
 Phase alternating line, 728-729
 Phase modulation, 110
 Photonic sublayer, SONET, 129
 Physical layer, 77-174
 cellular radio, 155-163
 communication satellites, 163-170
 OSI, 29-30
 telephone system, 102-163
 transmission media, 82-94
 wireless transmission, 94-102
 Physical medium, 18
 Piggybacking, 202-203
 PIM (*see* Protocol Independent Multicast)
 Pipelining, 209
 Pixel, 729
 Plain old telephone service, 142
 Plaintext, 580
 PMD sublayer, ATM, 64, 147, 235-239
 Point of presence, 107
 Point-to-point network, 8
 Point-to-point protocol, 231-235, 685
 Point-to-point subnet, 12
 Policy certification authority, 668
 Politics, 43
 Polling, 328
 Polymorphism, Java, 715
 Polynomial code, 187
 POP (*see* Point of Presence)

- POP3 (*see* Post Office Protocol-3)
- Port, TCP, 523
 - well-known, 523
- Portapotty, 15
- Post office protocol-3, 662
- Post Telegraph and Telephone Administration, 67
- POTS (*see* Plain Old Telephone Service)
- PPP (*see* Point-to-Point Protocol)
- Predictive encoding, 124
- Presentation layer, 33
- Pretty good privacy, 664-667
 - compared to PEM, 669-670
- Primary rate, ISDN, 142-143,
- Primitives, service, 25-27
- Principal, 601
- Privacy enhanced mail, 667-669
 - compared to PGP, 669-670
- Private branch exchange, 142
- Private key ring, PGP, 666
- Process server, 491
- Program stream, MPEG, 743
- Promiscuous mode, 306
- Protocol, 17
 - 1-bit, 205-207
 - 802.5, 296-298
 - AAL, 545-555
 - ADCCP, 226
 - ARP, 420-423
 - ARQ, 200-202
 - ATM AAL, 547-554
 - authentication, 601-613
 - BGP, 429-431
 - binary countdown, 255-256
 - bit-map, 254-255
 - BOOTP, 424
 - challenge-response, 602-604
 - collision-free, 254-256
 - CSMA, 250-254
 - DMSP, 662
 - DSMA, 270-271
 - DVMRP, 758-759
 - elementary data link, 190-202
 - exterior gateway, 405-406
 - gigabit network, 568-572
 - go back n, 207-213
 - HDLC, 225-228
 - HTTP, 689-691
- Protocol (*cont.*)
 - ICMP, 419-420
 - IEEE 802.3, 280-283
 - IEEE 802.4, 288-290
 - IEEE 802.5, 296-298
 - IGMP, 431-432
 - IMAP, 662
 - interior gateway, 405-406
 - IP, 36, 412-419
 - IPX, 46
 - LAP, 226
 - LCP, 231
 - limited contention, 256-259
 - MACA, 264-265
 - MACAW, 265
 - multiple access, 246-275
 - NCP (Network Control Protocol), 231
 - NCP (Network Core Protocol), 46
 - Needham-Schroeder, 608-609
 - NNTP, 677-680
 - noisy channel, 197-200
 - Otway-Rees, 609-610
 - PAR, 200-202
 - PIM, 760
 - POP3, 662
 - PPP, 231-235
 - Q.2931, 453
 - RARP, 423-424
 - RSVP, 394-396
 - SSCOP, 555
 - SDLC, 226-227
 - selective repeat, 213-219
 - sliding window, 202-219
 - SLIP, 229-230
 - SMTP, 658-660
 - SNMP, 642-643
 - TCP, 36-37, 521-542
 - tree walk, 258-259
 - UDP, 37, 542-544
 - unrestricted simplex, 195-197
 - WDMA, 260-262
 - wireless LAN, 262-265
- Protocol data unit, 22-23
- Protocol hierarchy, 17-20
- Protocol independent multicast, 760
- Protocol stack, 18
- Protocol verification, 219-224
- Proxy ARP, 423

- Proxy server, 688
- Pruning, 760
- PSTN (*see* Public Switched Telephone Network)
- PTT (*see* Post, Telegraph, and Telephone)
- Public key ring, 667
- Public switched telephone network, 102
- Public-key cryptography, 597-601
- Pulse code modulation, 121
- Push-to-talk system, 157
- PVC (*see* Permanent Virtual Circuit)

- Q**
- Q.2931, 453
- QAM (*see* Quadrature Amplitude Modulation)
- QoS (*see* Quality of Service)
- Quadrature amplitude modulation, 111
- Quality of service, 23, 460-463, 481-483
 - ATM, 460-463
- Quantization noise, 725
- Quoted printable encoding, 654

- R**
- Radio transmission, 97-98
- RAID (*see* Redundant Array of Inexpensive Disks)
- Random access channel, 243
- RARP (*see* Reverse Address Resolution Protocol)
- Rate-based congestion control, 469-471
- Reachability analysis, 20
- Realm, Kerberos, 611
- Receiving window, 203
- Recursive query, 630
- Redundant array of inexpensive disks, 748
- Reference model, 28-44
 - B-ISDN, 63-65
 - comparison of OSI and TCP/IP, 38-39
 - OSI 28-35
 - TCP/IP, 35-38
- Reference point, ISDN, 142
- Reference station, 329
- Reflection attack, 603
- Releasing a connection, 498-502
 - TCP, 530-533
- Remote login, 53
- Repeater, 91-94, 279, 398
- Replay attack, 608
- Request, 25-27
- Request for comment, 71
- Request-reply service, 24-25
- Resolver, DNS, 622
- Resource management cell, 470
- Resource record, 624-628
- Resource reservation, 468-469
- Resource reservation protocol, 394-395
- Response, 25-27
- Retransmission timer, TCP, 539-540
- Reverse address resolution protocol, 423-424
- Reverse path forwarding, 371-372
- RFC (*see also* Request For Comment)
 - RFC 768, 542
 - RFC 792, 420
 - RFC 793, 522
 - RFC 821, 651, 659, 761, 644
 - RFC 822 644, 650, 651-653, 655, 660, 661, 665, 667, 676, 677, 688, 689 690, 691, 761
 - RFC 826, 422
 - RFC 903, 423
 - RFC 951, 424
 - RFC 977, 677
 - RFC 1028, 630
 - RFC 1034, 622
 - RFC 1035, 622
 - RFC 1036, 676
 - RFC 1048, 424
 - RFC 1055, 229
 - RFC 1056, 662
 - RFC 1064, 662
 - RFC 1067, 630
 - RFC 1084, 424
 - RFC 1106, 528, 529
 - RFC 1112, 432
 - RFC 1122, 522
 - RFC 1144, 230
 - RFC 1155, 630
 - RFC 1157, 630
 - RFC 1213, 642
 - RFC 1225, 662
 - RFC 1247, 424

- RFC (*cont.*)
 - RFC 1268, 431
 - RFC 1323, 522, 528
 - RFC 1421, 667
 - RFC 1422, 667
 - RFC 1423, 667
 - RFC 1424, 667
 - RFC 1425, 659
 - RFC 1441, 630
 - RFC 1442, 630, 639
 - RFC 1443, 630
 - RFC 1444, 630
 - RFC 1445, 630
 - RFC 1446, 630
 - RFC 1447, 630
 - RFC 1448, 630, 643
 - RFC 1449, 630
 - RFC 1450, 630
 - RFC 1451, 630
 - RFC 1452, 630
 - RFC 1483, 473, 554
 - RFC 1519, 435
 - RFC 1521, 653, 654, 655
 - RFC 1550, 437
 - RFC 1577, 342, 473, 554
 - RFC 1654, 431
 - RFC 1661, 231, 234
 - RFC 1662, 231
 - RFC 1663, 231
 - RFC 1700, 415, 523
 - RFC 1715, 443
 - RFC 1883, 438
 - RFC 1884, 438
 - RFC 1885, 438
 - RFC 1886, 438
 - RFC 1887, 438
 - Ring, star-shaped, 295
 - RM cell (*see* Resource Management cell)
 - Rock 'n roll, signal-to-noise ratio, 739
 - Routing algorithm, 345-374
 - adaptive, 347
 - broadcast, 370-372
 - distance vector, 355-359
 - flooding, 351
 - flow-based, 353-355
 - hierarchical, 365-367
 - internetwork, 405-406
 - link state, 359-365
 - Routing algorithm (*cont.*)
 - Mobile host, 367-370
 - multicast, 372-374
 - nonadaptive, 347
 - reverse path forwarding, 371-372
 - shortest path
 - RS-232, 114-116
 - RS-422-A, 115
 - RS-423-A, 115
 - RS-449, 115-116
 - RSA algorithm, 598-600, 665-666
 - RSVP (*see* Resource reSerVation Protocol)
 - Run-length encoding, 731
- S**
- S-box, 587-588
 - SABME (*see* Set Asynchronous Balanced Mode Extended)
 - SAP (*see* Service Access Point)
 - SAR sublayer, ATM, 65, 546
 - Satellite network, 327-333
 - communication, 163-179
 - compared to fiber, 168-170
 - geosynchronous, 164-167
 - low-orbit, 167-70
 - SCR (*see* Sustained Cell Rate)
 - SDH (*see* Synchronous Digital Hierarchy)
 - SDLC (*see* Synchronous Data Link Control)
 - SDU (*see* Service Data Unit)
 - SEAL (*see* Simple Efficient Adaptation Layer)
 - Search engine, World Wide Web, 720-723
 - SECAM (*see* SEquentiel Couleur Avec Memoire)
 - SECBR (*see* Severely-Errored Cell Block Ratio)
 - Secret-key cryptography, 587-597
 - Section, SONET, 126
 - Section sublayer, SONET, 129-130
 - Secure hash algorithm, 618
 - Security
 - cellular telephone, 161
 - Java, 718-720
 - Segment, TCP, 525
 - Selective flooding, 351

- Selective repeat, 209
- Selective repeat protocol, 213-219
- Sending window, 203
- Sequentiel couleur avec memoire, 728-729
- Serial line IP, 229-230, 685
- Service
 - connection-oriented, 23-25
 - connectionless, 23-25
 - datagram, 24-25
 - request-reply, 24-25
- Service access point, 22
- Service data unit, 22
- Service primitive, 25-27
 - example, 510-512
- Service provider, 22
- Service user, 22
- Service-specific connection-oriented protocol, 555
- Session key, 602
- Session layer, OSI, 32-33
- Session routing, 346
- Set-top box, 754-756
- Set asynchronous balanced mode extended, 228
- Set normal response mode extended, 228
- Severely-errored cell block ratio, 463
- SGML (*see* Standard Generalized Markup Language)
- SHA (*see* Secure Hash Algorithm)
- Shannon limit, 81-82
- Shell account, 229
- Shortest path routing, 348-352
- Signal-to-noise ratio, 81
- Signature, digital, 613-620
- Silly window syndrome, 534-535
- Simple Efficient Adaptation Layer, 552-554
- Simple internet protocol plus, 438
- Simple mail transfer protocol, 658-660
- Simple network management protocol, 632-643
- Simplex communication, 21
- Single-mode fiber, 88
- SIPP (*see* Simple Internet Protocol Plus)
- Sliding window protocol, 202-219
 - 1-bit, 205-207
- SLIP (*see* Serial Line IP)
- Slow start algorithm, 538-539
- SMDS (*see* Switched Multimegabit Data Service)
- SMI (*see* Structure of Management Information)
- Smiley, 674
- SMTP (*see* Simple Mail Transfer Protocol)
- SNA (*see* Systems Network Architecture)
- SNMP (*see* Simple Network Management Protocol)
- SNMP agent, 631
- SNMP protocol, 642-643
- SNRME (*see* Set Normal Response Mode Extended)
- Social issues, 6-7
 - related to cryptography, 620-622
- Socket, 486-487
- Software, network, 16-28
- Soliton, 89
- SONET (*see* Synchronous Optical Network)
- Source encoding, 732-734
- Source routing, 415-416
- Source routing bridge, 314-316
- Space division switch, 136-138
- SPADE, 330
- SPAN, 53
- Spanning tree, 371
- SPE (*see* Synchronous Payload Envelope)
- Speed of light, 94
- Spider, 720
- Split horizon, 358-359
- Spot beam, 165
- Spread spectrum, 96
 - direct sequence, 96
- SPX, 46
- SSCOP (*see* Service-Specific Connection-Oriented Protocol)
- Standard generalized markup language, 695
- Standardization
 - Internet, 70-72
 - ISO, 69-70
 - network, 66-72
 - telecommunications, 67-69
 - video on demand, 756-757
- Star-shaped ring, 295
- Statistical encoding, 731-732
- Storage hierarchy, 746-747
- Store-and-forward subnet, 12

- Store-and-forward switching, 133
 - Striping, 748
 - Structure of management information, 639-641
 - STS-1 (*see* Synchronous Transport Signal-1)
 - Stub network, 430
 - Style sheet, 698
 - Subclass, Java, 713
 - Subnet, 11
 - Internet, 417-419
 - Subnet mask, 419
 - Subsplit cable, 85
 - Substitution cipher, 582-583
 - Superclass, Java, 713
 - Supergroup, 119
 - Supervisory frame, 226-228
 - Sustained cell rate, 461
 - Switch
 - crossbar, 135-138
 - space division, 136-138
 - time division, 138-139
 - Switch hierarchy, telephone, 134-135
 - Switched Ethernet, 285-287
 - Switched Multimegabit Data Service, 57-59
 - Switched virtual circuit, 60, 145-146
 - Switching
 - circuit, 130-134
 - message, 131-133
 - packet, 133-134
 - store-and-forward, 133
 - telephone, 130-139
 - Switching fabric, ATM, 148
 - Symmetric key cryptography, 598
 - Synchronization, 33
 - Synchronous data link control, 226-227
 - Synchronous digital hierarchy, 125-130
 - Synchronous optical network, 125-130
 - Synchronous payload envelope, 126-127
 - Synchronous transport signal-1, 126
 - Systems Network Architecture, 41
- T**
- T1 carrier, 121-122
 - T2 carrier, 124
 - T3 carrier, 124
 - T4 carrier, 124
 - Tag, HTML, 696-699
 - Tandem office, 104
 - Tariff, 67
 - TC sublayer, ATM, 64-65, 235-239
 - TCP (*see* Transmission Control Protocol)
 - TCP/IP reference model, 35-38, 43-44
 - compared to OSI, 38-39
 - TDM (*see* Time Division Multiplexing)
 - Telecommunications standardization, 67-69
 - Telephone, cellular, 157-163
 - Telephone system, 102-163
 - local loop, 108-118
 - politics, 106-108
 - SONET, 125-130
 - switching, 130-139
 - T1 carrier, 121-122
 - trunks and multiplexing, 118-130
 - Television
 - analog, 727-729
 - digital, 729-730
 - Telnet, 686-687, 693, 694
 - Terminal interface processor, 48
 - Thin Ethernet, 277
 - Three-way handshake, 496-498
 - Time division multiplexing, 118, 121-124, 330-333
 - Time division switch, 138-139
 - Time domain reflectometry, 277
 - Timer, token, 321
 - Timing wheel, 567-568
 - TIP (*see* Terminal Interface Processor)
 - Toaster-on-a-pole, 168
 - Token, 287-288, 293
 - Token bucket algorithm, 381-384
 - Token bus LAN (*see* IEEE 802.4)
 - Token management, 32
 - Token ring LAN, 292-299
 - Token-holding time, 296
 - Toll connecting trunk, 104
 - Toll office, 104
 - Torn tape office, 133
 - TPDU (*see* Transport Protocol Data Unit)
 - Traffic descriptor, 461
 - Traffic policing, 379-380
 - Traffic shaping, 379-380, 463-468
 - Transceiver, 277
 - Transformation encoding, 732-733
 - Transit network, 430

- Transmission control protocol, 36-37, 521-542, 658, 678, 685
 - congestion control, 536-539
 - connection management, 529-533
 - Karn's algorithm, 541
 - Nagle's algorithm, 534-535
 - segment header, 526-529
 - service model, 523-524
 - silly window syndrome, 534-535
 - timer management, 539-542
 - transmission policy, 533-536
 - wireless networks, 543-545
 - Transmission
 - infrared, 100
 - lightwave, 100-102
 - Transmission media, 82-94
 - Transponder, satellite, 164
 - Transport entity, 480
 - Transport gateway, 398
 - Transport layer, 479-576
 - ATM AAL, 545-555
 - example, 510-521
 - Internet, 521-545
 - network performance, 555-572
 - OSI, 31
 - protocol elements, 488-510
 - service provided, 479-487
 - Transport protocol, 488
 - addressing, 489-492
 - elements, 488-510
 - flow control, 502-506
 - Internet, 521-545
 - multiplexing, 506-508
 - Transport protocol data unit, 484
 - Transport service access point, 489
 - Transport service primitives, 483-486
 - Transport service provider, 481
 - Transport service user, 481
 - Transport stream, 743
 - Transposition cipher, 583-585
 - Trap, SNMP, 632
 - Trap directed polling, 632
 - Tree walk protocol, 258-259
 - Trellis coding, 112
 - Tributary, SONET, 127
 - Trigram, 583
 - Triple X, 60
 - Trunk, 11, 118-130
 - TSAP (*see* Transport Service Access Point)
 - Tunneling, 404-405
 - Twisted pair, 83-84
 - category 3, 83
 - category 5, 84
 - Two-army problem, 499-500
- U**
- UBR (*see* Unspecified Bit Rate service)
 - UDP (*see* User Datagram Protocol)
 - Unbalanced transmission, 115
 - Unconfirmed service, 26-27
 - Uniform resource locator, 692-695
 - schemes, 692-693
 - Universal resource identifier, 695
 - Unnumbered frame, 226-228
 - Unshielded twisted pair, 84
 - Unspecified bit rate service, 460
 - Urgent data, 524
 - URI (*see* Universal Resource Identifier)
 - URL (*see* Uniform Resource Locator)
 - USENET, 669-680, 693
 - implementation, 675-680
 - relationship to the Internet, 669
 - user view, 670-675
 - USENET hierarchies, 671
 - User agent, 645, 646-648
 - User datagram protocol, 37, 542-544
 - User plane, ATM, 64
 - User profile, 648
 - UTP (*see* Unshielded Twisted Pair)
 - UUCP, 669
- V**
- V.24, 114
 - V.32 111
 - V.32 bis, 111
 - V.34 111
 - V.42 bis, 112
 - Vacation daemon, 663
 - Variable bit rate service, 459
 - VBR (*see* Variable Bit Rate service)
 - Vector quantization, 733

Very high frequency band, 95, 97-98
 Very low frequency band, 95, 97
 Very small aperture terminal, 165
 VHF band (*see* Very High Frequency band)
 Video, 727-730
 analog, 727-729
 digital, 729-730
 interlaced, 728
 progressive, 728
 Video on demand, 744-756
 distribution network, 750-754
 server, 745-750
 set-top box, 754-756
 Video server, 745-750
 software, 747-749
 Videoconference, 5
 Virtual channel, ATM, 450
 Virtual circuit, 342-345
 compared to datagram, 344-345
 Virtual path, ATM, 450
 Virtual scheduling algorithm, 466
 VISTAnet, 56
 VLF band (*see* Very low frequency band)
 Voice-grade line, 79
 VSAT (*see* Very Small Aperture Terminal)
 VTMP, 572

W

WAN (*see* Wide Area Network)
 WARC (*see* World Administrative Radio
 Conference)
 Wavelength, 94
 Wavelength division multiple access, 260-262
 Wavelength division multiplexing, 119-121
 WDM (*see* Wavelength Division
 Multiplexing)
 WDMA (*see* Wavelength Division Multiple
 Access)
 Web (*see* World Wide Web)
 Web page, 682, 683, 697
 Weighted fair queueing, 388-389
 Well-known port, 523
 Wide area network, 11-13
 Wine policy, 390

Wireless networking, 13-15
 analog radio, 155-163
 digital radio, 266-275
 electromagnetic waves, 94-101
 mobile hosts, 367-370, 432-434
 wireless LANs, 262-265
 wireless TCP, 543-545
 Wireline carrier, 160
 Wiring closet, 83
 Work factor, 581
 World administrative radio conference, 95
 World Wide Web, 54, 681-723
 browser, 682
 CGI, 705-706
 external viewer
 fetching a page, 685-687
 HTML language, 691-706
 HTTP protocol, 689-691
 hyperlink, 682
 hypermedia, 684
 hypertext, 682
 Java, 706-720
 search engine, 720-723
 server, 685-689
 URL, 692-695
 Worm, 720
 WWV, 494
 WWW (*see* World Wide Web)
 WYSIWYG, 695

X

X.3, 60
 X.21, 59
 X.25, 59-60
 X.28, 60
 X.29, 60
 X.400, 644, 661
 X.509, 668-669
 XTP, 572

Z

Zipf's law, 746
 Zone, DNS, 628

About the Author

Andrew S. Tanenbaum has an S.B. degree from M.I.T. and a Ph.D. from the University of California at Berkeley. He is currently a Professor of Computer Science at the Vrije Universiteit in Amsterdam, The Netherlands, where he heads the Computer Systems Group. He is also Dean of the Advanced School for Computing and Imaging, an interuniversity graduate school doing research on advanced parallel systems, distributed systems, and imaging systems. Nevertheless, he is trying very hard to avoid turning into a bureaucrat.

In the past, he has done research on compilers, operating systems, networking, and local-area distributed systems. His current research focuses primarily on the design of wide-area distributed systems that scale to millions of users. These research projects have led to over 70 refereed papers in journals and conference proceedings. He is also the author of five books (see page ii).

Prof. Tanenbaum has also produced a considerable volume of software. He was the principal architect of the Amsterdam Compiler Kit, a widely-used toolkit for writing portable compilers, and MINIX, a small UNIX-like operating system for operating systems courses. Together with his Ph.D. students and programmers, he helped design the Amoeba distributed operating system, a high-performance microkernel-based distributed operating system. MINIX and Amoeba are now available for free for education and research via the Internet.

His Ph.D. students have gone on to greater glory after getting their degrees. He is very proud of them. In this respect he resembles a mother hen.

Prof. Tanenbaum is a Fellow of the ACM, a Senior Member of the IEEE, a member of the Royal Netherlands Academy of Arts and Sciences, and winner of the 1994 ACM Karl V. Karlstrom Outstanding Educator Award. He is also listed in *Who's Who in the World*. His home page on the World Wide Web is located at <http://www.cs.vu.nl/~ast/>.

COMPUTER NETWORKING

THIRD EDITION



COMPUTER NETWORKS

ANDREW S. TANENBAUM

Computer Networks is the ideal introduction to today's and tomorrow's networks. This classic best-seller has been totally rewritten to reflect the networks of the late 1990s and beyond.

Author, educator, and researcher Andrew S. Tanenbaum, winner of the ACM Karl V. Karlstrom Outstanding Educator Award, carefully explains how networks work inside, from the hardware technology up through the most popular network applications. The book takes a structured approach to networking, starting at the bottom (the physical layer) and gradually working up to the top (the application layer). The topics covered include:

- Physical layer (*e.g.*, copper, fiber, radio, and satellite communication)
- Data link layer (*e.g.*, protocol principles, HDLC, SLIP, and PPP)
- MAC Sublayer (*e.g.*, IEEE 802 LANs, bridges, new high-speed LANs)
- Network layer (*e.g.*, routing, congestion control, internetworking, IPv6)
- Transport layer (*e.g.*, transport protocol principles, TCP, network performance)
- Application layer (*e.g.*, cryptography, email, news, the Web, Java, multimedia)

In each chapter, the necessary principles are described in detail, followed by extensive examples taken from the Internet, ATM networks, and wireless networks.

Other bestselling titles by Andrew S. Tanenbaum:

Operating Systems: Design and Implementation, 2nd edition

Modern Operating Systems

Distributed Operating Systems

Structured Computer Organization, 3rd edition

PRENTICE HALL
Upper Saddle River, NJ 07458



REB120/107/

THIRD EDITION

COMPUTER NETWORKS

ANDREW S. TANENBAUM



Computer Networks

Third Edition

Andrew S. Tanenbaum

*Vrije Universiteit
Amsterdam, The Netherlands*

For book and bookstore information



<http://www.prenhall.com>



*Prentice Hall PTR
Upper Saddle River, New Jersey 07458*

ADDL for
ENGL

Library of Congress Cataloging in Publication Data

Tanenbaum, Andrew S. 1944-

Computer networks / Andrew S. Tanenbaum. -- 3rd ed.
p. cm.

Includes bibliographical references and index.

ISBN 0-13-349945-6

I. Computer networks. I. Title.

TK5105.5.T36 1996

004.6--dc20

96-4121

CIP

Editorial/production manager: *Camille Trentacoste*

Interior design and composition: *Andrew S. Tanenbaum*

Cover design director: *Jerry Votta*

Cover designer: *Don Martinetti, DM Graphics, Inc.*

Cover concept: *Andrew S. Tanenbaum, from an idea by Marilyn Tremaine*

Interior graphics: *Hadel Studio*

Manufacturing manager: *Alexis R. Heydt*

Acquisitions editor: *Mary Franz*

Editorial Assistant: *Noreen Regina*



© 1996 by Prentice Hall PTR

Prentice-Hall, Inc.

A Simon & Schuster Company

Upper Saddle River, New Jersey 07458

The publisher offers discounts on this book when ordered in bulk quantities. For more information, contact:

Corporate Sales Department, Prentice Hall PTR, One Lake Street, Upper Saddle River, NJ 07458.
Phone: (800) 382-3419; Fax: (201) 236-7141. E-mail: corpsales@prenhall.com

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

All product names mentioned herein are the trademarks of their respective owners.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-349945-6

Prentice-Hall International (UK) Limited, *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall Canada Inc., *Toronto*

Prentice-Hall Hispanoamericana, S.A., *Mexico*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Simon & Schuster Asia Pte. Ltd., *Singapore*

Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

CONTENTS

PREFACE

xv

1 INTRODUCTION

1

- 1.1 USES OF COMPUTER NETWORKS 3
 - 1.1.1 Networks for Companies 3
 - 1.1.2 Networks for People 4
 - 1.1.3 Social Issues 6
- ✓ 1.2 NETWORK HARDWARE 7 *copy 12*
 - 1.2.1 Local Area Networks 9
 - 1.2.2 Metropolitan Area Networks 10
 - 1.2.3 Wide Area Networks 11
 - 1.2.4 Wireless Networks 13
 - 1.2.5 Internetworks 16
- ✓ 1.3 NETWORK SOFTWARE 16
 - 1.3.1 Protocol Hierarchies 17
 - 1.3.2 Design Issues for the Layers 21
 - 1.3.3 Interfaces and Services 22
 - 1.3.4 Connection-Oriented and Connectionless Services 23
 - 1.3.5 Service Primitives 25
 - 1.3.6 The Relationship of Services to Protocols 27
- ✓ 1.4 REFERENCE MODELS 28
 - 1.4.1 The OSI Reference Model 28
 - 1.4.2 The TCP/IP Reference Model 35
 - 1.4.3 A Comparison of the OSI and TCP Reference Models 38
 - 1.4.4 A Critique of the OSI Model and Protocols 40
 - 1.4.5 A Critique of the TCP/IP Reference Model 43
- 1.5 EXAMPLE NETWORKS 44
 - 1.5.1 Novell Netware 45
 - 1.5.2 The ARPANET 47
 - 1.5.3 NSFNET 50
 - 1.5.4 The Internet 52
 - 1.5.5 Gigabit Testbeds 54

- 1.6 EXAMPLE DATA COMMUNICATION SERVICES 56
 - 1.6.1 SMDS—Switched Multimegabit Data Service 57
 - 1.6.2 X.25 Networks 59
 - 1.6.3 Frame Relay 60
 - 1.6.4 Broadband ISDN and ATM 61
 - 1.6.5 Comparison of Services 66
- 1.7 NETWORK STANDARDIZATION 66
 - 1.7.1 Who's Who in the Telecommunications World 67
 - 1.7.2 Who's Who in the International Standards World 69
 - 1.7.3 Who's Who in the Internet Standards World 70
- 1.8 OUTLINE OF THE REST OF THE BOOK 72
- 1.9 SUMMARY 73

2 THE PHYSICAL LAYER

77

- 2.1 THE THEORETICAL BASIS FOR DATA COMMUNICATION 77
 - 2.1.1 Fourier Analysis 78
 - 2.1.2 Bandwidth-Limited Signals 78
 - 2.1.3 The Maximum Data Rate of a Channel 81
- 2.2 TRANSMISSION MEDIA 82
 - 2.2.1 Magnetic Media 82
 - 2.2.2 Twisted Pair 83
 - 2.2.3 Baseband Coaxial Cable 84
 - 2.2.4 Broadband Coaxial Cable 85
 - 2.2.5 Fiber Optics 87
- 2.3 WIRELESS TRANSMISSION 94
 - 2.3.1 The Electromagnetic Spectrum 94
 - 2.3.2 Radio Transmission 97
 - 2.3.3 Microwave Transmission 98
 - 2.3.4 Infrared and Millimeter Waves 100
 - 2.3.5 Lightwave Transmission 100
- 2.4 THE TELEPHONE SYSTEM 102
 - 2.4.1 Structure of the Telephone System 103
 - 2.4.2 The Politics of Telephones 106
 - 2.4.3 The Local Loop 108
 - 2.4.4 Trunks and Multiplexing 118
 - 2.4.5 Switching 130

- 2.5 NARROWBAND ISDN 139
 - 2.5.1 ISDN Services 140
 - 2.5.2 ISDN System Architecture 140
 - 2.5.3 The ISDN Interface 142
 - 2.5.4 Perspective on N-ISDN 143
- 2.6 BROADBAND ISDN AND ATM 144
 - 2.6.1 Virtual Circuits versus Circuit Switching 145
 - 2.6.2 Transmission in ATM Networks 146
 - 2.6.3 ATM Switches 147
- 2.7 CELLULAR RADIO 155
 - 2.7.1 Paging Systems 155
 - 2.7.2 Cordless Telephones 157
 - 2.7.3 Analog Cellular Telephones 157
 - 2.7.4 Digital Cellular Telephones 162
 - 2.7.5 Personal Communications Services 162
- 2.8 COMMUNICATION SATELLITES 163
 - 2.8.1 Geosynchronous Satellites 164
 - 2.8.2 Low-Orbit Satellites 167
 - 2.8.3 Satellites versus Fiber 168
- 2.9 SUMMARY 170

3 THE DATA LINK LAYER

175

- 3.1 DATA LINK LAYER DESIGN ISSUES 176
 - 3.1.1 Services Provided to the Network Layer 176
 - 3.1.2 Framing 179
 - 3.1.3 Error Control 182
 - 3.1.4 Flow Control 183
- 3.2 ERROR DETECTION AND CORRECTION 183
 - 3.2.1 Error-Correcting Codes 184
 - 3.2.2 Error-Detecting Codes 186
- 3.3 ELEMENTARY DATA LINK PROTOCOLS 190
 - 3.3.1 An Unrestricted Simplex Protocol 195
 - 3.3.2 A Simplex Stop-and-Wait Protocol 195
 - 3.3.3 A Simplex Protocol for a Noisy Channel 197

- 3.4 SLIDING WINDOW PROTOCOLS 202
 - 3.4.1 A One Bit Sliding Window Protocol 206
 - 3.4.2 A Protocol Using Go Back n 207
 - 3.4.3 A Protocol Using Selective Repeat 213
- 3.5 PROTOCOL SPECIFICATION AND VERIFICATION 219
 - 3.5.1 Finite State Machine Models 219
 - 3.5.2 Petri Net Models 223
- 3.6 EXAMPLE DATA LINK PROTOCOLS 225
 - 3.6.1 HDLC—High-level Data Link Control 225
 - 3.6.2 The Data Link Layer in the Internet 229
 - 3.6.3 The Data Link Layer in ATM 235
- 3.7. SUMMARY 239

4 THE MEDIUM ACCESS SUBLAYER

243

- 4.1 THE CHANNEL ALLOCATION PROBLEM 244
 - 4.1.1 Static Channel Allocation in LANs and MANs 244
 - 4.1.2 Dynamic Channel Allocation in LANs and MANs 245
- 4.2 MULTIPLE ACCESS PROTOCOLS 246
 - 4.2.1 ALOHA 246
 - 4.2.2 Carrier Sense Multiple Access Protocols 250
 - 4.2.3 Collision-Free Protocols 254
 - 4.2.4 Limited-Contention Protocols 256
 - 4.2.5 Wavelength Division Multiple Access Protocols 260
 - 4.2.6 Wireless LAN Protocols 262
 - 4.2.7 Digital Cellular Radio 266
- 4.3 IEEE STANDARD 802 FOR LANS AND MANS 275
 - 4.3.1 IEEE Standard 802.3 and Ethernet 276
 - 4.3.2 IEEE Standard 802.4: Token Bus 287
 - 4.3.3 IEEE Standard 802.5: Token Ring 292
 - 4.3.4 Comparison of 802.3, 802.4, and 802.5 299
 - 4.3.5 IEEE Standard 802.6: Distributed Queue Dual Bus 301
 - 4.3.6 IEEE Standard 802.2: Logical Link Control 302

- 4.4 BRIDGES 304
 - 4.4.1 Bridges from 802.x to 802.y 307
 - 4.4.2 Transparent Bridges 310
 - 4.4.3 Source Routing Bridges 314
 - 4.4.4 Comparison of 802 Bridges 316
 - 4.4.5 Remote Bridges 317
- 4.5 HIGH-SPEED LANS 318
 - 4.5.1 FDDI 319
 - 4.5.2 Fast Ethernet 322
 - 4.5.3 HIPPI—High-Performance Parallel Interface 325
 - 4.5.4 Fibre Channel 326
- 4.6 SATELLITE NETWORKS 327
 - 4.6.1 Polling 328
 - 4.6.2 ALOHA 329
 - 4.6.3 FDM 330
 - 4.6.4 TDM 330
 - 4.6.5 CDMA 333
- 4.7 SUMMARY 333

5 THE NETWORK LAYER

339

- 5.1 NETWORK LAYER DESIGN ISSUES 339
 - 5.1.1 Services Provided to the Transport Layer 340
 - 5.1.2 Internal Organization of the Network Layer 342
 - 5.1.3 Comparison of Virtual Circuit and Datagram Subnets 344
- 5.2 ROUTING ALGORITHMS 345
 - 5.2.1 The Optimality Principle 347
 - 5.2.2 Shortest Path Routing 349
 - 5.2.3 Flooding 351
 - 5.2.4 Flow-Based Routing 353
 - 5.2.5 Distance Vector Routing 355
 - 5.2.6 Link State Routing 359
 - 5.2.7 Hierarchical Routing 365
 - 5.2.8 Routing for Mobile Hosts 367
 - 5.2.9 Broadcast Routing 370
 - 5.2.10 Multicast Routing 372

- 5.3 CONGESTION CONTROL ALGORITHMS 374
 - 5.3.1 General Principles of Congestion Control 376
 - 5.3.2 Congestion Prevention Policies 378
 - 5.3.3 Traffic Shaping 379
 - 5.3.4 Flow Specifications 384
 - 5.3.5 Congestion Control in Virtual Circuit Subnets 386
 - 5.3.6 Choke Packets 387
 - 5.3.7 Load Shedding 390
 - 5.3.8 Jitter Control 392
 - 5.3.9 Congestion Control for Multicasting 393
- 5.4 INTERNETWORKING 396
 - 5.4.1 How Networks Differ 399
 - 5.4.2 Concatenated Virtual Circuits 401 5.5.3 422
 - 5.4.3 Connectionless Internetworking 402
 - 5.4.4 Tunneling 404
 - 5.4.5 Internetwork Routing 405
 - 5.4.6 Fragmentation 406
 - 5.4.7 Firewalls 410
- 5.5 THE NETWORK LAYER IN THE INTERNET 412
 - 5.5.1 The IP Protocol 413
 - 5.5.2 IP Addresses 416
 - 5.5.3 Subnets 417
 - 5.5.4 Internet Control Protocols 419
 - 5.5.5 The Interior Gateway Routing Protocol: OSPF 424
 - 5.5.6 The Exterior Gateway Routing Protocol: BGP 429
 - 5.5.7 Internet Multicasting 431
 - 5.5.8 Mobile IP 432
 - 5.5.9 CIDR—Classless InterDomain Routing 434
 - 5.5.10 IPv6 437
- 5.6 THE NETWORK LAYER IN ATM NETWORKS 449
 - 5.6.1 Cell Formats 450
 - 5.6.2 Connection Setup 452
 - 5.6.3 Routing and Switching 455
 - 5.6.4 Service Categories 458
 - 5.6.5 Quality of Service 460
 - 5.6.6 Traffic Shaping and Policing 463
 - 5.6.7 Congestion Control 467
 - 5.6.8 ATM LANs 471
- 5.7 SUMMARY 473

6 THE TRANSPORT LAYER**479**

- 6.1 THE TRANSPORT SERVICE 479
 - 6.1.1 Services Provided to the Upper Layers 479
 - 6.1.2 Quality of Service 481
 - 6.1.3 Transport Service Primitives 483
- 6.2 ELEMENTS OF TRANSPORT PROTOCOLS 488
 - 6.2.1 Addressing 489
 - 6.2.2 Establishing a Connection 493
 - 6.2.3 Releasing a Connection 498
 - 6.2.4 Flow Control and Buffering 502
 - 6.2.5 Multiplexing 506
 - 6.2.6 Crash Recovery 508
- 6.3 A SIMPLE TRANSPORT PROTOCOL 510
 - 6.3.1 The Example Service Primitives 510
 - 6.3.2 The Example Transport Entity 512
 - 6.3.3 The Example as a Finite State Machine 519
- 6.4 THE INTERNET TRANSPORT PROTOCOLS (TCP AND UDP) 521
 - 6.4.1 The TCP Service Model 523
 - 6.4.2 The TCP Protocol 524
 - 6.4.3 The TCP Segment Header 526
 - 6.4.4 TCP Connection Management 529
 - 6.4.5 TCP Transmission Policy 533
 - 6.4.6 TCP Congestion Control 536
 - 6.4.7 TCP Timer Management 539
 - 6.4.8 UDP 542
 - 6.4.9 Wireless TCP and UDP 543
- 6.5 THE ATM AAL LAYER PROTOCOLS 545
 - 6.5.1 Structure of the ATM Adaptation Layer 546
 - 6.5.2 AAL 1 547
 - 6.5.3 AAL 2 549
 - 6.5.4 AAL 3/4 550
 - 6.5.5 AAL 5 552
 - 6.5.6 Comparison of AAL Protocols 554
 - 6.5.7 SSCOP—Service Specific Connection-Oriented Protocol 555
- 6.6 PERFORMANCE ISSUES 555
 - 6.6.1 Performance Problems in Computer Networks 556
 - 6.6.2 Measuring Network Performance 559

- 6.6.3 System Design for Better Performance 561
- 6.6.4 Fast TPDU Processing 565
- 6.6.5 Protocols for Gigabit Networks 568
- 6.7 SUMMARY 572

7 THE APPLICATION LAYER

577

- 7.1 NETWORK SECURITY 577
 - 7.1.1 Traditional Cryptography 580
 - 7.1.2 Two Fundamental Cryptographic Principles 585
 - 7.1.3 Secret-Key Algorithms 587
 - 7.1.4 Public-Key Algorithms 597
 - 7.1.5 Authentication Protocols 601
 - 7.1.6 Digital Signatures 613
 - 7.1.7 Social Issues 620
- 7.2 DNS—DOMAIN NAME SYSTEM 622
 - 7.2.1 The DNS Name Space 622
 - 7.2.2 Resource Records 624
 - 7.2.3 Name Servers 628
- 7.3 SNMP—SIMPLE NETWORK MANAGEMENT PROTOCOL 630
 - 7.3.1 The SNMP Model 631
 - 7.3.2 ASN.1—Abstract Syntax Notation 1 633
 - 7.3.3 SMI—Structure of Management Information 639
 - 7.3.4 The MIB—Management Information Base 641
 - 7.3.5 The SNMP Protocol 642
- 7.4 ELECTRONIC MAIL 643
 - 7.4.1 Architecture and Services 645
 - 7.4.2 The User Agent 646
 - 7.4.3 Message Formats 650
 - 7.4.4 Message Transfer 657
 - 7.4.5 Email Privacy 663
- 7.5 USENET NEWS 669
 - 7.5.1 The User View of USENET 670
 - 7.5.2 How USENET is Implemented 675

- 7.6 THE WORLD WIDE WEB 681
 - 7.6.1 The Client Side 682
 - 7.6.2 The Server Side 685
 - 7.6.3 Writing a Web Page in HTML 691
 - 7.6.4 Java 706
 - 7.6.5 Locating Information on the Web 720
- 7.7 MULTIMEDIA 723
 - 7.7.1 Audio 724
 - 7.7.2 Video 727
 - 7.7.3 Data Compression 730
 - 7.7.4 Video on Demand 744
 - 7.7.5 MBone—Multicast Backbone 756
- 7.8 SUMMARY 760

8 READING LIST AND BIBLIOGRAPHY

767

- 8.1 SUGGESTIONS FOR FURTHER READING 767
 - 8.1.1 Introduction and General Works 768
 - 8.1.2 The Physical Layer 769
 - 8.1.3 The Data Link Layer 770
 - 8.1.4 The Medium Access Control Sublayer 770
 - 8.1.5 The Network Layer 771
 - 8.1.6 The Transport Layer 772
 - 8.1.7 The Application Layer 772
- 8.2 ALPHABETICAL BIBLIOGRAPHY 775

INDEX 795