

Chapter 5

Switched Networks

The techniques we have considered thus far only support the interconnection of a relatively small number of computers that are physically located relatively near to one another. The Internet provides interconnections for a vast number of computers physically disbursed throughout the entire world. Accordingly, the next step is to consider how interconnections can be provided between such a large collection of computers.

In some ways, this is an old problem. Many of the issues that must be addressed to construct a large network of computers have arisen previously in the construction of the telephone network and even the telegraph network. So, we will begin by discussing the structure of the telephone network. Then, we will examine the important features that distinguish modern computer networks from earlier communications networks.

5.1 Switched Networks

Who invented the telephone?

Almost everyone asked this question quickly answers “Alexander Graham Bell”.

Who invented the Internet?

The only response this question will draw from most people is a blank stare.

For a long time, this struck me as a great injustice. Why should Bell deserve enduring fame while almost no one knows who deserves credit for the development of the Internet?

At first, I rationalized that the problem might be that one cannot give any one individual credit for all the technologies that have made the Internet possible. The inventors of all the systems that make the Internet possible — computers, modems, Ethernets, fiber optics cable, etc. — all deserve some credit.


Surprisingly, however, the more one examines the development of the phone system, the more one sees parallels with the development of the Internet. Just as the Internet required many technological developments for which many individuals deserve credit, the telephone system also depends on the inventions of many individuals other than Bell. In a history of the development of the telephone published in 1910, Herbert Casson describes the task of the others who contributed to the development of the phone system:

“All that these young men had received from Bell and Watson was that part of the telephone that we call the receiver. ... There were no switchboards of any account,

no cables of any value, no wires that were in any sense adequate, no theory of tests or signals, no exchanges, NO TELEPHONE SYSTEM OF ANY SORT WHATEVER.”

What Bell had invented was a device capable of encoding a human voice as an electrical current and decoding such a current to reproduce a close approximation to the original voice. He had not, however, developed the technology required to transport the electrical signals his device produced with the flexibility required for a phone system as we know it. The signals his device produced were quite weak. Long distance transmission of such signals would depend on the development of devices that could amplify such signals. In addition, there was no notion of how to interconnect thousands of telephones so that any pair could be interconnected.

The first commercial uses of telephones were on two-party private lines. A company with two offices would purchase a pair of telephones and arrange for the installation of wires between the locations of their two offices. The telephones would then act more like what we would call an intercom, providing communications only between those two locations. If such a company wanted to interconnect more than two offices, it would have to install wires between every pair of offices.



The diagram on the left illustrates the wiring needed from the point of just one office in such a system. The number of wires required for such an arrangement grows very large as the number of offices involved increases. Assuming there are 10 offices to be interconnected, each phone must be attached to the end of 9 wires connecting it to the other phones. Each of the 10 offices will require similar wiring. As a result, there will be a total of $9 \times 10 = 90$ connections between some phone and the end of a wire. Each wire has two ends, so there must be 45 wires total.

In general, if there are N offices, each office will have to be connected to the end of $N - 1$ wires. The total number of ends of wires connected to offices will be $N(N - 1)$. Again, since each wire has two ends this means that the total number of wires needed will be $\frac{N(N-1)}{2}$.

Unfortunately, $\frac{N(N-1)}{2}$ becomes a very big number very quickly. If you tried to connect just 100 offices in this way, $\frac{100 \times 99}{2} = 4950$ wires would be required. For 1000 offices, almost 500,000 wires would be needed. Clearly, it would be completely infeasible to build the modern phone system using this approach.


It didn't take long for someone to find a better approach. Rather than installing wires connecting phones to phones, the phone company began connecting phones to a central office or *exchange*. There a switchboard provided the means to interconnect the ends of the lines from any two phones connected to the switchboard. With this arrangement, any phone can be connected to any other phone as desired, but only N wires are required for N phones.



Interestingly, according to Casson, the first switchboard was devised by E. T. Holmes, a man whose principle business was providing burglar alarm systems. For his burglar alarm business, Holmes had already installed wires from his clients' premises to his central office. These were connected to devices to detect intruders at the clients' premises, allowing one employee at Holmes office to serve as watchman for many locations. The wires installed for Holmes' burglar alarm systems were in just the right configuration for a central telephone switchboard. Apparently, Holmes realized this and added telephones to the services he provided.

Before long, demand for telephone service expanded greatly and the size and complexity of the switchboards and exchanges grew with it. In particular, as demand grew for what we now call long-distance service, the simple model of many telephones connected to a single central exchange became insufficient. Instead, each local region had its own central exchange and it became necessary to provide connections between exchanges. When a call was placed to a phone connected to a remote exchange, the operator at the local exchange would make a connection from the caller's line to a line connecting to the desired exchange. The operator at the remote exchange would then connect this line to the line connected to the desired telephone.

Even though there would be far fewer exchanges than telephones in such a system, it was still not practical to arrange for a direct wire between every pair of exchanges. The number of wires required for such an arrangement would still be $\frac{N(N-1)}{2}$ if there were N exchanges. The alternative used is to ensure that there are enough wires connecting exchanges to make it possible to find a path of wires from any exchange to any other exchange. The diagram below illustrates such a system.



It shows switchboards for five exchanges labeled A, B, C, D, E. A number of phones are shown with direct connections to each switchboard. These would belong to customers located geographically near the associated switchboard's location. The diagram also shows connections between the switchboards. A is connected to D, B and C, but not to E. Switchboard E is connected to C and B but not to either A or D.

Despite the fact that there are not connections between every pair of switchboards in the diagram, it is possible to connect every pair of phones. As an example, consider how one could connect a phone attached to switchboard C (like the phone shaded in green near the lower left of the diagram) and a phone that is connected to switchboard D (like the one shaded in red). There is

no direct connection between switchboards C and D, but there are pathways through the available connections that are sufficient to make a connection between the green and red phone. One such path is highlighted in red in the diagram. Starting at the green phone, it connects to switchboard C, then to switchboard A, then to switchboard B, then to switchboard D and finally to the red phone. This is not the only such path in the diagram, but it is enough to ensure that there is some way that a caller at the green phone could reach the pink phone.




In the original phone system, a connection through several exchanges would be arranged manually by operators. The owner of the green phone would pick up the phone and ask the operator at switchboard C for a connection to the red phone. The operator at switchboard C would in turn “call” the operator at switchboard A and ask that operator to provide a connection to the switchboard for the pink phone’s exchange, D. The modern phone system works on similar principles, but the process of finding pathways through the network of wires

between phones and exchanges is now entirely automated.¹

The desire to interconnect many computers quickly leads to the same problem facing the pioneers of the phone system. It is not practical to make a direct connection between every pair of computers that might want to communicate. Instead, as in the phone system, a better approach is to connect computers to a switching system that provides pathways between computers. Unlike the evolution of the phone system, there is no point in using human operators or electromechanical systems to arrange for the connections. Instead, the switches in a computer network are themselves computers. These machines spend most or all of their time devoted to the task of providing interconnections between other computers, but they themselves are computers. Unlike the phone network, where the devices that serve as the sources and destinations for information, the phones, are distinctly different from the devices that handle switching, there is no strict dichotomy between the equipment at the heart of a computer network and those at its periphery. They are all computers!

This is a very important difference. A telephone is a very limited device. It can convert sound waves into electrical signals, but it can do little else. By comparison, a computer is a very flexible device. Like a telephone, it can convert information into forms more suitable for transmission, but it is not limited to a single kind of information, sound, or a single kind of encoding. A computer can be programmed to encode sound, images, text, video and many other forms of information. In addition, a computer, even one on the periphery of a computer network, has the capability of playing a significant role in the process of information transmission. All a telephone can do is send a voice signal into the phone network relying on the switching systems to ensure its delivery. A computer attached to a network can interact with the network to plan the route for its information, to double-check that information has been delivered correctly and in other ways. It is really part of the network, rather than just a device that depends on the network.

A diagram suggesting the organization of such a switched computer network is shown below.



The “wires” interconnecting the computers in such a network could use any of the techniques for transmitting data presented when we discussed the problem of connecting just a pair of computers.

¹The first automatic switching system was devised by Almon Strowger, an undertaker who is said to have been motivated by the fear that local telephone operators were deliberately redirecting calls made to his business to one of his competitors.

If the computers are close enough together, they could be specially installed cables. If the network is more widespread geographically, the interconnections could be made through modems or dedicated higher speed connections leased from a phone company. In fact, they need not be wires at all. They could be optical fiber or even micro-wave links. The key ideas are that each computer is connected to only a small number of other computers, that while not all computers are directly linked pathways exist between all pairs of computers and that a subset of the computers play the role of switches when such pathways are used to support communications.

5.2 Circuit Switching vs. Message Switching

In our introduction to the idea of a computer network based on switching, we have emphasized the similarities between the techniques used in the telephone network and in a switched computer network. There are, however, major differences in the way such networks operate. In this and the next section we will discuss three distinct approaches to switching — circuit switching, message switching and packet switching — and examine their relative advantages.

Any of the three switching techniques described in the next two sections could be used to build a computer network. In fact, however, the Internet is based on packet switching and many believe the decision to use packet switching was fundamental to the Internet's success. What is so special about packet switching? The answer is an economic one. While packet switching does not make anything totally new possible, it does make it possible to do an old thing, transmit information, more efficiently and therefore less expensively — MUCH less expensively.

Consider the functioning of an early phone system in which connections were made manually by human switchboard operators. For the sake of simplicity, assume that a call is made between two phones connected to the same exchange. Before the call is made, there is no direct electrical connection between the calling phone and the phone which is being called. Each phone is connected to the switchboard in the central exchange, but they are not directly connected to one another. The caller contacts the operator and asks to be connected. The operator takes a wire that is connected to the caller's phone on one end and has a jack on the other end and plugs it into a socket which is connected to the wire leading to the called party's phone. At this point, there is a physical path of wires capable of conducting electric current between the two phones. As the parties converse, this path carries the varying electric current encoding their voices back and forth without further intervention from the operator.

A physical path capable of conducting electrical current is called a circuit. Accordingly, the task performed by an operator in a phone system is called circuit switching. In the modern phone system, the task remains the same even though it is accomplished automatically. When a call is made, the result of the process of handling the call is the establishment of a path that can carry continuous stream of audio signals between the two phones. The path may not be easily described as a wire. It may include fiber optic cables on which the audio signal is multiplexed with hundreds or thousands of other signals. All components of the path, however, provide functionality equivalent to a dedicated wire in that the signal that emerges at the receiving phone is essentially identical to that which would be delivered by a physical path of simple wires. Once a connection has been established, the call proceeds without the involvement of the switching system.

In computer networks, an alternate approach to switching is possible. Just as it is easier to appreciate circuit switching by imagining the behavior of a turn of the century telephone operator than by imagining the operation of a modern digital telephone switching center, it is easier to

understand the technique used in computer networks by imagining how it might have been used in a much older network, the telegraph system.

We have already discussed the Morse code system used to encode messages. While this system seems incredibly primitive by today's standards, the introduction of the telegraph revolutionized communication in the 19th century. Before the telegraph, the fastest means of sending a message across a continent or across the ocean took days. Suddenly, such communication could take place instantaneously. Imagine how revolutionary this change must have been. In its day, Western Union was as powerful and profitable a corporation as AT&T was for most of the 20th century.

As with computer networks and the phone system, the telegraph was not very interesting by itself. Two telegraph instruments at the end of a wire did not constitute a communications system. Like the phone system and the Internet, the importance of the telegraph rested on the establishment of a network of telegraph links that spanned the whole world. Also, it was not practical to construct a telegraph network with direct connections between every telegraph office in every town and city in the country. Instead, the telegraph network relied on a network of connections between telegraph offices that provided paths of several links between any two offices, and on the ability of telegraph offices to act as intermediate switching centers between offices that were not directly connected.

Imagine you are part of a start-up company providing telegraph service along the east coast of the United States. You have established offices in Philadelphia, New York and Boston with plans of providing delivery service between those cities. Customers in each city will bring messages for delivery to your office. You will send the message electrically to your office in the destination city where it will be typed and hand delivered to the designated recipient.

Given the expense of installing cable, your company would probably not install three lines providing direct connections between all three cities. Since New York is located between Philadelphia and New York, a more economical approach would be to install a line between New York and Philadelphia and another line between New York and Boston. Then, when someone wants to send a message from Boston to Philadelphia, they could do so by sending the message through New York. The interesting question is exactly how to have the New York office fulfill its role as a "switch" between Philadelphia and Boston.

Let's imagine how things might be done in New York. The office might have three employees: a) Bob who sends and receives message from Boston, b) Phil who sends and receives messages from Philadelphia, and c) Ned who mans the front desk, taking messages from customers and giving them to Bob or Phil for transmission depending on their destination.

When the Boston office needs to send a message to Philadelphia, they will have to start by sending some message to Bob through the line to New York. If the telegraph system worked like the early phone system, they would send a message telling Bob that they need to be connected to the Philadelphia line. In this case, Bob would have to first check to see if Phil was busy sending or receiving a message. If so, he would send a message back to Boston telling them they have to wait. Once the line to Philadelphia became idle, Bob would tell Phil that Boston needed to use the line. Then, they would physically connect the line from Boston to the line from Philadelphia and sit back and wait until the Boston office finished sending its message. At that point, they would disconnect the lines and return to normal. This procedure is just the telegraph equivalent of circuit switching.

There is an alternative approach which requires no physical connection of the lines from Boston and Philadelphia. The simplicity of this alternate method depends on two simple practices that would be fairly standard in the operation of a telegraph office. First, think about what Ned, the

front desk operator does as customers give him messages. Depending on each message's destination, he simply gives the message to be sent to either Bob or Phil. Bob or Phil may be busy sending or receiving a message when this happens. If so, the new messages simply get placed in a pile of messages awaiting transmission. Second, note that whenever a message is sent, the information transmitted must include the address to which the message should be delivered. If this information were not included, the receiving office would not be able to dispatch a messenger to deliver the telegram.

Now, suppose that the operator in Boston simply sends all outgoing messages, whether destined to New York or Philadelphia out on the New York line. Bob, the operator in New York, can decide what to do with them after they arrive by examining the destination addresses. If a message arrives from Boston destined for an address in New York, Bob will simply give the message to Ned at the front desk who will assign it to one of the delivery boys. If Bob receives a message destined for Philadelphia, he can simply add it to Phil's pile of message awaiting transmission to Philadelphia in the same way Ned places messages from New York to Philadelphia on Phil's pile. Eventually, Phil will transmit all the messages, whether from Boston or New York, to Philadelphia as desired.

This approach is very different from physically connecting the line from Boston to the line from Philadelphia. There is never a physical, electrical connection all the way from Boston to Philadelphia in this scheme. Nevertheless, messages will get delivered from Boston to Philadelphia. The New York office is still functioning as a switching center. In this case, however, the technique employed is referred to as message switching rather than circuit switching.

Message switching would not be at all acceptable for use within the phone system. Imagine if when you made a call you actually only talked directly to an operator who repeated everything you said to the person with whom you really wanted to talk and also told you everything they wanted to say to you. Telephone calls would be a very different experience. On the other hand, message switching works quite naturally in the context of a telegraph system. There is no loss of "personal touch" when messages are finally delivered. The tasks required of the employees in the telegraph office are actually simplified. They only need to know how to send and receive messages and to determine for each message on which line it should be sent. They no longer need to know how to physically connect lines.

Just as it is a natural approach in the context of a telegraph system, message switching fits in well in a computer network. Again, there is no loss of "personal touch" as long as the intended sequence of 0's and 1's gets delivered to its intended destination. To participate in any computer network, computers have to be designed so that they can send and receive messages. This is true even for the computer on the periphery of the network that are not even serving as switches of any kind. Given a computer that knows how to send and receive messages, it is fairly simple to program the machine to forward messages from an incoming network connection to an outgoing connection to achieve message switching.

Compared to circuit switching, message switching simplifies the process of sharing communication lines. Suppose for a moment that our telegraph office used circuit switching. Recall that in this case when the Boston office requests a connection through New York to Philadelphia, the New York operator has to check to see if the line from New York to Philadelphia is free. If a message is currently being sent or received on the New York to Philadelphia line, the Boston office will have to wait until the Philadelphia line becomes free before beginning its transmission. This is a waste of resources. During the waiting period, the Boston to New York line will remain idle even though the Boston office may have other messages that need to be sent directly to New York.

With message switching the available lines in the system are used more efficiently. If a message that needs to be sent out on a given line arrives while that line is busy, it just is added to the collection of messages waiting to be sent out on that line later. Messages wait for lines, but lines are not left idle waiting for connections to other lines.

5.3 Packet Switching

With just one slight change, the message switching scheme we have described becomes packet switching, the switching technique used in the Internet. The change is quite simple. Rather than treating messages as indivisible units, a packet switched network takes the liberty of dividing user messages into smaller units called packets as it sees fit. It is then these smaller units that are forwarded through the network in much the same way that messages are handled under message switching.

A major motivation behind the use of this technique is the wide variation in the sizes of “messages” sent through a computer network. Traditionally, telegraph messages were short and concise. Such conciseness is not the norm in the Internet. Certainly, the economics of Internet usage and access do nothing to encourage brevity. Most fees paid for Internet access vary with the speed of the line through which access is provided, but not with the amount of data actually transmitted through the line. As a result, while some email messages are as short as a telegram, other “messages” sent through the Internet are quite long. A one minute video clip can require the transmission of several million bytes of binary data.

Recall that in the telegraph station analogy, the operator of each outgoing line maintained a pile of messages waiting to be sent on the line. When many customers arrived at the same time with messages for a particular destination, the pile for the line to that destination would grow. When business was light, the operator would catch up and might manage to run out of messages to be sent eventually leaving the line idle.

A similar process occurs in a computer network based on message switching. The computer functioning as a switch maintains a collection of messages waiting to be sent on each outgoing line. Rather than calling the collection of messages a “pile”, it is usually called a *queue*. An outgoing line’s queue of messages will grow and shrink, occasionally to nothing, as the demand to send messages on the line varies.

Given that messages sent through the Internet vary widely in size, the messages stored in a line’s queue will also vary in size. There may be some very large messages waiting to be sent on a line at the same time that there are short messages waiting to be sent. If the computer controlling the switching process decides to start sending a large message, the line will be busy for a relatively long time. During this time, all the short messages will be kept waiting.

Such delays will be insignificant if the short items are email messages. Even a message several megabytes long is unlikely to occupy any link in a modern network for more than a few minutes. But, for someone waiting for a simple web page to load, an added delay of a minute or two can seem like eternity.


A packet switching system solves this problem by breaking big messages up into smaller units called packets. One of the nice properties of digitally encoded information is that it can easily be broken into many pieces and reconstituted later to reproduce the original precisely. To make the process precise, imagine that you wanted to send a large text document (several hundred pages) through the postal system using only standard business envelopes. The first step is to make sure

that all the pages in the document are numbered. Then, you could stuff small groups of pages in separate envelopes, address and stamp all the envelopes and send them on their way. The receiver would obviously have to reassemble all the pages received, using the page numbers to place them in the right order.

In packet switched computer networks, the sending computer typically takes each logical message sent and breaks it into small pieces called packets for transmission. A sequence number is added to each piece to facilitate reassembly and each piece is separately addressed. In the case of a small message, of course, a single packet containing the entire message may be sufficient. These packets are then passed into the network. Switches view these packets rather than the original messages as the fundamental units they are responsible for transporting.

The first advantage of this approach is that large messages will no longer prevent the timely delivery of short messages. If a large message is sent through a particular switch, what will actually arrive at the switch is a large collection of packets. These packets, together with packets for other long or short messages trying to make their way through the switch, will form the collection of items awaiting transmission on an outgoing line. Now the computer controlling the lines can handle the short messages more quickly by sneaking their transmissions in between pieces of the large message.

Even if only a single message is being sent through a network, packet switching has several potential advantages over message switching. Consider the simple network shown below:



Suppose that computer A wants to send a large message to computer E using B, C, and D as intermediate switches. The network shown provides two paths from A to E. One path goes B and C and the other goes through D. While the network itself is unrealistically simple, the presence of such multiple paths is not unrealistic. Multiple paths are desirable because they provide reliability in the case that a link is damaged.

If message switching is used to send a large message through such a network, computer A will have to choose one of the two available paths and send its entire message through that path. In the case of packet switching, computer A can break a large message into packets and then send half the packet on the route through D while the other half are sent on the route through B and C. With any luck, this will cut the total time required to send the message almost by half.

Both packet switching and message switching are examples of what are called “store and forward” networks. This means that in both approaches, each switch waits until it has received an incoming packet or message before deciding how to forward the message along the next step toward its final destination. During the possibly brief period between the completion of the receipt of a message and its forwarding, the switch stores a copy of the message. By contrast, in a circuit switched network, the bits of a message flow through the switch but are not retained. In a circuit switched networks, a message is forwarded as it is being received. In message and packet switched systems, the forwarding of a message does not begin until it has been completely received.

As a result of this fact, even if only a single path exists from source to destination, packet switching can speed delivery compared to message switching. For example, suppose that computer A wanted to send a large message to computer E through computer D. Suppose that as a single unit, the message would take one minute to transmit, while if packet switching is used the message would be broken up into 61 packets, each of which would take just one second to transmit. (It is reasonable to assume that the total amount of data sent in packet form will be greater than when the message is sent as a single message since each packet will have to include a sequence number and information indicating to which message it belongs.) Finally, since the data will travel through the wires between A, D and E at the speed of light, we will assume that anything sent is received essentially instantaneously.

Now, if A starts sending its message to E at 12:00 noon using message switching, it will finish at 12:01. At this point, D will have received a complete copy of the message so it can begin forwarding the message to E. This will take another minute, so E will receive the last bit of the message at 12:02, 2 minutes after A started its transmission.

If, on the other hand, A starts sending its message using packet switching at 12:00, it will take it 61 seconds to complete transmission. It will finish sending the first packet at 12:00:01, the second packet at 12:00:02, and so on until the last piece is sent at 12:01:01, one second after it would have been finished in the message switching scenario.

Things are very different, however, if we consider D and E. D will have completely received the first packet by 12:00:01 and will begin to forward it to E at that point. It will take D one second to forward it, so E will receive it at 12:00:02. The same process will repeat for each of the packets. E will receive each packet one second after D. In particular, E will receive the last packet, and therefore the entire message, by 12:01:02, 58 second earlier than it would have been received in the message switching scenario.

The secret behind this magic is that the packet switched version keeps more data lines active more of the time. In the message switching scenario, nothing was being sent from D to E during the first minute while A was sending its message to D. Similarly, no use is made of the line from A to D while D is sending the message to E. In the packet switched approach, once the first packet is received by D, both the line from A to D and the line from D to E are actively and simultaneously used until the last packet is received by D. While the Nth packet is being sent from A to D, packet N-1 is being sent from D to E.

One could argue that this claim is founded entirely on the restriction that the network receive a complete copy of an incoming message before beginning to forward the message. If this restriction were relaxed, the scenarios just described would come out differently. There are, however, factors that make this restriction essential. For example, in many networks, a single switch is connected to lines that operate at different transmission speeds. If in our example, the line from D to E sent more bits per second than the line from A to D, then if the switch started to send a message to E before it had all arrived from A, it might run out of bits received from A while it was sending the packet to E. Waiting until a complete packet has been received provides a simple way to handle such mismatches in transmission speeds.

Also, as we will see later when we discuss how computer networks deal with errors, packet switching has a significant advantage in dealing with transmission errors. If one has to retransmit data because an error occurred, it is less costly to retransmit a small packet than a large message. Together such factors give packet switching a tremendous economic advantage over circuit switching.