Lecture 3

In this lecture, we will try understand how hub and switch operates, and what is the difference between them?

Ethernet is best understood by first considering the early *10Base5* and *10Base2* specifications. These two Ethernet specifications defined the details of the physical layer of early Ethernet networks. With these two specifications, the engineer installs a series of coaxial cables to each device on the Ethernet network—there is no hub, switch, or wiring panel. The series of cables creates an electrical bus that is shared among all devices on the Ethernet. Because it is a single bus, only one electrical signal can flow at a single time. If two or more signals were sent, the two would overlap (in technical terms this is called as collision), making both signals unintelligible.

Ethernet also defined a specification for how to ensure that only one device sends traffic on the Ethernet at one time—otherwise, the Ethernet would have been unusable. The algorithm, known as the carrier sense multiple access collision detect (CSMA/CD) algorithm, defines how the bus is accessed. In human terms, CSMA/CD is similar to what happens in a meeting room with many attendees. Some people talk much of the time. Some do not talk, but they listen. Others talk occasionally. Being humans, it's hard to understand what two people are saying at the same time, so generally, one person is talking and the rest are listening. Imagine that Bob and Larry both want to reply to the current speaker's comments. As soon as the speaker takes a breath, Bob and Larry might both try to speak. If Larry hears Bob's voice before Larry actually makes a noise, Larry might stop and let Bob speak. Or, maybe they both start at almost the same time, so they talk over each other, and many others in the room can't hear what was said. Then there's the proverbial "Excuse me, you talk next … " and eventually Larry or Bob talks. Or, in some cases, another person jumps in and talks while Larry and Bob are both backing off. These "rules" are based on your culture; CSMA/CD is based on Ethernet protocol specifications and achieves the same type of goal.

The figure below shows the basic logic of an old Ethernet 10Base2 network, which literally uses a single electrical bus, created with coaxial cable.

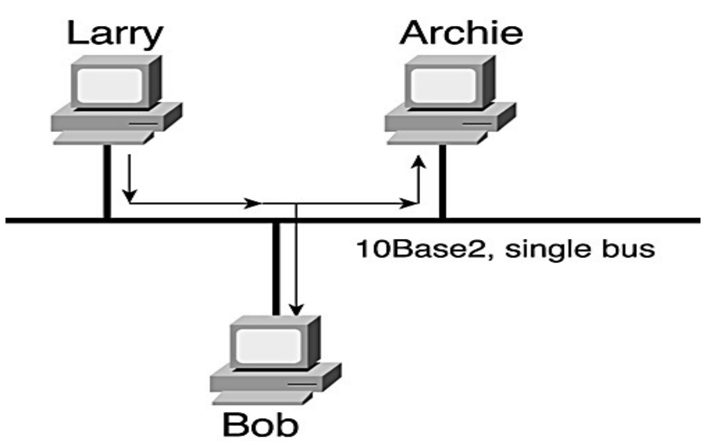

The CSMA/CD algorithm works like this:

1. A device with a frame to send listens until the Ethernet is not busy.
2. When the Ethernet is not busy, the sender begins sending the frame.
3. The sender listens to make sure that no collision occurred.
4. If there was a collision, the sender randomizes a timer and waits that long.
5. When the timer expires, the process starts over with Step 1.

10BaseT solved several problems with the early Ethernet specifications. 10BaseT allowed the use of telephone cabling that was already installed, or simply allowed the use of cheaper, easier-to-install cabling when new cabling was required.

The use of 10BaseT *hubs* gave Ethernet much higher availability because a single cable problem does not affect the other users. These hubs are essentially multiport repeaters; they extend the bus concept of 10Base2 and 10Base5 by regenerating the same electrical signal sent by the original sender of a frame out every other port. Therefore, collisions can still occur, so CSMA/CD access rules continue to be used. Figure below shows the same network as above, after migration to using hubs and twisted-pair cabling instead of co-ax cabling.

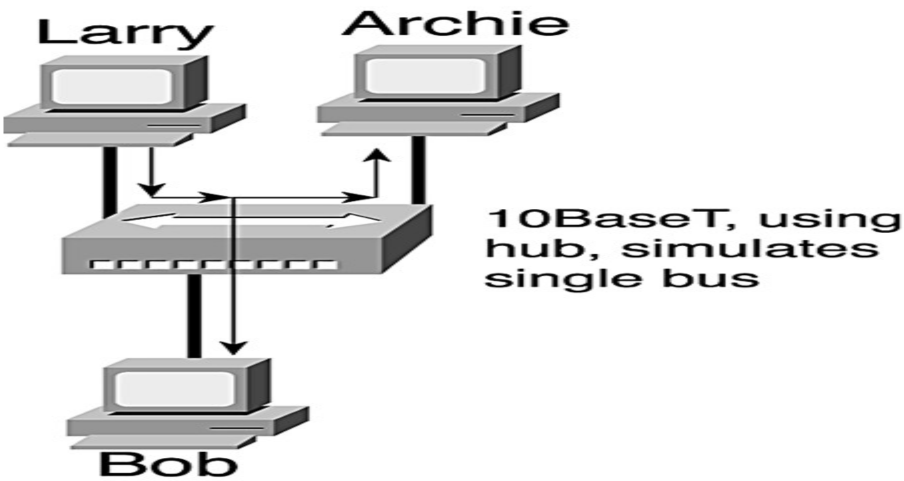

The concept of cabling each device to a central hub, with that hub creating the same electrical bus as in the older types of Ethernet, was a core fact of 10BaseT Ethernet. Because hubs continued the concept and physical reality of a single electrical path that is shared by all devices, we call this as *shared Ethernet:* All devices are sharing a single 10-Mbps bus.

However, the degrading performance of CSMA/CD when the Ethernet utilization increased was still not addressed. The next step was to make the hub smart enough to ensure that collisions simply did not happen—which means that CSMA/CD would no longer be needed. First, you need a deeper knowledge of 10BaseT hubs before the solution to the congestion problem becomes obvious.

Figure below outlines the operation of half-duplex 10BaseT with hubs.

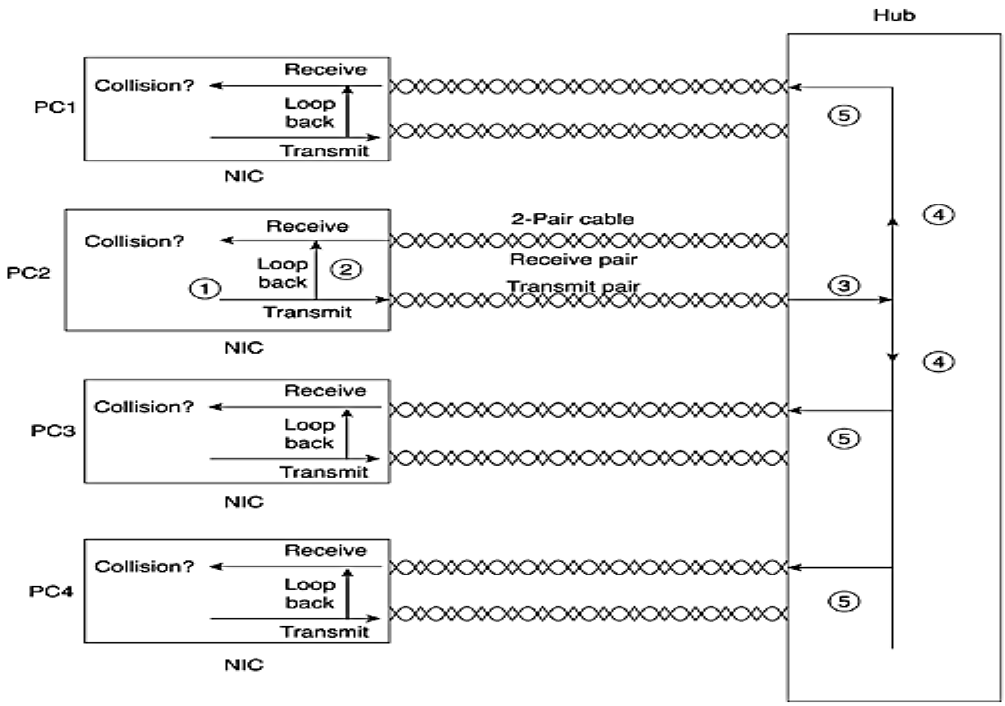

The operation of above hub in chronological steps is explained below

1. The network interface card (NIC) sends a frame.
2. The NIC loops the sent frame onto its receive pair.
3. The hub receives the frame.
4. The hub sends the frame across an internal bus so that all other NICs can receive the electrical signal. It does not forward the electrical signal back out to the device that sent the original signal.
5. The hub repeats the signal to each receive pair to all other devices. In other words, the hub sends so that the attached stations receive on their receive pair. (Similarly, the hub listens on the transmit pair because that is the pair used by each station for transmissions.)

The figure details how the hub works, with one device sending and no collision. If PC1 and PC2 sent a frame at the same time, a collision would occur. At Step 4, the hub would forward both electrical signals, which would cause the overlapping signals to be sent to all the NICs. CSMA/CD logic is still needed to have PC1 and PC2 wait and try again.

The term ***collision domain*** defines the set of devices for which their frames could collide. The original Ethernet specifications implied half-duplex behavior. Just like having a single speaker in a meeting room, there is a single sender on an Ethernet (so far).

LAN switches overcome the problems created by collisions and the CSMA/CD algorithm. Switches do not create a single shared bus, but rather they treat each individual physical port as a separate bus. Switches use memory buffers to hold incoming frames as well. Figure below explain the operation of switch.

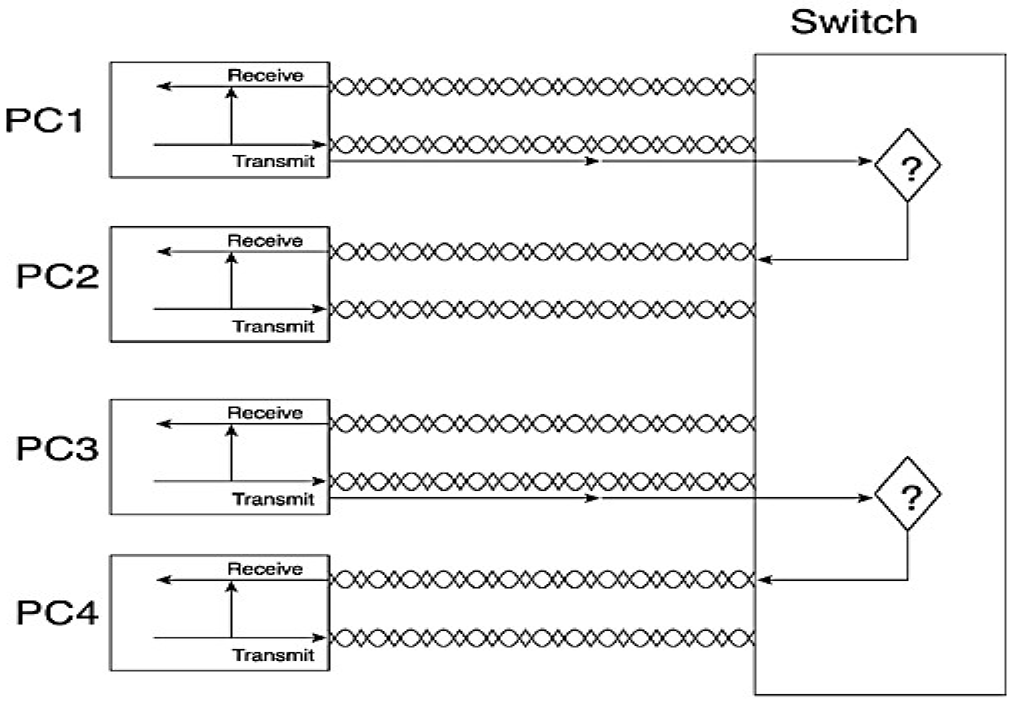

In above figure, both PC1 and PC3 are sending at the same time. The switch looks at the destination Ethernet address, and sends only the frame from PC1 to PC2 and the frame sent by PC3 to PC4. The big difference between the hub and the switch is that the switch interpreted the electrical signal as an Ethernet frame and processed the frame to make a decision. A hub simply repeats the electrical signal and makes no attempt to interpret the electrical signal (Layer 1) as a LAN frame (Layer 2).

Buffering also helps prevent collisions. Imagine that PC1 had sent a broadcast, which needs to be sent out all ports in above figure. PC3 is sending another frame to PC4. The switch, knowing that forwarding both frames to PC4 would cause a collision, buffers one frame until the first one has been completely sent to PC4.

10BaseT switches added forwarding logic based on the address of the frame, which avoids collisions. Switches also added buffering so that additional collisions were avoided. In short:

> If only one device is cabled to each port of a switch, no collisions occur.

The final concept in this section requires the assumption that no collision can occur. If no collision can occur, CSMA/CD is not needed—at least, not the collision detection and recovery part. If CSMA/CD is not needed, the implied half-duplex operation no longer is needed. In short, LAN switches with only one device cabled to the switch allow the use of *full-duplex* operation.