Network Security

Professor Gaurav S. Kasbekar
Department of Electrical Engineering
Indian Institute of Technology, Bombay

Week - 01

Lecture - 06

Review of Basic Concepts and Terminology in Communication Networks: Part 4

Hello, in this lecture we discuss the structure of the internet and various access networks, which connect end users to the internet. We start with our discussion of the structure of the internet, which is a popular network used currently as we know. The network is a network of many interconnected networks. These component networks are owned by different internet service providers. And the internet has a loosely hierarchical structure.

In this example, three networks are connected together. This is one network. This is another network and this is the third network. This network is owned by a local internet service provider. This network is owned by a regional internet service provider and this network is company's network.

This example illustrates three different networks connected together. The internet is a network of thousands of such interconnected networks. These networks are owned by different ISPs. And hence, the name internet, which stands for internetwork. It's a network of many networks.

What is an internet service provider? An internet service provider is an organization, which owns a network that connects end users to the rest of the internet and/or routes data among other ISPs. So, an ISP owns a network of many routers and communication links, such as this one shown here. In this figure, we see that there are several routers and these routers are connected together by communication links, which are shown here. So, an ISP owns such a network of routers and communication links, and using such a network, it connects end users to the rest of the internet.

For example, there might be some end users over here, which connect to the rest of the internet using this ISP network. And/or the ISP routes data among other ISPs. There may be one ISP network connected here. There may be one ISP network connected here. This particular ISP routes data between this ISP and this ISP via these routers and communication links.

This shows the structure of the internet. Each cloud is an ISP network like the one shown on the previous slide. So, we see that the internet is a hierarchical structure of ISP networks. So, it has several tiers. This is tier 1, this is tier 2, this is tier 3.

At the edge of this hierarchy that is over here, there are university and corporate campus networks, such as IIT Bombay's network. And there are internet service providers that connect end users to the internet. For example, BSNL, Tata Indicom, Airtel, and so on. So, these university and corporate campus networks and ISPs that connect end users to the internet, these are at the edge of this hierarchy over here. So, these ISPs use various technologies, such as DSL, cable internet, Ethernet, Wi-Fi, and so on, which we'll discuss later, to connect these end users to the rest of the internet.

And at the center of the hierarchy, there are larger ISPs that connect the access networks in one, and smaller ISPs among themselves. For example, ISPs in different states and countries are connected by larger ISPs to each other. As an example, consider one end user here, and one end user here. They are connected together by this higher-tier ISP. Consider an end user connected to this ISP, and an end user connected to this ISP.

Data is transported between this user and this user via all these ISPs. So, for example, this user's data might be sent to this ISP, which is sent to this ISP, and so on and so forth. It is then sent to this ISP, then to this ISP, then to this ISP, and then to the other user. So, via this hierarchical interconnection of many ISP networks, data is transported between end users. The financial aspects of this structure are as follows.

A lower-tier ISP pays the higher-tier ISP it is connected to for connectivity to the rest of the internet. For example, a tier-2 ISP pays the tier-1 ISP to which it is connected, and tier-3 ISP pays the tier-2 ISP to which it is connected. For example, this ISP pays this ISP because this ISP provides connectivity to the rest of the internet to this ISP. Similarly, this ISP pays this ISP for providing it connectivity to the rest of the internet. So, tier-1 ISPs typically don't pay each other for connectivity.

Tier-1 ISPs are shown in the center of the hierarchy. They're all large scale ISPs, such as for example, Tata Communications, Sprint, AT&T, and so on. Each operates over a wide area and has a large network of many routers and communication links. So, the relation among different tier-1 ISPs is symmetric. Each one provides connections to a large part of the internet to the other tier-1 ISPs.

On the other hand, the relation between a tier-3 ISP and a tier-2 ISP is asymmetric. A tier-2 ISP provides a wider coverage to a tier-3 ISP than the tier-3 ISP provides to the tier-2 ISP. So, hence, a tier-3 ISP pays a tier-2 ISP. So, this payment may be on a monthly basis or it may be on the basis of the amount of data transferred. Packets exchanged between end systems pass through many networks, as illustrated in this figure.

So, consider this desktop communicating with this desktop. The data sent by this desktop first goes through its local ISP, then to the tier-3 ISP, up the hierarchy, then to the tier-2 ISP above it, then to the tier-1 ISP, then to another tier-1 ISP, and then down the hierarchy to this tier-2 ISP, then to this local ISP, and then to this other end system. So, packets are exchanged through many networks doing their transit from one end system to another. Another thing to note is that this is a loosely hierarchical interconnection, not a strictly hierarchical interconnection. Most of the connections are between ISPs at one tier and ISP of the next higher or lower tier.

But there are some connections between ISPs of the same tier. For example, here, two tier-2 ISPs are connected directly. They don't have to go through this higher-tier ISP. So, for example, if there is a lot of traffic that flows between these two tier-2 ISPs, then they can use this direct connection. They don't have to send their traffic through this tier-1 ISP.

So, hence, it's a loosely hierarchical interconnection of ISPs. It's not a strictly hierarchical interconnection of ISPs. We now discuss access networks. In this hierarchy, we saw that access networks are at the edge of the hierarchy, and they connect end users to the rest of the internet. There are different types of access networks.

They connect end systems to the edge router. So, in this picture, the access networks are highlighted in red. So here, this access network connects these end systems to the edge router which is this one. The edge router is the router at the edge of this access network and it is connected to the rest of the internet. Similarly, this is the edge router for this network.

It is connected to the rest of the internet, and the end systems of this network are connected to this edge router via these access networks, which are shown by the red links. So, there are different types of access networks. One type of access networks are residential access networks, such as dial-up, DSL, and cable. DSL stands for Digital Subscriber Line. So, these networks connect users in homes to the rest of the internet.

Then there are institutional access networks, which are deployed in universities and company campuses. Examples are Ethernet and Wi-Fi. Ethernet is a wired local area network technology, and Wi-Fi is a wireless local area network technology. Then, there are wide-area wireless access networks, which provide connectivity over a wide area, such as an entire city or a town or a state and country, and so on. Examples are 2G, 3G, 4G, and 5G networks.

So, these are different types of access networks. So, how do we evaluate the performance of an access network? The key characteristics are as follows. One is how much bandwidth can the access network provide? Can it, for example, provide 1 Mbps or 1 Gbps or 10 Gbps?

So, the higher the bandwidth it can provide, the better it is. Another characteristic is, is it shared or dedicated? That means, by dedicated, we mean that the access network provides connectivity to an end user via a dedicated connection. For example, a cable between the end user and the access network. By shared connection, we mean that the bandwidth of the access network is shared among many users via a shared cable, for example.

So, we'll discuss examples of such shared and dedicated access networks. So, one widely used access network is digital subscriber line (DSL), which is used for residential connectivity. Some background: in the early days of the internet, for example, in the 1990s, the traditional analog telephone system was already widely deployed. It was connected to many homes. So, in this analog telephone system, telephone lines which are made of copper were connected to a large number of homes.

After the advent of the internet, it was costly to construct new cables for internet access to homes. It was easier to just reuse the existing telephone lines for internet access as well. So, to construct new cables, you would require alteration of buildings, digging of roads, and so on. So, to avoid that, it was easier to just use the existing telephone lines. DSL provides access to the internet in addition to telephone network or telephone lines.

This shows the architecture of a DSL network. This is a telephone cable which connects a user in a home, which is this one, to the nearest telephone exchange. So, this is the telephone exchange, and the telephone exchange is connected to the internet and the telephone network. This cable that connects the telephone office to the user's home, the bandwidth of this cable is divided into three parts. So, this cable has bandwidth of about 1 MHz.

This bandwidth is divided into three parts. 0 to 4 KHz is used for ordinary telephone calls. 4 KHz to 50 KHz is used for upstream data. That is, the traffic that is sent from the user to the internet. And 50 KHz to 1 MHz is used for downstream data.

That is, the data that is sent from the internet to the user. So, correspondingly, this DSL technology can support typical speeds of about 1 Mbps upstream and 8 Mbps downstream. So, these are speeds for the old version of DSL. Now, modern versions support much higher speeds. One observation here is that the bandwidth that is reserved for the downstream data, that is, the data that flows from the internet to the user, that is, this bandwidth is 950 KHz.

This is much more than the bandwidth reserved for the upstream data, that is, only 46 KHz. So, the reason is that most residential users download more data than they upload. So, hence, a higher bandwidth is provided for the downstream data than for the upstream data. This is also the case for other residential access networks. Now, we discuss cable internet which is another access network used for residential users.

Cable TV was widely deployed before the advent of the internet. And in cable TV, a shared cable broadcasts TV channels to hundreds of homes. For example, these TV channels might be Star Sports or Times Now, and so on. The cable internet reuses this infrastructure just as DSL uses the telephone network infrastructure. This shows the architecture of a cable network.

There's a cable headend, and there is a shared cable that connects this cable headend to a large number of homes, typically 500 to 5,000 homes. So, this headend is connected to the networks of different distributors, different TV channels, and it sends the TV channels over this shared cable. And when it sends the information of a TV channel, for example, Star Sports, when it sends this information over this cable, that information reaches every home connected to this cable. So, in cable internet, we use frequency division multiplexing on this cable and divide the bandwidth of this cable into three parts. 54 to 550 MHz is used for TV.

550 MHz to 750 MHz is used for downstream data, that is, the data from the internet to the homes, and 5 to 42 MHz is used for upstream data, that is from the homes to the internet. So, again note that more bandwidth is reserved for downstream data than for upstream data.

A feature to note here is that this cable is a shared cable because in the case of TV, it was designed for TV, and in the case of TV, the same information is to be transmitted to all the homes connected to the cable. So, for example, when we say Star Sports, the same information is to be transmitted to all these homes. Similarly, Times Now, the same information is to be transmitted to all these homes.

But the data exchange over the internet is not like that. A user in this home might be connected to Google. A user in this home might be connected to, say, India Times, and so on. So, there may be the data that is sent and received by different homes is different. How do we use the shared cable for providing internet connectivity in addition to TV access?

So, this cable is a shared broadcast medium since it was originally intended for cable TV, and TV is a broadcast kind of traffic. In the downstream, the signal sent by the headend is received by every home. So, whatever information is sent by this cable headend reaches every home because it's a shared cable. But each customer only uses the packets that are destined for itself. So, for example, if this user is downloading a file from the internet, then this file will reach every home, but only this user will use the file.

The other users will ignore the file. If some confidential information needs to be sent, then that can be encrypted as we'll discuss later. Now, consider the upstream, that is, the direction from the users to the internet. So, if two users transmit simultaneously to some other users in the internet, then the transmissions will collide over this cable. So, we require a medium access control protocol to share the bandwidth of this cable among the different users who are connected to this cable.

In the context of cable internet, a protocol known as DOCSIS is used as the medium access control protocol. So, DOCSIS is the protocol used for medium access control in cable internet. So, typical speeds provided by cable internet are; it provides up to 30 Mbps downstream and 2 Mbps upstream. Again these are for the basic version of cable internet. There are high speed versions available now.

But there are no guarantees on the speed, unlike DSL, because the actual speed depends on the number of active users since the medium is shared. So, in this case, this 30 Mbps is to be shared to all the users in all the homes. So, depending on how many users are active at a particular instant, the bandwidth that is available to one particular user will change. On the other hand, in the case of DSL, there was a dedicated cable from every user to the telephone office. So, although the speeds were lower, namely 8 Mbps and 1 Mbps, but those speeds were for the dedicated use of one user only, whereas these speeds, 30 Mbps and 2 Mbps, are to be shared among many users.

So, thus, we should be careful while comparing the speed of a shared medium technology like cable internet and a dedicated link technology, such as DSL. Just the raw numbers don't tell us which one will be faster. Also, we have to consider whether this bandwidth is shared among many users or it is for the exclusive use of one user. So, next we discuss Ethernet, which is a local area network often used in university and corporate local area networks. So, it is widely used in company and university lands.

It operates over short distances of a few kilometers, and it is available in speeds of 10 Mbps to 10 Gbps. In the original version of Ethernet, there was a shared cable, which we discussed in the previous lecture, and different end systems were connected to a shared cable. So, for example, the different computers in a university environment, they were all connected to the shared cable. It's a broadcast medium because whatever is transmitted by one node, that information reaches all the nodes connected to the cable. So, we require a MAC protocol for providing access to different nodes, and a MAC protocol called CSMA-CD.

It stands for Carrier Sense Multiple Access Collision Detection. This is the MAC protocol that was used for the original version of Ethernet. We won't discuss the details. In modern Ethernet, switches are used. So, the previous architecture, where there was a shared cable, that architecture doesn't operate very well in the case of high-speed local area networks, such as 1 Gbps or 10 Gbps networks.

So, in modern Ethernet, switches are used. There are dedicated connections between switches and end systems. A switched Ethernet network is shown in this picture here. These are the end systems, and they are connected via these switches. Switches are conceptually similar to routers.

They transport the data exchanged by different end systems. So, for example, consider this end system and suppose it wants to communicate with this end system. Then its data is sent via this dedicated connection to this switch, and then from this switch via a dedicated connection to this switch, and then from this switch to this switch, from this switch to this switch, and finally to this end system. So, there are dedicated connections between the end systems and the switches, and between different switches. And data is transported between end systems via these switches. And the backbone switch is connected to the rest of the internet.

Next we discuss Wi-Fi. Its technical name is IEEE 802.11. It's an IEEE standard. It is widely deployed in corporate and university lands, airports, shopping centers, homes, and so on. Your mobile hosts, such as laptops and smartphones, they are connected to a router via base station. Often the router and the base station are combined into one device.

And the base station is known as access point in Wi-Fi terminology. So, the access point is this one. It provides speeds of 1 Mbps to several Gbps. And it operates over a range of a few hundreds of meters. So, one Wi-Fi access point might operate over a range of say 100 meters.

And you can use many Wi-Fi access points to provide connectivity over a wide area, such as an entire campus of several kilometers in length and several kilometers in width. So, this is a wireless technology. So, wireless medium is a shared medium. Whatever information is transmitted by one node, it travels in all the directions around it. So, for example, the wireless transmissions of this node, they reach this node as well as this node as well as these nodes.

So, we require a MAC protocol to share the bandwidth among different nodes in a vicinity. So, the MAC protocol that is used in the case of Wi-Fi is CSMA-CA, which is Carrier Sense Multiple Access Collision Avoidance. We won't discuss the details of this. Wi-Fi uses unlicensed spectrum around 2.4 GHz and 5 GHz. By unlicensed spectrum, we mean that to use the spectrum, we don't have to acquire a license from the regulator.

Anyone can use certified Wi-Fi equipment or Bluetooth equipment and so on and operate on these frequency bands. So, these are in contrast to license bands, which are used by cellular operators, such as Airtel and Vi, and so on. These bands are inexpensive, but they are subject to interference from neighboring Wi-Fi hotspots, Bluetooth devices, microwave ovens, and so on. So, for example, the radiation emitted by microwave ovens is in this range. Yeah, the radiation emitted by microwave ovens is in these frequency bands, 2.4 GHz and 5 GHz, and so on.

And anyone can set up a Wi-Fi hotspot or a Bluetooth network, and so on. So, there can be interference among different networks using the unlicensed spectrum. Another limitation of Wi-Fi is that Wi-Fi connectivity is only available when the user has access to a hotspot in the vicinity. There may be no hotspot in the vicinity, or even if there is a hotspot, a user may not have the necessary login credentials to use that Wi-Fi network. So, hence, Wi-Fi provides connectivity over a limited range.

Another wireless technology, that is, a popular access network is wide-area wireless access. So, this is comprised of cellular networks, such as 2G, 3G, 4G, and 5G. Internet and telephone connectivity is provided in mobile phones throughout a large region, such as a city. It provides speeds of a few 10s of Mbps. The latest versions, such as 4G and 5G can provide much higher speeds, for example, hundreds of Mbps or a few Gbps.

Again, the medium used is a wireless medium, which is a shared medium. Whatever one node transmits reaches every other node in the vicinity. So, we require a MAC protocol to share out the medium among the different users in a vicinity. So, different MAC protocols are used depending on the version of the cellular network, whether it's 2G, 3G, 4G, or 5G. It is deployed by cellular operators, such as Airtel and Vi. So, we discussed that Wi-Fi uses unlicensed spectrum.

In contrast, cellular networks use licensed spectrum. So, they participate in an auction and license out the spectrum from the regulator. The spectrum is expensive because the operators need to pay the regulator to acquire the spectrum. But once it is acquired, it provides exclusive access. That is, only the operator who owns a particular frequency band can use that frequency band.

No one else is allowed to transmit on that frequency band. So, the architecture of a cellular network is shown in this picture. A region, such as a city is divided into small areas called cells. Each hexagon is a cell in this picture. And there is a base station serving the users in each cell.

For example, this base station serves the users in this cell. These cells are all connected to a switching center. This shows a 2G architecture. This is a mobile switching center and it is connected to the telephone network. So, in the case of modern networks, this is another switching center, which is connected to the internet as well as the telephone network.

Each cell is served by one base station. And when a user moves from the range of one base station to the range of another base station, then it is transferred from this base station to this base station. So now, what is the reason for dividing the region into cells? Why not have only one base station serving all the users in the entire city? So, the reason is to use a concept called frequency reuse.

That is, the same set of frequencies are used at far apart cells without mutual interference. This increases the capacity of the system. For example, consider a cellular operator, such as Airtel or Vi, who has licensed out a spectrum band. This is the frequency axis. The operator has licensed out this much spectrum from the regulator.

It divides the spectrum into different parts. Let's call these: these bands are labeled as A, B, C, D, and so on. And in this cell, band A is used. Band A is also used in this cell and in this cell. Similarly, band B is used in this cell.

It is also used in this cell and in this cell. So, the same bands are used over and over again in different cells, so by repeatedly using the same frequency bands for accessing different cells, the capacity of the system can be shown to increase. So, the idea is that by the time the transmissions of this base station reach this point, where the same frequencies are being used, the transmissions attenuate, so they don't interfere with the transmissions of this base station. So, by using this concept of frequency reuse, the capacity of the system is improved. Next we discuss communication links.

They connect end systems and routers together. So, in this picture, these are end systems and these are routers, and they are connected together by different communication links, wired and wireless. These communication links are made of various physical media, such as copper cable and wireless. For example, digital subscriber line and cable internet, they use copper cable. And Wi-Fi and cellular networks use the wireless medium.

We discuss one particular type of communication link in some more detail, fiber optic cable. Here, a thin glass fiber is used to transmit light signals. This shows a collection of fibers. And light signals are transmitted through these fibers. Fiber optic cables can provide very high speed operation.

A single fiber optic cable can provide hundreds of Gbps. It also has the desirable property of very low signal attenuation. What this means is that suppose there is a transmission from point A to point B. If this is a long distance, for example, say 1,000 kilometers, then the signal attenuates by the time it reaches the other end. So, we use what are known as repeaters on the path from A to B. Each repeater takes the incoming signal, and attenuates it, and filters it, and then transmits it to the next repeater, and so on and so forth until it reaches B. So, we require repeaters on long distance links. But these repeaters are required only once every 50 km or so in the case of fiber optic as opposed to once every 5 km for copper cable because the attenuation is lower in the case of fiber optic signals.

But the downside of fiber optic technology is that optical transmitters and receivers are expensive. Fiber optic technology has been the dominant long-haul medium. That is, for long distance communication, mostly fiber optic communication is used. And it is widely deployed in the core of the internet, including overseas links. For example, there are a large number of underwater links connecting different continents.

They are made up of fiber optic technology. But now, even DSL and cable connections to residences, companies, and so on are being upgraded to fiber links. And this technology is known as FTTX. X can be either P or H or B. These stand for fiber to the premises, or fiber to the home, fiber to the building, and so on. So, here are the connections to residences, companies, and so on.

They are also being upgraded to fiber optic links. in contrast to DSL or cable, which were used earlier. So, these fiber optic links provide much higher speeds to users in residences, companies, and so on. So, even the links at the edges of the network are being upgraded to fiber optic cables. This concludes our discussion of access networks.

Thank you.