

Computer Networks (CN)

Rupak R. Gupta (Siouka)

Aaditya Joil (Jojo)

Advait Desai (Drunkenstein)

2025

Contents

1	Introduction	4
1.1	Data Communications	4
1.1.1	Components of Data Flow	4
1.1.2	Types of Data Flow	5
1.2	Networks	6
1.2.1	Network Types	6
1.3	Networking Topologies	8
1.3.1	Mesh	8
1.3.2	Star	9
1.3.3	Bus	9
1.3.4	Ring	10
1.3.5	Hybrid	10
1.4	Internet	11
1.5	Network Models	11
1.5.1	OSI Model	11
1.5.2	TCP/IP Protocol Suite	13
1.6	Addressing	14
1.6.1	Link-layer Addresses	15
1.6.2	Logical Addresses	15
1.6.3	Port Addresses	15
1.6.4	Specific Addresses	15
2	Physical Layer	16
2.1	Transmission Media	16
2.2	Guided Media	16

2.2.1	Twisted-Pair Cable	16
2.2.2	Coaxial Cable	17
2.2.3	Fibre-Optic Cable	18
2.3	Unguided Media	19
2.3.1	Propagation Methods	19
2.3.2	Wireless Transmission Waves	20
3	MAC Layer	22
3.1	Design Issues	22
3.1.1	Addressing	22
3.2	Error Detection and Correction	23
3.2.1	Types of Errors	23
3.2.2	Redundancy	23
3.2.3	Detection versus Correction	23
3.2.4	Block Coding	23
3.2.5	Checksum	23
3.3	Network Performance	23
3.3.1	Latency	23
3.3.2	Bandwidth-Delay Product	23
3.3.3	Round Trip Time	23
3.4	Data Link Control	23
3.4.1	Noiseless Channels	23
3.4.2	Noisy Channels	23
4	Network Layer	24
4.1	IPv4 Addressing	24
4.1.1	Address Space	24
4.1.2	Notation	24
4.1.3	Classful Addressing	25
4.1.4	Classless Addressing	26
4.2	Network Address Translation	27
4.2.1	Address Translation	28
4.2.2	Translation Table	28
4.3	Internetworking	29
4.3.1	Need for Network Layer	29
4.3.2	Internet as a Connectionless Network	29

4.4	IPv4	29
4.4.1	Datagram	30
4.4.2	Fragmentation	32
4.4.3	Checksum	34
4.4.4	Options	34
4.5	Address Mapping	34
4.5.1	ARP	34
4.5.2	ICMP	34
4.6	Delivery	34
4.7	Forwarding	34
4.8	Routing	34
4.8.1	Unicast Routing Protocols	34
4.8.2	Multicast Routing Protocols	34
4.8.3	Routing Algorithms	34
5	Transport Layer	35
5.1	Process-to-Process Delivery	35
5.2	Addressing	36
5.2.1	IANA Ranges	36
5.2.2	Socket Addresses	36
5.3	User Datagram Protocol	36
5.3.1	Well-Known Ports for UDP	36
5.3.2	UDP Datagram	37
5.4	Transmission Control Protocol	38
6	Applications	39
6.1	Traditional Applications	39
6.1.1	World Wide Web	39
6.1.2	Hypertext Transfer Protocol	39
6.1.3	File Transfer Protocol	39
6.1.4	Email	39
6.1.5	Telnet and SSH	39
6.1.6	Domain Name Service	39
6.2	Peer-to-Peer Networks	39
6.3	Advances in the domain	39
	Acronyms	40

Chapter 1

Introduction

1.1 Data Communications

The term *telecommunication* means communication at a distance. The word *data* refers to information presented in whatever form is agreed upon by the parties creating and using the data. *Data communications* are the exchange of data between two devices via some form of transmission medium, such as a wire cable.

1.1.1 Components of Data Flow

There are five main components of data flow as shown in fig. 1.1.




Figure 1.1: Components of Data Flow

Message

Message is the information (data) to be communicated. Popular forms of information include text, numbers, pictures, audio and video.

Sender

The sender is the device that sends the data message. It can be a computer, workstation, telephone handset, video camera, and so on.

Receiver

The receiver is the device that receives the data message. It can be a computer, workstation, telephone handset, video camera, and so on.

Transmission Medium

The transmission media is the physical path by which a message travels from a sender to a receiver. Some examples of transmission media include twisted-pair, coaxial cable, fibre-optic cable, and radio waves.

Protocol


A protocol is a set of rules that govern data communications. It represents an agreement between the communicating parties. Without a protocol, two devices may be *connected but not communicating*. For example, two people speaking to each other in different languages.

1.1.2 Types of Data Flow

Data flow refers to the way data moves from one communicating party to the other. Depending on the flow of data, data communication is of three types:


Simplex

In simplex mode, the communication is *unidirectional*. Only one of the two devices on a link can transmit; the other can only receive. Keyboards and traditional monitors are examples of simplex devices.




Half-Duplex

In half-duplex mode, the communication is *partially bidirectional*. Both the devices can transmit and receive data, but not at the same time. When one device is acting as a transmitter, the other must act as a receiver, and vice versa. Walkie-talkies and Citizens Band (CB) radios are both half-duplex systems.



Full-Duplex

In full-duplex mode, the communication is *bidirectional*. Both the devices can transmit and receive data simultaneously. Telephone systems are examples of full-duplex systems.



1.2 Networks

A *network* is a set of devices (often referred to as *nodes*) connected by communication *links*. A node can be a computer, printer, or any other device capable of sending and/or receiving data generated by other nodes on the network. We usually divide nodes into two types: *host* and *communicating device*.

Host

A host (or an end system) is a device which appears at the either ends of networks. It is usually the initial source and final destination of the data. Hosts are usually large servers or mainframes, desktops, laptops, workstations, cellular phones, or security systems.

Communicating Devices

Communicating devices are devices which interconnect the various parts of the networks. These are devices such as routers, which connect one network to another, switches, which connect devices together, a Modulator-Demodulator (MoDem), which changes the form of data, etc.

1.2.1 Network Types


The criteria of distinguishing one network from the other is difficult and sometimes confusing. A few criteria such as size, geographical coverage, and ownership are used to make this distinction.

Local Area Networks

A Local Area Network (LAN) is usually privately owned and connects some hosts in a single office, building, or a campus. Depending on the needs of an organisation, a LAN can be as simple as two PCs and a printer in a home office, or it can extend throughout a company and include audio and video devices.

In the past, all hosts in the network were connected through a common cable, which meant that a packet sent from one host to another was received by all hosts. Only the intended host kept the packets, all other dropped it. In the present, most LANs use a smart connecting switch, which is able to recognise the destination address of the packet and guide the packet to its destination without sending it to all other hosts. This alleviates the traffic on the network and allows more than one pair of computers to communicate simultaneously if they do not share common sources and destinations.

When LANs are used in isolation, they are used to allow resources to be shared between the hosts (for example, printers). Most LANs today are connected to each other and to WANs to create communication at a wider level.




Wide Area Networks

A Wide Area Network (WAN) is also an interconnection of devices capable of communication. However, there are some differences between a LAN and a WAN. A LAN is normally limited in size, spanning an office, a building or a campus; a WAN has a wider geographical span, spanning a state, a country, or even the world.

A WAN interconnects communication devices as compared to hosts in a LAN. A LAN is normally privately owned by the organisation that creates it; whereas a WAN is normally created and run by communication companies and is leased to organisations that want to use it.


Point-to-Point WAN

A point-to-point WAN is a network that connects two communicating devices through a transmission media.



Switched WAN

A switched WAN is a network with more than two ends. A switched WAN is used in the backbone of global communication today. A switched WAN can be said to be a combination of several point-to-point WANs that are connected by switches.



Metropolitan Area Networks


A Metropolitan Area Network (MAN) is a computer network that connects computers and other devices across a larger geographical area than a LAN but smaller than a WAN. It usually covers metropolitan areas such as towns to small cities. It is formed by interconnecting multiple LANs together.

MANs are built for high-speed data transfer and connectivity, often using technologies like fibre optics and wireless communication. MANs can be owned and operated by a single organization, a consortium of users, or a network provider.

Internetwork

An internetwork, or an internet (note the lowercase *i*), is a network which is formed by connecting two different types of networks together. An example of an internetwork is a large national


company which has two LANs at two different office locations across the country and has leased a WAN from a service provider in order to connect the two LANs together.



1.3 Networking Topologies

1.3.1 Mesh

In a Mesh Topology, all the hosts have dedicated links to all other hosts. This setup requires $n(n - 1)/2$ cables (if the cables allow duplex data flow) or $n(n - 1)$ (if the cables allow only simplex data flow).



Advantages of mesh topology:


- Use of dedicated links guarantees that each connection can carry its own data load, thus eliminating the traffic problems that occur when links must be shared by multiple devices.
- One link becoming unusable does not incapacitate the entire system.
- Increased Privacy and Security, the message is only forwarded to it's intended recipient.
- Easier fault identification and isolation.

Disadvantages of mesh topology:

- Number of ports and cabling required.
- Amount of required cabling cannot be housed in walls or ceilings.
- Cost of hardware.

1.3.2 Star

In a Star Topology, each device has a dedicated point-to-point link only to a central controller, usually called a *hub*. The devices are not directly linked to each other. If one device wants to send data to another, it sends the data to the controller, which then relays the data to the other connected device.



Advantages of star topology:


- It is less expensive than a mesh topology.
- Far less cabling needs to be housed.
- Easier to install and reconfigure (additions, moves, deletions involve only one connection).
- Robustness of topology, if one link fails, only that link is affected.

Disadvantages of star topology:

- Main disadvantage is the dependency of the whole topology on the hub. If the hub goes down, the whole system is dead.
- Each device in the network must be linked to the hub. This may lead to unnecessary cabling.

1.3.3 Bus

A bus topology makes use of multipoint connections. One long cable acts as a backbone to link all the devices in a network. Nodes are connected to the bus cable by drop lines and taps.



Advantages of bus topology:


- Ease of installation. Backbone cable can be laid along the most efficient path, then connected to the nodes by drop lines of various lengths.

Disadvantages of bus topology:

- Difficult reconnection and fault isolation.
- It is difficult to reconfigure the network. Adding or removing devices requires the replacement of the entire backbone.
- A fault or a break stops all transmissions, even on the same side of the fault as the signal is reflected back as noise to the point of origin.

1.3.4 Ring

In a ring topology, each device has a dedicated point-to-point connection with only the two devices on either side of it. A signal passed along the ring in one direction, from device to device, until it reaches its destination. Each device in the ring incorporates a repeater. When a device receives a signal intended for another device, its repeater regenerates the bits and passes them along.



Advantages of ring topology:

- A ring is relatively easy to install and reconfigure.
- Each device is only connected to two devices, so to add or remove the device only two connections need to be modified.

Disadvantages of ring topology:

- Unidirectional traffic can be a disadvantage.
- In a simple ring, a break in the ring can disable the entire network. This weakness can be solved by using a dual ring or a switch capable of closing off the loop.

1.3.5 Hybrid

A single type of topology is not prevalent in the present, in most cases a hybrid sort of topology is used. For example, an office network may use a singular hub with three backbones sprawling across the entire office with drop lines attaching the devices to the closest one.

1.4 Internet

The Interconnected Network (Internet) is the world's largest internetwork. It is composed of thousands of interconnected networks.

The Internet is composed of several backbones, provider networks and customer networks. At the top levels, the backbones are owned by some communication companies. They are connected to each other by some complex switching systems known as *peering points*. At the second level, there are smaller networks, called the provider networks, that make use of the backbones for a fee. The provider networks are connected to the backbones as well as other provider networks. The customer networks are at the edge of the Internet and these actually use the services provided by the Internet.

Backbones and provider networks are also called Internet Service Provider (ISP). The backbones are referred to as international ISP and the provider networks are referred to as national or regional ISP.

1.5 Network Models

1.5.1 OSI Model

Established in 1947, the International Standards Organization (ISO) is a multinational body dedicated to worldwide agreement on international standards. An ISO standard that covers all aspects of network communications is the Open Systems Interconnection (OSI) model. It was first introduced in the late 1970s.

Application Layer
Presentation Layer
Session Layer
Transport Layer
Network Layer
Data-link Layer
Physical Layer

Physical Layer

The physical layer is responsible for movements of individual bits from one hop (node) to the next. Some of its responsibilities are:

Physical characteristics The interface between the devices and the transmission media.

Representation of bits Bits must be encoded into signals: electrical or optical.

Data rate The number of bits sent each second.

Synchronization of bits The sender and receiver must not only use the same bit rate, but must also be synchronized at the bit level.

Data Link Layer

The data link layer is responsible for moving frames from one hop (node) to the next. Some of its responsibilities are:

Framing Data link layer divides the stream of bits received from the network layer into manageable data units called *frames*.

Physical addressing The data link layer adds a header to the frame to define the sender and/or receiver of the frame.

Flow control The rate at which the data is absorbed by the receiver.

Error control To detect and retransmit damaged or lost frames, duplicate frames using trailer added to the end of the frame.

Access control Connection link is controlled by one device at any given time.

Network Layer

The network layer is responsible for the delivery of individual packets from the source host to the destination host. Some of its responsibilities are:

Logical addressing If a packet passes the network boundary, we need another addressing system to help distinguish the source and destination systems.

Routing When independent networks or links are connected together to create *internetworks* (network of networks) or a large network, the connecting devices route or switch the packets to their final destination.

Transport Layer

The transport layer is responsible for the delivery of a message from one process to another. Some of its responsibilities are:

Service point addressing Source-to-destination delivery means delivery not only from one computer to the next, but also from a specific process (running program) on one computer to a specific process (running program).

Segmentation and reassembly A message is divided into transmittable segments, with each segment containing a sequence number.

Connection control The transport layer can be either connectionless or connection-oriented.

Flow control Flow control at this layer is performed end to end rather than across a single link, unlike in the Data Link Layer.

Error control Error control at this layer is performed process-to-process rather than across a single link. Error correction is usually achieved through retransmission.

Session Layer

The session layer is responsible for dialog control and synchronization. It establishes, maintains, and synchronizes the interaction between communicating systems. Some of its responsibilities are:

Dialog control Allows the communication between two processes to take place in either half-duplex (one way at a time) or full-duplex (two ways at a time) mode.

Synchronization The session layer allows a process to add *checkpoints*(synchronization points) into a stream of data.

Presentation Layer

The presentation layer is concerned with the syntax and semantics of the information exchanged between two systems. Some of its responsibilities are:

Translation The processes (running programs) in two systems are usually exchanging information in the form of character strings, numbers, and so on.

Encryption *Encryption* means that the sender transforms the original information to another form and sends the resulting message out over the network.

Compression Data compression reduces the number of bits contained in the information such as text, audio, and video.

Application Layer

The application layer is responsible for providing services to the user.

Encapsulation/Decapsulation at each layer

In each layer of the OSI model, encapsulation and decapsulation of data occurs. At the source host, each layer receives the message from the upper layer, adds its own header and passes it on to the lower layers (thus encapsulating the message). This can be seen in fig. 1.2.




Figure 1.2: Encapsulation at Source Host

At the destination host, each layer receives the encapsulated message from the lower layers, removes the corresponding headers and passes it on to the upper layers (thus decapsulating the message). This can be seen in fig. 1.3.




Figure 1.3: Decapsulation at Destination Host

1.5.2 TCP/IP Protocol Suite

The layers in the Transmission Control Protocol/Internet Protocol (TCP/IP) suite do not exactly match those in the OSI model. The TCP/IP protocol suite was defined as having four layers: Network Access, Internet, Transport, and Application.

The transport layer in the TCP/IP protocol suite, encompasses a few elements of the the session layer from OSI and the remaining elements of the session layer and the entirety of the presentation layer are handled by the application layer in the TCP/IP protocol suite.

Application Layer
Transport Layer
Internet Layer
Network Access Layer
Physical Layer

The TCP/IP protocol suite was defined and implemented before the OSI model was published. The OSI model gained initial popularity but due to a few reasons such as the Session and Presentation Layers protocols not being defined as well as due to the TCP/IP model already being implemented, it would cost a lot to replace it. Apart from this, one implementation of the OSI model was actually completed but it did not show much performance. Hence, the OSI model is only a theoretical model whereas the TCP/IP is the one upon which the Internet actually works.

Network Access Layer

The Network Access layer is equivalent to TCP/IP protocol suite is equivalent to the Data-Link Layer in the OSI model.

Internet Layer

The Internet layer is equivalent to TCP/IP protocol suite is equivalent to the Network Layer in the OSI model.

Transport Layer

The Transport layer is equivalent to TCP/IP protocol suite is equivalent to the Transport Layer in the OSI model. It also incorporates some elements of the Session Layer as well.

Application Layer

The Application layer is equivalent to TCP/IP protocol suite is equivalent to the Session, Presentation and Application Layers in the OSI model.

1.6 Addressing

Any communication that involves two parties requires a pair of addresses. Since we have five layers in TCP/IP, it is logical to assume that we will require five pairs of addresses but it is not so. The unit of data exchanged below the physical layer are the bits, and we cannot address individual bits. Four levels of addresses are used in an internet employing the TCP/IP protocols: *link-layer*, *logical*, *port*, and *specific*.

1.6.1 Link-layer Addresses

The link-layer addresses (also called the MAC addresses), are locally defined addresses, each of which defines a specific host or router in a network.

1.6.2 Logical Addresses

The logical addresses (also known as IP addresses) are globally defined addresses with the whole Internet as their scope. These are defined by the IP itself in the third layer.

1.6.3 Port Addresses

Port numbers are local addresses running on a specific machine which differentiate between several programs running at the same time.

1.6.4 Specific Addresses

At the application layer, the names of the programs or application which provide services to the user are used as addresses themselves (*rrjo.vercel.app* or *example@email.com*).

Chapter 2

Physical Layer

The physical layer handles communication between two nodes (computers or routers) using individual bits. TCP/IP does not define a specific physical layer protocol.

Although the unit of exchange in this layer is the bits, the data is actually converted to *signals* to be transmitted along the media and reconverted back to data at the receiving end.

2.1 Transmission Media

A transmission medium can be broadly defined as anything that can carry information from a source to a destination. The information is usually a signal that is the result of a conversion of data from another form.

2.2 Guided Media

Guided media, which are those that provide a conduit from one device to another, include twisted-pair cable, coaxial cable, and fiber-optic cable.

2.2.1 Twisted-Pair Cable

A twisted pair consists of two conductors (usually copper), each with its own plastic insulation, *twisted together*, as shown in fig.

figure

One of the wires is used to carry the signal to the receiver and the other is used as a ground reference. The receiver uses the difference of the two signals to find out whether a one or a zero is being sent.

The wires may be subject to interference and crosstalk from other sources and this may distort the signal. If the wires are parallel, there is a high likelihood that the source of noise is not equidistant from them leading to an unequal change in the electrical signal in the wires. Since the wires are twisted together, they are both equally affected by the distorting signal and at the receiving end, due to the difference being taken, the noise effectively cancels out.

Types

The most common twisted-pair cables used in communication industries is the Unshielded Twisted-Pair (UTP). It is the standard twisted-pair wrapped in a plastic cover with no major modifications.

figure

IBM introduced another type of twisted-pair cable called the Shielded Twisted-Pair (STP). They have a metallic foil or braiding encasing the two cables. This improves the quality of the cable by preventing noise penetration and crosstalk. Due to the metal braiding, the cable becomes bulkier and expensive. STPs are seldom used outside of IBM.

Connectors

The most common UTP connector is the Registered Jack (RJ)45 connector. This is a keyed connector, meaning the connector can be inserted in only one way.

figure

Application

- Telephone lines to provide voice and data channels. The local loop which connects subscribers to the central telephone office.
- The DSL lines used by telephone companies to provide high-data-rate capacities.
- Local-area networks, such as the 10Base-T and the 100Base-T networks.

2.2.2 Coaxial Cable

Coaxial (Coax) cable are cables which are *coaxial* in nature. This type of cable consists of a singular stranded or solid wire core conductor, which is encased in an insulating sheath, which is further encased in a metallic braid (which acts as a shield as well as completes the circuit). The outer braid is also covered in an insulating sheath before being covered by the plastic cover of the cable.

figure

Due to this construction, the coaxial cables are able to carry frequencies in much higher ranges than the twisted-pair cables.

Types

Coaxial cables are categorized by their Radio Government ratings. An RG number denotes a unique set of physical specifications, such as the wire gauge of the inner conductors, the thickness, and type of insulators, the construction of the shield, etc.

<i>Category</i>	<i>Impedance</i>	<i>Use</i>
RG-59	75Ω	Cable TV
RG-58	50Ω	Thin Ethernet
RG-11	50Ω	Thick Ethernet

Table 2.1: Coax categories

Connectors

The most popular type of Coax connectors are the Bayonet Neill-Concelman (BNC) connectors.

figure

- The BNC connectors are used to connect the end of the cable to a device like a Television set.
- The BNC T-connector is used in Ethernet networks to branch out to a connection to a computer or other device.
- The BNC terminator is used at the end of a cable to prevent the reflection of the signal.

Applications

- Telephone networks for high bandwidth and data rate facilities.
- Television networks, near the homes of the consumers. The rest used fibre optics.
- Traditional Ethernet LANs, such as 10Base-2 (Thin Ethernet) and 10Base-5 (Thick Ethernet).

2.2.3 Fibre-Optic Cable

A fibre-optic cable is made up of glass or plastic and *transmits signals in the form of light*.

figure

An optically dense glass or plastic *core* is covered by a less dense *cladding*. An optical fibre makes use of the concept of *Total Internal Reflection* in order to transmit data.

Propagation Modes

Two major modes of light propagation are used in optical fibres, multimode and single-mode.

In multimode, multiple beams of light move through the cable at the same time. The ways in which these beams move depends on the index of the cable.

In single-mode, a single beam of light which is highly focused is used. The material used is such that the critical angle is around 90° .

figure

Construction

The main fibre optic cable is surrounded by several layers of different material to preserve the core and cladding. First a plastic buffer used to stabilise the fibre is present, then strands of Kevlar are present to strengthen the cable, the entire cable is wrapped in an outer jacket made from Teflon or PVC.

figure

Connectors

There are three major types of connectors used for optical fibres. Subscriber-Channel (SC), Straight-Tip (ST) and MT-RJ.

The SC connector is used for cable TV. It uses a push/pull locking system. The ST connector is used for connecting cable to networking devices. It uses a bayonet locking system. MT-RJ is a connector that is the same size as RJ45.

figure

Advantages

- Higher bandwidth leading to dramatically high data rates.
- Less signal attenuation and less repeaters required.
- Immunity to electromagnetic interference.
- Resistant to corrosive materials.
- Light in weight.
- Greater immunity to tapping.

Disadvantages

- Installation and maintenance is difficult and requires expertise.
- Unidirectional light propagation. If bidirectional communication is required then two cables are required.
- More expensive than the other forms of guided media.

2.3 Unguided Media

Unguided media transport electromagnetic waves without using a physical conductor. This type of communication is often referred to as wireless communication. Signals are often propagated through free space and are thus available to everyone.

2.3.1 Propagation Methods

Unguided signals can travel from the source to the destination in several ways, namely ground propagation, sky propagation and line-of-sight propagation.

Ground Propagation

In ground propagation, the waves travel through the lowest portion of the atmosphere closely hugging the ground. The low frequency signals used emanate in all directions from the transmitter and follow the curvature of the planet. The distance travelled is directly proportional to the power of the signal.

Sky Propagation

Higher frequency waves radiate upward into the ionosphere where they are reflected back to earth. This type of transmission allows for greater distances with lower output power.

Line-of-Sight Propagation

Very high frequency signals are transmitted in straight lines directly between antennas. The antennas must be directional, facing each other, either tall enough or close to each other so that they are not affected by the curvature of the earth.

2.3.2 Wireless Transmission Waves

Wireless transmission involves sending and receiving signals in the form of electromagnetic waves in free space. Broadly, three broad groups of waves are used: Radio, Microwaves and Infrared

Radio Waves

Electromagnetic waves ranging in the frequencies between 3 kHz and 1 GHz are normally called *radio waves*.

Radio waves are usually omnidirectional waves which are propagated in all directions starting from a point. The antennas need not be aligned. The omnidirectional property has a disadvantage that they are susceptible to interference by another antenna that may send signals using the same frequency or band.

Radio waves, especially those that propagate in the sky mode are able to travel long distances. This makes radio waves a good candidate for long-distance broadcasting such as AM radio.

Radio waves (of low and medium frequencies) can also penetrate through walls, this can both be an advantage and disadvantage as devices such as an AM radio may be able to receive a signal inside a building but we are not able to isolate wireless communication to inside a building.

Omnidirectional antenna with a point transmitter and receiver are usually used.

Applications

The omnidirectional characteristics of radio waves make them useful for multi-casting. Where there are many receivers but one transmitter. FM radio, AM radio, television, maritime radio, cordless phones, and paging.

Microwaves

Electromagnetic waves ranging in the frequencies between 1 and 300 GHz are called *microwaves*.

Microwaves are unidirectional in nature, i.e. when an antenna transmits, the signal can be focused to a narrow beam. The transmitting and receiving antennas need to be aligned. This has the advantage that one pair of antennas do not interfere with another pair and the same band of frequency can be used for multiple simultaneous communication.

Microwave propagation is line-of-sight. Since the antenna need to be facing each other, the towers with the antenna which are far apart need to be tall as well to overcome the curvature of the earth. Repeaters are often required for long range communication.

Very-high frequency microwaves cannot penetrate walls and hence communication inside a building is not possible.

Unidirectional antennas such as *parabolic dish* and *horn* antenna are used for sending and transmitting signals.

Applications

Due to unidirectional properties are very useful in situations where unicast communication is

required. They are used in cellular phones, satellite networks, and wireless LAN.

Infrared

Electromagnetic waves with frequencies from 300 GHz to 400 THz are known as *infrared waves*. These can be used for short range communications.

Due to high frequencies, these cannot penetrate walls. This is helpful because we can isolate one communication system from another.

Infrared signals cannot be used for long range communication. Furthermore, they cannot be used outside a building because the sun's rays contain infrared waves that can interfere with the communication.

Applications

The infrared band has an excellent potential for data transmission due to its high bandwidth (almost 400 THz). Very high data rates can be achieved.

Standards for data communication using infrared rays have been decided by the *Infrared Data Association*. Many computers provide a special IrDA port that allows a wireless keyboard to communicate with a the computer provided that the keyboard points to the port.

Apart from that almost all remote controlled devices such as television, air conditioner and remote controlled cars make use of Infrared communication.

Chapter 3

MAC Layer

The data-link layer (also sometimes called the MAC layer) is the second layer in the TCP/IP protocol suite.

This layer handles the communication of *frames* between two different nodes in a point-to-point or a multipoint link.

3.1 Design Issues

To better understand the working of the protocols at this layer, we can think of this layer as two independent sublayers, the Data Link Control (DLC) and the Media Access Control (MAC) layers.




Figure 3.1: Sublayers of data-link layers for broadcast and point-to-point links

The DLC layer, deals with issues regarding both point-to-point links and broadcast links, while the MAC layer deals with issues regarding broadcast links.

3.1.1 Addressing

The Internet is a connectionless network. We use IP addresses in the Internet, but the IP packet header only contains the source and destination addresses. The source and destination IP addresses are useless if we want to transfer packets across a non-mesh network.

We cannot change source and destination addresses in the IP packet header as they are required for error and flow control using ICMP. For a connectionless network, such as the Internet, we need to make use of another addressing mechanism. The *link-layer address* of the two nodes in a point-to-point connection is also known as the

3.2 Error Detection and Correction

3.2.1 Types of Errors

3.2.2 Redundancy

3.2.3 Detection versus Correction

3.2.4 Block Coding

Hamming Distance

3.2.5 Checksum

3.3 Network Performance

3.3.1 Latency

$$\text{Throughput} = \frac{\text{Number of bits carried}}{\text{Time taken}}$$

Latency = propagation time + transmission time + queuing time + processing delay

$$\text{Propagation time} = \frac{\text{Distance}}{\text{Propagation speed}}$$

$$\text{Transmission time} = \frac{\text{Message size}}{\text{Bandwidth}}$$

3.3.2 Bandwidth-Delay Product

3.3.3 Round Trip Time

3.4 Data Link Control

3.4.1 Noiseless Channels

Simplest Protocol

Stop-and-Wait Protocol

3.4.2 Noisy Channels

Go-Back-N ARQ

Stop-and-Wait ARQ

Selective Repeat ARQ

Chapter 4

Network Layer

4.1 IPv4 Addressing

An Internet Protocol (IP)v4 address is a 32-bit address that uniquely and universally defines the connection of a device (for example, a computer or a router) to the Internet.

IPv4 addresses are unique and universal. Two devices on the Internet can never have the same address at the same time.

4.1.1 Address Space

A protocol such as IPv4 that defines addresses, has an **address space**. An address space is the total number of addresses used by the protocol. If a protocol uses N bits to define the address, the address space is 2^N .

IPv4 uses 32-bit addresses, which means that the address space is 2^{32} or 4,294,967,296. This implies that, theoretically, more than 4 billion devices could be connected to the internet. In practice, however, this number is much less because of the restrictions imposed on the addresses.

4.1.2 Notation

There are two prevalent notations to show an IPv4 address: **binary notation** and **dotted-decimal notation**.

Binary notation

In binary notation, the IPv4 address is displayed as 32 bits. Each octet is often referred to as a byte. So it is common to hear an IPv4 address referred to as a 4-byte address. The following is an example of an IPv4 address in binary notation:

01110101 10010101 00011101 00000010

Dotted decimal notation

To make the IPv4 address more compact and easier to read, Internet addresses are usually written in decimal form, with a decimal point (dot) separating the bytes. The following is the

dotted-decimal notation of the above address:

117.149.29.2

4.1.3 Classful Addressing

IPv4 addressing, at its inception, used the concept of classes. This architecture is called **classful addressing**. Although this scheme is becoming obsolete, we briefly discuss it here to show the rationale behind classless addressing.

In classful addressing, the address space is divided into five classes: A, B, C, D, and E. Each class occupies some part of the address space.

We can find the class of an address when given the address in either notation.

- If the address is given in binary notation, the first few bits can immediately tell us the class of the address.
- If the address is given in dotted-decimal notation, the first byte defines the class.

Both methods are shown in fig. 4.1.




Figure 4.1: Finding the classes in binary and dotted-decimal notation

Classes and Blocks

One problem with classful addressing is that each class is divided into a fixed number of blocks with each block having a fixed size as shown in table 4.1.

<i>Class</i>	<i>Number of Blocks</i>	<i>Block Size</i>	<i>Application</i>
A	128	16,777,216	Unicast
B	16,384	65,536	Unicast
C	2,097,152	256	Unicast
D	1	268,435,456	Multicast
E	1	268,435,456	Reserved

Table 4.1: Number of blocks and block size in IPv4 addressing

We can observe the flaw in the above table. A block in class A addresses is too big for any organization, while a block in class C addresses is probably too small for many organizations, with class B being somewhere in the middle. Lastly, class D and class E addresses cannot be used for addressing in the typical sense; resulting in another waste of addresses.

In classful addressing, a large part of the available addresses were wasted.

Netid and Hostid

An IP address in class A, B or C is divided into **netid** and **hostid**. These parts are of varying lengths, depending on the class of the address. Note that this concept does not apply to classes D and E.

Mask

Although the length of the netid and hostid (in bits) is predetermined in classful addressing, we can also use a **mask**, a 32-bit number made of contiguous 1s followed by contiguous 0s. The concept does not apply to classes D and E.

<i>Class</i>	<i>Binary</i>	<i>Dotted-Decimal</i>	<i>CIDR</i>
A	11111111 00000000 00000000 00000000	255.0.0.0	/8
B	11111111 11111111 00000000 00000000	255.255.0.0	/16
C	11111111 11111111 11111111 00000000	255.255.255.0	/24

Table 4.2: Default masks for classful addressing

The mask can help us find the netid and the hostid. For example, the first 8 bits of any address in class A define the netid; the next 24 bits define the hostid.

The last column of table 4.2 shows the mask in the form $/n$, where n can be 8, 16 or 24. This notation is called the **Classless Inter-Domain Routing (CIDR)** notation. It will be important later for classless addressing.

Subnetting

If an organization was granted a large block in class A or B, it could divide the addresses into several contiguous groups and assign each group to smaller networks (called **subnets**).

Supernetting

In supernetting, an organization can combine several class C blocks to create a larger range of addresses. In other words, several networks are combined to create a **supernet**.

Address Depletion

So turns out we ran out of IP addresses. How can there be more than 4 billion devices using the internet at once?!?! One solution that has alleviated the problem is the idea of classless addressing.

Classful addressing, which is almost obsolete, is replaced with classless addressing.

4.1.4 Classless Addressing

To overcome address depletion and give more organizations access to the Internet, **classless addressing** was designed and implemented. In this scheme, there are no classes, but the addresses are still granted in blocks.

Address Blocks

In classless addressing, when an entity needs to be connected to the Internet, it is granted a block (range) of addresses. The size of the block (number of addresses) varies based on the

nature and size of the entity.

Restriction To simplify the handling of addresses, the Internet authorities impose three restrictions on classless addressing blocks:

1. The addresses in a block must be contiguous.
2. The number of addresses in a block must be a power of 2.
3. The first address must be evenly divisible by the number of addresses.

Mask

A better way to define a block of addresses is to select any address in the block and the mask. A *mask* is a 32-bit number in which the n leftmost bits are 1s and the $32 - n$ rightmost bits are 0s. Unlike classful addressing, n can take any value from 0 to 32 in classless addressing.

It is very convenient to represent the mask with just the value of n preceded by a slash (CIDR notation).

In IPv4 addressing, a block of addresses can be defined as

$$x.y.z.t/n$$

in which $x.y.z.t$ defines one of the addresses and the $/n$ defines the mask.

First address The first address in the block can be found by setting the $32 - n$ rightmost bits in the binary notation of the address to 0s.

Last address The last address in the block can be found by setting the $32 - n$ rightmost bits in the binary notation of the address to 1s.

Number of addresses The number of addresses is the difference between the last and first address. It can be found using the formula 2^{32-n} .

Network Addresses

The first address of a block is treated as a special address. It is called the **network address**, and it defines the organization network to the rest of the world.

4.2 Network Address Translation

Network Address Translation (NAT) enables a user to have a large set of addresses internally and one address, or a small set of addresses, externally.

To separate the addresses used inside the home or business and the ones used of the Internet, the Internet authorities have reserved three sets of addresses as private addresses, as shown in table 4.3.

<i>First address</i>	<i>Last address</i>	<i>Total</i>
10.0.0.0	10.255.255.255	2^{24}
172.16.0.0	172.31.255.255	2^{20}
192.168.0.0	192.168.255.255	2^{16}

Table 4.3: Addresses for private networks

Figure 4.2 shows a simple implementation of NAT. As fig. 4.2 shows, the private network uses




Figure 4.2: A NAT implementation

private addresses. The router that connects the network to the global address uses one private address and one global address.

4.2.1 Address Translation

All the outgoing packets go through the NAT router, which replaces the *source address* in the packet with the global NAT address. All incoming packets also pass through the NAT router, which replaces the *destination address* in the packet with the appropriate private address. An example of address translation is shown in fig. 4.3.




Figure 4.3: Addresses in a NAT

4.2.2 Translation Table

A NAT can determine the destination address for a packet coming from the Internet using a **translation table**. The process is illustrated in fig. 4.4.

A translation table may be defined on one IP address or a pool of IP addresses, but we use IP addresses as well as port numbers to allow a many-to-many relationship between private-network hosts and external server programs. An example five-column translation table is shown in table 4.4.

<i>Private Address</i>	<i>Private Port</i>	<i>External Address</i>	<i>External Port</i>	<i>Transport Protocol</i>
172.18.3.1	1400	25.8.3.2	80	Transmission Control Protocol (TCP)
172.18.3.2	1401	25.8.3.2	80	
...	...	...	...	...

Table 4.4: Five-column translation table




Figure 4.4: NAT address translation

4.3 Internetworking

4.3.1 Need for Network Layer

4.3.2 Internet as a Connectionless Network

4.4 IPv4

The **Internet Protocol version 4 (IPv4)** is the delivery mechanism used by the TCP/IP protocols. Figure 4.5 shows the position of IPv4 in the suite.




Figure 4.5: Position of IPv4 in TCP/IP protocol suite

IPv4 is an unreliable and connectionless datagram protocol: a **best-effort** delivery service. The term *best-effort* means that IPv4 provides no error control or flow control.

IPv4 is also a connectionless protocol for a packet-switching network that uses the datagram approach.

4.4.1 Datagram

Packets in the IPv4 layer are called **datagrams**. The IPv4 datagram format is shown in fig. 4.6.




Figure 4.6: IPv4 datagram format

A datagram is a variable-length packet consisting of two parts: header and data. The header is 20 to 60 bytes in length and contains information essential to routing and delivery.

A brief description of each field is in order.

Version (VER) This 4-bit field defines the version of the IPv4 protocol. IPv6 may totally replace version 4 in the future. If the machine is using some other version of IPv4, the datagram is discarded rather than interpreted incorrectly.

Header length (HLEN) This 4-bit field defines the total length of the datagram header in 4-byte words. This field is needed because the length of the header is variable (20 to 60 bytes) due to the options field. The minimum allowed value for this field is 5, which corresponds to 20 bytes ($5 \times 4 = 20$) and a maximum of 15, which corresponds to 60 bytes ($15 \times 4 = 60$).

Services This 8-bit field was previously called **service type**, but is now called **differentiated services**. Both interpretations are shown in fig. 4.7.




Figure 4.7: Service type or differentiated services

1. Service Type

In this interpretation, the first 3 bits are called precedence bits. The next 4 bits are called **type of service (TOS) bits**, and the last bit is unused.

- (a) **Precedence:** It defines the priority of the datagram in issues such as congestion. The datagrams with lower precedence are discarded first.
- (b) **TOS bits:** At most only one bit in this 4-bit subfield can be a 1 in each datagram. The bit patterns and their interpretations are given in table 4.5.

<i>TOS Bits</i>	<i>Description</i>
0000	Normal (default)
0001	Minimize cost
0010	Maximize reliability
0100	Maximize throughput
1000	Minimize delay

Table 4.5: Types of service

Application programs can request a specific type of service. The defaults for some applications are shown in table 4.6.

<i>Protocol</i>	<i>TOS Bits</i>	<i>Description</i>
ICMP	0000	Normal
BOOTP	0000	Normal
NNTP	0001	Minimize cost
IGP	0010	Maximize reliability
SNMP	0010	Maximize reliability
TELNET	1000	Minimize delay
FTP (data)	0100	Maximize throughput
FTP (control)	1000	Minimize delay
TFTP	1000	Minimize delay
SMTP (command)	1000	Minimize delay
SMTP (data)	0100	Maximize throughput
DNS (UDP query)	1000	Minimize delay
DNS (TCP query)	0000	Normal
DNS (zone)	0100	Maximize throughput

Table 4.6: Default types of service

(c) **Differentiated Services**

In this interpretation, the first 6 bits make up the **codepoint** subfield, and the last 2 bits are unused.

- i. When the rightmost 3 bits are 0s, the 3 leftmost bits are interpreted the same as the precedence bits in the service type interpretation.
- ii. When the rightmost 3 bits are not all 0s the 6 bits define 64 services based on the priority assignment by the Internet or local authorities according to table 4.7.

<i>Category</i>	<i>Codepoint</i>	<i>Assigning Authority</i>
1	XXXXX0	Internet
2	XXXX11	Local
3	XXXX01	Temporary or experimental

Table 4.7: Values for codepoints

Total length This 16-bit field defines the total length (**header plus data**) of the IPv4 datagram in bytes. To find the length of the data coming from the upper layer, subtract the

header length from the total length.

$$\text{Length of data} = \text{total length} - \text{header length}$$

Identification This field is used in fragmentation (refer §4.4.2).

Flags This field is used in fragmentation (refer §4.4.2).

Fragmentation offset This field is used in fragmentation (refer §4.4.2).

Time to live This field denotes the maximum number of hops (routers) the datagram can visit.

Protocol This 8-bit field defines the higher-level protocol (e.g. TCP, UDP, ICMP and IGMP) that uses the services of the IPv4 layer. It specifies the final destination protocol to which the IPv4 datagram is delivered (fig. 4.8).




Figure 4.8: Protocol field and encapsulated data

The value of this field for each highre-level protocol is shown in table 4.8.

<i>Value</i>	<i>Protocol</i>
1	ICMP
2	IGMP
6	TCP
17	UDP
89	OSPFI

Table 4.8: Protocol values

Checksum Refer §4.4.3.

Source address This 32-bit field defines the IPv4 address of the source.

Destination address This 32-bit field defines the IPv4 address of the destination.

The source address and destination address fields **must remain unchanged** during the time the IPv4 datagram travels from the source host to the destination host.

4.4.2 Fragmentation

A datagram can travel through different networks. Each router decapsulates the IPv4 datagram from the frame it receives, processes it, and then encapsulates it in another frame. The format and size of the received frame depend on the protocol used by the physical network through which the frame has just traveled.

Maximum Transfer Unit

Each data link layer protocol has its own frame format in most protocols. One of the fields defined in the format is the maximum size of the data field. In other words, when a datagram is encapsulated in a frame, the total size of the datagram must be less than the maximum size, which is defined by the restrictions imposed by the hardware and software used in the network (see fig. 4.9).




Figure 4.9: Maximum transfer unit (MTU)

The maximum length of the IPv4 is equal to 65,535 bytes, which makes transmission more efficient if we use an MTU protocol of that size. However, for other physical networks, we must divide the datagram to make it pass through these networks. This is called **fragmentation**.

A datagram may be fragmented by the source host or any router in the path. The reassembly of the datagram is done only by the destination host.

When a datagram is fragmented, required parts of the header must be copied by all fragments.

Fields Related to Fragmentation

The fields that are related to fragmentation and reassembly of an IPv4 datagram are the identification, flags, and fragmentation offset fields.

Identification

Flags

Offset

4.4.3 Checksum

4.4.4 Options

4.5 Address Mapping

4.5.1 ARP

4.5.2 ICMP

4.6 Delivery

4.7 Forwarding

4.8 Routing

4.8.1 Unicast Routing Protocols

4.8.2 Multicast Routing Protocols

4.8.3 Routing Algorithms

Transport Layer

5.1 Process-to-Process Delivery

The data link layer is responsible for delivery of frames between two neighboring nodes over a link. This is called *node-to-node delivery*.

The network layer is responsible for delivery of datagrams between two hosts. This is called *host-to-host delivery*.

Real communication takes place between two processes (application programs). We need **process-to-process delivery**.

However, at any moment, several processes may be running on the source host and several on the destination host. To complete this delivery, we need a mechanism to deliver data from one of these processes running on the source host to the corresponding process running on the destination host.

The transport layer is responsible for process-to-process delivery: the delivery of a packet, part of a message, from one process to another. These three types of deliveries and their domains are shown in fig. 5.1.




Figure 5.1: Types of data deliveries

5.2 Addressing

Just like the MAC address at data link and the IP address at network layer, we have a transport layer address called a **port number**.

In the Internet model, the port numbers are 16-bit integers between 0 and 65,535. The client program defines itself with a port number, chosen randomly by the transport layer software running on the client host. This is the **ephemeral port number**.

The server cannot choose a port number randomly, though, as the client would not have a way to know the destination port without incurring extra overhead. Therefore, the Internet has decided to use universal port numbers for servers; these are called **well-known port numbers**.

5.2.1 IANA Ranges

The Internet Assigned Number Authority (IANA) has divided the port numbers into three ranges:

Well-known ports The ports ranging from 0 to 1023 are assigned and controlled by IANA. These are well-known ports.

Registered ports The ports ranging from 1024 to 49,151 are not assigned or controlled by IANA. They can only be registered with IANA to prevent duplication.

Dynamic ports The ports ranging from 49,152 to 65,535 are neither controlled nor registered. They can be used for any process. These are ephemeral ports.

5.2.2 Socket Addresses

Process-to-process needs two identifiers, IP address and the port number, at each end to make a connection. The combination of an IP address and a port number is called a **socket address**.

For example, the IP 200.23.56.8 and the port number 69 can be combined to form the socket address 200.23.56.8:69.

Transport Layer Protocols

The position of the protocols used in the transport layer is shown in fig. 5.2.

5.3 User Datagram Protocol

The **User Datagram Protocol (UDP)** is called a **connectionless, unreliable transport protocol**.

UDP is a very simple protocol using a minimum of overhead. If a process wants to send a small message and does not care about reliability, it can use UDP.

5.3.1 Well-Known Ports for UDP

Some well-known port numbers used by UDP are shown in table 5.1. Some port numbers can be used by both UDP and TCP.

In UNIX, the well-known ports are stored in a file called `/etc/services`. Each line in this file




Figure 5.2: Position of User Datagram Protocol (UDP) and TCP in TCP/IP suite

<i>Port</i>	<i>Protocol</i>	<i>Description</i>
7	Echo	Echoes a received datagram back to the sender
9	Discard	Discards any datagram that is received
11	Users	Active users
13	Daytime	Returns the date and the time
17	Quote	Returns a quote of the day
19	Chargen	Returns a string of characters
53	Nameserver	Domain Name Service
67	BOOTPs	Server port to download bootstrap information
68	BOOTPc	Client port to download bootstrap information
69	TFTP	Trivial File Transfer Protocol
111	RPC	Remote Procedure Call
123	NTP	Network Time Protocol
161	SNMP	Simple Network Management Protocol
162	SNMP	Simple Network Management Protocol (trap)

Table 5.1: Well-known ports used with UDP

gives the name of the server and the well-known port number. The line corresponding to the desired application can be extracted using the **grep** utility.

```
$ grep ftp /etc/services
ftp 21/tcp
ftp 21/udp
```

5.3.2 UDP Datagram

UDP packets, called **datagrams**, have a fixed-size header of 8 bytes. The format of a user datagram is shown in fig. 5.3.

The fields are as follows:

Source port number The port number used by the process running on the source




Figure 5.3: User datagram format

5.4 Transmission Control Protocol

Aw hell nah GATE Smashers

Chapter 6

Applications

6.1 Traditional Applications

6.1.1 World Wide Web

6.1.2 Hypertext Transfer Protocol

6.1.3 File Transfer Protocol

6.1.4 Email

6.1.5 Telnet and SSH

6.1.6 Domain Name Service

6.2 Peer-to-Peer Networks

6.3 Advances in the domain

domain?

Acronyms

BNC Bayonet Neill-Concelman 18

CIDR Classless Inter-Domain Routing 26, 27

CN Computer Networks 1

Coax Coaxial 17, 18

DLC Data Link Control 22

Drunkenstein Advait Desai 1

IANA Internet Assigned Number Authority 36

Internet Interconnected Network 11, 13, 14, 22, 24, 26–28, 31, 36, 40

IP Internet Protocol 24–33, 36

ISO International Standards Organization 11

ISP Internet Service Provider 11

Jojo Aaditya Joil 1

LAN Local Area Network 6–8, 18, 21

MAC Media Access Control 22, 36

MAN Metropolitan Area Network 7

MoDem Modulator-Demodulator 6

NAT Network Address Translation 27–29

OSI Open Systems Interconnection 11, 13, 14

RJ Registered Jack 17–19

Siouka Rupak R. Gupta 1

STP Shielded Twisted-Pair 17

TCP Transmission Control Protocol 28, 36, 37

TCP/IP Transmission Control Protocol/Internet Protocol 13, 14, 16, 22, 29, 37

UDP User Datagram Protocol 36, 37

UTP Unshielded Twisted-Pair 17

WAN Wide Area Network 7, 8

$$\left[\begin{array}{cc} \mathbb{R} & \Re \\ g & \infty \end{array} \right] \begin{pmatrix} 0 \\ \mathbb{J} \end{pmatrix}^2$$

The End.