Access Networks

Visual 23

Access networks and physical media

Q: How to connect an end system to an edge router?

- Residential access nets
- Institutional access networks (school, company)
- Mobile access networks

Keep in mind:

- Bandwidth (bits per second) of access network?
- Shared or dedicated?

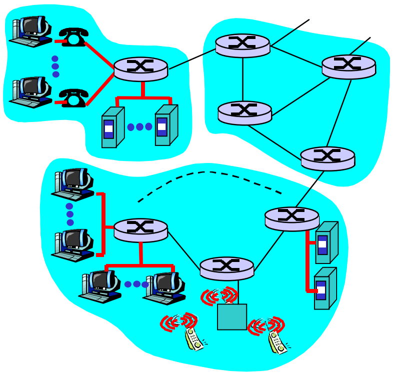

We have examined the roles of end systems and routers in network architecture. In this section we consider the access network – the physical link(s) that connect an end system to its edge router – that is, to the first router on a path from the end system to any other distant end system. Since access network technology is closely tied to physical media technology (fibre, coaxial pair, twisted-pair telephone wire, radio spectrum), we consider these two topics together in this section.

Access networks can be loosely divided into three categories:

- **Residential access networks**, connecting a home end system into the network.
- **Institutional access networks**, connecting an end system in a business or educational institution into the network.
- **Mobile access networks**, connecting a mobile end system into the network.

These categories are not hard and fast; some corporate end systems may well use the access network technology that we ascribe to residential access networks, and vice versa. The following descriptions are meant to hold for the common (if not every) case.

Residential Access Networks

A residential access network connects a home end system (typically a PC, but perhaps a Web TV or other residential system) to an edge router. Probably the most common form of home access is by use of a modem over a POTS (Plain Old Telephone System) dialup line to an Internet Service Provider (ISP). The home modem converts the digital output of the PC into analogue format for transmission over the analogue phone line. A modem in the ISP converts the analogue signal back into digital form for input to the ISP router. In this case, the access network is simply a point-to-point dialup link into an edge router. The point-to-point link is your ordinary twisted-pair phone line.

Visual 24

Residential access: point to point access

- Dialup via modem
  - up to 56 Kbps direct access to router (conceptually)
- ISDN: integrated services digital network: 128 Kbps all-digital connect to router
- ADSL: asymmetric digital subscriber line
  - up to 1 Mbps home-to-router
  - up to 8 Mbps router-to-home
  - ADSL deployment: happening

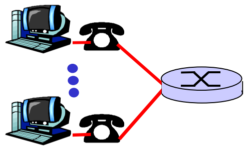

Today's modem speeds allow dialup access at rates up to 56 Kbps. However, due to the poor quality of twisted-pair line between many homes and ISPs, many users get an effective rate significantly less than 56 Kbps.

Whereas dialup modems require conversion of the end system's digital data into analogue form for transmission, so-called narrowband ISDN technology (Integrated Services Digital Network) allows for all-digital transmission of data from a home end system over ISDN 'telephone' lines to a phone company central office. Although ISDN was originally conceived as a way to carry digital data from one end of the phone system to another, it is also an important network access technology that provides higher speed access (for example, 128 Kbps) from the home into a data network such as the Internet. In this case, ISDN can simply be thought of as a 'better modem'.

Dialup modems and narrowband ISDN are already widely deployed technologies. Two new technologies, Asymmetric Digital Subscriber Line (ADSL) and Hybrid Fibre Coaxial cable (HFC) are currently being deployed.

ADSL is conceptually similar to dialup modems: It is a new modem technology again running over existing twisted-pair telephone lines, but it can transmit at rates of up to about 8 Mbps from the ISP router to a home end system. The data rate in the reverse direction, from the home end system to the central office router, is less than 1 Mbps. The asymmetry in the access speeds gives rise to the term asymmetric in ADSL. The asymmetry in the data rates reflects the belief that a home user is more likely to be a consumer of information (bringing data into the home) than a producer of information.

ADSL uses frequency division multiplexing, as described in the previous section. In particular, ADSL divides the communication link between the home and the ISP into three non-overlapping frequency bands.

Visual 25

Residential access: cable modems (1)

- HFC: hybrid fibre coax
  - asymmetric: up to 10Mbps upstream, 1 Mbps downstream
- Network of cable and fibre attaches homes to ISP router
  - shared access to router among homes
  - issues: congestion, dimensioning
- Deployment: available via cable companies, e.g., MediaOne

One of the features of ADSL is that the service allows the user to make an ordinary telephone call, using the POTS channel, while simultaneously surfing the Web. This feature is not available with standard dialup modems. The actual amount of downstream and upstream bandwidth available to the user is a function of the distance between the home modem and the ISP modem, the gauge of the twisted-pair line, and the degree of electrical interference. For a high-quality line with negligible electrical interference, an 8 Mbps downstream transmission rate is possible if the distance between the home and the ISP is less than 3,000 meters; the downstream transmission rate drops to about 2 Mbps for a distance of 6,000 meters. The upstream rate ranges from 16 Kbps to 1 Mbps.

ADSL, ISDN, and dialup modems all use ordinary phone lines, but HFC access networks are extensions of the current cable network used for broadcasting cable television. In a traditional cable system, a cable head end station broadcasts through a distribution of coaxial cable and amplifiers to residences, fibre optics (also to be discussed soon) connect the cable head end to neighbourhood-level junctions, from which traditional coaxial cable is then used to reach individual houses and apartments. Each neighbourhood juncture typically supports 500 to 5,000 homes.

Visual 26

Residential access: cable modems (2)

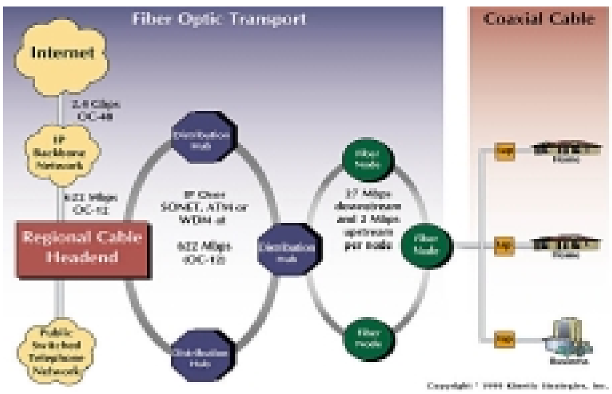

Diagram: http://www.cabledatacomnews.com/cmic/diagram.html

A Hybrid Fibre-Coax Access Network

As with ADSL, HFC requires special modems, called cable modems. Companies that provide cable Internet access require their customers to either purchase or lease a modem. One such company is Cyber Cable, which uses Motorola's Cyber Surfer Cable Modem and provides high-speed Internet access to most of the neighbourhoods in Paris. Typically, the cable modem is an external device and connects to the home PC through a 10-BaseT Ethernet port. Cable modems divide the HFC network into two channels, a downstream and an upstream channel. As with ADSL, the downstream channel is typically allocated more bandwidth and hence a larger transmission rate. For example, the downstream rate of the Cyber Cable system is 10 Mbps and the upstream rate is 768 Kbps. However, with HFC (and not with ADSL), these rates are shared among the homes, as we discuss next.

One important characteristic of HFC is that it is a shared broadcast medium. In particular, every packet sent by the head end travels downstream on every link to every home; and every packet sent by a home travels on the upstream channel to the head end. For this reason, if several users are receiving different Internet videos on the downstream channel, the actual rate at which each user receives its video will be significantly less than the downstream rate.

On the other hand, if all the active users are Web surfing, then each of the users may actually receive Web pages at the full downstream rate, as a small collection of users will rarely request a Web page at exactly the same time. Because the upstream channel is also shared, packets sent by two different homes at the same time will collide, which further decreases the effective upstream bandwidth.

Advocates of ADSL are quick to point out that ADSL is a point-to-point connection between the home and ISP, and therefore all the ADSL bandwidth is dedicated rather than shared. Cable advocates, however, argue that a reasonably dimensioned HFC network provides higher bandwidths than ADSL. The battle between ADSL and HFC for high speed residential access has clearly begun.

Company Access Networks

Visual 27

Institutional access: local area networks

- Company/univ local area network (LAN) connects end system to edge router
- Ethernet:
  - shared or dedicated cable connects end system and router
  - 10 Mbs, 100 Mbps, Gigabit Ethernet
- Deployment: institutions, home LANs happening now
- LANs: chapter 5

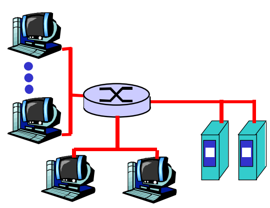

In company access networks, a local area network (LAN) is used to connect an end system to an edge router. As we will see in a later chapter, there are many different types of LAN technology. However, Ethernet technology is currently by far the most prevalent access technology in company networks.

Ethernet operates at 10 Mbps or 100 Mbps (and now even at 1 Gbps). It uses either twisted-pair copper wire or coaxial cable to connect a number of end systems with each other and with an edge router. The edge router is responsible for routing packets that have destinations outside of that LAN. Like HFC, Ethernet uses a shared medium, so that end users share the transmission rate of the LAN.

More recently, shared Ethernet technology has been migrating towards switched Ethernet technology. Switched Ethernet uses multiple coaxial cable or twisted-pair Ethernet segments connected at a 'switch' to allow the full bandwidth of an Ethernet to be delivered to different users on the same LAN simultaneously. We will explore shared and switched Ethernet in some detail in a later chapter.

Mobile Access Networks

Visual 28

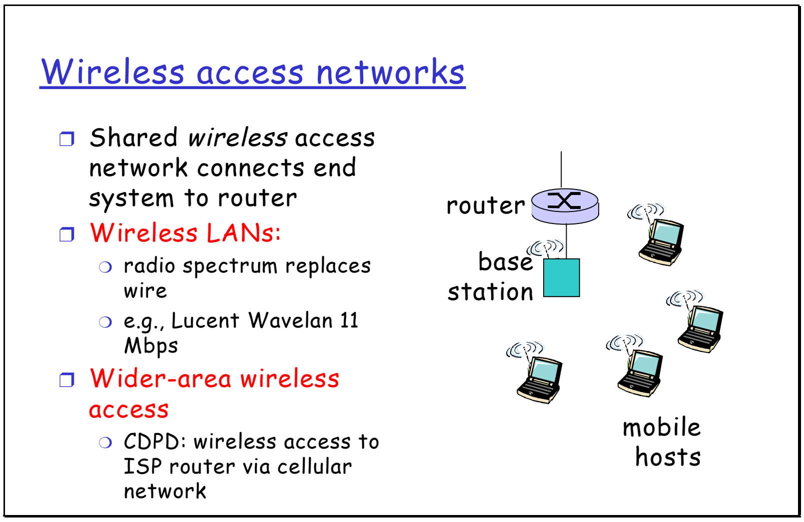

Mobile access networks use the radio spectrum to connect a mobile end system (for example, a laptop PC or a PDA with a wireless modem) to a base station. This base station, in turn, is connected to an edge router of a data network.

An emerging standard for wireless data networking is Cellular Digital Packet Data (CDPD). As the name suggests, a CDPD network operates as an overlay network (that is, as a separate, smaller virtual network, as a piece of the larger network) within the cellular telephone network. A CDPD network thus uses the same radio spectrum as the cellular phone system, and operates at speeds in the tens of Kbits per second. As with cable-based access networks and shared Ethernet, CDPD end systems must share the transmission media with other CDPD end systems within the cell covered by a base station.

A Media Access Control (MAC) protocol is used to arbitrate channel sharing among the CDPD end systems.

The CDPD system supports the IP protocol and thus allows an IP end system to exchange IP packets over the wireless channel with an IP base station. CDPD does not provide for any protocols above the network layer. From an Internet perspective, CDPD can be viewed as extending the Internet dial tone (that is, the ability to transfer IP packets) across a wireless link between a mobile end system and an Internet router.

Home Networks

Visual 29

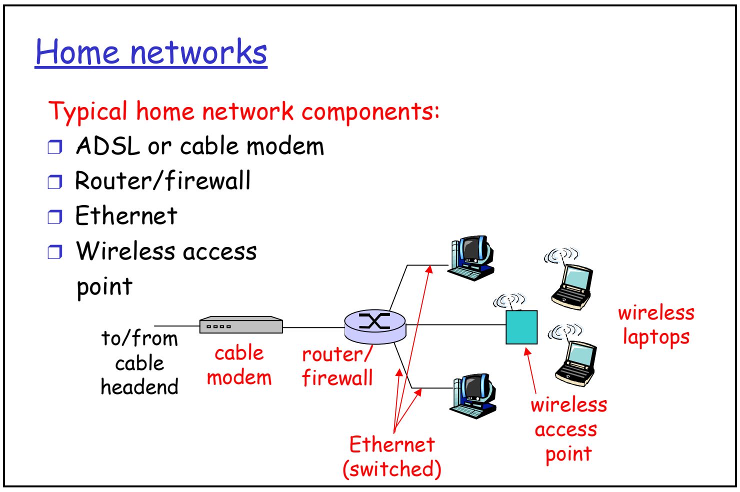

The above diagram shows typical home network components.

Physical Media

Visual 30

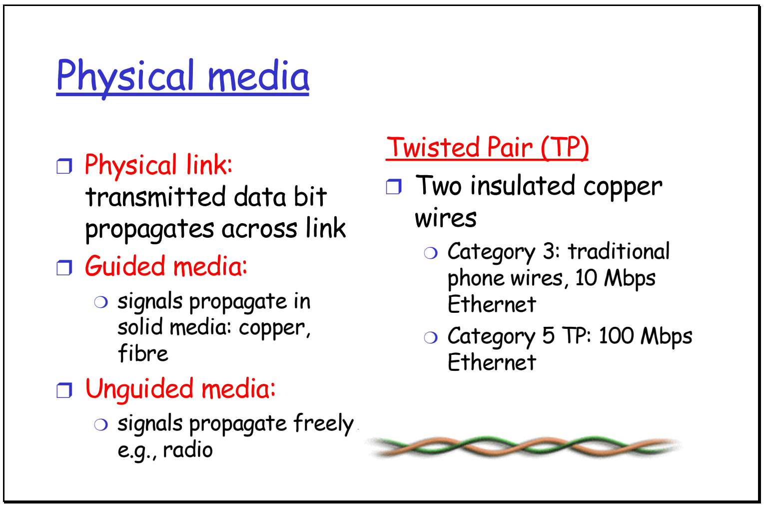

In the previous subsection, we gave an overview of some of the most important access network technologies in the Internet. As we described these technologies, we also indicated the physical media used. For example, we said that HFC uses a combination of fibre cable and coaxial cable. We said that ordinary modems, ISDN, and ADSL use twisted-pair copper wire. And we said that mobile access networks use the radio spectrum. In this subsection we provide a brief overview of these and other transmission media that are commonly employed in the Internet.

In order to define what is meant by a physical medium; let us reflect on the brief life of a bit. Consider a bit travelling from one end system, through a series of links and routers, to another end system. This poor bit gets transmitted many, many times! The source end system first transmits the bit, and shortly thereafter the first router in the series receives the bit; the first router then transmits the bit, and shortly afterwards the second router receives the bit, and so on.

Thus our bit, when travelling from source to destination, passes through a series of transmitter-receiver pairs. For each transmitter-receiver pair, the bit is sent by propagating electromagnetic waves or optical pulses across a physical medium. The physical medium can take many shapes and forms and does not have to be of the same type for each transmitter-receiver pair along the path.

Examples of physical media include twisted-pair copper wire, coaxial cable, multimode fibre-optic cable, terrestrial radio spectrum, and satellite radio spectrum.

Physical media fall into two categories: guided media and unguided media. With guided media, the waves are guided along a solid medium, such as a fibre-optic cable, a twisted-pair copper wire, or a coaxial cable. With unguided media, the waves propagate in the atmosphere and in outer space, such as in a digital satellite channel or in a CDPD system.

Some Popular Physical Media

Suppose you want to wire a building to allow computers to access the Internet or an intranet. Should you use twisted-pair copper wire, coaxial cable, or fibre optics? Which of these media gives the highest bit rates over the longest distances? We shall address these questions below.

However, before we get into the characteristics of the various guided medium types, let us say a few words about their costs. The actual cost of the physical link (copper wire, fibre-optic cable, and so on) is often relatively minor compared with the other networking costs. In particular, the labour cost associated with the installation of the physical link can be orders of magnitude higher than the cost of the material.

For this reason, many builders install twisted pair, optical fibre, and coaxial cable to every room in a building. Even if only one medium is initially used, there is a good chance that another medium could be used in the near future, and so money is saved by not having to lay additional wires.

Twisted-Pair Copper Wire

The least-expensive and most commonly used transmission medium is twisted-pair copper wire. For over one-hundred years it has been used by telephone networks. In fact, more than 99 percent of the wired connections from the telephone handset to the local telephone switch use twisted-pair copper wire. Most of us have seen twisted-pair in our homes and work environments. Twisted-pair consists of two insulated copper wires, each about 1 mm thick, arranged in a regular spiral pattern. The wires are twisted together to reduce the electrical interference from similar pairs close by. Typically, a number of pairs are bundled together in a cable by wrapping the pairs in a protective shield. A wire pair constitutes a single communication link.

Twisted Pair

Unshielded Twisted Pair (UTP) is commonly used for computer networks within a building, that is, for local area networks (LANs). Data rates for LANs using twisted pair today range from 10 Mbps to 100 Mbps. The data rates that can be achieved depend on the thickness of the wire and the distance between transmitter and receiver.

Two types of UTP are common in LANs – category 3 and category 5:

- Category 3 corresponds to voice-grade twisted pair, commonly found in office buildings. Office buildings are often pre-wired with two or more parallel pairs of category 3 twisted pair; one pair is used for telephone communication, and the additional pairs can be used for additional telephone lines or for LAN networking. 10 Mbps Ethernet, one of the most prevalent LAN types, can use category 3 UTP.
- Category 5, with its more twists per centimetre and TeflonTM insulation, can handle higher bit rates. 100 Mbps Ethernet running on category 5 UTP has become very popular in recent years. In recent years, category 5 UTP has become common for pre-installation in new office buildings.

When fibre-optic technology emerged in the 1980s, many people disparaged twisted pair because of its relatively low bit rates. Some people even felt that fibre-optic technology would completely replace twisted pair. But twisted pair did not give up so easily. Modern twisted-pair technology, such as category 5 UTP, can achieve data rates of 100 Mbps for distances up to a few hundred meters. Even higher rates are possible over shorter distances. In the end, twisted pair has emerged as the dominant solution for high-speed LAN networking.

As discussed in the section on access networks, twisted pair is also commonly used for residential Internet access. We saw that dialup modem technology enables access at rates of up to 56 Kbps over twisted pair. We also saw that ISDN is available in many communities, providing access rates of about 128 Kbps over twisted pair. We also saw that ADSL (asymmetric digital subscriber loop) technology has enabled residential users to access the Internet at rates in excess of 6 Mbps over twisted pair.

Visual 31

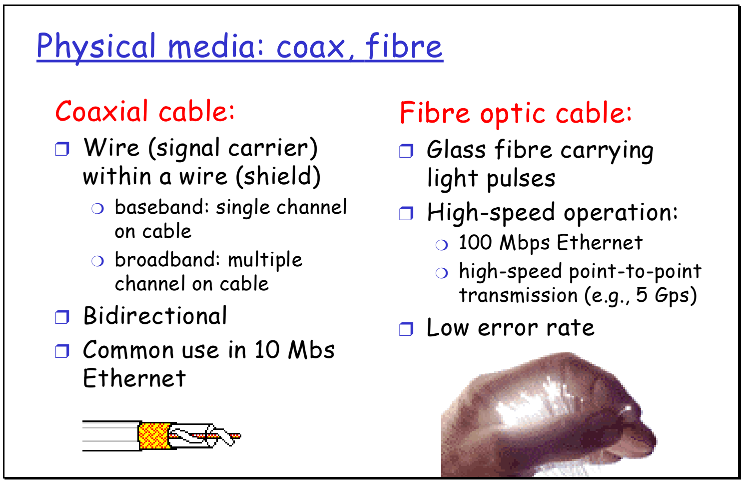

Coaxial Cable

Like twisted pair, coaxial cable consists of two copper conductors, but the two conductors are concentric rather than parallel. With this construction and a special insulation and shielding, coaxial cable can have higher bit rates than twisted pair. Coaxial cable comes in two varieties: baseband coaxial cable and broadband coaxial cable.

Baseband coaxial cable, also called 50-ohm cable, is about a centimetre thick, lightweight, and easy to bend. It is commonly used in LANs; in fact, the computer you use at work or at school is probably connected to a LAN with either baseband coaxial cable or with UTP.

Take a look at the connection to your computer's interface card:

- If you see a telephone-like jack and some wire that resembles telephone wire, you are using UTP.
- If you see a T-connector and a cable running out of both sides of the T-connector, you are using baseband coaxial cable.

The term baseband comes from the fact that the stream of bits is dumped directly into the cable, without shifting the signal to a different frequency band; 10 Mbps Ethernets can use either UTP or baseband coaxial cable.

Broadband Coaxial Cable

Broadband coaxial cable, also called 75-ohm cable, is quite a bit thicker, heavier, and stiffer than the baseband variety. It was once commonly used in LANs and can still be found in some older installations. For LANs, baseband cable is now preferable since it is less expensive, easier to physically handle, and does not require attachment cables.

Broadband cable, however, is quite common in cable television systems. As we saw earlier, cable television systems have recently been coupled with cable modems to provide residential users with Web access at rates of 10 Mbps or higher. With broadband coaxial cable, the transmitter shifts the digital signal to a specific frequency band, and the resulting analogue signal is sent from the transmitter to one or more receivers.

Both baseband and broadband coaxial cable can be used as a guided shared medium. Specifically, a number of end systems can be connected directly to the cable, and all the end systems receive whatever any one of the computers transmits.

Fibre Optics

An optical fibre is a thin, flexible medium that conducts pulses of light, with each pulse representing a bit. A single optical fibre can support tremendous bit rates, up to tens or even hundreds of gigabits per second. They are immune to electromagnetic interference, have very low signal attenuation up to 100 kilometres, and are very hard to tap. These characteristics have made fibre optics the preferred long-haul guided transmission media, particularly for overseas links.

Many of the long-distance telephone networks in the United States and elsewhere now use fibre optics exclusively. Fibre optics is also prevalent in the backbone of the Internet. However, the high cost of optical devices – such as transmitters, receivers, and switches – has hindered their deployment for short-haul transport, such as in a LAN or into the home in a residential access network.

Terrestrial and Satellite Radio Channels

Visual 32

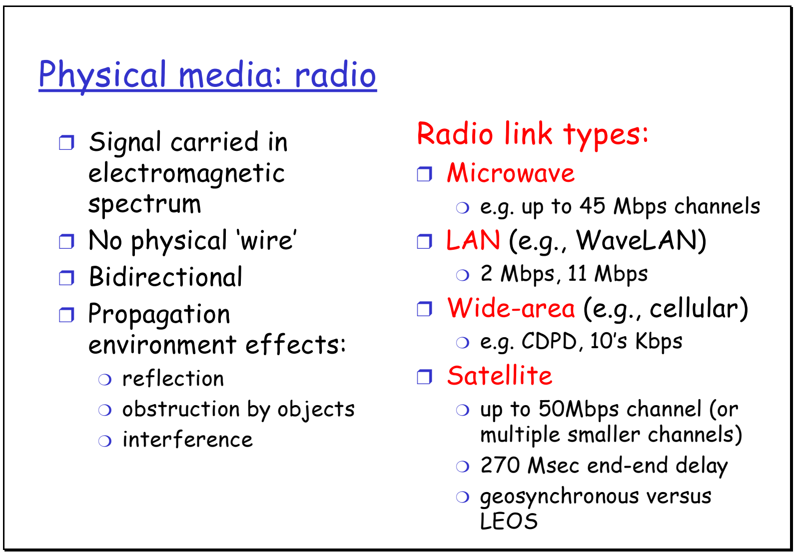

Radio channels carry signals in the electromagnetic spectrum. They are an attractive media because they require no physical wire to be installed, can penetrate walls, provide connectivity to a mobile user, and can potentially carry a signal for long distances. The characteristics of a radio channel depend significantly on the propagation environment and the distance over which a signal is to be carried. Environmental considerations determine path loss and shadow fading (which decrease in signal strength as the signal travels over a distance, and around/through obstructing objects), multipath fading (due to signal reflection off interfering objects), and interference (due to other radio channels or electromagnetic signals).

Terrestrial radio channels can be broadly classified into two groups: those that operate as local area networks (typically spanning from ten to a few hundred meters) and wide-area radio channels that are used for mobile data services (typically operating within a metropolitan region). A number of wireless LAN products are on the market, operating in the range of one to tens of Mbps. Mobile data services typically provide channels that operate at tens of Kbps.

A communication satellite links two or more Earth-based microwave transmitter/receivers, known as ground stations. The satellite receives transmissions on one frequency band, regenerates the signal using a repeater and transmits the signal on another frequency. Satellites can provide bandwidths in the gigabit per second range.

Two types of satellites are used in communications: geostationary satellites and low-altitude satellites.

- **Geostationary satellites** permanently remain above the same spot on Earth. This stationary presence is achieved by placing the satellite in orbit at 36,000 kilometres above Earth’s surface. This huge distance from ground station through satellite back to ground station introduces a substantial signal propagation delay of 250 milliseconds. Nevertheless, satellite links, which can operate at speeds of hundreds of Mbps, are often used in telephone networks and in the backbone of the Internet.
- **Low-altitude satellites** are placed much closer to Earth and do not remain permanently above one spot on Earth. They rotate around Earth just as the Moon does. To provide continuous coverage to an area, many satellites need to be placed in orbit. There are currently many low-altitude communication systems in development. The Iridium system, for example, consists of 66 low-altitude satellites. Lloyd’s Satellite Constellation Systems’ Web page provides and collects information on Iridium as well as other satellite constellation systems. The low-altitude satellite technology may be used for Internet access sometime in the future.

Delay and Loss in Packet-Switched Networks

Having briefly considered the major pieces of the Internet architecture – the applications, end systems, end-to-end transport protocols, routers, and links – let us now consider what can happen to a packet as it travels from its source to its destination.

Recall that a packet starts in a host (the source), passes through a series of routers, and ends its journey in another host (the destination). As a packet travels from one node (host or router) to the subsequent node (host or router) along this path, the packet suffers from several different types of delays at each node along the path. The most important of these delays are the nodal processing delay, queuing delay, transmission delay, and propagation delay. Together, these delays accumulate to give a total nodal delay.

In order to acquire a deep understanding of packet switching and computer networks, we must understand the nature and importance of these delays.

Visual 33

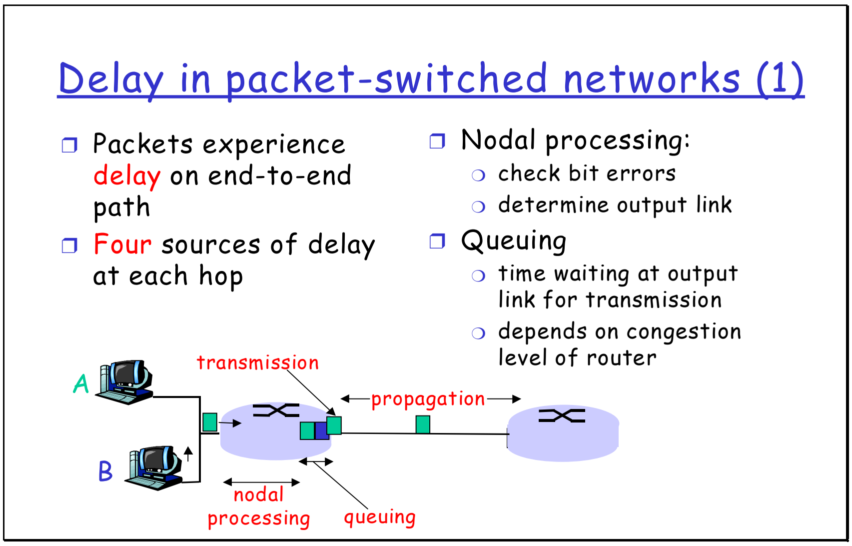

Types of Delay

- **Processing delay** – the time required to examine the packet's header and determine where to direct the packet is part of the processing delay. The processing delay can also include other factors, such as the time needed to check for bit-level errors in the packet that occurred in transmitting the packet's bits from the upstream router to router A. Processing delays in high-speed routers are typically on the order of microseconds or less. After this nodal processing, the router directs the packet to the queue that precedes the link to router B.
- **Queuing delay** – at the queue, the packet experiences a queuing delay as it waits to be transmitted onto the link. The queuing delay of a specific packet will depend on the number of other, earlier-arriving packets that are queued and waiting for transmission across the link. The delay of a given packet can vary significantly from packet to packet. If the queue is empty and no other packet is currently being transmitted, then our packet's queuing delay is zero. On the other hand, if the traffic is heavy and many other packets are also

waiting to be transmitted, the queuing delay will be long. We will see shortly that the number of packets that an arriving packet might expect to find on arrival is a function of the intensity and nature of the traffic arriving to the queue. Queuing delays can be on the order of milliseconds to microseconds in practice.

- **Transmission delay** – assuming that packets are transmitted in first-come-first-serve manner, as is common in the Internet, our packet can be transmitted once all the packets that have arrived before it have been transmitted. Denote the length of the packet by L bits, and denote the transmission rate of the link from router A to router B by R bits/sec. The rate R is determined by transmission rate of the link to router B. For example, for a 10-Mbps Ethernet link, the rate is R = 10 Mbps; for a 100-Mbps Ethernet link, the rate is R = 100 Mbps. The transmission delay is also known as the store-and-forward delay. This is the amount of time required to transmit all of the packet's bits into the link. Transmission delays are typically on the order of microseconds or less in practice.

Visual 34

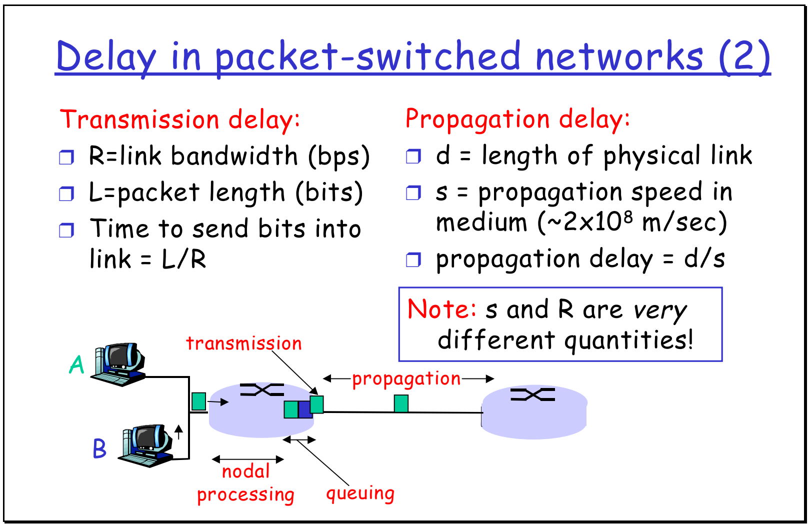

- **Propagation delay** – once a bit is pushed onto the link, it needs to propagate to router B. The time required to propagate from the beginning of the link to router B is the propagation delay. The bit propagates at the propagation speed of the link. The propagation speed depends on the physical medium of the link (that is, multimode fibre, twisted-pair copper wire, and so on) and is in the range of:

$$2 * 10^8 \text{ metres/sec to } 3 * 10^8 \text{ metres/sec}$$

which is equal to, or a little less than, the speed of light. The propagation delay is the distance between two routers divided by the propagation speed. That is, the propagation delay is d/s, where d is the distance between router A and router B and s is the propagation speed of the link. Once the last bit of the packet propagates to node B, it and all the preceding bits of the packet are stored in router B. The whole process then continues with router B now performing the forwarding. In wide-area networks, propagation delays are on the order of milliseconds.

Comparing Transmission and Propagation Delay

Newcomers to the field of computer networking sometimes have difficulty understanding the difference between transmission delay and propagation delay. The difference is subtle but important. The transmission delay is the amount of time required for the router to push out the packet; it is a function of the packet's length and the transmission rate of the link, but has nothing to do with the distance between the two routers. The propagation delay, on the other hand, is the time it takes a bit to propagate from one router to the next; it is a function of the distance between the two routers, but has nothing to do with the packet's length or the transmission rate of the link.

Queuing Delay (Revisited)

Visual 35

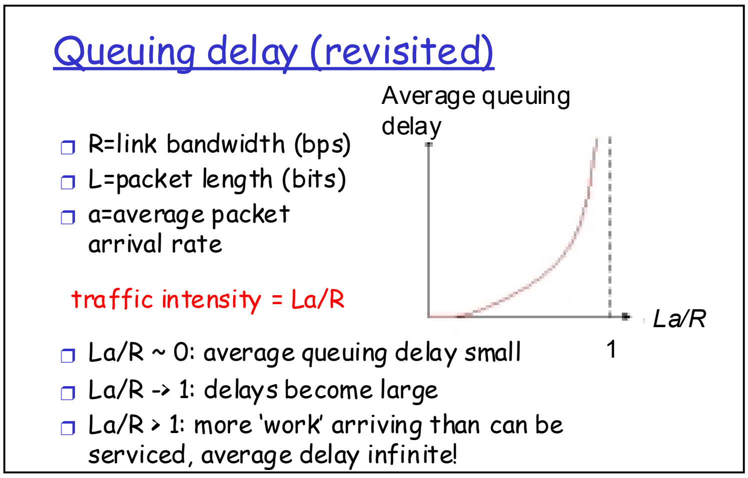

Real Internet Delays and Routes

Visual 36

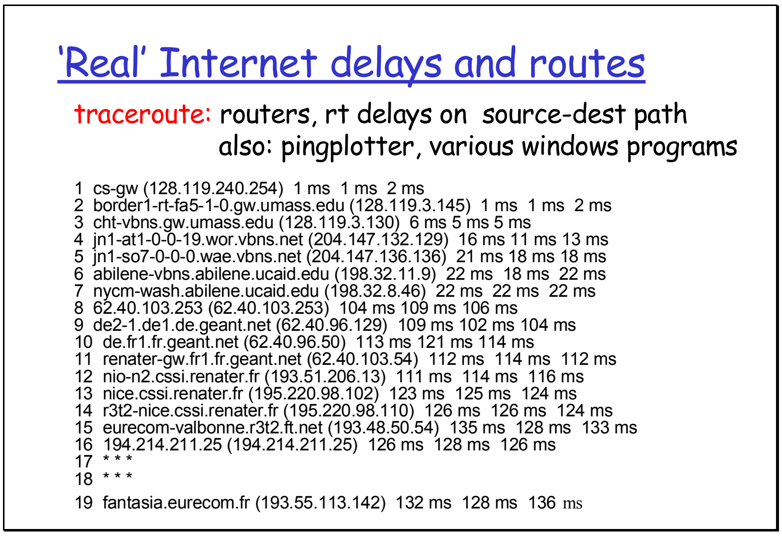

- **Exercise** – Try traceroute command if you are using Linux or tracert if you are using windows and note the result.

Syntax: traceroute <domainname)

Example: traceroute nccedu.com

Protocol Layers and their Service Models

Visual 37

Protocol 'Layers'

Networks are complex!

- Many 'pieces':
  - hosts
  - routers
  - links of various media
  - applications
  - protocols
  - hardware, software

Question:

Is there any hope of *organizing* the structure of a network?

Or at least our discussion of networks?

From our discussion so far, it is apparent that the Internet is an extremely complicated system. We have seen that there are many pieces to the Internet: numerous applications and protocols, various types of end systems and connections between end systems, routers, and various types of link-level media. Given this enormous complexity, is there any hope of organizing network architecture, or at least our discussion of network architecture? Fortunately, the answer to both questions is yes.

To reduce design complexity, network designers organize protocols – and the network hardware and software that implement the protocols – in layers. With a layered protocol architecture, each protocol belongs to one of the layers. It is important to realize that a protocol in layer n is distributed among the network entities (including end systems and packet switches) that implement that protocol, just as the functions in our layered airline architecture were distributed between the departing and arriving airports. In other words, there's a piece of layer n in each of the network entities. These pieces communicate with each other by exchanging layer-n messages. These messages are called layer-n protocol data units, or more commonly n-PDUs. The contents and format of an n-PDU, as well as the manner in which the n-PDUs are exchanged among the network elements, are defined by a layer-n protocol. When taken together, the protocols of the various layers are called the protocol stack.

Interestingly enough, this notion of relying on lower-layer services is prevalent in many other forms of communication. For example, consider ordinary postal mail. When you write a letter, you include envelope information such as the destination address and the return address with the letter. The letter, along with the address information, can be considered a PDU at the highest layer of the protocol stack. You then drop the PDU in a mailbox. At this point, the letter is out of your hands. The postal service may then add some of its own internal information onto your letter, essentially adding a header to your letter. For example, in the United States a barcode is often printed on your letter.

Once you drop your envelope into a mailbox, you rely on the postal service to deliver the letter to the correct destination in a timely manner. For example, you don't worry about whether a postal truck will break down while carrying the letter. Instead the postal service takes care of this, presumably with well-defined plans to recover from such failures. Furthermore, within the postal service itself there are layers, and the protocols at one layer rely on and use the services of the layer below.

In order for one layer to interoperate with the layer below it, the interfaces between the two layers must be precisely defined. Standards bodies define precisely the interfaces between adjacent layers (for example, the format of the PDUs passed between the layers) and permit the developers of networking software and hardware to implement the interior of the layers as they please. Therefore, if a new and improved implementation of a layer is released, the new implementation can replace the old implementation and, in theory, the layers will continue to interoperate.

Layer Functions

In a computer network, each layer may perform one or more of the following generic set of tasks:

- Error control, which makes the logical channel between the layers in two peer network elements more reliable.
- Flow control, which avoids overwhelming a slower peer with PDUs.
- Segmentation and reassembly, which at the transmitting side divides large data chunks into smaller pieces and at the receiving side reassembles the smaller pieces into the original large chunk.
- Multiplexing, which allows several higher-level sessions to share a single lower-level connection.

The Internet Protocol Stack, and Protocol Data Units

A protocol layer can be implemented in software, in hardware, or using a combination of the two. Application-layer protocols – such as HTTP and SMTP – are almost always implemented in software in the end systems; so are transport-layer protocols. Because the physical layer and data link layers are responsible for handling communication over a specific link, they are typically implemented in a network interface card (for example, Ethernet or ATM interface cards) associated with a given link. The network layer is often a mixed implementation of hardware and software. We now summarize the Internet layers and the services they provide:

Application Layer

The application layer is responsible for supporting network applications. The application layer includes many protocols, including HTTP to support the Web, SMTP to support electronic mail, and FTP to support file transfer. We shall see in Chapter 2 that it is very easy to create our own new application-layer protocols.

Transport Layer

The transport layer provides the service of transporting application-layer messages between the client and server sides of an application. In the Internet there are two transport protocols, TCP and UDP, either of which can transport application-layer messages. TCP provides a connection-oriented service to its applications. This service includes guaranteed delivery of application-layer messages to the destination and flow control (that is, sender/receiver speed matching). TCP also segments long messages into shorter segments and provides a congestion control mechanism, so that a source throttles its transmission rate when the network is congested. The UDP protocol provides its applications a connectionless service, which (as we saw in Section 1.3) is very much a no-frills service.

Network Layer

The network layer is responsible for routing datagrams from one host to another. The Internet's network layer has two principle components. It has a protocol that defines the fields in the IP datagram as well as how the end systems and routers act on these fields. This protocol is the celebrated IP protocol. There is only one IP protocol, and all Internet components that have a network layer must run the IP protocol. The Internet's network layer also contains routing protocols that determine the routes that datagrams take between sources and destinations. The Internet has many routing protocols. As we saw in Section 1.4, the Internet is a network of networks, and within a network, the network administrator can run any routing protocol desired. Although the network layer contains both the IP protocol and numerous routing protocols, it is often simply referred to as the IP layer, reflecting the fact that IP is the glue that binds the Internet together.

The Internet transport layer protocols (TCP and UDP) in a source host passes a transport-layer segment and a destination address to the IP layer, just as you give the postal service a letter with a destination address. The IP layer then provides the service of routing the segment to its destination. When the packet arrives at the destination, IP passes the segment to the transport layer within the destination.

Link Layer

The network layer routes a packet through a series of packet switches (called routers, in the Internet) between the source and destination. To move a packet from one node (host or packet switch) to the next node in the route, the network layer must rely on the services of the link layer. In particular, at each node IP passes the datagram to the link layer, which delivers the datagram to the next node along the route. At this next node, the link layer passes the IP datagram to the network layer. The process is analogous to the postal worker at a mailing centre who puts a letter into a plane that will deliver the letter to the next postal centre along the route. The services provided at the link layer depend on the specific link-layer protocol that is employed over the link. For example, some protocols provide reliable delivery on a link basis, that is, from transmitting node, over one link, to receiving node. Note that this reliable delivery service is different from the reliable delivery service of TCP, which provides reliable delivery from one end system to another. Examples of link layers include Ethernet and Point-to-Point Protocol (PPP); in some contexts, ATM and frame relay can be considered link layers. As datagrams typically need to traverse several links to travel from source to destination, a datagram may be handled by different link-layer protocols at different links along its route. For example, a datagram may be handled by Ethernet on one link and then PPP on the next link. IP will receive a different service from each of the different link-layer protocols.

Physical Layer

While the job of the link layer is to move entire frames from one network element to an adjacent network element, the job of the physical layer is to move the individual bits within the frame from one node to the next. The protocols in this layer are again link dependent, and further depend on the actual transmission medium of the link (for example, twisted-pair copper wire, single-mode fibre optics). For example, Ethernet has many physical layer protocols: one for twisted-pair copper wire, another for coaxial cable, another for fibre, and so on. In each case, a bit is moved across the link in a different way.

Internet History

Development and Demonstration of Early Packet Switching Principles: 1961-1972

Visual 38

Internet history

1961-1972: Early packet-switching principles

- 1961: Kleinrock - queueing theory shows effectiveness of packet-switching
- 1964: Baran - packet-switching in military nets
- 1967: ARPAnet conceived by Advanced Research Projects Agency
- 1969: first ARPAnet node operational
- 1972:
  - ARPAnet demonstrated publicly
  - NCP (Network Control Protocol) first host-host protocol
  - First e-mail program
  - ARPAnet has 15 nodes

The fields of computer networking and today's Internet trace their beginnings back to the early 1960s, a time at which the telephone network was the world's dominant communication network. Given the increasing importance (and great expense) of computers in the early 1960s and the advent of timeshared computers, it was perhaps natural (at least with perfect hindsight!) to consider the question of how to hook computers together so that they could be shared among geographically distributed users. The traffic generated by such users was likely to be 'bursty' – intervals of activity, such as the sending of a command to a remote computer, followed by periods of inactivity while waiting for a reply or while contemplating the received response.

Three research groups around the world, all unaware of the others' work, began inventing the notion of packet switching as an efficient and robust alternative to circuit switching. The first published work on packet-switching techniques was that of Leonard Kleinrock, at that time a graduate student at MIT. Using queuing theory, Kleinrock's work elegantly demonstrated the effectiveness of the packet-switching approach for bursty traffic sources. In 1964, Paul Baran at the Rand Institute had begun investigating the use of packet switching for secure voice over military networks, and at the National Physical Laboratory in England, Donald Davies and Roger Scantlebury were also developing their ideas on packet switching.

The work at MIT, Rand, and NPL laid the foundations for today's Internet. But the Internet also has a long history of a let's-build-it-and-demonstrate-it attitude that also dates back to the early 1960s. J C R Licklider and Lawrence Roberts, both colleagues of Kleinrock's at MIT, went on to lead the computer science program at the Advanced Research Projects Agency (ARPA) in the United States. Roberts published an overall plan for the so-called ARPAnet, the first packet-switched computer network and a direct ancestor of today's public Internet.

The early packet switches were known as interface message processors (IMPs) and the contract to build these switches was awarded to the BBN company. On Labour Day in 1969, the first IMP was installed at UCLA under Kleinrock's supervision, with three additional IMPs being installed shortly thereafter at the Stanford Research Institute (SRI), UC Santa Barbara, and the University of Utah (Figure 1.28). The fledgling precursor to the

Internet was four nodes large by the end of 1969. Kleinrock recalls the very first use of the network to perform a remote login from UCLA to SRI, crashing the system.

By 1972, ARPAnet had grown to approximately 15 nodes, and was given its first public demonstration by Robert Kahn at the 1972 International Conference on Computer Communications. The first host-to-host protocol between ARPAnet end systems known as the network-control protocol (NCP) was completed [RFC 001]. With an end-to-end protocol available, applications could now be written. The first e-mail program was written by Ray Tomlinson at BBN in 1972.

Internetworking, and New and Proprietary Networks: 1972-1980

Visual 39

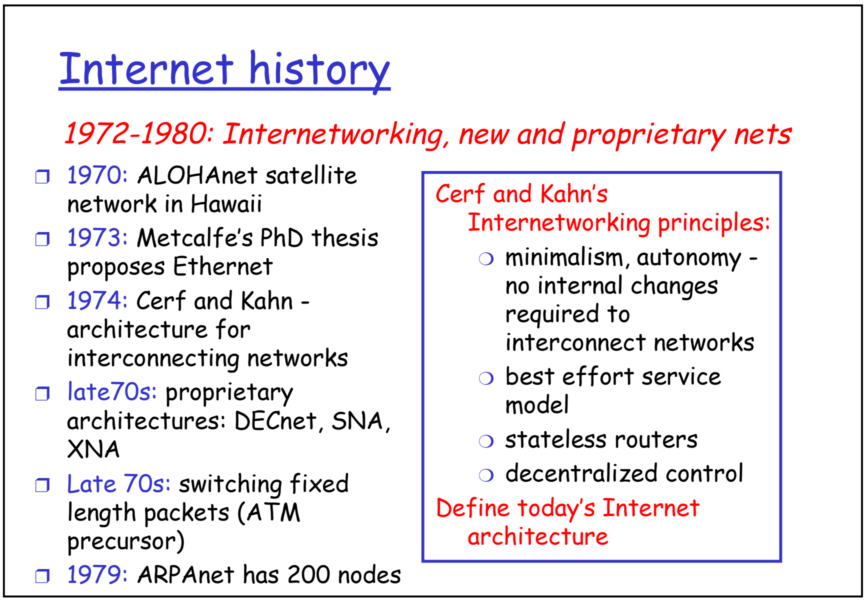

The initial ARPAnet was a single, closed network. In order to communicate with an ARPAnet host, one had to actually be attached to another ARPAnet IMP. In the early to mid 1970s, additional packet-switching networks besides ARPAnet came into being:

- ALOHAnet, a microwave network linking together universities on the Hawaiian islands;
- Telenet, a BBN commercial packet-switching network based on ARPAnet technology;
- Tymnet;
- Transpac, a French packet-switching network.

The number of networks was beginning to grow. In 1973, Robert Metcalfe's PhD thesis laid out the principle of Ethernet, which would later lead to a huge growth in so-called local area networks (LANs) that operated over a small distance based on the Ethernet protocol.

Once again, with perfect hindsight one might now see that the time was ripe for developing an encompassing architecture for connecting networks together. Pioneering work on interconnecting networks, once again under the sponsorship of DARPA (Defense Advanced Research Projects Agency) – in essence creating a network of networks – was done by Vinton Cerf and Robert Kahn. The term ‘internetting’ was coined to describe this work.

These architectural principles were embodied in the TCP protocol. The early versions of TCP, however, were quite different from today’s TCPs. The early versions of TCP combined a reliable in-sequence delivery of data via end-system retransmission (still part of today’s TCP) with forwarding functions (which today are performed by IP). Early experimentation with TCP, combined with the recognition of the importance of an unreliable, non-flow-controlled end-to-end transport service for applications such as packetised voice, led to the separation of IP out of TCP and the development of the UDP protocol. The three key Internet protocols that we see today – TCP, UDP, and IP – were conceptually in place by the end of the 1970s.

In addition to the DARPA Internet-related research, many other important networking activities were underway. In Hawaii, Norman Abramson was developing ALOHAnet, a packet-based radio network that allowed multiple remote sites on the Hawaiian Islands to communicate with each other.

The ALOHA protocol was the first so-called multiple-access protocol, allowing geographically distributed users to share a single broadcast communication medium (a radio frequency). Abramson’s work on multiple-access protocols was built upon by Metcalfe and Boggs in the development of the Ethernet protocol for wire-based shared broadcast networks; interestingly, Metcalfe and Boggs’ Ethernet protocol was motivated by the need to connect multiple PCs, printers, and shared disks together.

Twenty-five years ago, well before the PC revolution and the explosion of networks, Metcalfe and Boggs were laying the foundation for today’s PC LANs. Ethernet technology represented an important step for internetworking as well. Each Ethernet local area network was itself a network, and as the number of LANs proliferated, the need to internetwork these LANs together became increasingly important.

Metcalfe’s Original Conception of the Ethernet

In addition to the DARPA internetworking efforts and the Aloha/Ethernet multiple-access networks, a number of companies were developing their own proprietary network architectures. Digital Equipment Corporation (Digital) released the first version of the DECnet in 1975, allowing two PDP-11 minicomputers to communicate with each other. DECnet has continued to evolve since then, with significant portions of the OSI protocol suite being based on ideas pioneered in DECnet. Other important players during the 1970s were Xerox (with the XNS architecture) and IBM (with the SNA architecture). Each of these early networking efforts would contribute to the knowledge base that would drive networking in the 80s and 90s.

It is important to note here that in the 1980s (and even before), researchers such as [Fraser 1983, 1993] and [Turner 1986] were also developing a competitor technology to the Internet architecture. These efforts have contributed to the development of the ATM architecture, a connection-oriented approach based on the use of fixed-size packets, known as cells.

A Proliferation of Networks: 1980-1990

Visual 40

Internet history

1980-1990: new protocols, a proliferation of networks

- 1983: deployment of TCP/IP
- 1982: smtp e-mail protocol defined
- 1983: DNS defined for name-to-IP-address translation
- 1985: ftp protocol defined
- 1988: TCP congestion control
- New national networks: Csnet, BITnet, NSFnet, Minitel
- 100,000 hosts connected to confederation of networks

By the end of the 1970s, approximately 200 hosts were connected to the ARPAnet. By the end of the 1980s the number of hosts connected to the public Internet, a confederation of networks looking much like today's Internet, would reach 100,000. The 1980s would be a time of tremendous growth.

Much of the growth in the early 1980s resulted from several distinct efforts to create computer networks linking universities together. BITnet (because it's their network) provided e-mail and file transfers among several universities in the Northeast. CSNET (computer science network) was formed to link together university researchers without access to ARPAnet. In 1986, NSFNET was created to provide access to NSF-sponsored supercomputing centres. Starting with an initial backbone speed of 56 Kbps, NSFNET's backbone would be running at 1.5 Mbps by the end of the decade, and would be serving as a primary backbone linking together regional networks.

In the ARPAnet community, many of the final pieces of today's Internet architecture were falling into place. January 1, 1983, saw the official deployment of TCP/IP as the new standard host protocol for ARPAnet (replacing the NCP protocol). The transition [RFC 801] from NCP to TCP/IP was a 'flag day' type event – all hosts were required to transfer over to TCP/IP as of that day. In the late 1980s, important extensions were made to TCP to implement host-based congestion control. The Domain Name System, used to map between a human-readable Internet name (for example, gaia.cs.umass.edu) and its 32-bit IP address, was also developed [RFC 1034].

In parallel with this development of the ARPAnet (which was for the most part a United States effort), in the early 1980s the French launched the Minitel project, an ambitious plan to bring data networking into everyone's home. Sponsored by the French government, the Minitel system consisted of a public packet-switched network (based on the X.25 protocol suite, which uses virtual circuits), Minitel servers, and inexpensive terminals with built-in low speed modems.

The Minitel became a huge success in 1984 when the French government gave away a free Minitel terminal to each French household that wanted one. Minitel sites included free sites – such as a telephone directory site – as well as private sites, which collected a usage-based fee from each user. At its peak in the mid 1990s, it offered more than 20,000 different services, ranging from home banking to specialized research databases. It was used by over 20% of France's population, generated more than $1 billion each year, and created 10,000 jobs.

The Minitel was in a large proportion of French homes 10 years before most Americans had ever heard of the Internet. It still enjoys widespread use in France, but is increasingly facing stiff competition from the Internet.

Commercialization and the Web: the 1990s

The 1990s were ushered in with two events that symbolized the continued evolution and the soon-to-arrive commercialization of the Internet. First, ARPAnet, the progenitor of the Internet ceased to exist. MILNET and the Defence Data Network had grown in the 1980s to carry most of the US Department-of-Defence-related traffic and NSFnet had begun to serve as a backbone network connecting regional networks in the United States and national networks overseas. In 1991, NSFNET lifted its restrictions on use of NSFNET for commercial purposes. NSFNET itself would be decommissioned in 1995, with Internet backbone traffic being carried by commercial Internet service providers.

Visual 41

Internet history

1990s: commercialization, the WWW

- Early 1990s: ARPAnet decommissioned
- 1991: NSF lifts restrictions on commercial use of NSFnet (decommissioned, 1995)
- Early 1990s: WWW
  - hypertext [Bush 1945, Nelson 1960s]
  - HTML, http: Berners-Lee
  - 1994: Mosaic, later Netscape
  - late 1990s: commercialization of the WWW

Late 1990s:

- Estimated 50 million computers on Internet
- Estimated 100 million+ users
- Backbone links running at 1 Gbps

The main event of the 1990s, however, was to be the release of the World Wide Web, which brought the Internet into the homes and businesses of millions and millions of people worldwide. The Web also served as a platform for enabling and deploying hundreds of new applications, including online stock trading and banking, streamed multimedia services, and information retrieval services.

The Web was invented at CERN by Tim Berners-Lee in 1989-1991, based on ideas originating in earlier work on hypertext from the 1940s by Bush and since the 1960s by Ted Nelson. Berners-Lee and his associates developed initial versions of HTML, HTTP, a Web server, and a browser – the four key components of the Web. The original CERN browsers only provided a line-mode interface.

Around the end of 1992 there were about 200 Web servers in operation, this collection of servers being the tip of the iceberg for what was about to come. At about this time several researchers were developing Web browsers with GUI interfaces, including Marc Andreesen, who led the development of the popular GUI browser Mosaic for X. Andreesen and his colleagues released an alpha version of his browser in 1993, and in 1994 he and James Baker formed Mosaic Communications, which later became Netscape Communications Corporation. By 1995, university students were using Mosaic and Netscape browsers to surf the Web on a daily basis. At about this time companies – big and small – began to operate Web servers and transact commerce over the Web. In 1996, Microsoft got into the Web business in a big way though lately realising the huge business potential of the Internet.

During the 1990s, networking research and development also made significant advance. The technical community struggled with the problems of defining and implementing an Internet service model for traffic requiring real-time constraints, such as continuous media applications. The need to secure and manage Internet infrastructure also became of paramount importance as e-commerce applications proliferated and the Internet became a central component of the world's telecommunications infrastructure.

Summary

Visual 42