Network
Engineering and Architecture

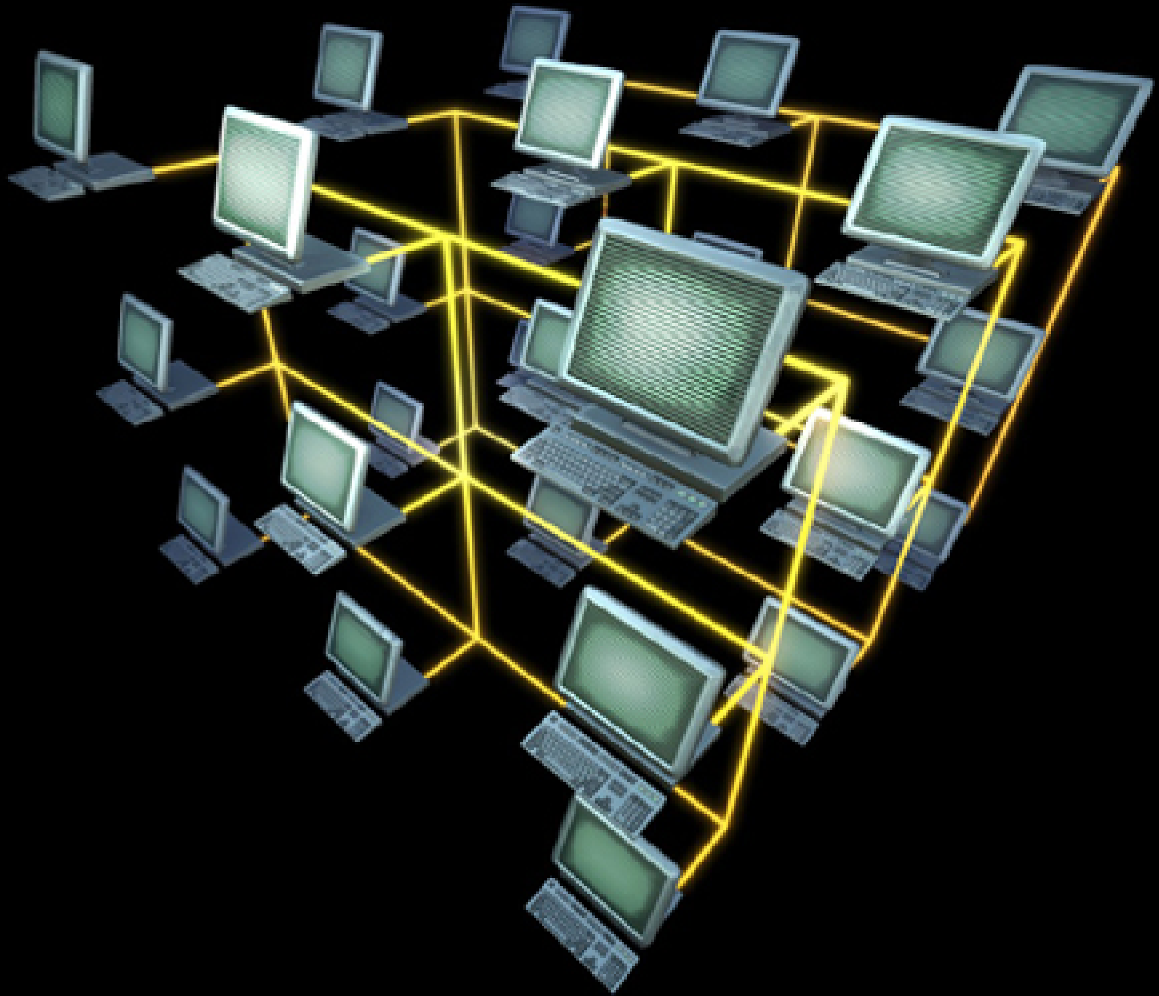

Celina Hardaway

Shameka Deal

First Edition, 2012

ISBN 978-81-323-0959-8

Published by:
Academic Studio
4735/22 Prakashdeep Bldg,
Ansari Road, Darya Ganj,
Delhi - 110002
Email: info@wtbooks.com

Table of Contents

Chapter- 1

Computer Network

A **computer network**, often simply referred to as a network, is a collection of computers and devices interconnected by communications channels that facilitate communications among users and allows users to share resources. Networks may be classified according to a wide variety of characteristics. A computer network allows sharing of resources and information among interconnected devices.

History

Early networks of communicating computers included the military radar system Semi-Automatic Ground Environment (SAGE) and its relative the commercial airline reservation system Semi-Automatic Business Research Environment (SABRE), started in the late 1950s. In the 1960s, the Advanced Research Projects Agency (ARPA) started funding the design of the Advanced Research Projects Agency Network (ARPANET) for the United States Department of Defense. Development of the network began in 1969, based on designs developed during the 1960s. The ARPANET evolved into the modern Internet.

Purpose

Computer networks can be used for a variety of purposes:

- *Facilitating communications.* Using a network, people can communicate efficiently and easily via email, instant messaging, chat rooms, telephone, video telephone calls, and video conferencing.
- *Sharing hardware.* In a networked environment, each computer on a network may access and use hardware resources on the network, such as printing a document on a shared network printer.
- *Sharing files, data, and information.* In a network environment, authorized user may access data and information stored on other computers on the network. The capability of providing access to data and information on shared storage devices is an important feature of many networks.
- *Sharing software.* Users connected to a network may run application programs on remote computers.
- *Information preservation.*

- *Easy communication*

Network classification

The following list presents categories used for classifying networks.

Connection method

Computer networks can be classified according to the hardware and software technology that is used to interconnect the individual devices in the network, such as optical fiber, Ethernet, wireless LAN, HomePNA, power line communication or G.hn.

Ethernet as it is defined by IEEE 802 utilizes various standards and mediums that enable communication between devices. Frequently deployed devices include hubs, switches, bridges, or routers. Wireless LAN technology is designed to connect devices without wiring. These devices use radio waves or infrared signals as a transmission medium. ITU-T G.hn technology uses existing home wiring (coaxial cable, phone lines and power lines) to create a high-speed (up to 1 Gigabit/s) local area network.

Wired technologies

- *Twisted pair wire* is the most widely used medium for telecommunication. Twisted-pair cabling consist of copper wires that are twisted into pairs. Ordinary telephone wires consist of two insulated copper wires twisted into pairs. Computer networking cabling consist of 4 pairs of copper cabling that can be utilized for both voice and data transmission. The use of two wires twisted together helps to reduce crosstalk and electromagnetic induction. The transmission speed ranges from 2 million bits per second to 100 million bits per second. Twisted pair cabling comes in two forms which are Unshielded Twisted Pair (UTP) and Shielded twisted-pair (STP) which are rated in categories which are manufactured in different increments for various scenarios.

- *Coaxial cable* is widely used for cable television systems, office buildings, and other work-sites for local area networks. The cables consist of copper or aluminum wire wrapped with insulating layer typically of a flexible material with a high dielectric constant, all of which are surrounded by a conductive layer. The layers of insulation help minimize interference and distortion. Transmission speed range from 200 million to more than 500 million bits per second.

- *Optical fiber cable* consists of one or more filaments of glass fiber wrapped in protective layers that carries a data by means of pulses of light. It transmits light which can travel over extended distances. Fiber-optic cables are not affected by electromagnetic radiation. Transmission speed may reach trillions of bits per second. The transmission speed of fiber optics is hundreds of times faster than for coaxial cables and thousands of times faster than a twisted-pair wire.A recent innovation in fiber-optic cable is the use of colored light.Instead of carrying one message in a stream of white light impulses, this technology can carry multiple signals in a single strand.

Wireless technologies

- *Terrestrial microwave* – Terrestrial microwaves use Earth-based transmitter and receiver. The equipment looks similar to satellite dishes. Terrestrial microwaves use low-gigahertz range, which limits all communications to line-of-sight. Path between relay stations spaced approx, 30 miles apart. Microwave antennas are usually placed on top of buildings, towers, hills, and mountain peaks.

- *Communications satellites* – The satellites use microwave radio as their telecommunications medium which are not deflected by the Earth's atmosphere. The satellites are stationed in space, typically 22,000 miles (for geosynchronous satellites) above the equator. These Earth-orbiting systems are capable of receiving and relaying voice, data, and TV signals.

- *Cellular and PCS systems* – Use several radio communications technologies. The systems are divided to different geographic areas. Each area has a low-power transmitter or radio relay antenna device to relay calls from one area to the next area.

- *Wireless LANs* – Wireless local area network use a high-frequency radio technology similar to digital cellular and a low-frequency radio technology. Wireless LANs use spread spectrum technology to enable communication between multiple devices in a limited area. An example of open-standards wireless radio-wave technology is IEEE.

- Infrared communication , which can transmit signals between devices within small distances not more than 10 meters peer to peer or (face to face) without any body in the line of transmitting.

Scale

Networks are often classified as local area network (LAN), wide area network (WAN), metropolitan area network (MAN), personal area network (PAN), virtual private network (VPN), campus area network (CAN), storage area network (SAN), and others, depending on their scale, scope and purpose, e.g., controller area network (CAN) usage, trust level, and access right often differ between these types of networks. LANs tend to be designed for internal use by an organization's internal systems and employees in individual physical locations, such as a building, while WANs may connect physically separate parts of an organization and may include connections to third parties.

Functional relationship (network architecture)

Computer networks may be classified according to the functional relationships which exist among the elements of the network, e.g., active networking, client–server, Wireless ad hoc network and peer-to-peer (workgroup) architecture.

Network topology

Computer networks may be classified according to the network topology upon which the network is based, such as bus network, star network, ring network, mesh network. Network topology is the coordination by which devices in the network are arranged in their logical relations to one another, independent of physical arrangement. Even if networked computers are physically placed in a linear arrangement and are connected to a hub, the network has a star topology, rather than a bus topology. In this regard the visual and operational characteristics of a network are distinct. Networks may be classified based on the method of data used to convey the data, these include digital and analog networks.

Types of networks based on physical scope

Common types of computer networks may be identified by their scale.

Local area network

A local area network (LAN) is a network that connects computers and devices in a limited geographical area such as home, school, computer laboratory, office building, or closely positioned group of buildings. Each computer or device on the network is a node. Current wired LANs are most likely to be based on Ethernet technology, although new standards like ITU-T G.hn also provide a way to create a wired LAN using existing home wires (coaxial cables, phone lines and power lines).

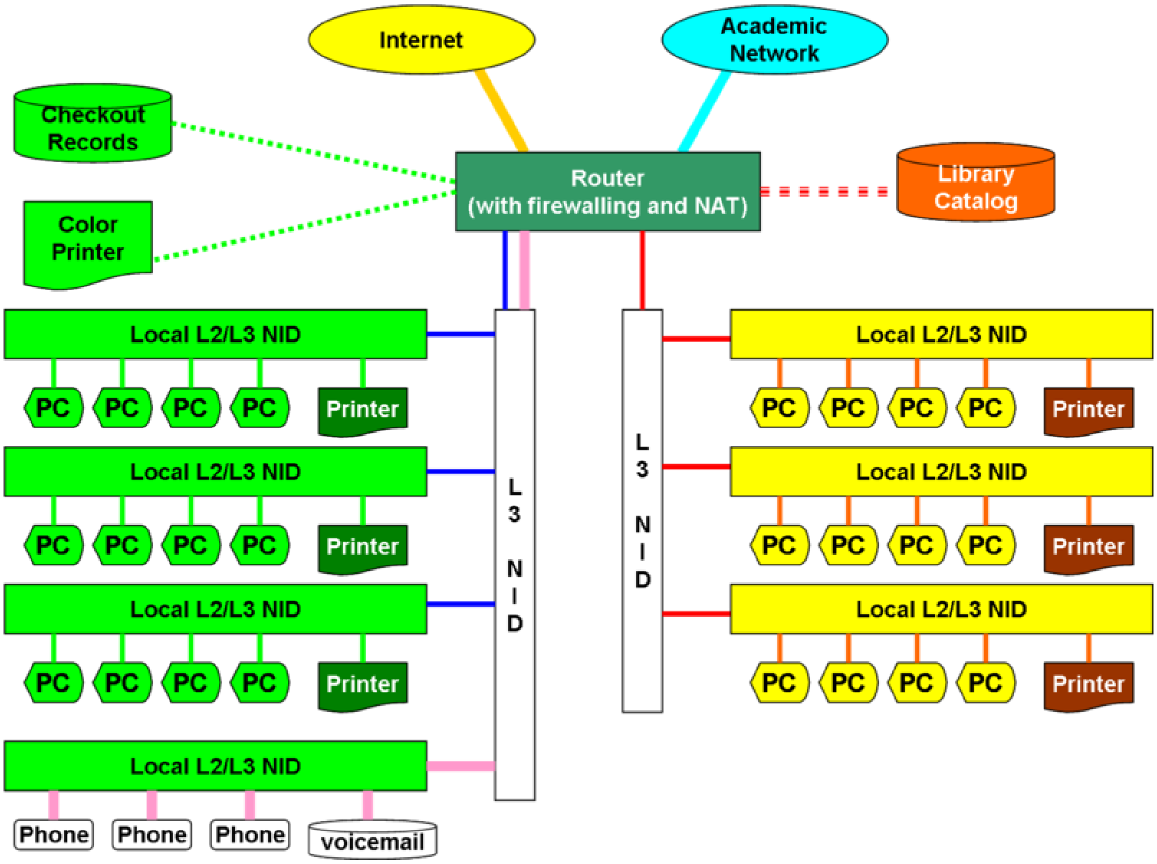

Typical library network, in a branching tree topology and controlled access to resources

All interconnected devices must understand the network layer (layer 3), because they are handling multiple subnets (the different colors). Those inside the library, which have only 10/100 Mbit/s Ethernet connections to the user device and a Gigabit Ethernet connection to the central router, could be called "layer 3 switches" because they only have Ethernet interfaces and must understand IP. It would be more correct to call them access routers, where the router at the top is a distribution router that connects to the Internet and academic networks' customer access routers.

The defining characteristics of LANs, in contrast to WANs (Wide Area Networks), include their higher data transfer rates, smaller geographic range, and no need for leased telecommunication lines. Current Ethernet or other IEEE 802.3 LAN technologies operate at speeds up to 10 Gbit/s. This is the data transfer rate. IEEE has projects investigating the standardization of 40 and 100 Gbit/s.

Personal area network

A personal area network (PAN) is a computer network used for communication among computer and different information technological devices close to one person. Some examples of devices that are used in a PAN are personal computers, printers, fax machines, telephones, PDAs, scanners, and even video game consoles. A PAN may include wired and wireless devices. The reach of a PAN typically extends to 10 meters. A wired PAN is usually constructed with USB and Firewire connections while technologies such as Bluetooth and infrared communication typically form a wireless PAN.

Home area network

A home area network (HAN) is a residential LAN which is used for communication between digital devices typically deployed in the home, usually a small number of personal computers and accessories, such as printers and mobile computing devices. An important function is the sharing of Internet access, often a broadband service through a CATV or Digital Subscriber Line (DSL) provider. It can also be referred to as an office area network (OAN).

Wide area network

A wide area network (WAN) is a computer network that covers a large geographic area such as a city, country, or spans even intercontinental distances, using a communications channel that combines many types of media such as telephone lines, cables, and air waves. A WAN often uses transmission facilities provided by common carriers, such as telephone companies. WAN technologies generally function at the lower three layers of the OSI reference model: the physical layer, the data link layer, and the network layer.

Campus network

A campus network is a computer network made up of an interconnection of local area networks (LAN's) within a limited geographical area. The networking equipments (switches, routers) and transmission media (optical fiber, copper plant, Cat5 cabling etc.) are almost entirely owned (by the campus tenant / owner: an enterprise, university, government etc.).

In the case of a university campus-based campus network, the network is likely to link a variety of campus buildings including; academic departments, the university library and student residence halls.

Metropolitan area network

A Metropolitan area network is a large computer network that usually spans a city or a large campus.

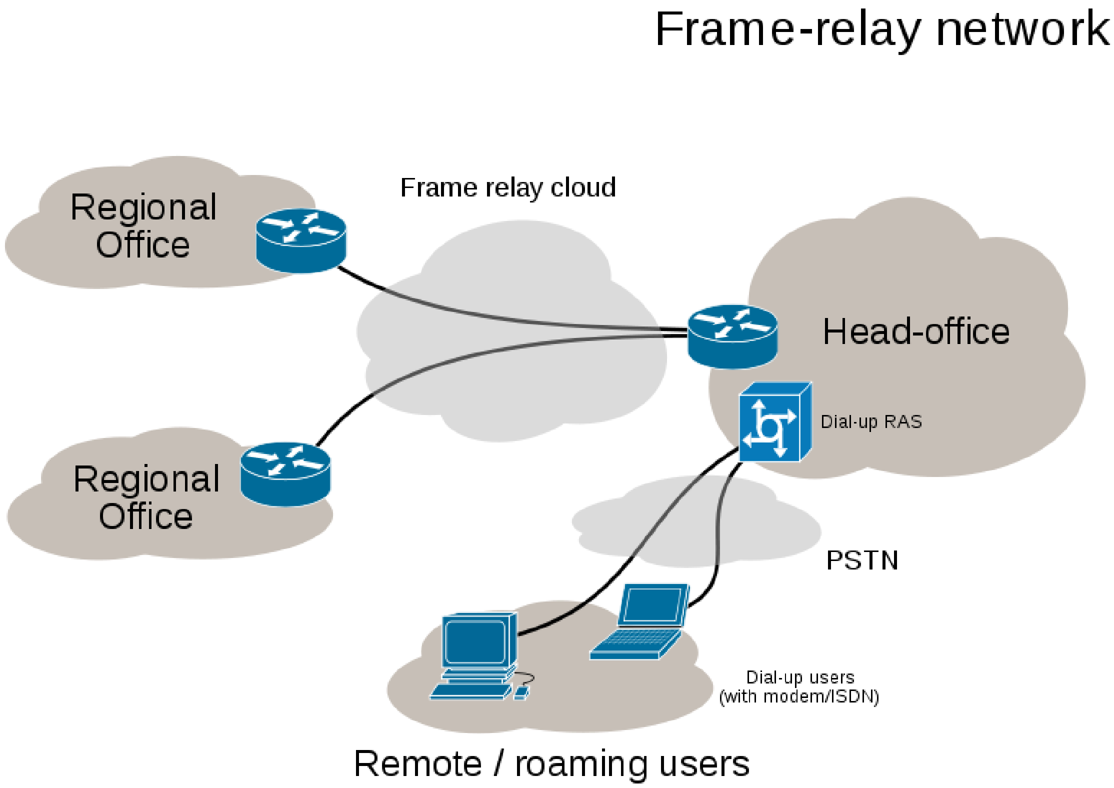

Sample EPN made of Frame relay WAN connections and dialup remote access.

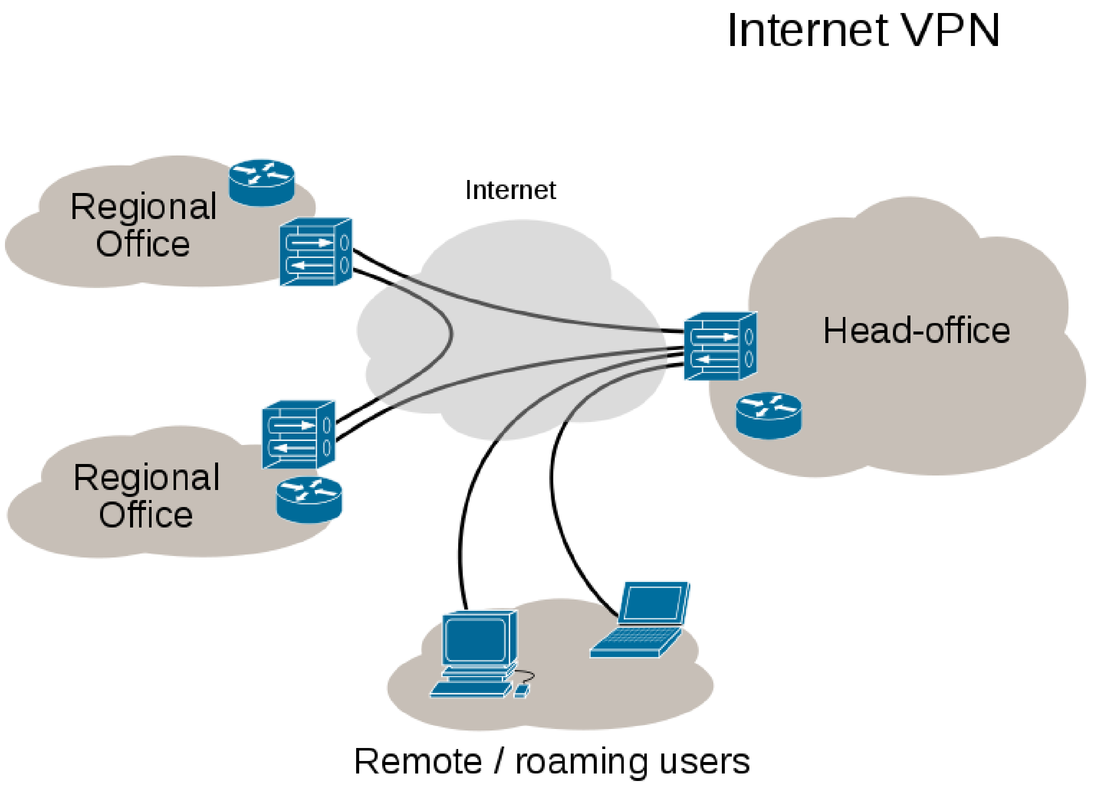

Sample VPN used to interconnect 3 offices and remote users

Enterprise private network

An enterprise private network is a network build by an enterprise to interconnect various company sites, e.g., production sites, head offices, remote offices, shops, in order to share computer resources.

Virtual private network

A virtual private network (VPN) is a computer network in which some of the links between nodes are carried by open connections or virtual circuits in some larger network (e.g., the Internet) instead of by physical wires. The data link layer protocols of the virtual network are said to be tunneled through the larger network when this is the case. One common application is secure communications through the public Internet, but a VPN need not have explicit security features, such as authentication or content encryption. VPNs, for example, can be used to separate the traffic of different user communities over an underlying network with strong security features.

VPN may have best-effort performance, or may have a defined service level agreement (SLA) between the VPN customer and the VPN service provider. Generally, a VPN has a topology more complex than point-to-point.

Internetwork

An internetwork is the connection of two or more private computer networks via a common routing technology (OSI Layer 3) using routers. The Internet is an aggregation of many internetworks, hence its name was shortened to Internet.

Backbone network

A Backbone network (BBN) A backbone network or network backbone is part of a computer network infrastructure that interconnects various pieces of network, providing a path for the exchange of information between different LANs or subnetworks. A backbone can tie together diverse networks in the same building, in different buildings in a campus environment, or over wide areas. Normally, the backbone's capacity is greater than the networks connected to it.

A large corporation that has many locations may have a backbone network that ties all of the locations together, for example, if a server cluster needs to be accessed by different departments of a company that are located at different geographical locations. The pieces of the network connections (for example: ethernet, wireless) that bring these departments together is often mentioned as network backbone. Network congestion is often taken into consideration while designing backbones.

Backbone networks should not be confused with the Internet backbone.

Global area network

A global area network (GAN) is a network used for supporting mobile communications across an arbitrary number of wireless LANs, satellite coverage areas, etc. The key challenge in mobile communications is handing off the user communications from one local coverage area to the next. In IEEE Project 802, this involves a succession of terrestrial wireless LANs.

Internet

The Internet is a global system of interconnected governmental, academic, corporate, public, and private computer networks. It is based on the networking technologies of the Internet Protocol Suite. It is the successor of the Advanced Research Projects Agency Network (ARPANET) developed by DARPA of the United States Department of Defense. The Internet is also the communications backbone underlying the World Wide Web (WWW).

Participants in the Internet use a diverse array of methods of several hundred documented, and often standardized, protocols compatible with the Internet Protocol Suite and an addressing system (IP addresses) administered by the Internet Assigned Numbers Authority and address registries. Service providers and large enterprises exchange information about the reachability of their address spaces through the Border Gateway Protocol (BGP), forming a redundant worldwide mesh of transmission paths.

Intranets and extranets

Intranets and extranets are parts or extensions of a computer network, usually a local area network.

An intranet is a set of networks, using the Internet Protocol and IP-based tools such as web browsers and file transfer applications, that is under the control of a single administrative entity. That administrative entity closes the intranet to all but specific, authorized users. Most commonly, an intranet is the internal network of an organization. A large intranet will typically have at least one web server to provide users with organizational information.

An extranet is a network that is limited in scope to a single organization or entity and also has limited connections to the networks of one or more other usually, but not necessarily, trusted organizations or entities—a company's customers may be given access to some part of its intranet—while at the same time the customers may not be considered *trusted* from a security standpoint. Technically, an extranet may also be categorized as a CAN, MAN, WAN, or other type of network, although an extranet cannot consist of a single LAN; it must have at least one connection with an external network.

Overlay network

An overlay network is a virtual computer network that is built on top of another network. Nodes in the overlay are connected by virtual or logical links, each of which corresponds to a path, perhaps through many physical links, in the underlying network.

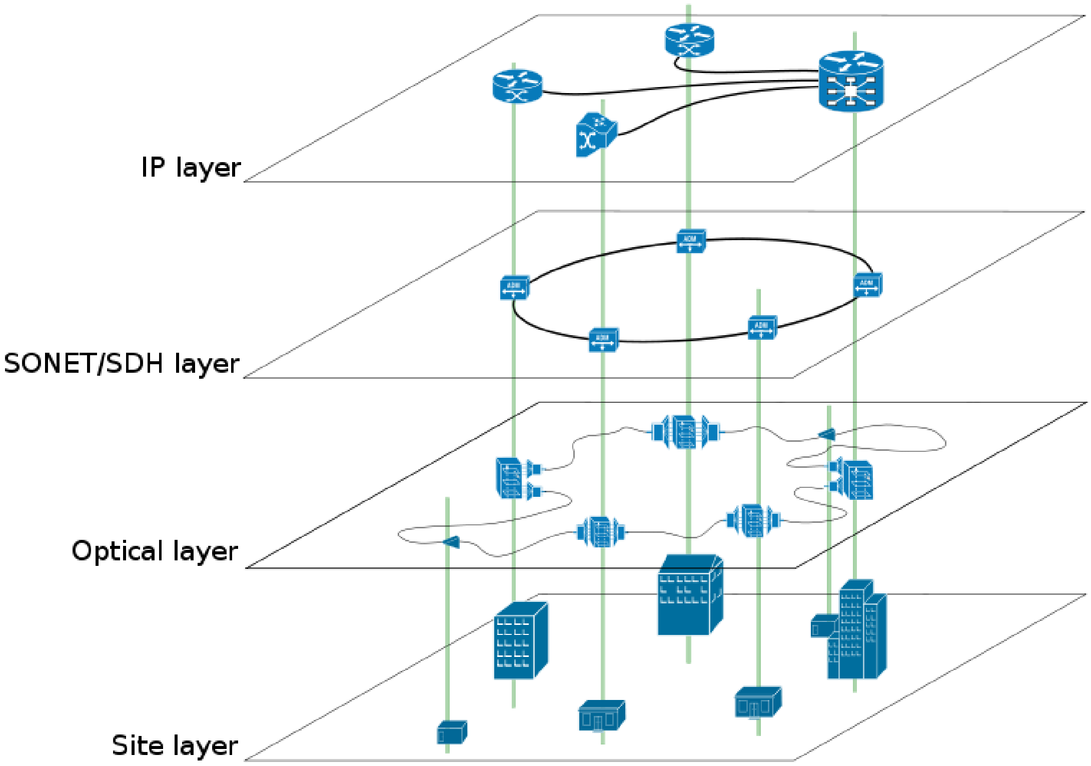

A sample overlay network: IP over SONET over Optical

For example, many peer-to-peer networks are overlay networks because they are organized as nodes of a virtual system of links run on top of the Internet. The Internet was initially built as an overlay on the telephone network.

Overlay networks have been around since the invention of networking when computer systems were connected over telephone lines using modem, before any data network existed.

Nowadays the Internet is the basis for many overlaid networks that can be constructed to permit routing of messages to destinations specified by an IP address. For example, distributed hash tables can be used to route messages to a node having a specific logical address, whose IP address is known in advance.

Overlay networks have also been proposed as a way to improve Internet routing, such as through quality of service guarantees to achieve higher-quality streaming media. Previous proposals such as IntServ, DiffServ, and IP Multicast have not seen wide acceptance largely because they require modification of all routers in the network. On the other hand, an overlay network can be incrementally deployed on end-hosts running the overlay protocol software, without cooperation from Internet service providers. The overlay has no control over how packets are routed in the underlying network between two overlay nodes, but it can control, for example, the sequence of overlay nodes a message traverses before reaching its destination.

For example, Akamai Technologies manages an overlay network that provides reliable, efficient content delivery (a kind of multicast). Academic research includes End System Multicast and Overcast for multicast; RON (Resilient Overlay Network) for resilient routing; and OverQoS for quality of service guarantees, among others. A backbone network or network backbone is a part of computer network infrastructure that interconnects various pieces of network, providing a path for the exchange of information between different LANs or subnetworks. A backbone can tie together diverse networks in the same building, in different buildings in a campus environment, or over wide areas. Normally, the backbone's capacity is greater than the networks connected to it.

Basic hardware components

All networks are made up of basic hardware building blocks to interconnect network nodes, such as Network Interface Cards (NICs), Bridges, Hubs, Switches, and Routers. In addition, some method of connecting these building blocks is required, usually in the form of galvanic cable (most commonly Category 5 cable). Less common are microwave links (as in IEEE 802.12) or optical cable ("optical fiber").

Network interface cards

A network card, network adapter, or NIC (network interface card) is a piece of computer hardware designed to allow computers to communicate over a computer network. It provides physical access to a networking medium and often provides a low-level addressing system through the use of MAC addresses.

Each network interface card has its unique id. This is written on a chip which is mounted on the card.

Repeaters

A repeater is an electronic device that receives a signal, cleans it of unnecessary noise, regenerates it, and retransmits it at a higher power level, or to the other side of an obstruction, so that the signal can cover longer distances without degradation. In most twisted pair Ethernet configurations, repeaters are required for cable that runs longer than 100 meters. A repeater with multiple ports is known as a hub. Repeaters work on the Physical Layer of the OSI model. Repeaters require a small amount of time to regenerate the signal. This can cause a propagation delay which can affect network communication when there are several repeaters in a row. Many network architectures limit the number of repeaters that can be used in a row (e.g. Ethernet's 5-4-3 rule).

Bridges

A network bridge connects multiple network segments at the data link layer (layer 2) of the OSI model. Bridges broadcast to all ports except the port on which the broadcast was received. However, bridges do not promiscuously copy traffic to all ports, as hubs do, but learn which MAC addresses are reachable through specific ports. Once the bridge associates a port and an address, it will send traffic for that address to that port only.

Bridges learn the association of ports and addresses by examining the source address of frames that it sees on various ports. Once a frame arrives through a port, its source address is stored and the bridge assumes that MAC address is associated with that port. The first time that a previously unknown destination address is seen, the bridge will forward the frame to all ports other than the one on which the frame arrived.

Bridges come in three basic types:

- Local bridges: Directly connect local area networks (LANs)
- Remote bridges: Can be used to create a wide area network (WAN) link between LANs. Remote bridges, where the connecting link is slower than the end networks, largely have been replaced with routers.
- Wireless bridges: Can be used to join LANs or connect remote stations to LANs.

Switches

A network switch is a device that forwards and filters OSI layer 2 datagrams (chunks of data communication) between ports (connected cables) based on the MAC addresses in the packets. A switch is distinct from a hub in that it only forwards the frames to the ports involved in the communication rather than all ports connected. A switch breaks the collision domain but represents itself as a broadcast domain. Switches make forwarding decisions of frames on the basis of MAC addresses. A switch normally has numerous ports, facilitating a star topology for devices, and cascading additional switches.] Some switches are capable of routing based on Layer 3 addressing or additional logical levels; these are called multi-layer switches. The term *switch* is used loosely in marketing to encompass devices including routers and bridges, as well as devices that may distribute traffic on load or by application content (e.g., a Web URL identifier).

Routers

A router is an internetworking device that forwards packets between networks by processing information found in the datagram or packet (Internet protocol information from Layer 3 of the OSI Model). In many situations, this information is processed in conjunction with the routing table (also known as forwarding table). Routers use routing tables to determine what interface to forward packets (this can include the "null" also known as the "black hole" interface because data can go into it, however, no further processing is done for said data).

Firewalls

Firewalls are the most important aspect of a network with respect to security. A firewalled system does not need every interaction or data transfer monitored by a human, as automated processes can be set up to assist in rejecting access requests from unsafe sources, and allowing actions from recognized ones. The vital role firewalls play in network security grows in parallel with the constant increase in 'cyber' attacks for the purpose of stealing/corrupting data, planting viruses, etc.

Chapter- 2

Internet

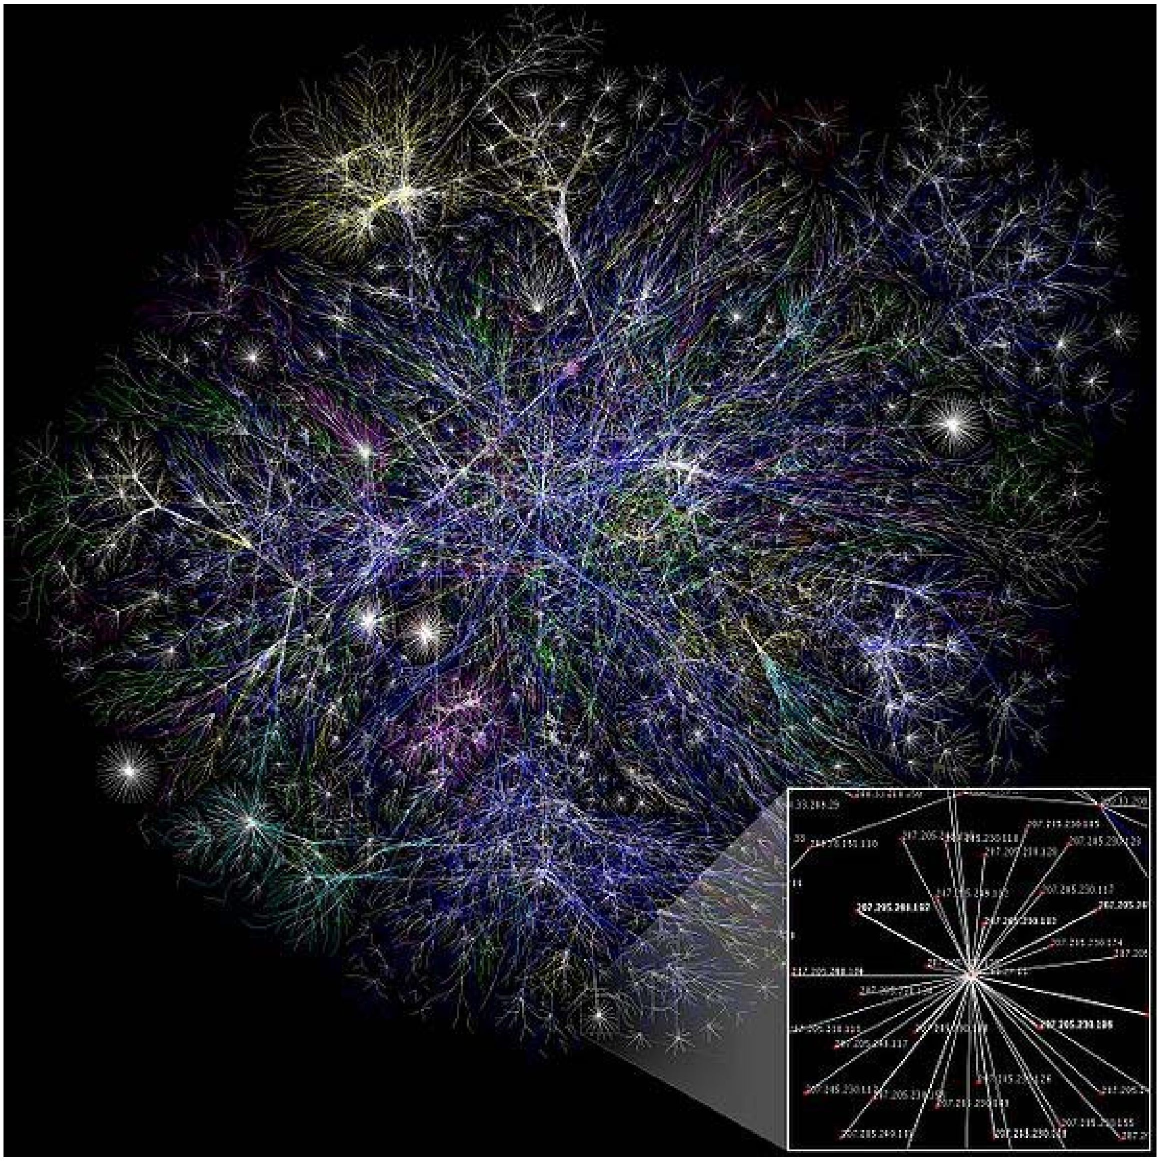

Visualization from the Opte Project of the various routes through a portion of the Internet

The **Internet** is a global system of interconnected computer networks that use the standard Internet Protocol Suite (TCP/IP) to serve billions of users worldwide. It is a *network of networks* that consists of millions of private, public, academic, business, and government networks, of local to global scope, that are linked by a broad array of electronic, wireless and optical networking technologies. The Internet carries a vast range of information resources and services, such as the inter-linked hypertext documents of the World Wide Web (WWW) and the infrastructure to support electronic mail.

Most traditional communications media including telephone, music, film, and television are reshaped or redefined by the Internet, giving birth to new services such as Voice over Internet Protocol (VoIP) and IPTV. Newspaper, book and other print publishing are adapting to Web site technology, or are reshaped into blogging and web feeds. The Internet has enabled or accelerated new forms of human interactions through instant messaging, Internet forums, and social networking. Online shopping has boomed both for major retail outlets and small artisans and traders. Business-to-business and financial services on the Internet affect supply chains across entire industries.

The origins of the Internet reach back to research of the 1960s, commissioned by the United States government in collaboration with private commercial interests to build robust, fault-tolerant, and distributed computer networks. The funding of a new U.S. backbone by the National Science Foundation in the 1980s, as well as private funding for other commercial backbones, led to worldwide participation in the development of new networking technologies, and the merger of many networks. The commercialization of what was by the 1990s an international network resulted in its popularization and incorporation into virtually every aspect of modern human life. As of 2009, an estimated quarter of Earth's population used the services of the Internet.

The Internet has no centralized governance in either technological implementation or policies for access and usage; each constituent network sets its own standards. Only the overreaching definitions of the two principal name spaces in the Internet, the Internet Protocol address space and the Domain Name System, are directed by a maintainer organization, the Internet Corporation for Assigned Names and Numbers (ICANN). The technical underpinning and standardization of the core protocols (IPv4 and IPv6) is an activity of the Internet Engineering Task Force (IETF), a non-profit organization of loosely affiliated international participants that anyone may associate with by contributing technical expertise.

Terminology

Internet is a short form of the technical term internetwork, the result of interconnecting computer networks with special gateways or routers. The Internet is also often referred to as *the Net*.

The term *the Internet*, when referring to the entire global system of IP networks, has been treated as a proper noun and written with an initial capital letter. Some guides specify that the word should be capitalized as a noun but not capitalized as an adjective.

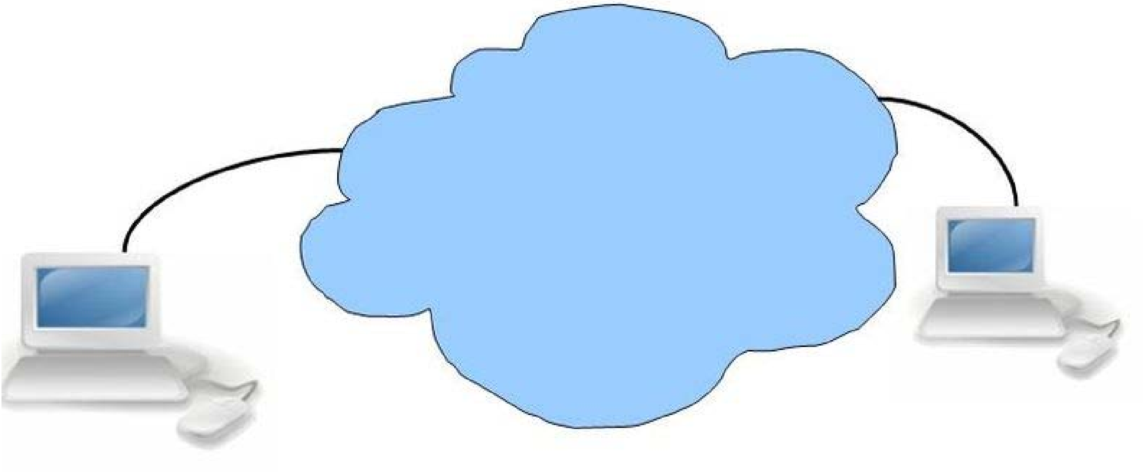

Depiction of the Internet as a *cloud* in network diagrams

The terms *Internet* and *World Wide Web* are often used in everyday speech without much distinction. However, the Internet and the World Wide Web are not one and the same. The Internet is a global data communications system. It is a hardware and software infrastructure that provides connectivity between computers. In contrast, the Web is one of the services communicated via the Internet. It is a collection of interconnected documents and other resources, linked by hyperlinks and URLs.

In many technical illustrations when the precise location or interrelation of Internet resources is not important, extended networks such as the Internet are often depicted as a cloud. The verbal image has been formalized in the newer concept of cloud computing.

History

The USSR's launch of Sputnik spurred the United States to create the Advanced Research Projects Agency (ARPA or DARPA) in February 1958 to regain a technological lead. ARPA created the Information Processing Technology Office (IPTO) to further the research of the Semi Automatic Ground Environment (SAGE) program, which had networked country-wide radar systems together for the first time. The IPTO's purpose was to find ways to address the US military's concern about survivability of their communications networks, and as a first step interconnect their computers at the Pentagon, Cheyenne Mountain, and Strategic Air Command headquarters (SAC). J. C. R. Licklider, a promoter of universal networking, was selected to head the IPTO. Licklider moved from the Psycho-Acoustic Laboratory at Harvard University to MIT in 1950, after becoming interested in information technology. At MIT, he served on a committee that established Lincoln Laboratory and worked on the SAGE project. In 1957 he became a Vice President at BBN, where he bought the first production PDP-1 computer and conducted the first public demonstration of time-sharing.

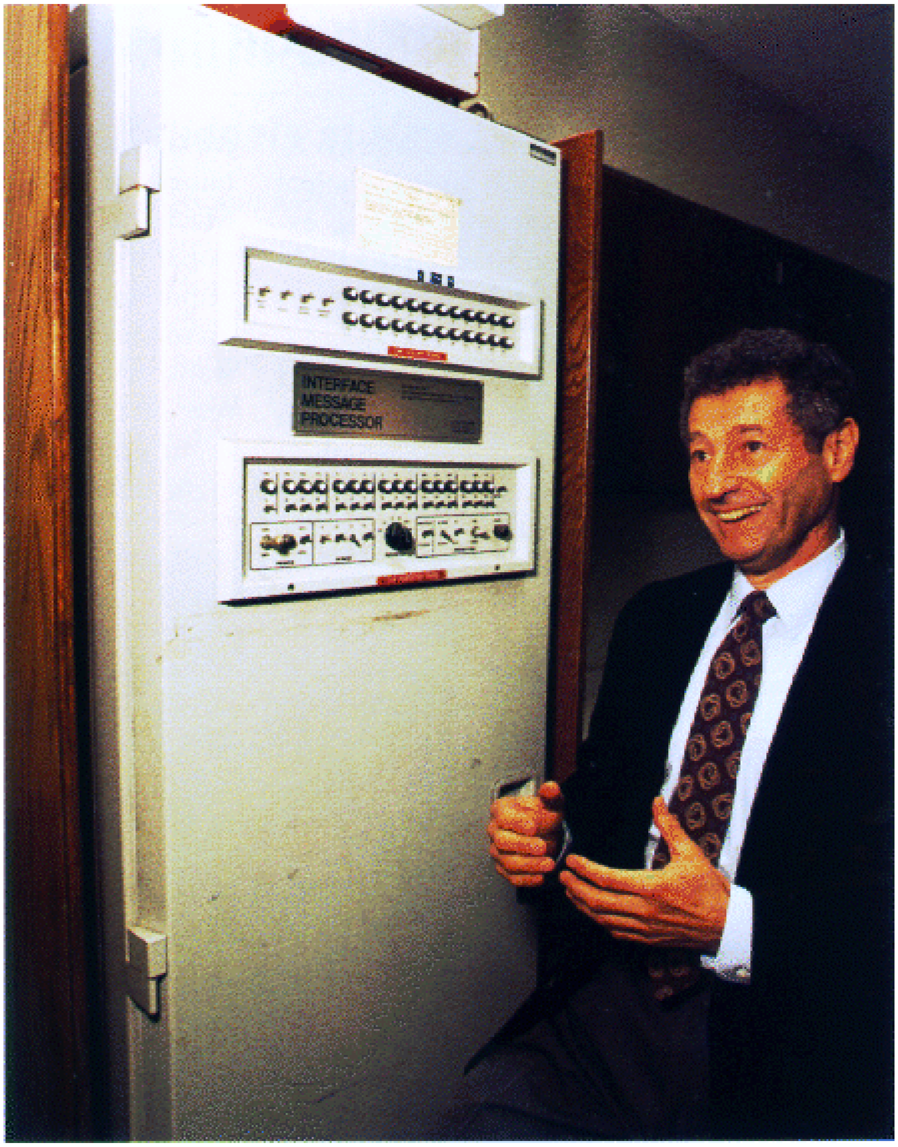

Professor Leonard Kleinrock with the first ARPANET Interface Message Processors at UCLA

A plaque commemorating the birth of the Internet at Stanford University

At the IPTO, Licklider's successor Ivan Sutherland in 1965 got Lawrence Roberts to start a project to make a network, and Roberts based the technology on the work of Paul Baran, who had written an exhaustive study for the United States Air Force that recommended packet switching (opposed to circuit switching) to achieve better network robustness and disaster survivability. Roberts had worked at the MIT Lincoln Laboratory originally established to work on the design of the SAGE system. UCLA professor Leonard Kleinrock had provided the theoretical foundations for packet networks in 1962, and later, in the 1970s, for hierarchical routing, concepts which have been the underpinning of the development towards today's Internet.

Sutherland's successor Robert Taylor convinced Roberts to build on his early packet switching successes and come and be the IPTO Chief Scientist. Once there, Roberts prepared a report called *Resource Sharing Computer Networks* which was approved by Taylor in June 1968 and laid the foundation for the launch of the working ARPANET the following year.

After much work, the first two nodes of what would become the ARPANET were interconnected between Kleinrock's Network Measurement Center at the UCLA's School of Engineering and Applied Science and Douglas Engelbart's NLS system at SRI International (SRI) in Menlo Park, California, on 29 October 1969. The third site on the ARPANET was the Culler-Fried Interactive Mathematics center at the University of California at Santa Barbara, and the fourth was the University of Utah Graphics Department. In an early sign of future growth, there were already fifteen sites connected to the young ARPANET by the end of 1971.

In an independent development, Donald Davies at the UK National Physical Laboratory developed the concept of packet switching in the early 1960s, first giving a talk on the subject in 1965, after which the teams in the new field from two sides of the Atlantic ocean first became acquainted. It was actually Davies' coinage of the wording *packet* and *packet switching* that was adopted as the standard terminology. Davies also built a packet-switched network in the UK, called the Mark I in 1970. Bolt, Beranek & Newman

(BBN), the private contractors for ARPANET, set out to create a separate commercial version after establishing "value added carriers" was legalized in the U.S. The network they established was called Telenet and began operation in 1975, installing free public dial-up access in cities throughout the U.S. Telenet was the first packet-switching network open to the general public.

Following the demonstration that packet switching worked on the ARPANET, the British Post Office, Telenet, DATAPAC and TRANSPAC collaborated to create the first international packet-switched network service. In the UK, this was referred to as the International Packet Switched Service (IPSS), in 1978. The collection of X.25-based networks grew from Europe and the US to cover Canada, Hong Kong and Australia by 1981. The X.25 packet switching standard was developed in the CCITT (now called ITU-T) around 1976. X.25 was independent of the TCP/IP protocols that arose from the experimental work of DARPA on the ARPANET, Packet Radio Net, and Packet Satellite Net during the same time period.

The early ARPANET ran on the Network Control Program (NCP), implementing the host-to-host connectivity and switching layers of the protocol stack, designed and first implemented in December 1970 by a team called the Network Working Group (NWG) led by Steve Crocker. To respond to the network's rapid growth as more and more locations connected, Vinton Cerf and Robert Kahn developed the first description of the now widely used TCP protocols during 1973 and published a paper on the subject in May 1974. Use of the term "Internet" to describe a single global TCP/IP network originated in December 1974 with the publication of RFC 675, the first full specification of TCP that was written by Vinton Cerf, Yogen Dalal and Carl Sunshine, then at Stanford University. During the next nine years, work proceeded to refine the protocols and to implement them on a wide range of operating systems. The first TCP/IP-based wide-area network was operational by 1 January 1983 when all hosts on the ARPANET were switched over from the older NCP protocols.

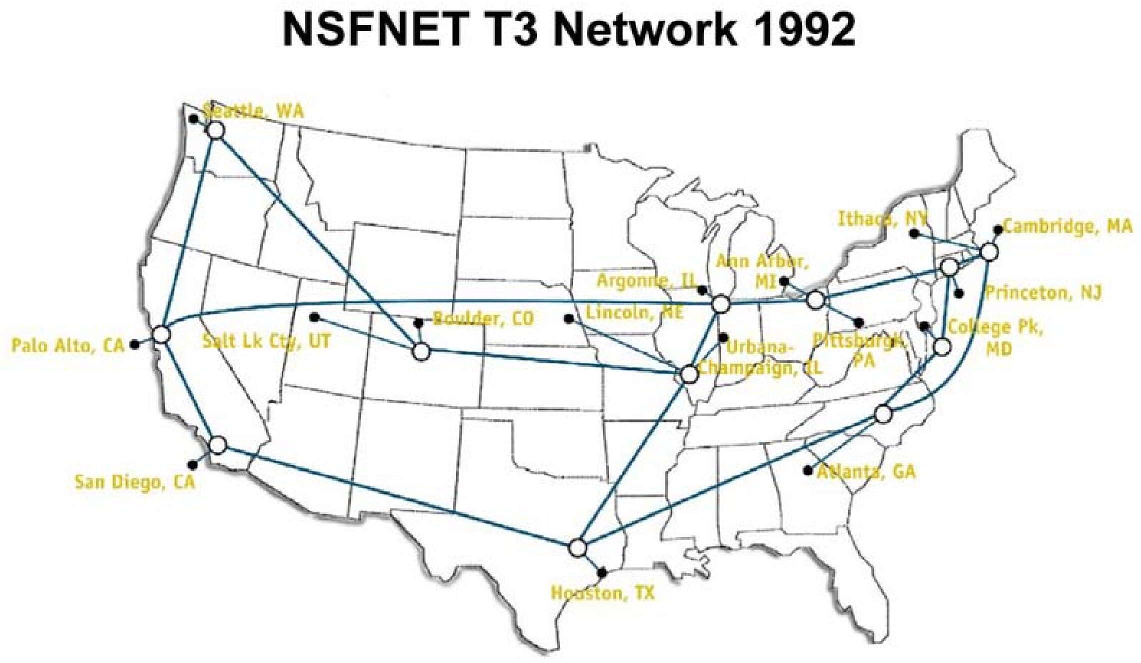

T3 NSFNET Backbone, c. 1992

In 1985, the United States' National Science Foundation (NSF) commissioned the construction of the NSFNET, a university 56 kilobit/second network backbone using computers called "fuzzballs" by their inventor, David L. Mills. The following year, NSF sponsored the conversion to a higher-speed 1.5 megabit/second network that became operational in 1988. A key decision to use the DARPA TCP/IP protocols was made by Dennis Jennings, then in charge of the Supercomputer program at NSF. The NSFNET backbone was upgraded to 45 Mbps in 1991 and decommissioned in 1995 when it was replaced by new backbone networks operated by commercial Internet Service Providers.

The opening of the NSFNET to other networks began in 1988.] The US Federal Networking Council approved the interconnection of the NSFNET to the commercial MCI Mail system in that year and the link was made in the summer of 1989. Other commercial electronic mail services were soon connected, including OnTyme, Telemail and Compuserve. In that same year, three commercial Internet service providers (ISPs) began operations: UUNET, PSINet, and CERFNET. Important, separate networks that offered gateways into, then later merged with, the Internet include Usenet and BITNET. Various other commercial and educational networks, such as Telenet (by that time renamed to Sprintnet), Tymnet, Compuserve and JANET were interconnected with the growing Internet in the 1980s as the TCP/IP protocol became increasingly popular. The adaptability of TCP/IP to existing communication networks allowed for rapid growth. The open availability of the specifications and reference code permitted commercial vendors to build interoperable network components, such as routers, making standardized network gear available from many companies. This aided in the rapid growth of the Internet and the proliferation of local-area networking. It seeded the widespread implementation and rigorous standardization of TCP/IP on UNIX and virtually every other common operating system.

This NeXT Computer was used by Sir Tim Berners-Lee at CERN and became the world's first Web server.

Although the basic applications and guidelines that make the Internet possible had existed for almost two decades, the network did not gain a public face until the 1990s. On 6 August 1991, CERN, a pan-European organization for particle research, publicized the new World Wide Web project. The Web was invented by British scientist Tim Berners-Lee in 1989. An early popular web browser was ViolaWWW, patterned after HyperCard and built using the X Window System. It was eventually replaced in popularity by the Mosaic web browser. In 1993, the National Center for Supercomputing Applications at the University of Illinois released version 1.0 of Mosaic, and by late 1994 there was growing public interest in the previously academic, technical Internet. By 1996 usage of the word *Internet* had become commonplace, and consequently, so had its use as a synecdoche in reference to the World Wide Web.

Meanwhile, over the course of the decade, the Internet successfully accommodated the majority of previously existing public computer networks (although some networks, such as FidoNet, have remained separate). During the late 1990s, it was estimated that traffic on the public Internet grew by 100 percent per year, while the mean annual growth in the number of Internet users was thought to be between 20% and 50%. This growth is often attributed to the lack of central administration, which allows organic growth of the network, as well as the non-proprietary open nature of the Internet protocols, which encourages vendor interoperability and prevents any one company from exerting too much control over the network. The estimated population of Internet users is 1.97 billion as of 30 June 2010.

From 2009 onward, the Internet is expected to grow significantly in Brazil, Russia, India, China, and Indonesia (BRICI countries). These countries have large populations and moderate to high economic growth, but still low Internet penetration rates. In 2009, the BRICI countries represented about 45 percent of the world's population and had approximately 610 million Internet users, but by 2015, Internet users in BRICI countries will double to 1.2 billion, and will triple in Indonesia.

Technology

Protocols

The complex communications infrastructure of the Internet consists of its hardware components and a system of software layers that control various aspects of the architecture. While the hardware can often be used to support other software systems, it is the design and the rigorous standardization process of the software architecture that characterizes the Internet and provides the foundation for its scalability and success. The responsibility for the architectural design of the Internet software systems has been delegated to the Internet Engineering Task Force (IETF). The IETF conducts standard-setting work groups, open to any individual, about the various aspects of Internet architecture. Resulting discussions and final standards are published in a series of publications, each called a Request for Comments (RFC), freely available on the IETF web site. The principal methods of networking that enable the Internet are contained in specially designated RFCs that constitute the Internet Standards. Other less rigorous documents are simply informative, experimental, or historical, or document the best current practices (BCP) when implementing Internet technologies.

The Internet Standards describe a framework known as the Internet Protocol Suite. This is a model architecture that divides methods into a layered system of protocols (RFC 1122, RFC 1123). The layers correspond to the environment or scope in which their services operate. At the top is the Application Layer, the space for the application-specific networking methods used in software applications, e.g., a web browser program. Below this top layer, the Transport Layer connects applications on *different hosts* via the network (e.g., client–server model) with appropriate data exchange methods. Underlying these layers are the core networking technologies, consisting of two layers. The Internet Layer enables computers to identify and locate each other via Internet Protocol (IP) addresses, and allows them to connect to one-another via intermediate (transit) networks. Lastly, at the bottom of the architecture, is a software layer, the Link Layer, that provides connectivity between hosts on the same local network link, such as a local area network (LAN) or a dial-up connection. The model, also known as TCP/IP, is designed to be independent of the underlying hardware which the model therefore does not concern itself with in any detail. Other models have been developed, such as the Open Systems Interconnection (OSI) model, but they are not compatible in the details of description, nor implementation, but many similarities exist and the TCP/IP protocols are usually included in the discussion of OSI networking.

The most prominent component of the Internet model is the Internet Protocol (IP) which provides addressing systems (IP addresses) for computers on the Internet. IP enables internetworking and essentially establishes the Internet itself. IP Version 4 (IPv4) is the initial version used on the first generation of the today's Internet and is still in dominant use. It was designed to address up to ~4.3 billion (10^9) Internet hosts. However, the explosive growth of the Internet has led to IPv4 address exhaustion which is estimated to enter its final stage in approximately 2011. A new protocol version, IPv6, was developed in the mid 1990s which provides vastly larger addressing capabilities and more efficient routing of Internet traffic. IPv6 is currently in commercial deployment phase around the world and Internet address registries (RIRs) have begun to urge all resource managers to plan rapid adoption and conversion.

IPv6 is not interoperable with IPv4. It essentially establishes a "parallel" version of the Internet not directly accessible with IPv4 software. This means software upgrades or translator facilities are necessary for every networking device that needs to communicate on the IPv6 Internet. Most modern computer operating systems are already converted to operate with both versions of the Internet Protocol. Network infrastructures, however, are still lagging in this development. Aside from the complex physical connections that make up its infrastructure, the Internet is facilitated by bi- or multi-lateral commercial contracts (e.g., peering agreements), and by technical specifications or protocols that describe how to exchange data over the network. Indeed, the Internet is defined by its interconnections and routing policies.

Structure

The Internet structure and its usage characteristics have been studied extensively. It has been determined that both the Internet IP routing structure and hypertext links of the World Wide Web are examples of scale-free networks. Similar to the way the commercial Internet providers connect via Internet exchange points, research networks tend to interconnect into large subnetworks such as GEANT, GLORIAD, Internet2 (successor of the Abilene Network), and the UK's national research and education network JANET. These in turn are built around smaller networks.

Many computer scientists describe the Internet as a "prime example of a large-scale, highly engineered, yet highly complex system". The Internet is extremely heterogeneous; for instance, data transfer rates and physical characteristics of connections vary widely. The Internet exhibits "emergent phenomena" that depend on its large-scale organization. For example, data transfer rates exhibit temporal self-similarity. The principles of the routing and addressing methods for traffic in the Internet reach back to their origins the

1960s when the eventual scale and popularity of the network could not be anticipated. Thus, the possibility of developing alternative structures is investigated.

Governance

ICANN headquarters in Marina Del Rey, California, United States

The Internet is a globally distributed network comprising many voluntarily interconnected autonomous networks. It operates without a central governing body. However, to maintain interoperability, all technical and policy aspects of the underlying core infrastructure and the principal name spaces are administered by the Internet Corporation for Assigned Names and Numbers (ICANN), headquartered in Marina del Rey, California. ICANN is the authority that coordinates the assignment of unique identifiers for use on the Internet, including domain names, Internet Protocol (IP) addresses, application port numbers in the transport protocols, and many other parameters. Globally unified name spaces, in which names and numbers are uniquely assigned, are essential for the global reach of the Internet. ICANN is governed by an international board of directors drawn from across the Internet technical, business, academic, and other non-commercial communities. The government of the United States continues to have the primary role in approving changes to the DNS root zone that lies at the heart of the domain name system. ICANN's role in coordinating the assignment of unique identifiers distinguishes it as perhaps the only central coordinating body on the global Internet. On 16 November 2005, the World Summit on the Information Society, held in Tunis, established the Internet Governance Forum (IGF) to discuss Internet-related issues.

Modern uses

The Internet is allowing greater flexibility in working hours and location, especially with the spread of unmetered high-speed connections and web applications.

The Internet can now be accessed almost anywhere by numerous means, especially through mobile Internet devices. Mobile phones, datacards, handheld game consoles and cellular routers allow users to connect to the Internet from anywhere there is a wireless network supporting that device's technology. Within the limitations imposed by small screens and other limited facilities of such pocket-sized devices, services of the Internet, including email and the web, may be available. Service providers may restrict the services offered and wireless data transmission charges may be significantly higher than other access methods.

Educational material at all levels from pre-school to post-doctoral is available from websites. Examples range from CBeebies, through school and high-school revision guides, virtual universities, to access to top-end scholarly literature through the likes of Google Scholar. In distance education, help with homework and other assignments, self-guided learning, whiling away spare time, or just looking up more detail on an interesting fact, it has never been easier for people to access educational information at any level from anywhere. The Internet in general and the World Wide Web in particular are important enablers of both formal and informal education.

The low cost and nearly instantaneous sharing of ideas, knowledge, and skills has made collaborative work dramatically easier, with the help of collaborative software. Not only can a group cheaply communicate and share ideas, but the wide reach of the Internet allows such groups to easily form in the first place. An example of this is the free software movement, which has produced, among other programs, Linux, Mozilla Firefox, and OpenOffice.org. Internet "chat", whether in the form of IRC chat rooms or channels, or via instant messaging systems, allow colleagues to stay in touch in a very convenient way when working at their computers during the day. Messages can be exchanged even more quickly and conveniently than via email. Extensions to these systems may allow files to be exchanged, "whiteboard" drawings to be shared or voice and video contact between team members.

Version control systems allow collaborating teams to work on shared sets of documents without either accidentally overwriting each other's work or having members wait until they get "sent" documents to be able to make their contributions. Business and project teams can share calendars as well as documents and other information. Such collaboration occurs in a wide variety of areas including scientific research, software development, conference planning, political activism and creative writing. Social and political collaboration is also becoming more widespread as both Internet access and computer literacy grow. From the flash mob 'events' of the early 2000s to the use of social networking in the 2009 Iranian election protests, the Internet allows people to work together more effectively and in many more ways than was possible without it.

The Internet allows computer users to remotely access other computers and information stores easily, wherever they may be across the world. They may do this with or without the use of security, authentication and encryption technologies, depending on the requirements. This is encouraging new ways of working from home, collaboration and information sharing in many industries. An accountant sitting at home can audit the books of a company based in another country, on a server situated in a third country that is remotely maintained by IT specialists in a fourth. These accounts could have been created by home-working bookkeepers, in other remote locations, based on information emailed to them from offices all over the world. Some of these things were possible before the widespread use of the Internet, but the cost of private leased lines would have made many of them infeasible in practice. An office worker away from their desk, perhaps on the other side of the world on a business trip or a holiday, can open a remote desktop session into his normal office PC using a secure Virtual Private Network (VPN) connection via the Internet. This gives the worker complete access to all of his or her normal files and data, including email and other applications, while away from the office. This concept has been referred to among system administrators as the Virtual Private Nightmare, because it extends the secure perimeter of a corporate network into its employees' homes.

Services

Information

Many people use the terms *Internet* and *World Wide Web*, or just the *Web*, interchangeably, but the two terms are not synonymous. The World Wide Web is a global set of documents, images and other resources, logically interrelated by hyperlinks and referenced with Uniform Resource Identifiers (URIs). URIs allow providers to symbolically identify services and clients to locate and address web servers, file servers, and other databases that store documents and provide resources and access them using the Hypertext Transfer Protocol (HTTP), the primary carrier protocol of the Web. HTTP is only one of the hundreds of communication protocols used on the Internet. Web services may also use HTTP to allow software systems to communicate in order to share and exchange business logic and data.

World Wide Web browser software, such as Microsoft's Internet Explorer, Mozilla Firefox, Opera, Apple's Safari, and Google Chrome, let users navigate from one web page to another via hyperlinks embedded in the documents. These documents may also contain any combination of computer data, including graphics, sounds, text, video, multimedia and interactive content including games, office applications and scientific demonstrations. Through keyword-driven Internet research using search engines like Yahoo! and Google, users worldwide have easy, instant access to a vast and diverse amount of online information. Compared to printed encyclopedias and traditional libraries, the World Wide Web has enabled the decentralization of information.

The Web has also enabled individuals and organizations to publish ideas and information to a potentially large audience online at greatly reduced expense and time delay. Publishing a web page, a blog, or building a website involves little initial cost and many cost-free services are available. Publishing and maintaining large, professional web sites with attractive, diverse and up-to-date information is still a difficult and expensive proposition, however. Many individuals and some companies and groups use *web logs* or blogs, which are largely used as easily updatable online diaries. Some commercial organizations encourage staff to communicate advice in their areas of specialization in the hope that visitors will be impressed by the expert knowledge and free information, and be attracted to the corporation as a result. One example of this practice is Microsoft, whose product developers publish their personal blogs in order to pique the public's interest in their work. Collections of personal web pages published by large service providers remain popular, and have become increasingly sophisticated. Whereas operations such as Angelfire and GeoCities have existed since the early days of the Web, newer offerings from, for example, Facebook and MySpace currently have large followings. These operations often brand themselves as social network services rather than simply as web page hosts.

Advertising on popular web pages can be lucrative, and e-commerce or the sale of products and services directly via the Web continues to grow.

When the Web began in the 1990s, a typical web page was stored in completed form on a web server, formatted with HTML, ready to be sent to a user's browser in response to a request. Over time, the process of creating and serving web pages has become more automated and more dynamic. Websites are often created using content management or software with, initially, very little content. Contributors to these systems, who may be paid staff, members of a club or other organization or members of the public, fill underlying databases with content using editing pages designed for that purpose, while casual visitors view and read this content in its final HTML form. There may or may not be editorial, approval and security systems built into the process of taking newly entered content and making it available to the target visitors.

Communication

Electronic mail, or email, is an important communications service available on the Internet. The concept of sending electronic text messages between parties in a way analogous to mailing letters or memos predates the creation of the Internet. Pictures, documents and other files are sent as email attachments. Emails can be cc-ed to multiple email addresses.

Internet telephony is another common communications service made possible by the creation of the Internet. VoIP stands for Voice-over-Internet Protocol, referring to the protocol that underlies all Internet communication. The idea began in the early 1990s with walkie-talkie-like voice applications for personal computers. In recent years many VoIP systems have become as easy to use and as convenient as a normal telephone. The benefit is that, as the Internet carries the voice traffic, VoIP can be free or cost much less than a traditional telephone call, especially over long distances and especially for those with always-on Internet connections such as cable or ADSL. VoIP is maturing into a competitive alternative to traditional telephone service. Interoperability between different providers has improved and the ability to call or receive a call from a traditional telephone is available. Simple, inexpensive VoIP network adapters are available that eliminate the need for a personal computer.

Voice quality can still vary from call to call but is often equal to and can even exceed that of traditional calls. Remaining problems for VoIP include emergency telephone number dialing and reliability. Currently, a few VoIP providers provide an emergency service, but it is not universally available. Traditional phones are line-powered and operate during a power failure; VoIP does not do so without a backup power source for the phone equipment and the Internet access devices. VoIP has also become increasingly popular for gaming applications, as a form of communication between players. Popular VoIP clients for gaming include Ventrilo and Teamspeak. Wii, PlayStation 3, and Xbox 360 also offer VoIP chat features.

Data transfer

File sharing is an example of transferring large amounts of data across the Internet. A computer file can be emailed to customers, colleagues and friends as an attachment. It can be uploaded to a website or FTP server for easy download by others. It can be put into a "shared location" or onto a file server for instant use by colleagues. The load of bulk downloads to many users can be eased by the use of "mirror" servers or peer-to-peer networks. In any of these cases, access to the file may be controlled by user authentication, the transit of the file over the Internet may be obscured by encryption, and money may change hands for access to the file. The price can be paid by the remote charging of funds from, for example, a credit card whose details are also passed—usually fully encrypted—across the Internet. The origin and authenticity of the file received may be checked by digital signatures or by MD5 or other message digests. These simple features of the Internet, over a worldwide basis, are changing the production, sale, and distribution of anything that can be reduced to a computer file for transmission. This includes all manner of print publications, software products, news, music, film, video, photography, graphics and the other arts. This in turn has caused seismic shifts in each of the existing industries that previously controlled the production and distribution of these products.

Streaming media is the real-time delivery of digital media for the immediate consumption or enjoyment by end users. Many radio and television broadcasters provide Internet feeds of their live audio and video productions. They may also allow time-shift viewing or listening such as Preview, Classic Clips and Listen Again features. These providers have been joined by a range of pure Internet "broadcasters" who never had on-air licenses. This means that an Internet-connected device, such as a computer or something more specific, can be used to access on-line media in much the same way as was previously possible only with a television or radio receiver. The range of available types of content is much wider, from specialized technical webcasts to on-demand popular multimedia services. Podcasting is a variation on this theme, where—usually audio—material is downloaded and played back on a computer or shifted to a portable media player to be listened to on the move. These techniques using simple equipment allow anybody, with little censorship or licensing control, to broadcast audio-visual material worldwide.

Digital media streaming increases the demand for network bandwidth. For example, standard image quality needs 1 Mbps link speed for SD 480p, HD 720p quality requires 2.5 Mbps, and the top-of-the-line HDX quality needs 4.5 Mbps for 1080p.

Webcams are a low-cost extension of this phenomenon. While some webcams can give full-frame-rate video, the picture is usually either small or updates slowly. Internet users can watch animals around an African waterhole, ships in the Panama Canal, traffic at a local roundabout or monitor their own premises, live and in real time. Video chat rooms and video conferencing are also popular with many uses being found for personal webcams, with and without two-way sound. YouTube was founded on 15 February 2005 and is now the leading website for free streaming video with a vast number of users. It uses a flash-based web player to stream and show video files. Registered users may upload an unlimited amount of video and build their own personal profile. YouTube claims that its users watch hundreds of millions, and upload hundreds of thousands of videos daily.

Access

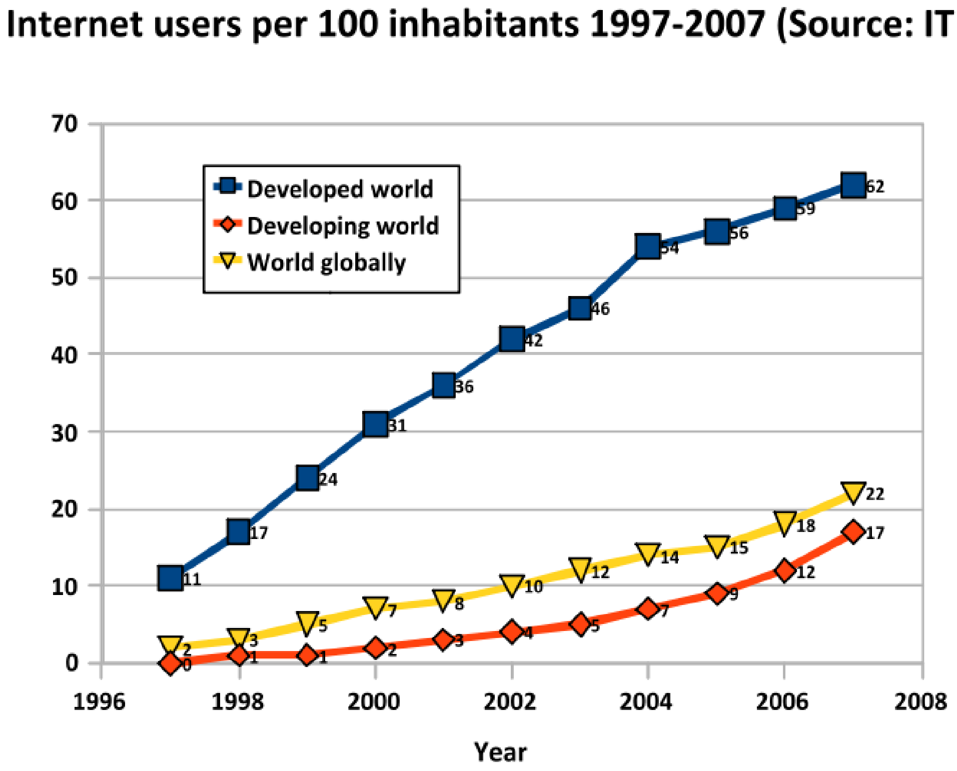

Graph of Internet users per 100 inhabitants between 1997 and 2007 by International Telecommunication Union

The prevalent language for communication on the Internet has been English. This may be a result of the origin of the Internet, as well as the language's role as a lingua franca. Early computer systems were limited to the characters in the American Standard Code for Information Interchange (ASCII), a subset of the Latin alphabet.

After English (27%), the most requested languages on the World Wide Web are Chinese (23%), Spanish (8%), Japanese (5%), Portuguese and German (4% each), Arabic, French and Russian (3% each), and Korean (2%). By region, 42% of the world's Internet users are based in Asia, 24% in Europe, 14% in North America, 10% in Latin America and the Caribbean taken together, 6% in Africa, 3% in the Middle East and 1% in Australia/Oceania. The Internet's technologies have developed enough in recent years, especially in the use of Unicode, that good facilities are available for development and communication in the world's widely used languages. However, some glitches such as *mojibake* (incorrect display of some languages' characters) still remain.

Common methods of Internet access in homes include dial-up, landline broadband (over coaxial cable, fiber optic or copper wires), Wi-Fi, satellite and 3G/4G technology cell phones. Public places to use the Internet include libraries and Internet cafes, where computers with Internet connections are available. There are also Internet access points in many public places such as airport halls and coffee shops, in some cases just for brief use while standing. Various terms are used, such as "public Internet kiosk", "public access terminal", and "Web payphone". Many hotels now also have public terminals, though these are usually fee-based. These terminals are widely accessed for various usage like ticket booking, bank deposit, online payment etc. Wi-Fi provides wireless access to computer networks, and therefore can do so to the Internet itself. Hotspots providing such access include Wi-Fi cafes, where would-be users need to bring their own wireless-enabled devices such as a laptop or PDA. These services may be free to all, free to customers only, or fee-based. A hotspot need not be limited to a confined location. A whole campus or park, or even an entire city can be enabled. Grassroots efforts have led to wireless community networks. Commercial Wi-Fi services covering large city areas are in place in London, Vienna, Toronto, San Francisco, Philadelphia, Chicago and Pittsburgh. The Internet can then be accessed from such places as a park bench. Apart from Wi-Fi, there have been experiments with proprietary mobile wireless networks like Ricochet, various high-speed data services over cellular phone networks, and fixed wireless services. High-end mobile phones such as smartphones generally come with Internet access through the phone network. Web browsers such as Opera are available on these advanced handsets, which can also run a wide variety of other Internet software. More mobile phones have Internet access than PCs, though this is not as widely used. An Internet access provider and protocol matrix differentiates the methods used to get online.

In contrast, an *Internet blackout* or *outage* can be caused by accidental local signaling interruptions. Disruptions of submarine communications cables may cause blackouts or slowdowns to large areas depending on them, such as in the 2008 submarine cable disruption. Internet blackouts of almost entire countries can be achieved by governments as Internet censorship, such as with the Internet in Egypt, where approximately 93% of networks were shut down in 2011 in an attempt to stop mobilisation for anti-government protests.

In an American study in 2005, the percentage of men using the Internet was very slightly ahead of the percentage of women, although this difference reversed in those under 30. Men logged on more often, spend more time online, and are more likely to be broadband users, whereas women tended to make more use of opportunities to communicate (such as email). Men were more likely to use the Internet to pay bills, participate in auctions, and for recreation such as downloading music and videos. Men and women were equally likely to use the Internet for shopping and banking. More recent studies indicate that in 2008, women significantly outnumbered men on most social networking sites, such as Facebook and Myspace, although the ratios varied with age. In addition, women watched more streaming content, whereas men downloaded more. In terms of blogs, men were more likely to blog in the first place; among those who blog, men were more likely to have a professional blog, whereas women were more likely to have a personal blog.

Overall Internet usage has seen tremendous growth. From 2000 to 2009, the number of Internet users globally rose from 394 million to 1.858 billion.

Social impact

The Internet has enabled entirely new forms of social interaction, activities, and organizing, thanks to its basic features such as widespread usability and access. Social networking websites such as Facebook, Twitter and MySpace have created new ways to socialize and interact. Users of these sites are able to add a wide variety of information to pages, to pursue common interests, and to connect with others. It is also possible to find existing acquaintances, to allow communication among existing groups of people. Sites like LinkedIn foster commercial and business connections. YouTube and Flickr specialize in users' videos and photographs.

In the first decade of the 21st century the first generation is raised with widespread availability of Internet connectivity, bringing consequences and concerns in areas such as personal privacy and identity, and distribution of copyrighted materials. These "digital natives" face a variety of challenges that were not present for prior generations.

The Internet has achieved new relevance as a political tool, leading to Internet censorship by some states. The presidential campaign of Howard Dean in 2004 in the United States was notable for its success in soliciting donation via the Internet. Many political groups use the Internet to achieve a new method of organizing in order to carry out their mission, having given rise to Internet activism. Some governments, such as those of Iran, North Korea, Myanmar, the People's Republic of China, and Saudi Arabia, restrict what people in their countries can access on the Internet, especially political and religious content. This is accomplished through software that filters domains and content so that they may not be easily accessed or obtained without elaborate circumvention.

In Norway, Denmark, Finland and Sweden, major Internet service providers have voluntarily, possibly to avoid such an arrangement being turned into law, agreed to restrict access to sites listed by authorities. While this list of forbidden URLs is only supposed to contain addresses of known child pornography sites, the content of the list is secret. Many countries, including the United States, have enacted laws against the possession or distribution of certain material, such as child pornography, via the Internet, but do not mandate filtering software. There are many free and commercially available software programs, called content-control software, with which a user can choose to block offensive websites on individual computers or networks, in order to limit a child's access to pornographic materials or depiction of violence.

The Internet has been a major outlet for leisure activity since its inception, with entertaining social experiments such as MUDs and MOOs being conducted on university servers, and humor-related Usenet groups receiving much traffic. Today, many Internet forums have sections devoted to games and funny videos; short cartoons in the form of Flash movies are also popular. Over 6 million people use blogs or message boards as a means of communication and for the sharing of ideas. The pornography and gambling industries have taken advantage of the World Wide Web, and often provide a significant source of advertising revenue for other websites. Although many governments have attempted to restrict both industries' use of the Internet, this has generally failed to stop their widespread popularity.

One main area of leisure activity on the Internet is multiplayer gaming. This form of recreation creates communities, where people of all ages and origins enjoy the fast-paced world of multiplayer games. These range from MMORPG to first-person shooters, from role-playing video games to online gambling. This has revolutionized the way many people interact while spending their free time on the Internet. While online gaming has been around since the 1970s, modern modes of online gaming began with subscription services such as GameSpy and MPlayer. Non-subscribers were limited to certain types of game play or certain games. Many people use the Internet to access and download music, movies and other works for their enjoyment and relaxation. Free and fee-based services exist for all of these activities, using centralized servers and distributed peer-to-peer technologies. Some of these sources exercise more care with respect to the original artists' copyrights than others.

Many people use the World Wide Web to access news, weather and sports reports, to plan and book vacations and to find out more about their interests. People use chat, messaging and email to make and stay in touch with friends worldwide, sometimes in the same way as some previously had pen pals. The Internet has seen a growing number of Web desktops, where users can access their files and settings via the Internet.

Cyberslacking can become a drain on corporate resources; the average UK employee spent 57 minutes a day surfing the Web while at work, according to a 2003 study by Peninsula Business Services. Internet addiction disorder is excessive computer use that interferes with daily life. Some psychologists believe that Internet use has other effects on individuals for instance interfering with the deep thinking that leads to true creativity.

Internet usage has been correlated to users' loneliness.] Lonely people tend to use the Internet as an outlet for their feelings and to share their stories with others, such as in the "I am lonely will anyone speak to me" thread.

Chapter- 3

Adaptive Bit Rate

Adaptive Bitrate Streaming (or Adaptive Streaming) is a technique used in streaming multimedia over computer networks. While in the past most video streaming technologies utilized streaming protocols such RTSP, today's adaptive streaming technologies are almost exclusively based on HTTP and designed to work efficiently over large distributed HTTP networks such as the Internet.

It works by detecting a user's bandwidth and CPU capacity in real time and adjusting the quality of a video stream accordingly. It requires the use of an encoder which can encode a single source video at multiple bit rates. The player client switches between streaming the different encodings depending on available resources. "The result: very little buffering, fast start time and a good experience for both high-end and low-end connections."

The acronym ABR can be confusing because its most common meaning is "Average bitrate," so it has not been widely adopted by video compression professionals.

Current uses

Post-production houses, content delivery networks and studios use adaptive bit rate technology in order to provide consumers with higher quality video using less manpower and fewer resources. When all is said and done, the creation of multiple video outputs, particularly for adaptive bit rate streaming, adds great value to consumers. If the technology is working as designed, the end user or consumer should be completely unaware of it. Therefore, even though media companies have been actively using adaptive bit rate technology for many years now and it has essentially become a standard practice for high-end streaming providers, mainstream consumers are relatively ignorant of its necessity.

Benefits of adaptive bit rate streaming

Consumers of streaming media experience the highest quality material when adaptive bit rate streaming is used because the user's network and playback conditions are automatically adapted to at any given time under changing conditions.

The media and entertainment industry are the main beneficiaries of adaptive bit rate streaming. As the video space grows exponentially, content delivery networks and video providers can provide customers with a superior viewing experience. Adaptive bit rate technology requires less encoding which simplifies overall workflow and creates better results.

The use of a CDN to deliver media streaming to an Internet audience is often used, as it allows scalability. The CDN received the stream from the source at its Origin server, then replicates it to many or all of its Edge cache servers. The end-user requests the stream and is redirected to the "closest" Edge server. The use of HTTP-base adaptive streaming allows the Edge server to run a simple HTTP server software, whose licence cost is cheap or free, reducing software licencing cost, compared to costly media server licences (eg. Adobe Flash Media Streaming Server). The CDN cost for HTTP streaming media is then similar to HTTP web caching CDN cost.

Implementations

Adaptive bit rate streaming was conceptualized by Move Networks and is now being developed and utilized by Adobe Systems, Apple, Microsoft and Octoshape. In September 2010, Move Networks was awarded a patent for their adaptive bit rate streaming.

Adobe Dynamic Streaming for Flash

"Dynamic streaming is the process of efficiently delivering streaming video to users by dynamically switching among different streams of varying quality and size during playback. This provides users with the best possible viewing experience their bandwidth and local computer hardware (CPU) can support. Another major goal of dynamic streaming is to make this process smooth and seamless to users, so that if up-scaling or down-scaling the quality of the stream is necessary, it is a smooth and nearly unnoticeable switch without disrupting the continuous playback."

The latest versions of Flash Player and Flash Media Server support adaptive bit-rate streaming over the traditional RTMP protocol, as well as HTTP, similar to the HTTP-based solutions from Apple and Microsoft. HTTP-based streaming has the advantage of not requiring any firewall ports being opened outside of the normal ports used by web browsers. HTTP-based streaming also allows video fragments to be cached by browsers, proxies, and CDNs, drastically reducing the load on the source server.

Apple HTTP Adaptive Streaming for iPhone/iPad

"HTTP Live Streaming is an HTTP-based media streaming communications protocol implemented by Apple Inc. as part of their QuickTime X, and iPhone software systems." Apple's newly released iPad also provides HTTP Live Streaming capabilities. It works by breaking down streams into several small HTTP-based file downloads that load simultaneously at variable adaptive rates.

HTTP Live Streaming is a standard feature in the iPhone 3.0 and newer versions.

HTTP adaptive bit rate streaming is based on HTTP progressive download, but contrary to the previous approach, here the files are very small, so that they can be compared to the streaming of packets, much like the case of using RTSP and RTP.

While all adaptive bit-rate streaming solutions are proprietary offerings as of October 2010, Apple has submitted its solution to the IETF for consideration as an Internet standard.

Microsoft Smooth Streaming for Silverlight and Windows Phone 7

Smooth Streaming is an IIS Media Services extension that enables adaptive streaming of media to clients over HTTP.] The format specification is based on the ISO Base Media File Format and standardized by Microsoft as the Protected Interoperable File Format. Microsoft is actively involved with 3GPP, MPEG and DECE organizations' efforts to standardize adaptive bit-rate HTTP streaming. Microsoft provides Smooth Streaming Client software development kits for Silverlight and Windows Phone 7. IIS Media Services 4.0, released in November 2010, introduced a feature which enables Smooth Streaming H.264/AAC videos, both live and on-demand, to be dynamically repackaged into the Apple HTTP Adaptive Streaming format and delivered to iOS devices without the need for re-encoding. Microsoft has successfully demonstrated delivery of both live and on-demand 1080p HD video with Smooth Streaming to Silverlight clients. In 2010 Microsoft also partnered with NVIDIA to demonstrate live streaming of 1080p stereoscopic 3D video to PCs equipped with NVIDIA 3D Vision technology.

Octoshape Multi-BitRate

Octoshape supports automatic mutli-bit rate streaming using standard streaming formats like Flash, Windows and HLS inputs. Octoshape uses a unique throughput optimization technology and resilient coding schemes to maximize the throughput consistency of a video stream over the Internet. Octoshape supports shifting to the appropriate bit rate of the particular consumer of the video. However, the core transport provides for a stable throughput profile over the Internet unlike TCP based technologies like HTTP or RTMP that have variable throughput profiles based on packet loss and latency. The technology selects appropriate bit rates instantly on startup, but rarely makes use of the rate shifting technology during a viewing session, giving the consumer a consistent TV quality video experience. Octoshape is also the first technology to deploy automatic multi-bit rate technology along with Multicast transport over the public Internet.

Criticisms

HTTP based adaptive bit rate technologies are significantly more operationally complex than traditional streaming technologies. Some of the documented considerations are things such as additional storage and encoding costs, and challenges with maintaining quality globally. There have also been some interesting dynamics found around the interactions between complex adaptive bit rate logic competing with complex TCP flow control logic.

Chapter- 4

Router

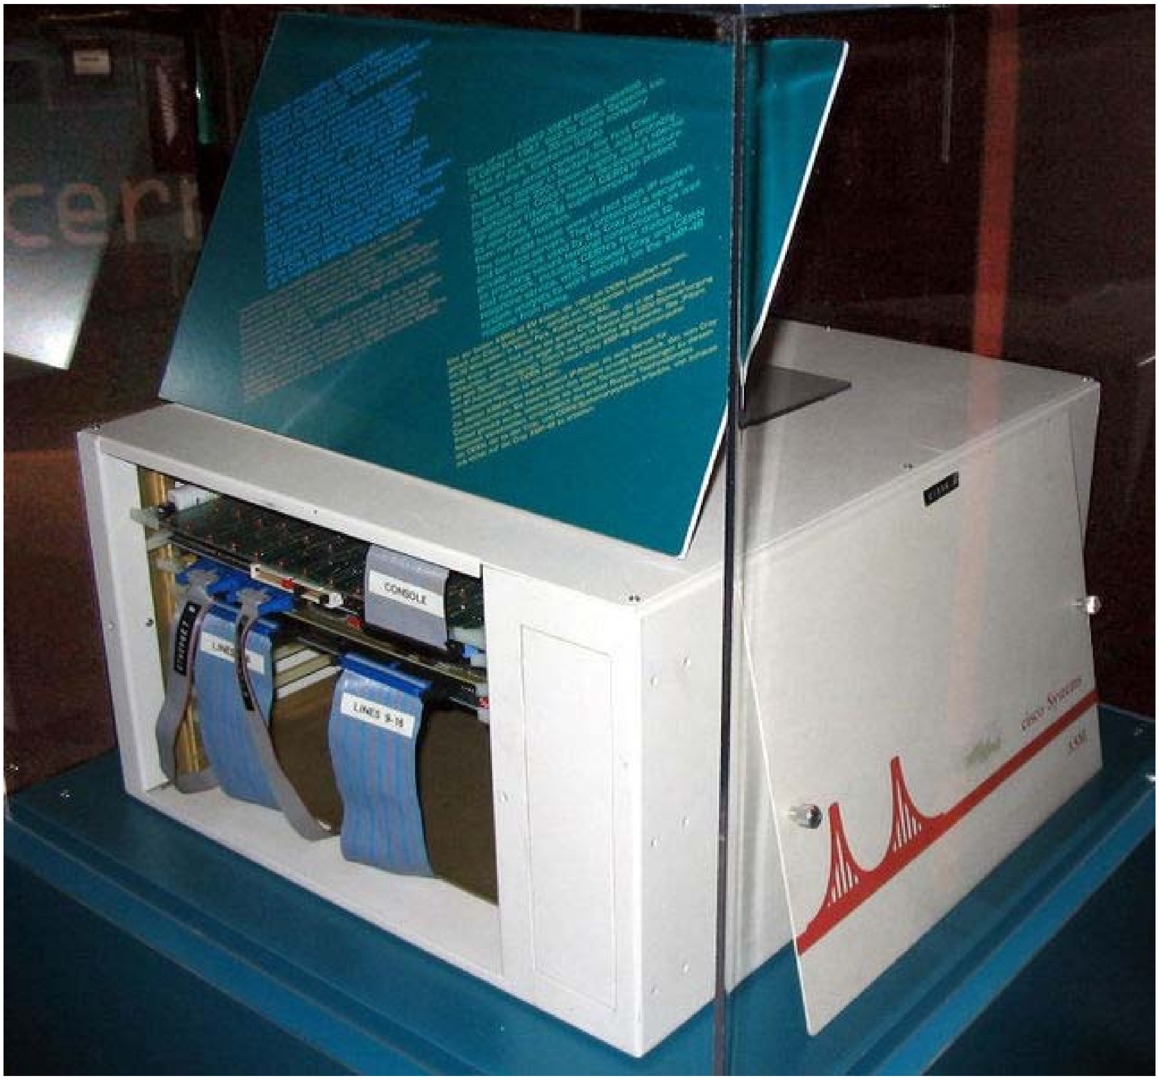

A Cisco ASM/2-32EM router deployed at CERN in 1987

Juniper SRX210 service gateway router

A **router** is a device that forwards data packets across computer networks. Routers perform the data "traffic directing" functions on the Internet. A router is a microprocessor-controlled device that is connected to two or more data lines from different networks. When a data packet comes in on one of the lines, the router reads the address information in the packet to determine its ultimate destination. Then, using information in its routing table, it directs the packet to the next network on its journey. A data packet is typically passed from router to router through the networks of the Internet until it gets to its destination computer. Routers also perform other tasks such as translating the data transmission protocol of the packet to the appropriate protocol of the next network, and preventing unauthorized access to a network by the use of a firewall.

The most familiar type of routers are home and small office routers that simply pass data, such as web pages and email, between the home computers and the owner's cable or DSL modem, which connects to the Internet (ISP). However more sophisticated routers range from enterprise routers, which connect large business or ISP networks up to the powerful core routers that forward data at high speed along the optical fiber lines of the Internet backbone.

Applications

When multiple routers are used in interconnected networks, the routers exchange information about destination addresses, using a dynamic routing protocol. Each router builds up a table listing the preferred routes between any two systems on the interconnected networks. A router has interfaces for different physical types of network connections, (such as copper cables, fiber optic, or wireless transmission). It also contains firmware for different networking protocol standards. Each network interface uses this specialized computer software to enable data packets to be forwarded from one protocol transmission system to another.

Routers may also be used to connect two or more logical groups of computer devices known as subnets, each with a different sub-network address. The subnets addresses recorded in the router do not necessarily map directly to the physical interface connections. A router has two stages of operation called planes:

- Control plane: A router records a routing table listing what route should be used to forward a data packet, and through which physical interface connection. It does this using internal pre-configured addresses, called static routes.

A typical home or small office router showing the ADSL telephone line and ETHERNET network cable connections.

- Forwarding plane: The router forwards data packets between incoming and outgoing interface connections. It routes it to the correct network type using information that the packet header contains. It uses data recorded in the routing table control plane.

Routers may provide connectivity within enterprises, between enterprises and the Internet, and between internet service providers (ISPs) networks. The largest routers (such as the Cisco CRS-1 or Juniper T1600) interconnect the various ISPs, or may be used in large enterprise networks. Smaller routers usually provide connectivity for typical home and office networks. Other networking solutions may be provided by a backbone Wireless Distribution System (WDS), which avoids the costs of introducing networking cables into buildings.

Enterprise routers

All sizes of routers may be found inside enterprises. The most powerful routers are usually found in ISPs, academic and research facilities. Large businesses may also need more powerful routers to cope with ever increasing demands of intranet data traffic. A three-layer model is in common use, not all of which need be present in smaller networks.

Access

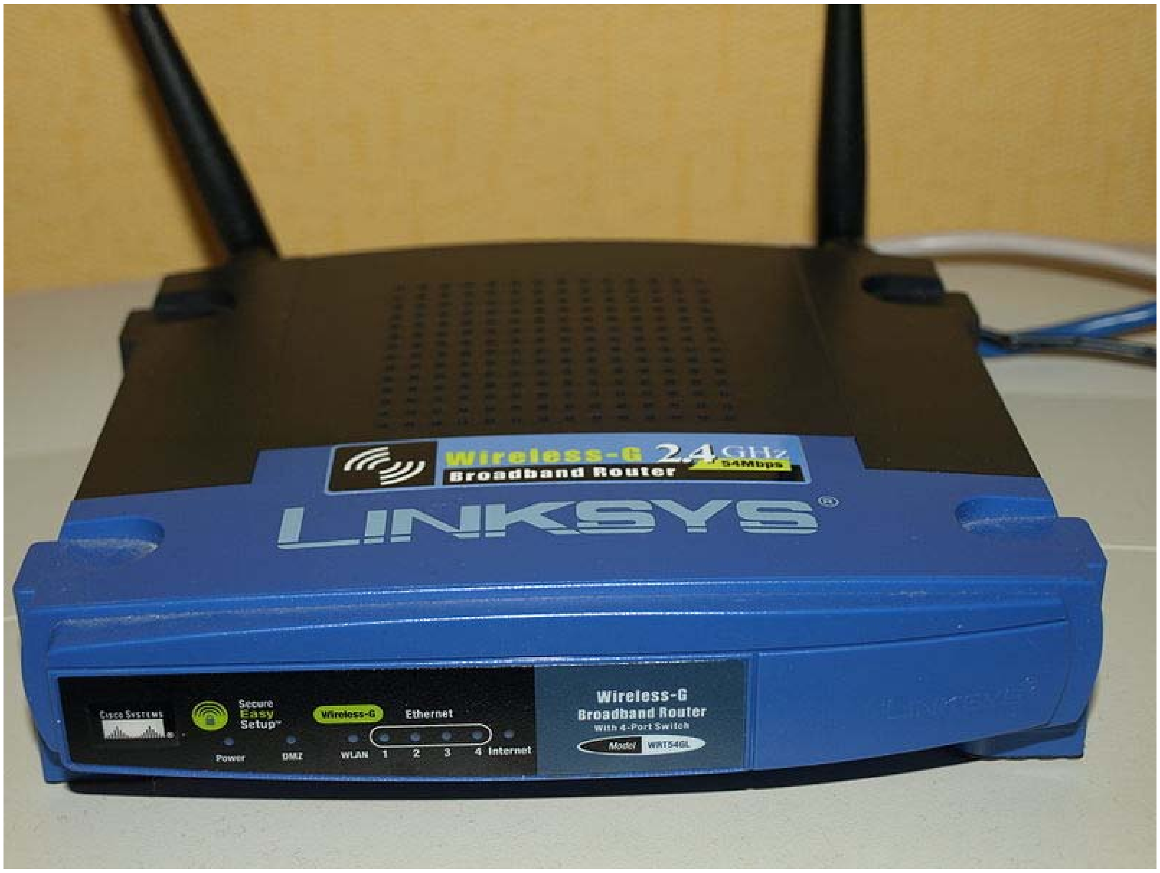

Linksys by Cisco WRT54GL SoHo Router

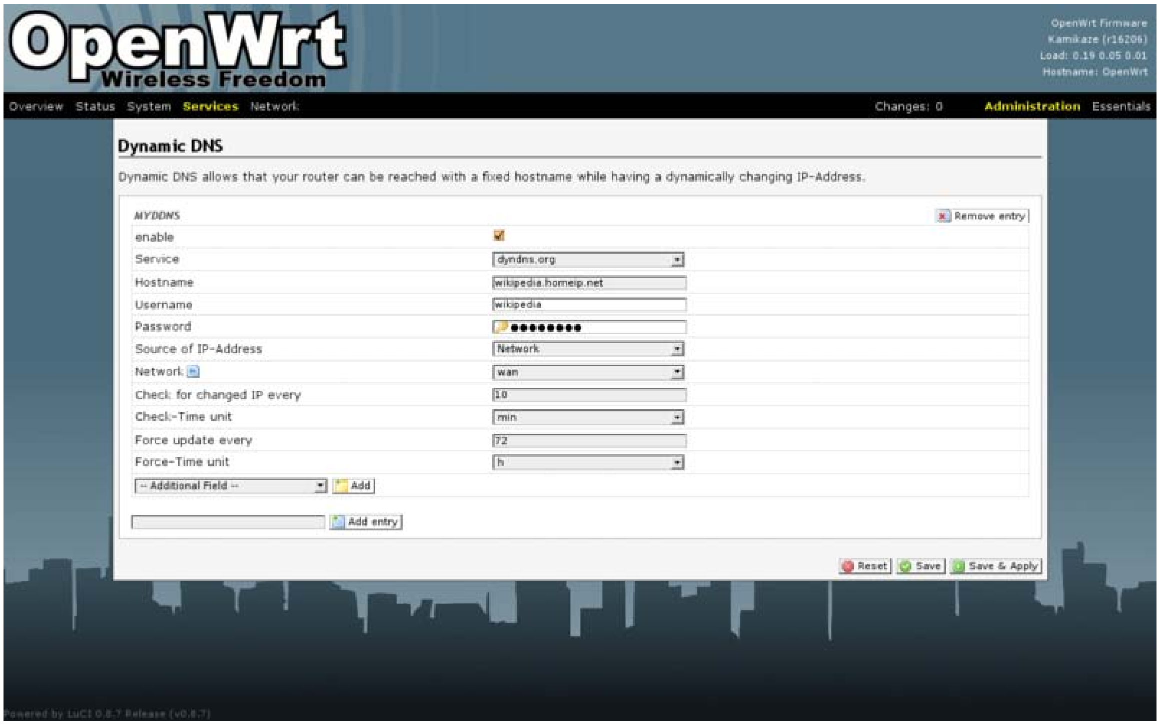

Dynamic DNS

Dynamic DNS allows that your router can be reached with a fixed hostname while having a dynamically changing IP-Address.

MYDDNS

enable	☑
Service	dyndns.org
Hostname	wikipedia.homeip.net
Username	wikipedia
Password	●●●●●●●●●
Source of IP-Address	Network
Network	wan
Check for changed IP every	10
Check-Time unit	min
Force update every	72
Force-Time unit	h

-- Additional Field --

A screenshot of the LuCI web interface used by OpenWrt. Here it is being used to configure Dynamic DNS.

Access routers, including 'small office/home office' (SOHO) models, are located at customer sites such as branch offices that do not need hierarchical routing of their own. Typically, they are optimized for low cost. Some SOHO routers are capable of running alternative free Linux-based firmwares like Tomato, OpenWrt or DD-WRT.

Distribution

Distribution routers aggregate traffic from multiple access routers, either at the same site, or to collect the data streams from multiple sites to a major enterprise location. Distribution routers are often responsible for enforcing quality of service across a WAN, so they may have considerable memory installed, multiple WAN interface connections, and substantial onboard data processing routines. They may also provide connectivity to groups of file servers or other external networks.

Security

External networks must be carefully considered as part of the overall security strategy. Separate from the router may be a firewall or VPN handling device, or the router may include these and other security functions. Many companies produced security-oriented routers, including Cisco Systems' PIX and ASA5500 series, Juniper's Netscreen, Watchguard's Firebox, Barracuda's variety of mail-oriented devices, and many others.

Core

In enterprises, a core router may provide a "collapsed backbone" interconnecting the distribution tier routers from multiple buildings of a campus, or large enterprise locations. They tend to be optimized for high bandwidth.

Internet connectivity and internal use

Routers intended for ISP and major enterprise connectivity usually exchange routing information using the Border Gateway Protocol (BGP). RFC 4098 standard defines the types of BGP-protocol routers according to the routers' functions:

- *Edge router*: Also called a Provider Edge router, is placed at the edge of an ISP network. The router uses External BGP to EBGP protocol routers in other ISPs, or a large enterprise Autonomous System.
- *Subscriber edge router*: Also called a Customer Edge router, is located at the edge of the subscriber's network, it also uses EBGP protocol to its provider's Autonomous System. It is typically used in an (enterprise) organization.
- *Inter-provider border router*: Interconnecting ISPs, is a BGP-protocol router that maintains BGP sessions with other BGP protocol routers in ISP Autonomous Systems.
- Core router: A *core router* resides within an Autonomous System as a back bone to carry traffic between edge routers.]

  Within an ISP: In the ISPs Autonomous System, a router uses internal BGP protocol to communicate with other ISP edge routers, other intranet core routers, or the ISPs intranet provider border routers.
  "Internet backbone:" The Internet no longer has a clearly identifiable backbone, unlike its predecessor networks. The major ISPs system routers make up what could be considered to be the current Internet backbone core. ISPs operate all four types of the BGP-protocol routers described here. An ISP "core" router is used to interconnect its edge and border routers. Core routers may also have specialized functions in virtual private networks based on a combination of BGP and Multi-Protocol Label Switching protocols.

- Port forwarding: Routers are also used for port forwarding between private internet connected servers.

- Voice/Data/Fax/Video Processing Routers: Commonly referred to as access servers or gateways, these devices are used to route and process voice, data, video, and fax traffic on the internet. Since 2005, most long-distance phone calls have been processed as IP traffic (VOIP) through a voice gateway,. Voice traffic that the traditional cable networks once carried. Use of access server type routers expanded with the advent of the internet, first with dial-up access, and another resurgence with voice phone service.

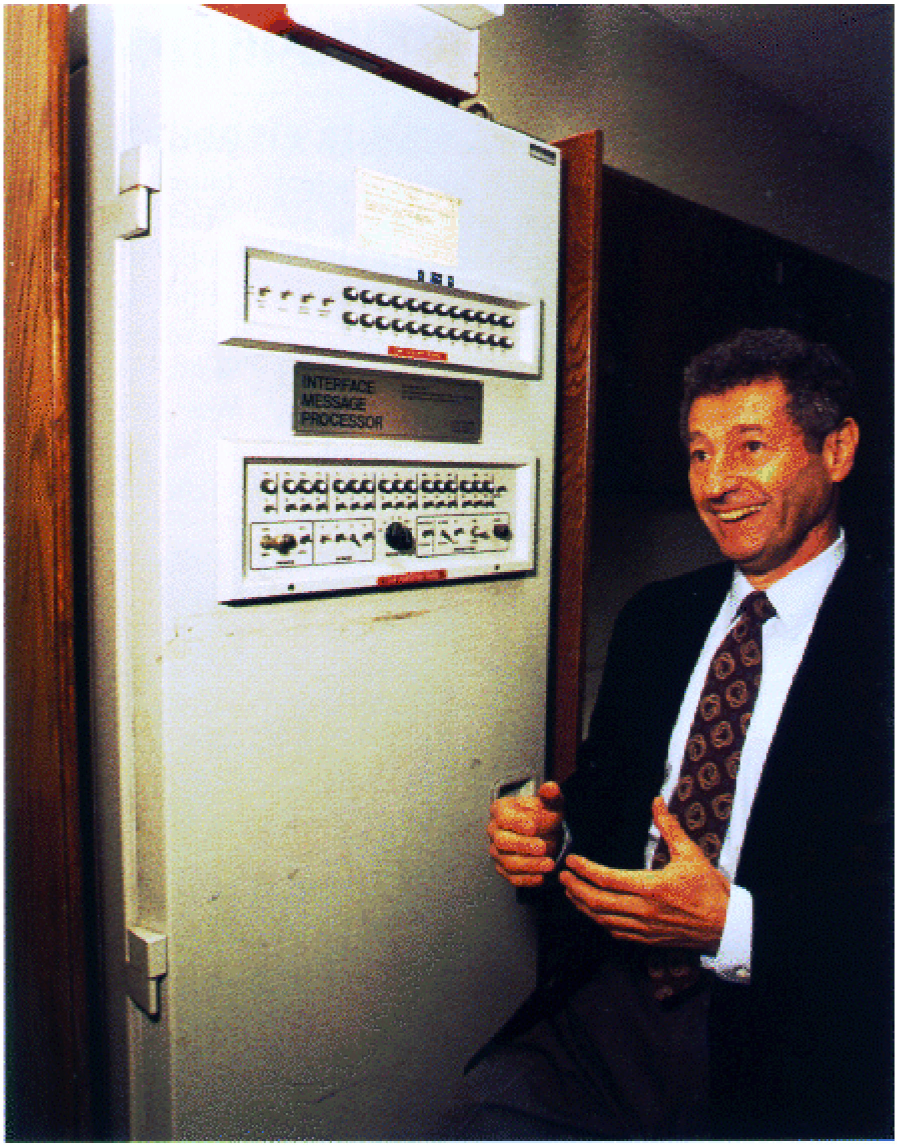

Leonard Kleinrock and the first IMP.

Avaya ERS 8600 (2010)

The very first device that had fundamentally the same functionality as a router does today, was the Interface Message Processor (IMP); IMPs were the devices that made up the ARPANET, the first packet network. The idea for a router (called "gateways" at the time) initially came about through an international group of computer networking researchers called the International Network Working Group (INWG). Set up in 1972 as an informal group to consider the technical issues involved in connecting different networks, later that year it became a subcommittee of the International Federation for Information Processing.

These devices were different from most previous packet networks in two ways. First, they connected dissimilar kinds of networks, such as serial lines and local area networks. Second, they were connectionless devices, which had no role in assuring that traffic was delivered reliably, leaving that entirely to the hosts (this particular idea had been previously pioneered in the CYCLADES network).

The idea was explored in more detail, with the intention to produce a prototype system, as part of two contemporaneous programs. One was the initial DARPA-initiated program, which created the TCP/IP architecture in use today. The other was a program at Xerox

PARC to explore new networking technologies, which produced the PARC Universal Packet system, due to corporate intellectual property concerns it received little attention outside Xerox for years.

Some time after early 1974 the first Xerox routers became operational. The first true IP router was developed by Virginia Strazisar at BBN, as part of that DARPA-initiated effort, during 1975-1976. By the end of 1976, three PDP-11-based routers were in service in the experimental prototype Internet.

The first multiprotocol routers were independently created by staff researchers at MIT and Stanford in 1981; the Stanford router was done by William Yeager, and the MIT one by Noel Chiappa; both were also based on PDP-11s.

Virtually all networking now uses TCP/IP, but multiprotocol routers are still manufactured. They were important in the early stages of the growth of computer networking, when protocols other than TCP/IP were in use. Modern Internet routers that handle both IPv4 and IPv6 are multiprotocol, but are simpler devices than routers processing AppleTalk, DECnet, IP, and Xerox protocols.

From the mid-1970s and in the 1980s, general-purpose mini-computers served as routers. Modern high-speed routers are highly specialized computers with extra hardware added to speed both common routing functions, such as packet forwarding, and specialised functions such as IPsec encryption.

There is substantial use of Linux and Unix software based machines, running open source routing code, for research and other applications. Cisco's operating system was independently designed. Major router operating systems, such as those from Juniper Networks and Extreme Networks, are extensively modified versions of Unix software.

Forwarding

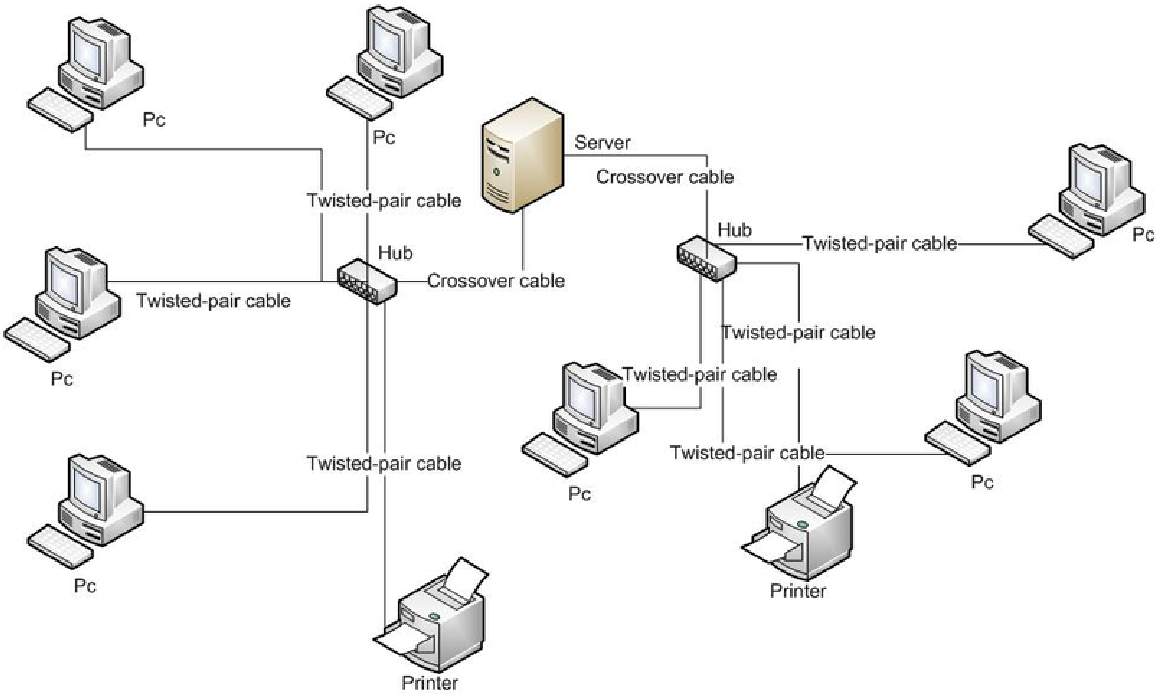

Computer network

For pure Internet Protocol (IP) forwarding function, a router is designed to minimize the state information associated with individual packets. The main purpose of a router is to connect multiple networks and forward packets destined either for its own networks or other networks. A router is considered a Layer 3 device because its primary forwarding decision is based on the information in the Layer 3 IP packet, specifically the destination IP address. This process is known as routing. When each router receives a packet, it searches its routing table to find the best match between the destination IP address of the packet and one of the network addresses in the routing table. Once a match is found, the packet is encapsulated in the Layer 2 data link frame for that outgoing interface. A router does not look into the actual data contents that the packet carries, but only at the layer 3 addresses to make a forwarding decision, plus optionally other information in the header for hint on, for example, QoS. Once a packet is forwarded, the router does not retain any historical information about the packet, but the forwarding action can be collected into the statistical data, if so configured.

Forwarding decisions can involve decisions at layers other than layer 3. A function that forwards based on layer 2 information, is properly called a bridge. This function is referred to as layer 2 bridging, as the addresses it uses to forward the traffic are layer 2 addresses (e.g. MAC addresses on Ethernet).

Besides making decision as which interface a packet is forwarded to, which is handled primarily via the routing table, a router also has to manage congestion, when packets arrive at a rate higher than the router can process. Three policies commonly used in the

Internet are tail drop, random early detection (RED), and weighted random early detection (WRED). Tail drop is the simplest and most easily implemented; the router simply drops packets once the length of the queue exceeds the size of the buffers in the router. RED probabilistically drops datagrams early when the queue is exceeds a pre-configured size of the queue until a pre-configured max when it becomes tail drop. WRED requires a weight on the average queue size to act upon when the traffic is about to exceed the pre-configured size, so that short bursts will not trigger random drops.

Another function a router performs is to decide which packet should be processed first when multiple queues exist. This is managed through quality of service (QoS), which is critical when Voice over IP is deployed, so that delays between packets do not exceed 150ms to maintain the quality of voice conversations.

Yet another function a router performs is called policy-based routing where special rules are constructed to override the rules derived from the routing table when a packet forwarding decision is made.

These functions may be performed through the same internal paths that the packets travel inside the router. Some of the functions may be performed through an application-specific integrated circuit (ASIC) to avoid overhead caused by multiple CPU cycles, and others may have to be performed through the CPU as these packets need special attention that cannot be handled by an ASIC.

Chapter- 5

Bridging (Networking)

Bridging is a forwarding technique used in packet-switched computer networks. Unlike routing, bridging makes no assumptions about where in a network a particular address is located. Instead, it depends on flooding and examination of source addresses in received packet headers to locate unknown devices. Once a device has been located, its location is recorded in a table where the MAC address is stored so as to preclude the need for further broadcasting. The utility of bridging is limited by its dependence on flooding, and is thus only used in local area networks.

Bridging generally refers to *Transparent bridging* or *Learning bridge* operation which predominates in Ethernet. Another form of bridging, Source route bridging, was developed for token ring networks.

A **network bridge** connects multiple network segments at the data link layer (Layer 2) of the OSI model. In Ethernet networks, the term *bridge* formally means a device that behaves according to the IEEE 802.1D standard. A bridge and switch are very much alike; a switch being a bridge with numerous ports. *Switch* or *Layer 2 switch* is often used interchangeably with *bridge*.

Bridges are similar to repeaters or network hubs, devices that connect network segments at the physical layer (Layer 1) of the OSI model; however, with bridging, traffic from one network is managed rather than simply rebroadcasted to adjacent network segments. Bridges are more complex than hubs or repeaters. Bridges can analyze incoming data packets to determine if the bridge is able to send the given packet to another segment of the network.

Transparent bridging operation

A bridge uses a forwarding database to send frames across network segments. The forwarding database is initially empty and entries in the database are built as the bridge receives frames. If an address entry is not found in the forwarding database, the frame is flooded to all other ports of the bridge, forwarding the frame to all segments except the source address. By means of these broadcast frames, the destination network will respond and forwarding database entry will be created.

As an example, consider three hosts, A, B and C and a bridge. The bridge has three ports. A is connected to bridge port 1, B is connected bridge port 2, C is connected to bridge port 3. A sends a frame addressed to B to the bridge. The bridge examines the source address of the frame and creates an address and port number entry for A in its forwarding table. The bridge examines the destination address of the frame and does not find it in its forwarding table so it floods it to all other ports: 2 and 3. The frame is received by hosts B and C. Host C examines the destination address and ignores the frame. Host B recognizes a destination address match and generates a response to A. On the return path, the bridge adds an address and port number entry for B to its forwarding table. The bridge already has A's address in its forwarding table so it forwards the response only to port 1. Host C or any other hosts on port 3 are not burdened with the response. Two-way communication is now possible between A and B without any further flooding.

Note that both source and destination addresses are used in this algorithm. Source addresses are recorded in entries in the table, while destination addresses are looked up in the table and matched to the proper segment to send the frame to.

The technology was originally developed by the Digital Equipment Corp. in the 1980s.

Filtering database

To translate between two segments, a bridge reads a frame's destination MAC address and decides to either forward or filter. If the bridge determines that the destination node is on another segment on the network, it forwards (retransmits) the packet to that segment. If the destination address belongs to the same segment as the source address, the bridge filters (discards) the frame. As nodes transmit data through the bridge, the bridge establishes a filtering database (also known as a forwarding table) of known MAC addresses and their locations on the network. The bridge uses its filtering database to determine whether a packet should be forwarded or filtered.

Destination lookup failure

Layer 2 (L2) Ethernet Switch is looking at the MAC Destination address of the Ethernet frame in order to switch it to the appropriate port/s. In case that MAC address exists in the Switch L2 Table, it transmits the Frame only to the port which is tied to that entry. In case that MAC address doesn’t exist in the Switch L2 Table, the frame is considered DLF and it been transmitted to all forwarding ports of that VLAN. (Also Broadcasts such as ARP Request messages are transmitted to the same ports)

Advantages of network bridges

- Self-configuring
- Simple bridges are inexpensive
- Isolate collision domain
- Reduce the size of collision domain by microsegmentation in non-switched networks

- Transparent to protocols above the MAC layer
- Allows the introduction of management/performance information and access control
- LANs interconnected are separate, and physical constraints such as number of stations, repeaters and segment length don't apply
- Helps minimize bandwidth usage

Disadvantages of network bridges

- Does not limit the scope of broadcasts [broadcast domain cannot be controlled]
- Does not scale to extremely large networks
- Buffering and processing introduces delays
- Bridges are more expensive than repeaters or hubs
- A complex network topology can pose a problem for transparent bridges. For example, multiple paths between transparent bridges and LANs can result in *bridge loops*. The spanning tree protocol helps to reduce problems with complex topologies.

Bridging versus routing

Bridging and routing are both ways of performing data control, but work through different methods. Bridging takes place at OSI Model Layer 2 (data-link layer) while routing takes place at the OSI Model Layer 3 (network layer). This difference means that a bridge directs frames according to hardware assigned MAC addresses while a router makes its decisions according to arbitrarily assigned IP Addresses. As a result of this, bridges are not concerned with and are unable to distinguish networks while routers can.

When designing a network, one can choose to put multiple segments into one bridged network or to divide it into different networks interconnected by routers. If a host is physically moved from one network area to another in a routed network, it has to get a new IP address; if this system is moved within a bridged network, it doesn't have to reconfigure anything.

Chapter- 6

Network Switch

Typical SOHO network switch.

Back view of Atlantis network switch with Ethernet ports.

A **network switch** or **switching hub** is a computer networking device that connects network segments.

The term commonly refers to a multi-port network bridge that processes and routes data at the data link layer (layer 2) of the OSI model. Switches that additionally process data at the network layer (Layer 3) and above are often referred to as Layer 3 switches or multilayer switches.

The term **network switch** does not generally encompass unintelligent or passive network devices such as hubs and repeaters.

The first Ethernet switch was introduced by Kalpana in 1990.

Function

The network switch plays an integral part in most modern Ethernet local area networks (LANs). Mid-to-large sized LANs contain a number of linked managed switches. Small office/home office (SOHO) applications typically use a single switch, or an all-purpose converged device such as a gateway to access small office/home broadband services such as DSL or cable internet. In most of these cases, the end-user device contains a router and components that interface to the particular physical broadband technology. User devices may also include a telephone interface for VoIP.

An Ethernet switch operates at the data link layer of the OSI model to create a separate collision domain for each switch port. With 4 computers (e.g., A, B, C, and D) on 4 switch ports, A and B can transfer data back and forth, while C and D also do so simultaneously, and the two conversations will not interfere with one another. In the case of a hub, they would all share the bandwidth and run in half duplex, resulting in collisions, which would then necessitate retransmissions. Using a switch is called microsegmentation. This allows computers to have dedicated bandwidth on a point-to-point connections to the network and to therefore run in Full duplex without collisions.

Role of switches in networks

Switches may operate at one or more layers of the OSI model, including data link, network, or transport (i.e., end-to-end). A device that operates simultaneously at more than one of these layers is known as a multilayer switch.

In switches intended for commercial use, built-in or modular interfaces make it possible to connect different types of networks, including Ethernet, Fibre Channel, ATM, ITU-T G.hn and 802.11. This connectivity can be at any of the layers mentioned. While Layer 2 functionality is adequate for bandwidth-shifting within one technology, interconnecting technologies such as Ethernet and token ring are easier at Layer 3.

Interconnection of different Layer 3 networks is done by routers. If there are any features that characterize "Layer-3 switches" as opposed to general-purpose routers, it tends to be that they are optimized, in larger switches, for high-density Ethernet connectivity.

In some service provider and other environments where there is a need for a great deal of analysis of network performance and security, switches may be connected between WAN routers as places for analytic modules. Some vendors provide firewall, network intrusion detection, and performance analysis modules that can plug into switch ports. Some of these functions may be on combined modules.

In other cases, the switch is used to create a mirror image of data that can go to an external device. Since most switch port mirroring provides only one mirrored stream, network hubs can be useful for fanning out data to several read-only analyzers, such as intrusion detection systems and packet sniffers.

Layer-specific functionality

A modular network switch with three network modules (a total of 24 Ethernet and 14 Fast Ethernet ports) and one power supply.

While switches may learn about topologies at many layers, and forward at one or more layers, they do tend to have common features. Other than for high-performance applications, modern commercial switches use primarily Ethernet interfaces, which can have different input and output bandwidths of 10, 100, 1000 or 10,000 megabits per second.

At any layer, a modern switch may implement power over Ethernet (PoE), which avoids the need for attached devices, such as an VoIP phone or wireless access point, to have a separate power supply. Since switches can have redundant power circuits connected to uninterruptible power supplies, the connected device can continue operating even when regular office power fails.

Layer 1 hubs versus higher-layer switches

A network hub, or repeater, is a simple network device. Hubs do not manage any of the traffic that comes through them. Any packet entering a port is broadcast out or "repeated" on every other port, except for the port of entry. Since every packet is repeated on every other port, packet collisions result, which slows down the network.

There are specialized applications where a hub can be useful, such as copying traffic to multiple network sensors. High end switches have a feature which does the same thing called port mirroring.

There is no longer any significant price difference between a hub and a low-end switch.

Layer 2

A network bridge, operating at the data link layer, may interconnect a small number of devices in a home or the office. This is a trivial case of bridging, in which the bridge learns the MAC address of each connected device.

Single bridges also can provide extremely high performance in specialized applications such as storage area networks.

Classic bridges may also interconnect using a spanning tree protocol that disables links so that the resulting local area network is a tree without loops. In contrast to routers, spanning tree bridges must have topologies with only one active path between two points. The older IEEE 802.1D spanning tree protocol could be quite slow, with forwarding stopping for 30 seconds while the spanning tree would reconverge. A Rapid Spanning Tree Protocol was introduced as IEEE 802.1w, but the newest edition of IEEE 802.1D adopts the 802.1w extensions as the base standard.

The IETF is specifying the TRILL protocol, which is the application of link-state routing technology to the layer-2 bridging problem. Devices which implement TRILL, called RBridges, combine the best features of both routers and bridges.

While "layer 2 switch" remains more of a marketing term than a technical term, the products that were introduced as "switches" tended to use microsegmentation and Full duplex to prevent collisions among devices connected to Ethernet. By using an internal forwarding plane much faster than any interface, they give the impression of simultaneous paths among multiple devices.

Once a bridge learns the topology through a spanning tree protocol, it forwards data link layer frames using a layer 2 forwarding method. There are four forwarding methods a bridge can use, of which the second through fourth method were performance-increasing methods when used on "switch" products with the same input and output port bandwidths:

1. Store and forward: The switch buffers and verifies each frame before forwarding it.
2. Cut through: The switch reads only up to the frame's hardware address before starting to forward it. Cut-through switches have to fall back to store and forward if the outgoing port is busy at the time the packet arrives. There is no error checking with this method.
3. Fragment free: A method that attempts to retain the benefits of both store and forward and cut through. Fragment free checks the first 64 bytes of the frame, where addressing information is stored. According to Ethernet specifications, collisions should be detected during the first 64 bytes of the frame, so frames that are in error because of a collision will not be forwarded. This way the frame will always reach its intended destination. Error checking of the actual data in the packet is left for the end device.
4. Adaptive switching: A method of automatically selecting between the other three modes.

While there are specialized applications, such as storage area networks, where the input and output interfaces are the same bandwidth, this is rarely the case in general LAN applications. In LANs, a switch used for end user access typically concentrates lower bandwidth (e.g., 10/100 Mbit/s) into a higher bandwidth (at least 1 Gbit/s). Alternatively, a switch that provides access to server ports usually connects to them at a much higher bandwidth than is used by end user devices.

Layer 3

Within the confines of the Ethernet physical layer, a layer 3 switch can perform some or all of the functions normally performed by a router. The most common layer-3 capability is awareness of IP multicast through IGMP snooping. With this awareness, a layer-3 switch can increase efficiency by delivering the traffic of a multicast group only to ports where the attached device has signaled that it wants to listen to that group.

Layer 4

While the exact meaning of the term Layer-4 switch is vendor-dependent, it almost always starts with a capability for network address translation, but then adds some type of load distribution based on TCP sessions.

The device may include a stateful firewall, a VPN concentrator, or be an IPSec security gateway.

Layer 7

Layer 7 switches may distribute loads based on URL or by some installation-specific technique to recognize application-level transactions. A Layer-7 switch may include a web cache and participate in a content delivery network.

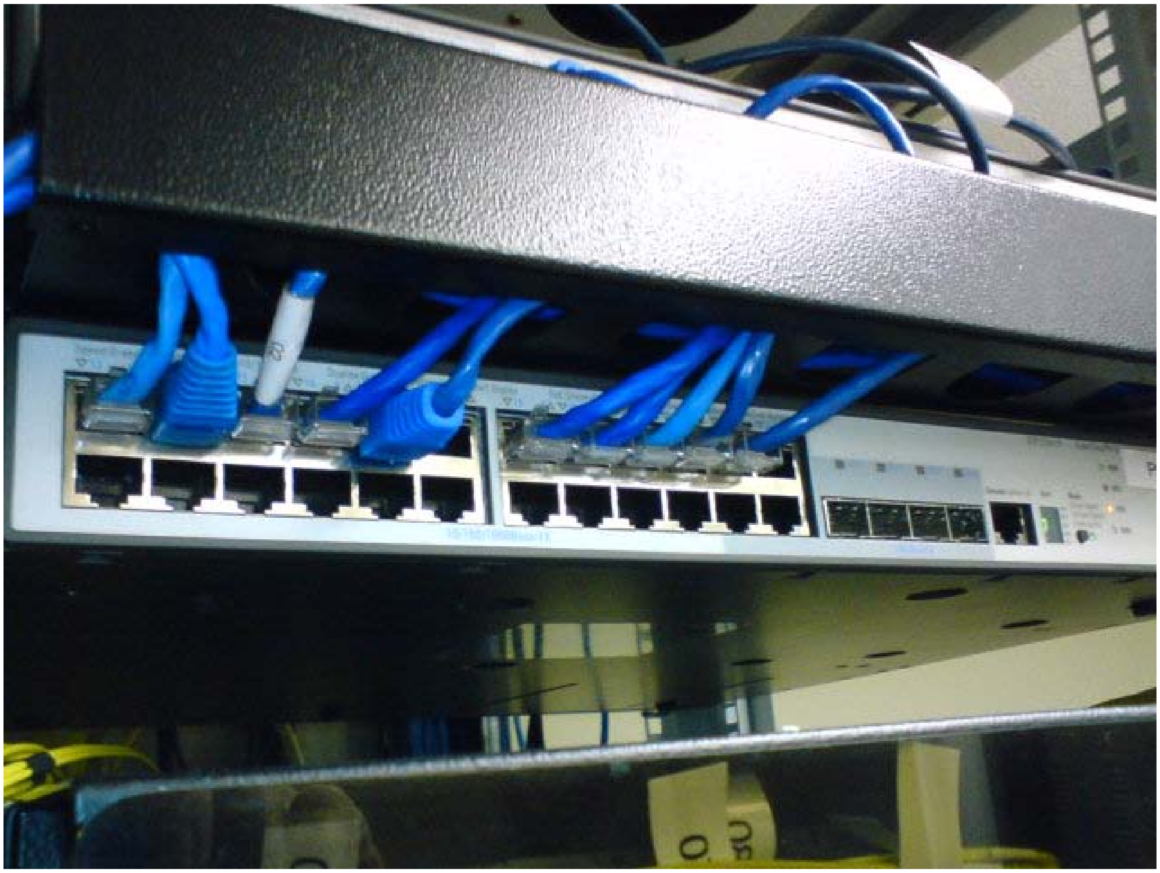

Rack-mounted 24-port 3Com switch

Types of switches

Form factor

- Desktop, not mounted in an enclosure, typically intended to be used in a home or office environment outside of a wiring closet
- Rack mounted
- Chassis — with swappable "switch module" cards. e.g. Alcatel's OmniSwitch 9000; Cisco Catalyst switch 4500 and 6500; 3Com 7700, 7900E, 8800.
- DIN rail mounted, normally seen in industrial environments or panels

Configuration options

- *Unmanaged* switches — These switches have no configuration interface or options. They are plug and play. They are typically the least expensive switches, found in home, SOHO, or small businesses. They can be desktop or rack mounted.
- *Managed* switches — These switches have one or more methods to modify the operation of the switch. Common management methods include: a command-line interface (CLI) accessed via serial console, telnet or Secure Shell, an embedded

Simple Network Management Protocol (SNMP) agent allowing management from a remote console or management station, or a web interface for management from a web browser. Examples of configuration changes that one can do from a managed switch include: enable features such as Spanning Tree Protocol, set port bandwidth, create or modify Virtual LANs (VLANs), etc. Two sub-classes of managed switches are marketed today:

  - *Smart* (or intelligent) switches — These are managed switches with a limited set of management features. Likewise "web-managed" switches are switches which fall in a market niche between unmanaged and managed. For a price much lower than a fully managed switch they provide a web interface (and usually no CLI access) and allow configuration of basic settings, such as VLANs, port-bandwidth and duplex.
  - *Enterprise Managed* (or fully managed) switches — These have a full set of management features, including CLI, SNMP agent, and web interface. They may have additional features to manipulate configurations, such as the ability to display, modify, backup and restore configurations. Compared with smart switches, enterprise switches have more features that can be customized or optimized, and are generally more expensive than smart switches. Enterprise switches are typically found in networks with larger number of switches and connections, where centralized management is a significant savings in administrative time and effort. A stackable switch is a version of enterprise-managed switch.

Traffic monitoring on a switched network

Unless port mirroring or other methods such as RMON or SMON are implemented in a switch,[1] it is difficult to monitor traffic that is bridged using a switch because only the sending and receiving ports can see the traffic. These monitoring features are rarely present on consumer-grade switches.

Two popular methods that are specifically designed to allow a network analyst to monitor traffic are:

- Port mirroring — the switch sends a copy of network packets to a monitoring network connection.
- SMON — "Switch Monitoring" is described by RFC 2613 and is a protocol for controlling facilities such as port mirroring.

Another method to monitor may be to connect a Layer-1 hub between the monitored device and its switch port. This will induce minor delay, but will provide multiple interfaces that can be used to monitor the individual switch port.

Typical switch management features

Linksys 48-port switch

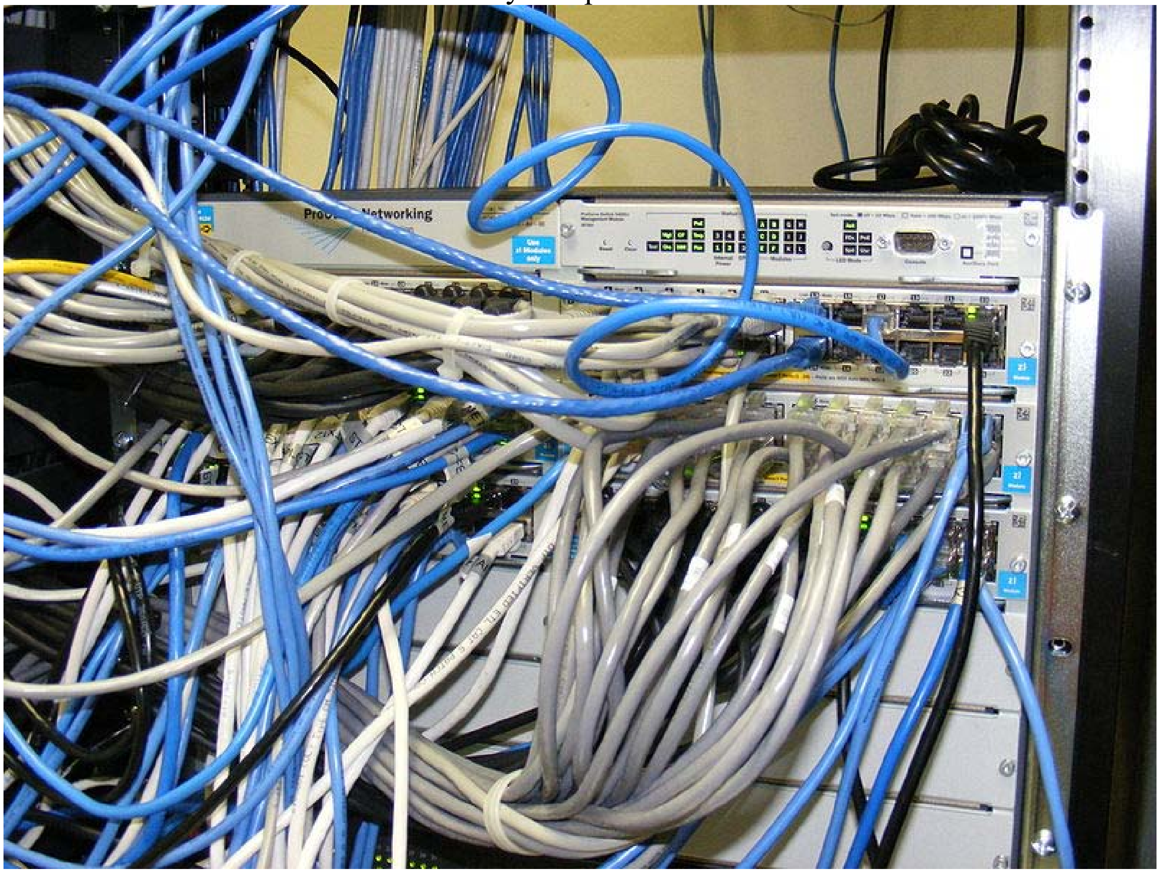

HP Procurve rack-mounted switches Mounted in a standard 19inch Telco Rack 19-inch rack with network cables

- Turn particular port range on or off
- Link bandwidth and duplex settings
- Priority settings for ports
- MAC filtering and other types of "port security" features which prevent MAC flooding
- Use of Spanning Tree Protocol
- SNMP monitoring of device and link health
- Port mirroring (also known as: port monitoring, spanning port, SPAN port, roving analysis port or link mode port)
- Link aggregation (also known as *bonding*, *trunking* or *teaming*)
- VLAN settings
- 802.1X network access control
- IGMP snooping

Link aggregation allows the use of multiple ports for the same connection achieving higher data transfer rates. Creating VLANs can serve security and performance goals by reducing the size of the broadcast domain.

Chapter- 7

Network Interface Controller & Ethernet Hub

Network Interface Controller

Network Interface Card (NIC)

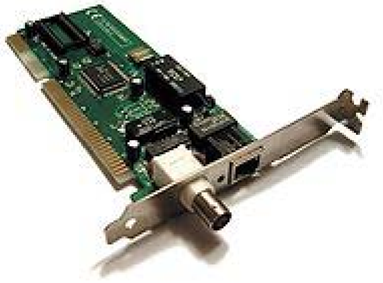

A 1990s Ethernet network interface controller card which connects to the motherboard via the now-obsolete ISA bus. This combination card features both a (now obsolete) bayonet cap BNC connector (left) for use in coaxial-based 10base2 networks and an RJ-45 connector (right) for use in twisted pair-based 10baseT networks. (The ports could not be used simultaneously.)

Connects to	Motherboard via one of: • Integrated • PCI Connector • ISA Connector • PCI-E • FireWire • USB Network via one of: • Fast Ethernet • Gigabit Ethernet

	- Optical fiber - Token ring
Speeds	10 Mbit/s 100 Mbit/s 1000 Mbit/s up to 160 Gbit/s
Common manufacturers	Novell Intel Realtek Others

A **network interface controller** (also known as a **network interface card**, **network adapter**, **LAN adapter** and by similar terms) is a computer hardware component that connects a computer to a computer network.

Whereas network interface controllers were commonly implemented on expansion cards that plug into a computer bus, the low cost and ubiquity of the Ethernet standard means that most newer computers have a network interface built into the motherboard.

Purpose

The network controller implements the electronic circuitry required to communicate using a specific physical layer and data link layer standard such as Ethernet, Wi-Fi, or Token Ring. This provides a base for a full network protocol stack, allowing communication among small groups of computers on the same LAN and large-scale network communications through routable protocols, such as IP.

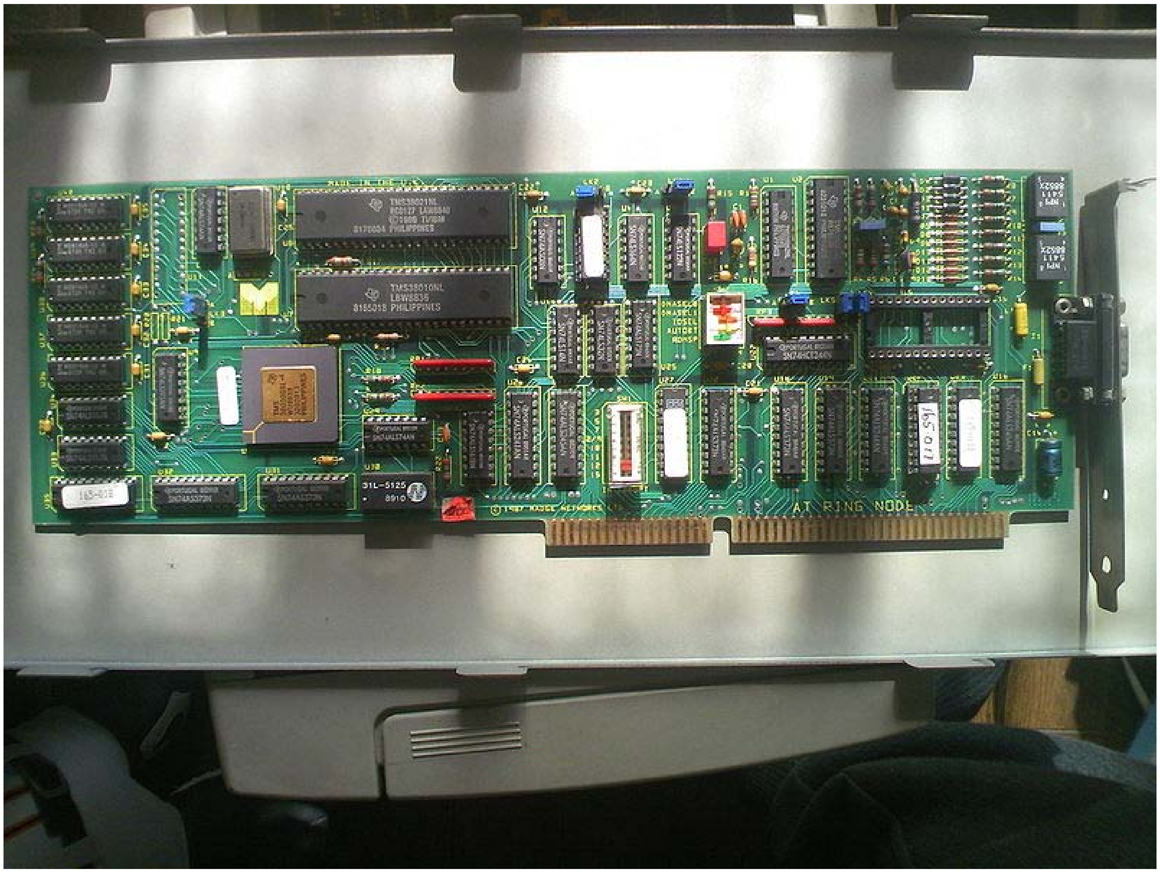

Madge 4/16Mbps TokenRing ISA NIC

Although other network technologies exist (e.g. token ring), Ethernet has achieved near-ubiquity since the mid-1990s.

Every Ethernet network controller has a unique 48-bit serial number called a MAC address, which is stored in read-only memory carried on the card for add-on cards. Every computer on an Ethernet network must have at least one controller. Each controller must have a unique MAC address. Normally it is safe to assume that no two network controllers will share the same address, because controller vendors purchase blocks of addresses from the Institute of Electrical and Electronics Engineers (IEEE) and assign a unique address to each controller at the time of manufacture.

The NIC allows computers to communicate over a computer network. It is both an OSI layer 1 (physical layer) and layer 2 (data link layer) device, as it provides physical access to a networking medium and provides a low-level addressing system through the use of MAC addresses. It allows users to connect to each other either by using cables or wirelessly.

Implementation

Whereas network controllers used to be expansion cards that plugged into a computer bus, the low cost and ubiquity of the Ethernet standard means that most newer computers have a network interface built into the motherboard. Newer server motherboards may even have dual network interfaces built-in. The Ethernet capabilities are either integrated into the motherboard chipset or implemented via a low cost dedicated Ethernet chip, connected through the PCI (or the newer PCI express) bus. A separate network card is not required unless additional interfaces are needed or some other type of network is used.

There are four techniques used to transfer data, the NIC may use one or more of these techniques.

- Polling is where the CPU examines the status of the peripheral under program control.
- Programmed I/O is where the microprocessor alerts the designated peripheral by applying its address to the system's address bus.
- Interrupt-driven I/O is where the peripheral alerts the microprocessor that it is ready to transfer data.
- Direct memory access is where an intelligent peripheral assumes control of the system bus to access memory directly. This removes load from the CPU but requires a separate processor on the card.

An Ethernet network controller typically has a RJ45 socket where the network cable is connected. Older NICs also supplied BNC, or AUI connections. A few LEDs inform the user of whether the network is active, and whether or not there is data being transmitted on it. Ethernet network controllers typically support 10 Mbit/s Ethernet, 100 Mbit/s Ethernet, and 1000 Mbit/s Ethernet varieties. Such controllers are designated *10/100/1000* and this means they can support a notional maximum transfer rate of 10, 100 or 1000 Megabits per second.

Ethernet Hub

4-port Ethernet hub

An **Ethernet hub**, **active hub**, **network hub**, **repeater hub** or **hub** is a device for connecting multiple twisted pair or fiber optic Ethernet devices together and making them act as a single network segment. Hubs work at the physical layer (layer 1) of the OSI model. The device is a form of multiport repeater. Repeater hubs also participate in collision detection, forwarding a jam signal to all ports if it detects a collision.

Hubs also often come with a BNC and/or Attachment Unit Interface (AUI) connector to allow connection to legacy 10BASE2 or 10BASE5 network segments. The availability of low-priced network switches has largely rendered hubs obsolete but they are still seen in older installations and more specialized applications.

Technical information

A network hub is a fairly unsophisticated broadcast device. Hubs do not manage any of the traffic that comes through them, and any packet entering any port is regenerated and broadcast out on all other ports. Since every packet is being sent out through all other ports, packet collisions result—which greatly impedes the smooth flow of traffic.

100 Mbit/s hubs/repeater come in two different speed grades: Class I delay the signal for a maximum of 140 bit times (enabling translation between 100Base-TX, 100Base-FX and 100Base-T4) and Class II ones delay the signal for a maximum of 92 bit times (enabling installation of two hubs in a single collision domain).

The need for hosts to be able to detect collisions limits the number of hubs and the total size of a network built using hubs (a network built using switches does not have these limitations). For 10 Mbit/s networks built using repeater hubs, the 5-4-3 rule must be followed: up to 5 segments (4 hubs) are allowed between any two end stations. For 100 Mbit/s networks, the limit is reduced to 3 segments (2 hubs) between any two end stations, and even that is only allowed if the hubs are of Class II. Some hubs have manufacturer specific stack ports allowing them to be combined in a way that allows more hubs than simple chaining through Ethernet cables, but even so, a large fast Ethernet network is likely to require switches to avoid the chaining limits of hubs.

Most hubs detect typical problems, such as excessive collisions and jabbering on individual ports, and *partition* the port, disconnecting it from the shared medium. Thus, hub-based Ethernet is generally more robust than coaxial cable-based Ethernet (e.g. 10BASE2), where a misbehaving device can adversely affect the entire collision domain. Even if not partitioned automatically, a hub simplifies troubleshooting because they remove the need to troubleshoot faults on a long cable with multiple taps; status lights on the hub can indicate the possible problem source or, as a last resort, devices can be disconnected from a hub one at a time much more easily than from a coaxial cable.

Hubs are classified as Layer 1 (physical layer) devices in the OSI model. At the physical layer, hubs support little in the way of sophisticated networking. Hubs do not read any of the data passing through them and are not aware of their source or destination. A hub simply receives incoming Ethernet frames, regenerates the electrical signal, and broadcasts these packets out to all other devices on the network.

To pass data through the repeater in a usable fashion from one segment to the next, the packets and the Logical Link Control (LLC) protocols must be the same on the each segment. This means that a repeater will not enable communication, for example, between an 802.3 segment (Ethernet) and an 802.5 segment (Token Ring). That is, they cannot translate an Ethernet packet into a Token Ring packet. In other words, repeaters do not translate anything.

Dual speed hubs

In the early days of fast Ethernet, Ethernet switches were relatively expensive devices. Hubs suffered from the problem that if there were any 10BASE-T devices connected then the whole network needed to run at 10 Mbit/s. Therefore a compromise between a hub and a switch was developed, known as a **dual-speed hub**. These devices consisted of an internal two-port switch, dividing the 10 Mbit/s and 100 Mbit/s segments. The device would typically consist of more than two physical ports. When a network device becomes active on any of the physical ports, the device attaches it to either the 10 Mbit/s segment or the 100 Mbit/s segment, as appropriate. This prevented the need for an all-or-nothing migration fast Ethernet networks. These devices are considered hubs because the traffic between devices connected at the same speed is not switched.

Uses

Historically, the main reason for purchasing hubs rather than switches was their price. This motivator has largely been eliminated by reductions in the price of switches, but hubs can still be useful in special circumstances:

- For inserting a protocol analyzer into a network connection, a hub is an alternative to a network tap or port mirroring.
- When a switch is accessible for end users to make connections, for example, in a conference room, an inexperienced or careless user (or saboteur) can bring down the network by connecting two ports together, causing a loop. This can be prevented by using a hub, where a loop will break other users on the hub, but not the rest of the network. This hazard can also be avoided by using switches that can detect and deal with loops, for example by implementing the spanning tree protocol.
- A hub with a 10BASE2 port can be used to connect devices that only support 10BASE2 to a modern network. The same goes for linking in an old 10BASE5 network segment using an AUI port on a hub (individual devices that were intended for thicknet can be linked to modern Ethernet by using an AUI-10BASE-T transceiver).

Chapter- 8

Modem

A **modem** (**mo**dulator-**dem**odulator) is a device that modulates an analog carrier signal to encode digital information, and also demodulates such a carrier signal to decode the transmitted information. The goal is to produce a signal that can be transmitted easily and decoded to reproduce the original digital data. Modems can be used over any means of transmitting analog signals, from driven diodes to radio.

The most familiar example is a voice band modem that turns the digital data of a personal computer into modulated electrical signals in the voice frequency range of a telephone channel. These signals can be transmitted over telephone lines and demodulated by another modem at the receiver side to recover the digital data.

Modems are generally classified by the amount of data they can send in a given time unit, normally measured in bits per second (bit/s, or bps). They can also be classified by the symbol rate measured in baud, the number of times the modem changes its signal state per second. For example, the ITU V.21 standard used audio frequency-shift keying, aka tones, to carry 300 bit/s using 300 baud, whereas the original ITU V.22 standard allowed 1,200 bit/s with 600 baud using phase shift keying.

History

News wire services in 1920s used multiplex equipment that met the definition, but the modem function was incidental to the multiplexing function, so they are not commonly included in the history of modems.

TeleGuide terminal

Modems grew out of the need to connect teletype machines over ordinary phone lines instead of more expensive leased lines which had previously been used for current loop-based teleprinters and automated telegraphs. George Stibitz connected a New Hampshire teletype to a computer in New York City by a subscriber telephone line in 1940.

In 1943, IBM adapted this technology to their unit record equipment and were able to transmit punched cards at 25 bits/second. Mass-produced modems in the United States began as part of the SAGE air-defense system in 1958, connecting terminals at various airbases, radar sites, and command-and-control centers to the SAGE director centers scattered around the U.S. and Canada. SAGE modems were described by AT&T's Bell Labs as conforming to their newly published Bell 101 dataset standard. While they ran on dedicated telephone lines, the devices at each end were no different from commercial acoustically coupled Bell 101, 110 baud modems.

In the summer of 1960, the name *Data-Phone* was introduced to replace the earlier term *digital subset*. The *202 Data-Phone* was a half-duplex asynchronous service that was marketed extensively in late 1960. In 1962, the *201A* and *201B Data-Phones* were introduced. They were synchronous modems using two-bit-per-baud phase-shift keying (PSK). The 201A operated half-duplex at 2,000 bit/s over normal phone lines, while the

201B provided full duplex 2,400 bit/s service on four-wire leased lines, the send and receive channels running on their own set of two wires each.

The famous *Bell 103A dataset* standard was also introduced by Bell Labs in 1962. It provided full-duplex service at 300 baud over normal phone lines. Frequency-shift keying was used with the call originator transmitting at 1,070 or 1,270 Hz and the answering modem transmitting at 2,025 or 2,225 Hz. The readily available 103A2 gave an important boost to the use of remote low-speed terminals such as the KSR33, the ASR33, and the IBM 2741. AT&T reduced modem costs by introducing the originate-only 113D and the answer-only 113B/C modems.

The Carterfone decision

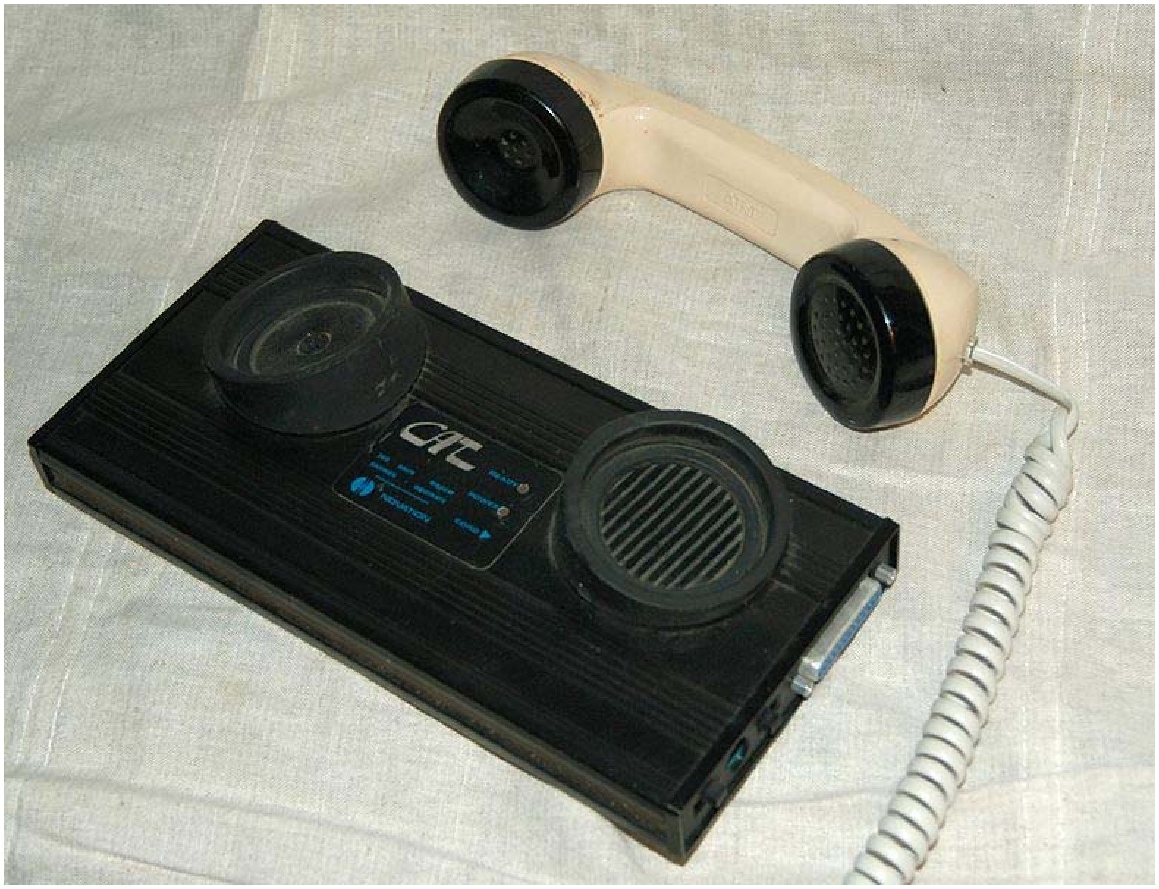

The *Novation CAT* acoustically coupled modem

For many years, the Bell System (AT&T) maintained a monopoly on the use of its phone lines, allowing only Bell-supplied devices to be attached to its network. Before 1968, AT&T maintained a monopoly on what devices could be *electrically* connected to its phone lines. This led to a market for 103A-compatible modems that were *mechanically* connected to the phone, through the handset, known as acoustically coupled modems. Particularly common models from the 1970s were the Novation CAT and the Anderson-Jacobson, spun off from an in-house project at Stanford Research Institute (now SRI

International). Hush-a-Phone v. FCC was a seminal ruling in United States telecommunications law decided by the DC Circuit Court of Appeals on November 8, 1956. The District Court found that it was within the FCC's authority to regulate the terms of use of AT&T's equipment. Subsequently, the FCC examiner found that as long as the device was not physically attached it would not threaten to degenerate the system. Later, in the Carterfone decision of 1968, the FCC passed a rule setting stringent AT&T-designed tests for electronically coupling a device to the phone lines. AT&T's tests were complex, making electronically-coupled modems expensive, so acoustically-coupled modems remained common into the early 1980s.

In December 1972, Vadic introduced the *VA3400*. This device was remarkable because it provided full duplex operation at 1,200 bit/s over the dial network, using methods similar to those of the 103A in that it used different frequency bands for transmit and receive. In November 1976, AT&T introduced the 212A modem to compete with Vadic. It was similar in design to Vadic's model, but used the lower frequency set for transmission. It was also possible to use the 212A with a 103A modem at 300 bit/s. According to Vadic, the change in frequency assignments made the 212 intentionally incompatible with acoustic coupling, thereby locking out many potential modem manufacturers. In 1977, Vadic responded with the VA3467 triple modem, an answer-only modem sold to computer center operators that supported Vadic's 1,200-bit/s mode, AT&T's 212A mode, and 103A operation.

The Smartmodem and the rise of BBSes

US Robotics Sportster 14,400 Fax modem (1994)

The next major advance in modems was the *Smartmodem*, introduced in 1981 by Hayes Communications. The Smartmodem was an otherwise standard 103A 300-bit/s modem, but was attached to a small controller that let the computer send commands to it and enable it to operate the phone line. The command set included instructions for picking up and hanging up the phone, dialing numbers, and answering calls. The basic Hayes command set remains the basis for computer control of most modern modems.

Prior to the Hayes *Smartmodem*, dial-up modems almost universally required a two-step process to activate a connection: first, the user had to manually dial the remote number on a standard phone handset, and then secondly, plug the handset into an acoustic coupler. Hardware add-ons, known simply as *dialers*, were used in special circumstances, and generally operated by emulating someone dialing a handset.

With the Smartmodem, the computer could dial the phone directly by sending the modem a command, thus eliminating the need for an associated phone instrument for dialing and the need for an acoustic coupler. The Smartmodem instead plugged directly into the phone line. This greatly simplified setup and operation. Terminal programs that maintained lists of phone numbers and sent the dialing commands became common.

The Smartmodem and its clones also aided the spread of bulletin board systems (BBSs). Modems had previously been typically either the call-only, acoustically coupled models used on the client side, or the much more expensive, answer-only models used on the server side. The Smartmodem could operate in either mode depending on the commands sent from the computer. There was now a low-cost server-side modem on the market, and the BBSs flourished.

Almost all modern modems can interoperate with fax machines. Digital faxes, introduced in the 1980s, are simply a particular image format sent over a high-speed (commonly 14.4 kbit/s) modem. Software running on the host computer can convert any image into fax-format, which can then be sent using the modem. Such software was at one time an add-on, but since has become largely universal.

Softmodem (dumb modem)

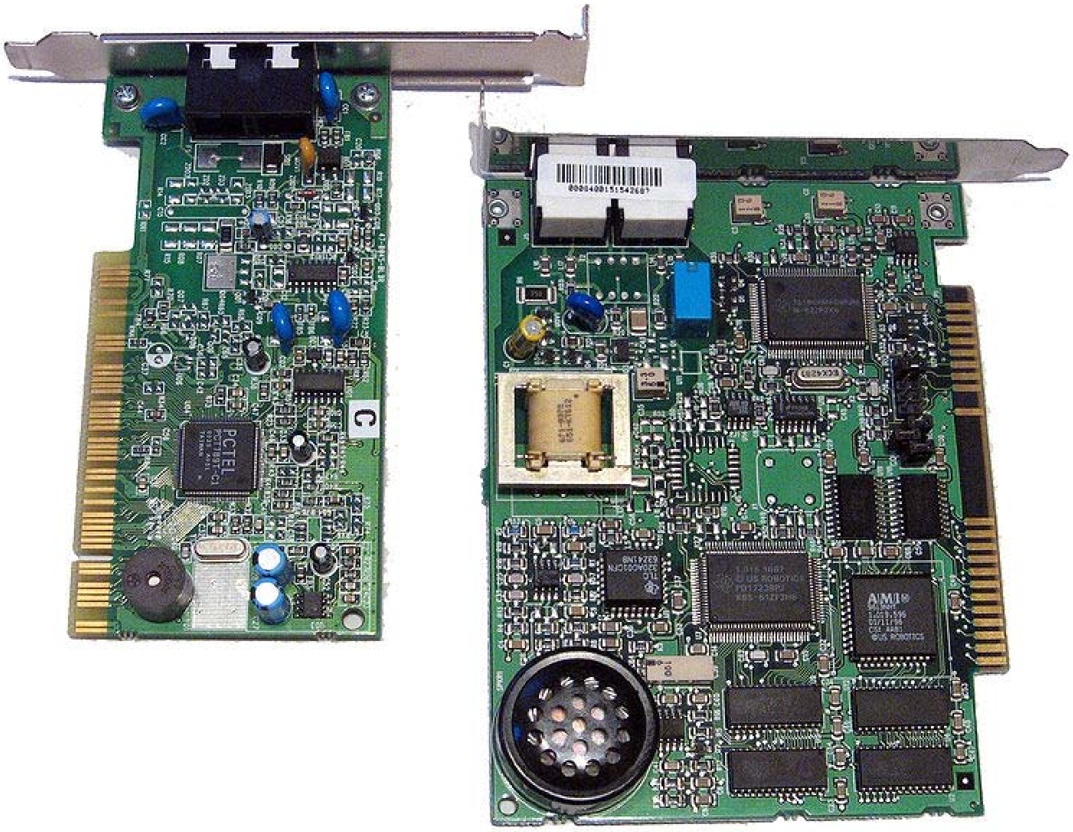

A PCI Winmodem/softmodem (on the left) next to a traditional ISA modem (on the right). Notice the less complex circuitry of the modem on the left.

A *Winmodem* or *softmodem* is a stripped-down modem that replaces tasks traditionally handled in hardware with software. In this case the modem is a simple interface designed to create voltage variations on the telephone line and to sample the line voltage levels (digital to analog and analog to digital converters). Softmodems are cheaper than traditional modems, since they have fewer hardware components. One downside is that the software generating and interpreting the modem tones is not simple (as most of the protocols are complex), and the performance of the computer as a whole often suffers when it is being used. For online gaming this can be a real concern. Another problem is lack of portability such that non-Windows operating systems (such as Linux) often do not have an equivalent driver to operate the modem.

Narrow-band/phone-line dialup modems

A standard modem of today contains two functional parts: an analog section for generating the signals and operating the phone, and a digital section for setup and control. This functionality is often incorporated into a single chip nowadays, but the division remains in theory. In operation the modem can be in one of two modes, *data mode* in which data is sent to and from the computer over the phone lines, and *command mode* in which the modem listens to the data from the computer for commands, and carries them out. A typical session consists of powering up the modem (often inside the computer itself) which automatically assumes command mode, then sending it the command for dialing a number. After the connection is established to the remote modem, the modem automatically goes into data mode, and the user can send and receive data. When the user is finished, the escape sequence, "+++" followed by a pause of about a second, may be sent to the modem to return it to command mode, then a command (e.g. "ATH") to hang up the phone is sent. Note that on many modem controllers it is possible to issue commands to disable the escape sequence so that it is not possible for data being exchanged to trigger the mode change inadvertently.

The commands themselves are typically from the Hayes command set, although that term is somewhat misleading. The original Hayes commands were useful for 300 bit/s operation only, and then extended for their 1,200 bit/s modems. Faster speeds required new commands, leading to a proliferation of command sets in the early 1990s. Things became considerably more standardized in the second half of the 1990s, when most modems were built from one of a very small number of chipsets. We call this the Hayes command set even today, although it has three or four times the numbers of commands as the actual standard.

Increasing speeds (V.21, V.22, V.22bis)

A 2,400 bit/s modem for a laptop.

The 300 bit/s modems used audio frequency-shift keying to send data. In this system the stream of 1s and 0s in computer data is translated into sounds which can be easily sent on the phone lines. In the Bell 103 system the *originating* modem sends 0s by playing a 1,070 Hz tone, and 1s at 1,270 Hz, with the *answering* modem putting its 0s on 2,025 Hz and 1s on 2,225 Hz. These frequencies were chosen carefully, they are in the range that suffer minimum distortion on the phone system, and also are not harmonics of each other.

In the 1,200 bit/s and faster systems, phase-shift keying was used. In this system the two tones for any one side of the connection are sent at the similar frequencies as in the

300 bit/s systems, but slightly out of phase. By comparing the phase of the two signals, 1s and 0s could be pulled back out, for instance if the signals were 90 degrees out of phase, this represented two digits, *1, 0*, at 180 degrees it was *1, 1*. In this way each cycle of the signal represents two digits instead of one. 1,200 bit/s modems were, in effect, 600 symbols per second modems (600 baud modems) with 2 bits per symbol.

Voiceband modems generally remained at 300 and 1,200 bit/s (V.21 and V.22) into the mid 1980s. A V.22bis 2,400-bit/s system similar in concept to the 1,200-bit/s Bell 212 signalling was introduced in the U.S., and a slightly different one in Europe. By the late 1980s, most modems could support all of these standards and 2,400-bit/s operation was becoming common.

Increasing speeds (one-way proprietary standards)

Many other standards were also introduced for special purposes, commonly using a high-speed channel for receiving, and a lower-speed channel for sending. One typical example was used in the French Minitel system, in which the user's terminals spent the majority of their time receiving information. The modem in the Minitel terminal thus operated at 1,200 bit/s for reception, and 75 bit/s for sending commands back to the servers.

Three U.S. companies became famous for high-speed versions of the same concept. Telebit introduced its *Trailblazer* modem in 1984, which used a large number of 36 bit/s channels to send data one-way at rates up to 18,432 bit/s. A single additional channel in the reverse direction allowed the two modems to communicate how much data was waiting at either end of the link, and the modems could change direction on the fly. The Trailblazer modems also supported a feature that allowed them to spoof the UUCP *g* protocol, commonly used on Unix systems to send e-mail, and thereby speed UUCP up by a tremendous amount. Trailblazers thus became extremely common on Unix systems, and maintained their dominance in this market well into the 1990s.

U.S. Robotics (USR) introduced a similar system, known as *HST*, although this supplied only 9,600 bit/s (in early versions at least) and provided for a larger backchannel. Rather than offer spoofing, USR instead created a large market among Fidonet users by offering its modems to BBS sysops at a much lower price, resulting in sales to end users who wanted faster file transfers. Hayes was forced to compete, and introduced its own 9,600-bit/s standard, *Express 96* (also known as *Ping-Pong*), which was generally similar to Telebit's PEP. Hayes, however, offered neither protocol spoofing nor sysop discounts, and its high-speed modems remained rare.

4,800 and 9,600 bit/s (V.27ter, V.32)

Echo cancellation was the next major advance in modem design. Local telephone lines use the same wires to send and receive, which results in a small amount of the outgoing signal bouncing back. This signal can confuse the modem, which was unable to distinguish between the echo and the signal from the remote modem. This was why earlier modems split the signal frequencies into 'answer' and 'originate'; the modem could then ignore its own transmitting frequencies. Even with improvements to the phone system allowing higher speeds, this splitting of available phone signal bandwidth still imposed a half-speed limit on modems.

Echo cancellation got around this problem. Measuring the echo delays and magnitudes allowed the modem to tell if the received signal was from itself or the remote modem, and create an equal and opposite signal to cancel its own. Modems were then able to send over the whole frequency spectrum in both directions at the same time, leading to the development of 4,800 and 9,600 bit/s modems.

Increases in speed have used increasingly complicated communications theory. 1,200 and 2,400 bit/s modems used the phase shift key (PSK) concept. This could transmit two or three bits per symbol. The next major advance encoded four bits into a combination of amplitude and phase, known as Quadrature Amplitude Modulation (QAM). Best visualized as a constellation diagram, the bits are mapped onto points on a graph with the x (real) and y (quadrature) coordinates transmitted over a single carrier.

The new V.27ter and V.32 standards were able to transmit 4 bits per symbol, at a rate of 1,200 or 2,400 baud, giving an effective bit rate of 4,800 or 9,600 bit/s. The carrier frequency was 1,650 Hz. For many years, most engineers considered this rate to be the limit of data communications over telephone networks.

Error correction and compression

Operations at these speeds pushed the limits of the phone lines, resulting in high error rates. This led to the introduction of error-correction systems built into the modems, made most famous with Microcom's MNP systems. A string of MNP standards came out in the 1980s, each increasing the effective data rate by minimizing overhead, from about 75% theoretical maximum in MNP 1, to 95% in MNP 4. The new method called MNP 5 took this a step further, adding data compression to the system, thereby increasing the data rate above the modem's rating. Generally the user could expect an MNP5 modem to transfer at about 130% the normal data rate of the modem. Details of MNP were later released and became popular on a series of 2,400-bit/s modems, and ultimately led to the development of V.42 and V.42bis ITU standards. V.42 and V.42bis were non-compatible with MNP but were similar in concept: Error correction and compression.

Another common feature of these high-speed modems was the concept of fallback, or *speed hunting*, allowing them to talk to less-capable modems. During the call initiation the modem would play a series of signals into the line and wait for the remote modem to respond to them. They would start at high speeds and progressively get slower and slower until they heard an answer. Thus, two USR modems would be able to connect at 9,600 bit/s, but, when a user with a 2,400-bit/s modem called in, the USR would fallback to the common 2,400-bit/s speed. This would also happen if a V.32 modem and a HST modem were connected. Because they used a different standard at 9,600 bit/s, they would fall back to their highest commonly supported standard at 2,400 bit/s. The same applies to

V.32bis and 14,400 bit/s HST modem, which would still be able to communicate with each other at only 2,400 bit/s.

Breaking the 9.6k barrier

In 1980, Gottfried Ungerboeck from IBM Zurich Research Laboratory applied powerful channel coding techniques to search for new ways to increase the speed of modems. His results were astonishing but only conveyed to a few colleagues. Finally in 1982, he agreed to publish what is now a landmark paper in the theory of information coding. By applying powerful parity check coding to the bits in each symbol, and mapping the encoded bits into a two-dimensional diamond pattern, Ungerboeck showed that it was possible to increase the speed by a factor of two with the same error rate. The new technique was called *mapping by set partitions* (now known as trellis modulation).

Error correcting codes, which encode code words (sets of bits) in such a way that they are far from each other, so that in case of error they are still closest to the original word (and not confused with another) can be thought of as analogous to sphere packing or packing pennies on a surface: the further two bit sequences are from one another, the easier it is to correct minor errors.

V.32bis was so successful that the older high-speed standards had little to recommend them. USR fought back with a 16,800 bit/s version of HST, while AT&T introduced a one-off 19,200 bit/s method they referred to as V.32ter (also known as V.32 terbo or *tertiary*), but neither non-standard modem sold well.

V.34/28.8k and 33.6k

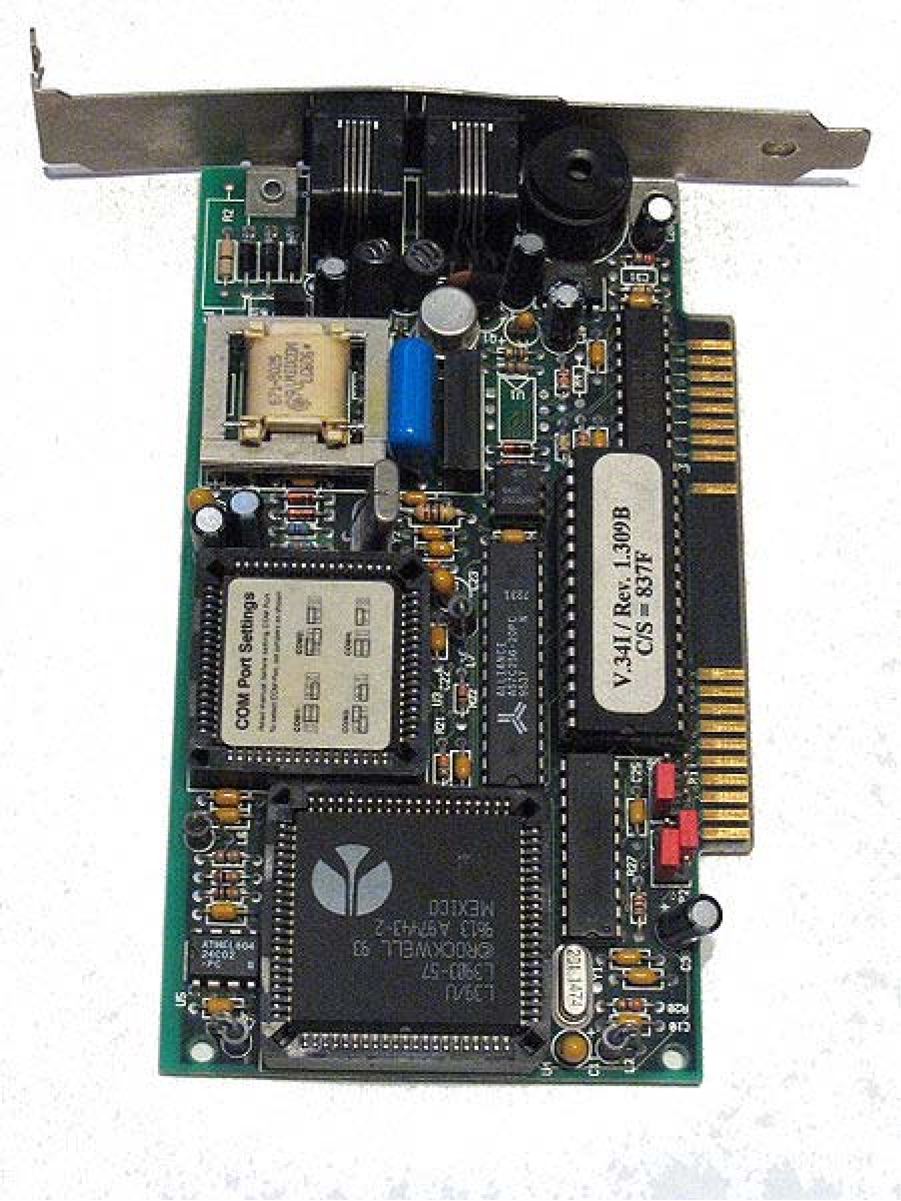

An ISA modem manufactured to conform to the V.34 protocol.

Any interest in these systems was destroyed during the lengthy introduction of the 28,800 bit/s V.34 standard. While waiting, several companies decided to release hardware and introduced modems they referred to as *V.FAST*. In order to guarantee compatibility with V.34 modems once the standard was ratified (1994), the manufacturers were forced to use more flexible parts, generally a DSP and microcontroller, as opposed to purpose-designed ASIC modem chips.

Today, the ITU standard V.34 represents the culmination of the joint efforts. It employs the most powerful coding techniques including channel encoding and shape encoding. From the mere 4 bits per symbol (9.6 kbit/s), the new standards used the functional equivalent of 6 to 10 bits per symbol, plus increasing baud rates from 2,400 to 3,429, to create 14.4, 28.8, and 33.6 kbit/s modems. This rate is near the theoretical Shannon limit. When calculated, the Shannon capacity of a narrowband line is $Bandwidth*log_2(1+P_u/P_n)$, with $P_u/P_n$the (linear) signal-to-noise ratio. Narrowband phone lines have a bandwidth from 300-4000 Hz, so using $P_u/P_n=1000$(SNR = 30dB): capacity is approximately 35 kbit/s.

Without the discovery and eventual application of trellis modulation, maximum telephone rates using voice-bandwidth channels would have been limited to 3,429 baud * 4 bit/symbol == approximately 14 kbit/s using traditional QAM. (DSL makes use of the bandwidth of traditional copper-wire twisted pairs between subscriber and the central office, which far exceeds that of analog voice circuitry.)

V.61/V.70 Analog/Digital Simultaneous Voice and Data

The V.61 Standard introduced Analog Simultaneous Voice and Data (ASVD). This technology allowed users of v.61 modems to engage in point-to-point voice conversations with each other while their respective modems communicated.

In 1995, the first DSVD (Digital Simultaneous Voice and Data) modems became available to consumers, and the standard was ratified as v.70 by the International Telecommunication Union (ITU) in 1996.

Two DSVD modems can establish a completely digital link between each other over standard phone lines. Sometimes referred to as "the poor man's ISDN," and employing a similar technology, v.70 compatible modems allow for a maximum speed of 33.6 kbps between peers. By using a majority of the bandwidth for data and reserving part for voice transmission, DSVD modems allow users to pick up a telephone handset interfaced with the modem, and initiate a call to the other peer.

One practical use for this technology was realized by early two player video gamers, who could hold voice communication with each other while in game over the PSTN.

Advocates of DSVD envisioned whiteboard sharing and other practical applications for the standard, however, with advent of cheaper 56kbps analog modems intended for Internet connectivity, peer-to-peer data transmission over the PSTN became quickly irrelevant. Also, the standard was never expanded to allow for the making or receiving of arbitrary phone calls while the modem was in use, due to the cost of infrastructure upgrades to telephone companies, and the advent of ISDN and DSL technologies which effectively accomplished the same goal.

Today, Multi-Tech is the only known company to continue to support a v.70 compatible modem. While their device also offers v.92 at 56kbps, it remains significantly more expensive than comparable modems sans v.70 support.

Using digital lines and PCM (V.90/92)

Modem bank at an ISP.

In the late 1990s Rockwell/Lucent and U.S. Robotics introduced new competing technologies based upon the digital transmission used in modern telephony networks. The standard digital transmission in modern networks is 64 kbit/s but some networks use a part of the bandwidth for remote office signaling (e.g., to hang up the phone), limiting the effective rate to 56 kbit/s DS0. This new technology was adopted into ITU standards V.90 and is common in modern computers. The 56 kbit/s rate is only possible from the central office to the user site (downlink). In the United States, government regulation limits the maximum power output, resulting in a maximum data rate of 53.3 kbit/s. The uplink (from the user to the central office) still uses V.34 technology at 33.6 kbit/s.

Later in V.92, the digital PCM technique was applied to increase the upload speed to a maximum of 48 kbit/s, but at the expense of download rates. For example a 48 kbit/s upstream rate would reduce the downstream as low as 40 kbit/s, due to echo on the telephone line. To avoid this problem, V.92 modems offer the option to turn off the digital upstream and instead use a 33.6 kbit/s analog connection, in order to maintain a high digital downstream of 50 kbit/s or higher. V.92 also adds two other features. The first is the ability for users who have call waiting to put their dial-up Internet connection on hold for extended periods of time while they answer a call. The second feature is the ability to quickly connect to one's ISP. This is achieved by remembering the analog and digital characteristics of the telephone line, and using this saved information to reconnect at a fast pace.

Using compression to exceed 56k

Today's V.42, V.42bis and V.44 standards allow the modem to transmit data faster than its basic rate would imply. For instance, a 53.3 kbit/s connection with V.44 can transmit up to 53.3*6 == 320 kbit/s using pure text. However, the compression ratio tends to vary due to noise on the line, or due to the transfer of already-compressed files (ZIP files, JPEG images, MP3 audio, MPEG video). At some points the modem will be sending compressed files at approximately 50 kbit/s, uncompressed files at 160 kbit/s, and pure text at 320 kbit/s, or any value in between.

In such situations a small amount of memory in the modem, a buffer, is used to hold the data while it is being compressed and sent across the phone line, but in order to prevent overflow of the buffer, it sometimes becomes necessary to tell the computer to pause the datastream. This is accomplished through *hardware flow control* using extra lines on the modem–computer connection. The computer is then set to supply the modem at some higher rate, such as 320 kbit/s, and the modem will tell the computer when to start or stop sending data.

Compression by the ISP

As telephone-based 56k modems began losing popularity, some Internet service providers such as Netzero and Juno started using pre-compression to increase the throughput and maintain their customer base. As example, the Netscape ISP uses a compression program that squeezes images, text, and other objects at the modem server, just prior to sending them across the phone line. Certain content using lossy compression (e.g., images) may be recompressed (transcoded) using different parameters to the compression algorithm, making the transmitted content smaller but of lower quality. The server-side compression operates much more efficiently than the on-the-fly compression of V.44-enabled modems due to the fact that V.44 is a generalized compression algorithm whereas other compression techniques are application-specific (JPEG, MPEG, Vorbis, etc.). Typically Website text is compacted to 4% thus increasing effective throughput to approximately 1,300 kbit/s. The accelerator also pre-compresses Flash executables and images to approximately 30% and 12%, respectively.

The drawback of this approach is a loss in quality, where the GIF and JPEG images are lossy compressed, which causes the content to become pixelated and smeared. However the speed is dramatically improved such that Web pages load in less than 5 seconds, and the user can manually choose to view the uncompressed images at any time. The ISPs employing this approach advertise it as "surf 5× faster" or simply "accelerated dial-up".

List of dialup speeds

Note that the values given are maximum values, and actual values may be slower under certain conditions (for example, noisy phone lines). A baud is one symbol per second; each symbol may encode one or more data bits.

Connection	**Bitrate (kbit/s)**	**Year Released**
110 baud Bell 101 modem	0.1	1958
300 baud (Bell 103 or V.21)	0.3	1962
1200 modem (1200 baud) (Bell 202)	1.2	
1200 Modem (600 baud) (Bell 212A or V.22)	1.2	
2400 Modem (600 baud) (V.22bis)	2.4	
2400 Modem (1200 baud) (V.26bis)	2.4	
4800 Modem (1600 baud) (V.27ter)	4.8	
9600 Modem (2400 baud) (V.32)	9.6	
14.4k Modem (2400 baud) (V.32bis)	14.4	
28.8k Modem (3200 baud) (V.34)	28.8	
33.6k Modem (3429 baud) (V.34)	33.6	
56k Modem (8000/3429 baud) (V.90)	56.0/33.6	
56k Modem (8000/8000 baud) (V.92)	56.0/48.0	
Bonding modem (two 56k modems)) (V.92)	112.0/96.0	
Hardware compression (variable) (V.90/V.42bis)	56.0-220.0	
Hardware compression (variable) (V.92/V.44)	56.0-320.0	
Server-side web compression (variable) (Netscape ISP)	100.0-1,000.0	

Radio modems

Direct broadcast satellite, WiFi, and mobile phones all use modems to communicate, as do most other wireless services today. Modern telecommunications and data networks also make extensive use of radio modems where long distance data links are required. Such systems are an important part of the PSTN, and are also in common use for high-speed computer network links to outlying areas where fibre is not economical.

Even where a cable is installed, it is often possible to get better performance or make other parts of the system simpler by using radio frequencies and modulation techniques through a cable. Coaxial cable has a very large bandwidth, however signal attenuation becomes a major problem at high data rates if a digital signal is used. By using a modem, a much larger amount of digital data can be transmitted through a single piece of wire. Digital cable television and cable Internet services use radio frequency modems to provide the increasing bandwidth needs of modern households. Using a modem also allows for frequency-division multiple access to be used, making full-duplex digital communication with many users possible using a single wire.

Wireless modems come in a variety of types, bandwidths, and speeds. Wireless modems are often referred to as transparent or smart. They transmit information that is modulated onto a carrier frequency to allow many simultaneous wireless communication links to work simultaneously on different frequencies.

Transparent modems operate in a manner similar to their phone line modem cousins. Typically, they were half duplex, meaning that they could not send and receive data at the same time. Typically transparent modems are polled in a round robin manner to collect small amounts of data from scattered locations that do not have easy access to wired infrastructure. Transparent modems are most commonly used by utility companies for data collection.

Smart modems come with a media access controller inside which prevents random data from colliding and resends data that is not correctly received. Smart modems typically require more bandwidth than transparent modems, and typically achieve higher data rates. The IEEE 802.11 standard defines a short range modulation scheme that is used on a large scale throughout the world.

WiFi and WiMax

Wireless data modems are used in the WiFi and WiMax standards, operating at microwave frequencies.

WiFi is principally used in laptops for Internet connections (wireless access point) and wireless application protocol (WAP).

Mobile modems and routers

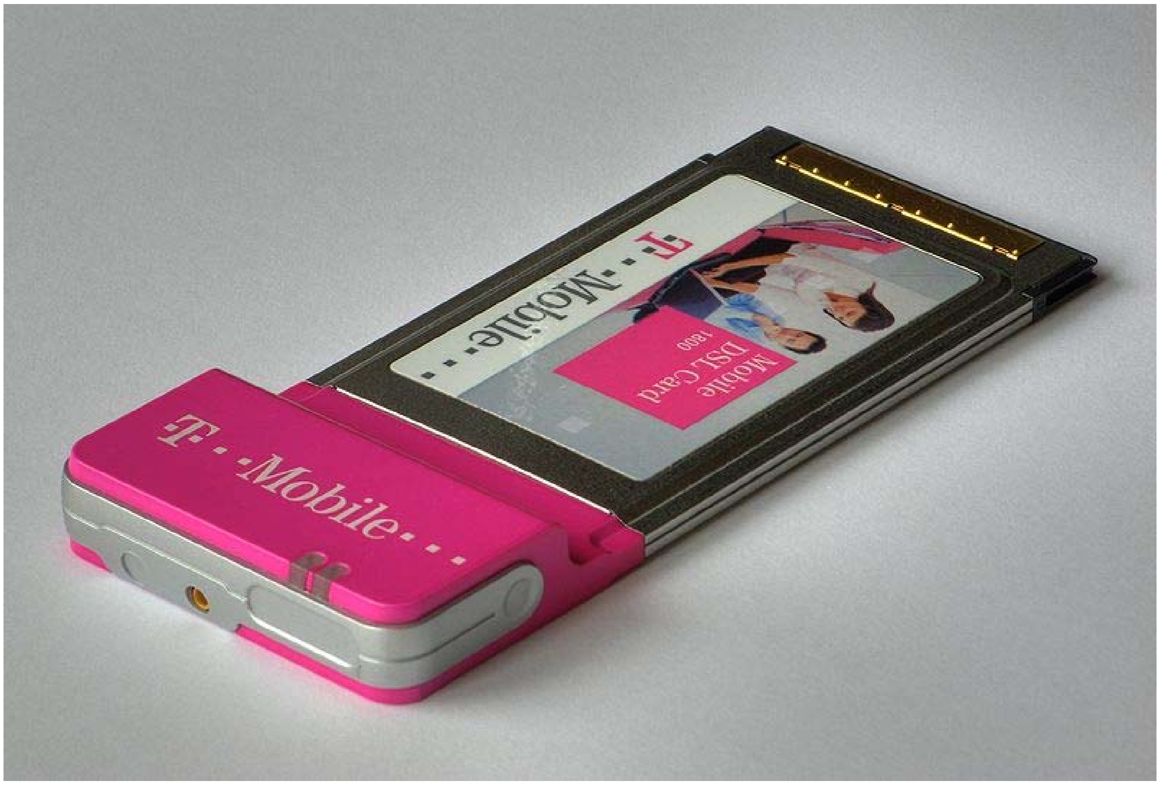

T-Mobile Universal Mobile Telecommunications System PC Card modem

Huawei CDMA2000 Evolution-Data Optimized USB wireless modem

Modems which use a mobile telephone system (GPRS, UMTS, HSPA, EVDO, WiMax, etc.), are known as wireless modems (sometimes also called cellular modems). Wireless modems can be embedded inside a laptop or appliance or external to it. External wireless modems are connect cards, usb modems for mobile broadband and cellular routers. A connect card is a PC card or ExpressCard which slides into a PCMCIA/PC card/ExpressCard slot on a computer. USB wireless modems use a USB port on the laptop instead of a PC card or ExpressCard slot. A cellular router may have an external datacard (*AirCard*) that slides into it. Most cellular routers do allow such datacards or USB modems. Cellular Routers may not be modems per se, but they contain modems or allow modems to be slid into them. The difference between a cellular router and a wireless modem is that a cellular router normally allows multiple people to connect to it (since it can route, or support multipoint to multipoint connections), while the modem is made for one connection.

Most of the GSM wireless modems come with an integrated SIM cardholder (i.e., Huawei E220, Sierra 881, etc.) and some models are also provided with a microSD memory slot and/or jack for additional external antenna such as Huawei E1762 and Sierra Wireless Compass 885.] The CDMA (EVDO) versions do not use R-UIM cards, but use Electronic Serial Number (ESN) instead.

The cost of using a wireless modem varies from country to country. Some carriers implement flat rate plans for unlimited data transfers. Some have caps (or maximum limits) on the amount of data that can be transferred per month. Other countries have plans that charge a fixed rate per data transferred—per megabyte or even kilobyte of data downloaded; this tends to add up quickly in today's content-filled world, which is why many people are pushing for flat data rates.

The faster data rates of the newest wireless modem technologies (UMTS, HSPA, EVDO, WiMax) are also considered to be *broadband wireless modems* and compete with other broadband modems below.

Broadband

DSL modem

ADSL modems, a more recent development, are not limited to the telephone's voiceband audio frequencies. Some ADSL modems use coded orthogonal frequency division modulation (DMT, for Discrete MultiTone; also called COFDM, for digital TV in much of the world).

Cable modems use a range of frequencies originally intended to carry RF television channels. Multiple cable modems attached to a single cable can use the same frequency band, using a low-level media access protocol to allow them to work together within the same channel. Typically, 'up' and 'down' signals are kept separate using frequency division multiple access.

New types of broadband modems are beginning to appear, such as doubleway satellite and power line modems.

Broadband modems should still be classed as modems, since they use complex waveforms to carry digital data. They are more advanced devices than traditional dial-up modems as they are capable of modulating/demodulating hundreds of channels simultaneously.

Many broadband modems include the functions of a router (with Ethernet and WiFi ports) and other features such as DHCP, NAT and firewall features.

When broadband technology was introduced, networking and routers were unfamiliar to consumers. However, many people knew what a modem was as most internet access was through dial-up. Due to this familiarity, companies started selling broadband modems using the familiar term *modem* rather than vaguer ones like *adapter* or *transceiver*, or even "bridge".

Many broadband modems must be configured in bridge mode before they can use a router.

Home networking

Although the name *modem* is seldom used in this case, modems are also used for high-speed home networking applications, specially those using existing home wiring. One example is the G.hn standard, developed by ITU-T, which provides a high-speed (up to 1 Gbit/s) Local area network using existing home wiring (power lines, phone lines and coaxial cables). G.hn devices use orthogonal frequency-division multiplexing (OFDM) to modulate a digital signal for transmission over the wire.

The phrase "Null modem" was used to describe attaching a specially wired cable between the serial ports of two personal computers. Basically, the transmit output of one computer was wired to the receive input of the other; this was true for both computers. The same software used with modems (such as Procomm or Minicom) could be used with the null modem connection.

Deep-space telecommunications

Many modern modems have their origin in deep space telecommunications systems of the 1960s.

Differences between deep space telecom modems and landline modems:

- digital modulation formats that have high doppler immunity are typically used
- waveform complexity tends to be low, typically binary phase shift keying
- error correction varies mission to mission, but is typically much stronger than most landline modems

Voice modem

Voice modems are regular modems that are capable of recording or playing audio over the telephone line. They are used for telephony applications. This type of modem can be used as an FXO card for Private branch exchange systems (compare V.92).

Popularity

A CEA study in 2006 found that dial-up Internet access is on a notable decline in the U.S. In 2000, dial-up Internet connections accounted for 74% of all U.S. residential Internet connections. The US demographic pattern for (dial-up modem users per capita) has been more or less mirrored in Canada and Australia for the past 20 years.

Dial-up modem use in the US had dropped to 60% by 2003, and in 2006 stood at 36%. Voiceband modems were once the most popular means of Internet access in the U.S., but with the advent of new ways of accessing the Internet, the traditional 56K modem is losing popularity.

Chapter- 9

Passive Optical Network

A **passive optical network** (**PON**) is a point-to-multipoint, fiber to the premises network architecture in which unpowered optical splitters utilizing Brewster's angle principles are used to enable a single optical fiber to serve multiple premises, typically 32-128. A PON consists of an optical line terminal (OLT) at the service provider's central office and a number of optical network units (ONUs) near end users. A PON configuration reduces the amount of fiber and central office equipment required compared with point to point architectures. A passive optical network is a form of fiber-optic access network.

Downstream signals are broadcast to each premises sharing a single fiber. Encryption is used to prevent eavesdropping.

Upstream signals are combined using a multiple access protocol, usually time division multiple access (TDMA). The OLTs "range" the ONUs in order to provide time slot assignments for upstream communication.

Standards

- IEEE 802.3
  - **EPON** (Ethernet PON) is part of IEEE standard Ethernet with options for 1/1 Gbit/s 10/1 Gbit/s and 10/10 Gbit/s. There are currently over 40 million installed EPON ports making it the most widely deployed PON technology globally. EPON is also the foundation for cable operators business services as part of the DOCSIS Provisioning of EPON (DPoE) specifications.

- ITU-T
  - G.983
    - **APON** (ATM PON). This was the first Passive optical network standard. It was used primarily for business applications, and was based on ATM.
    - **BPON** (Broadband PON) is a standard based on APON. It adds support for WDM, dynamic and higher upstream bandwidth allocation, and survivability. It also created a standard management interface, called OMCI, between the OLT and ONU/ONT, enabling mixed-vendor networks.

        - G.984
            - **G-PON** (Gigabit PON) is an evolution of the BPON standard. It supports higher rates, enhanced security, and choice of Layer 2 protocol (ATM, GEM, Ethernet). By mid-2008, Verizon had installed over 800 thousand lines. British Telecom, Mobily-SaudiArabia, Etisalat-UAE, and AT&T are in advanced trials. It is the successor to G.983.
        - G.987
            - **10G-PON** has 10 Gbit/s downstream and 2.5 Gbit/s upstream – framing is "G-PON like" and designed to coexist with GPON devices on the same network.
- SCTE IPS910
    - **RFoG** (RFoverGlass) is an SCTE Interface Practices Subcomittee standard in development for carrying HFC RF signals over a passive optical Network (PON).

History

Early work on efficient fiber to the home architectures was done in the 1990s by the Full Service Access Network (FSAN) working group, formed by major telecommunications service providers and system vendors. The International Telecommunications Union (ITU) did further work, and has since standardized on two generations of PON. The older ITU-T G.983 standard is based on Asynchronous Transfer Mode (ATM), and has therefore been referred to as APON (ATM PON). Further improvements to the original APON standard – as well as the gradual falling out of favor of ATM as a protocol – led to the full, final version of ITU-T G.983 being referred to more often as broadband PON, or BPON. A typical APON/BPON provides 622 megabits per second (Mbit/s) (OC-12) of downstream bandwidth and 155 Mbit/s (OC-3) of upstream traffic, although the standard accommodates higher rates.

The ITU-T G.984 (GPON) standard represents a boost, compared to BPON, in both the total bandwidth and bandwidth efficiency through the use of larger, variable-length packets. Again, the standards permit several choices of bit rate, but the industry has converged on 2.488 gigabits per second (Gbit/s) of downstream bandwidth, and 1.244 Gbit/s of upstream bandwidth. GPON Encapsulation Method (GEM) allows very efficient packaging of user traffic with frame segmentation.

The IEEE 802.3 Ethernet PON (EPON or GEPON) standard was completed in 2004, as part of the Ethernet First Mile project. EPON uses standard 802.3 Ethernet frames with symmetric 1 gigabit per second upstream and downstream rates. EPON is applicable for data-centric networks, as well as full-service voice, data and video networks. 10Gbit/s EPON or 10G-EPON was ratified as an amendment IEEE 802.3av to IEEE 802.3. 10G-EPON supports 10/1 Gbit/s. The downstream wavelength plan support simultaneous operation of 10 Gbit/s on one wavelength and 1 Gbit/s on a separate wavelength for operation of IEEE 802.3av and IEEE 802.3ah on the same PON concurrently. The upstream channel can support simultaneous operation of IEEE 802.3av and 1 Gbit/s 802.3ah simultaneously on a single shared (1,310 nm) channel.

Network elements

A PON takes advantage of wavelength division multiplexing (WDM), using one wavelength for downstream traffic and another for upstream traffic on a single nondispersion-shifted fiber (ITU-T G.652). BPON, EPON, GEPON, and GPON have the same basic wavelength plan and use the 1,490 nanometer (nm) wavelength for downstream traffic and 1310 nm wavelength for upstream traffic. 1550 nm is reserved for optional overlay services, typically RF (analog) video.

As with bit rate, the standards describe several optical budgets, most common is 28 dB of loss budget for both BPON and GPON, but products have been announced using less expensive optics as well. 28 dB corresponds to about 20 km with a 32-way split. Forward error correction (FEC) may provide another 2–3 dB of loss budget on GPON systems. As optics improve, the 28 dB budget will likely increase. Although both the GPON and EPON protocols permit large split ratios (up to 128 subscribers for GPON, up to 32,768 for EPON), in practice most PONs are deployed with a split ratio of 1x32 or smaller.

A PON consists of a central office node, called an optical line terminal (OLT), one or more user nodes, called optical network units (ONUs) or optical network terminals (ONTs), and the fibers and splitters between them, called the optical distribution network (ODN). ONT is an ITU-T term to describe a special, single-user case of an ONU. In Multiple Tenant Units, the ONU may be bridged to a customer premise device within the individual dwelling unit using technologies such as Ethernet over twisted pair, G.hn (a high-speed ITU-T standard that can operate over any existing home wiring - power lines, phone lines and coaxial cables) or DSL. An ONU is a device that terminates the PON and presents customer service interfaces to the user. Some ONUs implement a separate subscriber unit to provide services such as telephony, Ethernet data, or video.

The OLT provides the interface between the PON and the service providers network services. These typically include:

- Internet Protocol (IP) traffic over gigabit/s, 10 Gbit/s, or 100 Mbit/s Ethernet
- standard time division multiplexed (TDM) interfaces such as SONET or SDH
- ATM UNI at 155–622 Mbit/s

The ONT or ONU terminates the PON and presents the native service interfaces to the user. These services can include voice (plain old telephone service (POTS) or voice over IP (VoIP)), data (typically Ethernet or V.35), video, and/or telemetry (TTL, ECL, RS530, etc.). Often, the ONU functions are separated into two parts:

- the ONU, which terminates the PON and presents a converged interface – such as xDSL, coax, or multiservice Ethernet – toward the user, and

- network termination equipment (NTE), which provides the separate, native service interfaces directly to the user

A PON is a shared network, in that the OLT sends a single stream of downstream traffic that is seen by all ONUs. Each ONU only reads the content of those packets that are addressed to it. Encryption is used to prevent eavesdropping on downstream traffic.

Upstream bandwidth allocation

The OLT is responsible for allocating upstream bandwidth to the ONUs. Because the optical distribution network (ODN) is shared, ONU upstream transmissions could collide if they were transmitted at random times. ONUs can lie at varying distances from the OLT, meaning that the transmission delay from each ONU is unique. The OLT measures delay and sets a register in each ONU via PLOAM (physical layer operations and maintenance) messages to equalize its delay with respect to all of the other ONUs on the PON.

Once the delay of all ONUs has been set, the OLT transmits so-called grants to the individual ONUs. A grant is permission to use a defined interval of time for upstream transmission. The grant map is dynamically re-calculated every few milliseconds. The map allocates bandwidth to all ONUs, such that each ONU receives timely bandwidth for its service needs.

Some services – POTS, for example – require essentially constant upstream bandwidth, and the OLT may provide a fixed bandwidth allocation to each such service that has been provisioned. DS1 and some classes of data service may also require constant upstream bit rate. But much data traffic – internet surfing, for example – is bursty and highly variable. Through dynamic bandwidth allocation (DBA), a PON can be oversubscribed for upstream traffic, according to the traffic engineering concepts of statistical multiplexing. (Downstream traffic can also be oversubscribed, in the same way that any LAN can be oversubscribed. The only special feature in the PON architecture for downstream oversubscription is the fact that the ONU must be able to accept completely arbitrary downstream time slots, both in time and in size.)

In GPON there are two forms of DBA, status-reporting (SR) and non-status reporting (NSR).

In NSR DBA, the OLT continuously allocates a small amount of extra bandwidth to each ONU. If the ONU has no traffic to send, it transmits idle frames during its excess allocation. If the OLT observes that a given ONU is not sending idle frames, it increases the bandwidth allocation to that ONU. Once the ONU's burst has been transferred, the OLT observes a large number of idle frames from the given ONU, and reduces its allocation accordingly. NSR DBA has the advantage that it imposes no requirements on the ONU, and the disadvantage that there is no way for the OLT to know how best to assign bandwidth across several ONUs that need more.

In SR DBA, the OLT polls ONUs for their backlogs. A given ONU may have several so-called traffic containers (T-CONTs), each with its own priority or traffic class. The ONU reports each T-CONT separately to the OLT. The report message contains a logarithmic measure of the backlog in the T-CONT queue. By knowledge of the service level agreement for each T-CONT across the entire PON, as well as the size of each T-CONT's backlog, the OLT can optimize allocation of the spare bandwidth on the PON.

EPON systems use a DBA mechanism equivalent to GPON's SR DBA solution. The OLT polls ONUs for their queue status and grants bandwidth using the MPCP GATE message, while ONUs report their status using the MPCP REPORT message.

Current status

TDM-PON

Both APON/BPON and EPON/GEPON have been deployed widely, but most networks designed in 2008 use GPON or GEPON. GPON has fewer than 2 million installed ports. GEPON has approximately 30 million deployed ports. For TDM-PON, a passive power splitter is used as the remote terminal. Each ONUs (Optical network units) signals are multiplexed in the time domain. ONUs see their own data through the address labels embedded in the signal.

DOCSIS Provisioning of EPON or DPoE

Data Over Cable Service Interface Specification (DOCSIS) Provisioning of Ethernet Passive Optical Network, or DPoE, is a set of Cable Television Laboratory specifications that implement the DOCSIS service layer interface on existing Ethernet PON (EPON, GEPON or 10G-EPON) Media Access Control (MAC) and Physical layer (PHY) standards. In short it implements the DOCSIS Operations Administration Maintenance and Provisioning (OAMP) functionality on existing EPON equipment. It makes the EPON OLT look and act like a DOCSIS Cable Modem Termination Systems (CMTS) platform (which is called a DPoE System in DPoE terminology). In addition to the offering the same IP service capabilities as a CMTS, DPoE supports Metro Ethernet Forum (MEF) 9 and 14 services for the delivery of Ethernet services for business customers.

RFoG

Radio Frequency over Glass (RFoG) is a type of passive optical networking, that transports RF signals that are now transported over copper (principally over a hybrid fiber and coaxial cable) over PON. In the forward direction RFoGis either a stand alone P2MP system or an optical overlay for existing PON such as GEPON/EPON. The overlay for RFoG is based on Wave Division Multiplexing (WDM) -- the passive combination of wavelengths on a single strand of glass. Reverse RF support is provided by transporting the upstream or return RF into on a separate lambda from the PON return wavelength. The Society of Cable and Telecommunications Engineers (SCTE) Interface

Practices Subcomittee (IPS) Work Group 5, is currently working on IPS 910 RF over Glass. RFoG offers backwards compatibility with existing RF modulation technology, but offers no additional bandwidth for RF based services. Although not yet completed, the RFoG standard is actually a collection of standardized options which are not compatible with each other (they cannot be mixed on the same PON). Some of the standards may interoperate with other PONs, others may not. It offers a means to support RF technologies in locations where only fiber is available or where copper is not permitted or feasible. This technology is targeted towards Cable TV operators and their existing HFC networks. Some describe RFoG as "all of the cost of fiber to the home with none of the benefits."

WDM-PON

Wavelength Division Multiplexing PON, or WDM-PON, is a non-standard type of passive optical networking, being developed by some companies.

The multiple wavelengths of a WDM-PON can be used to separate Optical Network Units (ONUs) into several virtual PONs co-existing on the same physical infrastructure. Alternatively the wavelengths can be used collectively through statistical multiplexing to provide efficient wavelength utilization and lower delays experienced by the ONUs.

There is no common standard for WDM-PON nor any unanimously agreed upon definition of the term. By some definitions WDM-PON is a dedicated wavelength for each ONU. Other more liberal definitions suggest the use of more than one wavelength in any one direction on a PON is WDM-PON. It is difficult to point to an un-biased list of WDM-PON vendors when there is no such unanimous definition. PONs provide higher bandwidth than traditional copper based access networks. WDM-PON has better privacy and better scalability because of each ONU only receives its own wavelength.

Advantages: The MAC layer is simplified because the P2P connections between OLT and ONUs are realized in wavelength domain, so no P2MP media access control is needed. In WDM-PON each wavelength can run at a different speed and protocol so there is a easy pay-as-you-grow upgrade

Challenges: High cost of initial set-up, the cost of the WDM components. Temperature control is another challenge because of how wavelengths tend to drift with environmental temperatures.

Long-Reach Optical Access Networks

The concept of the Long-Reach Optical Access Network (LROAN) is to replace the optical/electrical/optical conversion that takes place at the local exchange with a continuous optical path that extends from the customer to the core of the network. Work by Davey and Payne at BT showed that significant cost savings could be made by reducing the electronic equipment and real-estate required at the local exchange or wire center. A proof of concept demonstrator showed that it was possible to serve 1024 at 10GBit/s with 100km reach.

This technology has sometimes been termed Long-Reach PON, however, many argue that the term PON is no longer applicable as, in most instances, only the distribution remains passive.

Enabling technologies

Due to the topology of PON, the transmission modes for downstream (i.e., from OLT to ONU) and upstream (i.e., from ONU to OLT) are different. For the downstream transmission, the OLT broadcasts optical signal to all the ONUs in continuous mode (CM), i.e., the downstream channel always has optical data signal. However, in the upstream channel, ONUs can not transmit optical data signal in CM. Use of CM would result in all of the signals transmitted from the ONUs converging (with attenuation) into one fiber by the power splitter (serving as power coupler), and overlapping. To solve this problem, burst mode (BM) transmission is adopted for upstream channel. The given ONU only transmits optical packet when it is allocated a time slot and it needs to transmit, and all the ONUs share the upstream channel in the time division multiplexing (TDM) mode. The phases of the BM optical packets received by the OLT are different from packet to packet, since the ONUs are not synchronized to transmit optical packet in the same phase, and the distance between OLT and given ONU are random. Since the distance between the OLT and ONUs are not uniform, the optical packets received by the OLT may have different amplitudes. In order to compensate the phase variation and amplitude variation in a short time (e.g., within 40 ns for GPON), burst mode clock and data recovery (BM-CDR) and burst mode amplifier (e.g., burst mode TIA) need to be employed, respectively. Furthermore, the BM transmission mode requires the transmitter to work in burst mode. Such a burst mode transmitter is able to turn on and off in short time. The above three kinds of circuitries in PON are quite different from their counterparts in the point-to-point continuous mode optical communication link.

Fiber to the premises

Passive optical networks do not use electrically powered components to split the signal. Instead, the signal is distributed using beam splitters. Each splitter typically splits the signal from a single fiber into 16, 32, or 64 fibers, depending on the manufacturer, and several splitters can be aggregated in a single cabinet. A beam splitter cannot provide any switching or buffering capabilities; the resulting connection is called a point-to-multipoint link. For such a connection, the optical network terminals on the customer's end must perform some special functions which would not otherwise be required. For example, due to the absence of switching capabilities, each signal leaving the central office must be broadcast to all users served by that splitter (including to those for whom the signal is not intended). It is therefore up to the optical network terminal to filter out any signals intended for other customers. In addition, since beam splitters cannot perform buffering, each individual optical network terminal must be coordinated in a multiplexing scheme to prevent signals leaving the customer from colliding at the intersection. Two types of multiplexing are possible for achieving this: wavelength-division multiplexing and time-division multiplexing. With wavelength-division multiplexing, each customer transmits their signal using a unique wavelength. With time-division multiplexing (TDM), the customers "take turns" transmitting information. TDM equipment has been on the market longest; WDM-PON equipment became available in 2005.

Passive optical networks have both advantages and disadvantages over active networks. They avoid the complexities involved in keeping electronic equipment operating outdoors. They also allow for analog broadcasts, which can simplify the delivery of analog television. However, because each signal must be pushed out to everyone served by the splitter (rather than to just a single switching device), the central office must be equipped with a particularly powerful piece of transmitting equipment called an optical line terminal (OLT). In addition, because each customer's optical network terminal must transmit all the way to the central office (rather than to just the nearest switching device), customers can't be as far from the central office as is possible with active optical networks.

Passive optical components

The drivers behind the modern passive optical network are the optical components that enable Quality of Service (QoS).

Single-mode, passive optical components include branching devices such as Wavelength-Division Multiplexer/Demultiplexers– (WDMs), isolators, circulators, and filters. These components are used in interoffice, loop feeder, Fiber In The Loop (FITL), Hybrid Fiber-Coaxial Cable (HFC), Synchronous Optical Network (SONET), and Synchronous Digital Hierarchy (SDH) systems; and other telecommunications networks employing optical communications systems that utilize Optical Fiber Amplifiers (OFAs) and Dense Wavelength Division Multiplexer (DWDM) systems. Industry proposed requirements for these components are detailed in GR-1209, Generic Requirements for Passive Optical Components.

The broad variety of passive optical components applications include multichannel transmission, distribution, optical taps for monitoring, pump combiners for fiber amplifiers, bit-rate limiters, optical connects, route diversity, polarization diversity, interferometers, and conherent communication.

WDMs are optical components in which power is split or combined based on the wavelength composition of the optical signal. Dense Wavelength Division Multiplexers (DWDMs) are optical components that split power over at least four wavelengths. Wavelength insensitive couplers are passive optical components in which power is split or combined independently of the wavelength composition of the optical signal. A given component may combine and divide optical signals simultaneously, as in bidirectional (duplex) transmission over a single fiber. Passive optical components are data format transparent, combining and dividing optical power in some predetermined ratio (coupling ratio) regardless of the information content of the signals. WDMs can be thought of as wavelength splitters and combiners. Wavelength insensitive couplers can be thought of as power splitters and combiners.

An optical isolator is a two-port passive component that allows light (in a given wavelength range) to pass through with low attenuation in one direction, while isolating (providing a high attenuation for) light propagating in the reverse direction. Isolators are used as both integral and in-line components in laser diode modules and optical amplifiers, and to reduce noise caused by multi-path reflection in highbit-rate and analog transmission systems.

An optical circulator operates in a similar way to an optical isolator, except that the reverse propagating lightwave is directed to a third port for output, instead of being lost. An optical circulator can be used for bidirectional transmission, as a type of branching component that distributes (and isolates) optical power among fibers, based on the direction of the lightwave propagation.

A fiber optic filter is a component with two or more ports that provides wavelength sensitive loss, isolation and/or return loss. Fiber optic filters are in-line, wavelength selective, components that allow a specific range of wavelengths to pass through (or reflect) with low attenuation for classification of filter types).

GR-1221-CORE, Generic Reliability Assurance Requirements for Passive Optical Components, addresses the long-term reliability of passive optical components.

Chapter- 10

Ambient Network and Bus Network

Ambient network

Ambient Networks is a network integration design that seeks to solve problems relating to switching between networks to maintain contact with the outside world. This project aims to develop a network software-driven infrastructure that will run on top of all current or future network physical infrastructures to provide a way for devices to connect to each other, and through each other to the outside world.

The concept of Ambient Networks comes from the IST Ambient Network project, which is a research project sponsored by the European Commission within the Sixth Framework Programme (FP6).

The Ambient Networks Project

Ambient Networks is a large-scale collaborative project within the European Union's Sixth Framework Programme that investigates future communications systems beyond today's fixed and 3rd generation mobile networks. It is part of the Wireless World Initiative. The project works at a new concept called Ambient Networking, to provide suitable mobile networking technology for the future mobile and wireless communications environment. Ambient Networks aims to provide a unified networking concept that can adapt to the very heterogeneous environment of different radio technologies and service and network environments. Special focus is put on facilitating both competition and cooperation of various market players by defining interfaces, which allow the instant negotiation of agreements. This approach goes clearly beyond interworking of well-defined protocols and is expected to have a long-term effect on the business landscape in the Wireless World. Central to the project is the concept of composition of networks, which is an approach to address the dynamic nature of the target environment. The approach is based on an open framework for network control functionality, which can be extended with new capabilities as well as operating over existing connectivity infrastructure.

- **Phase 1** of the project (2004-2005) has laid the conceptual foundations. The Deliverable D1-5 "Ambient Networks Framework Architecture" summarizes the work from phase 1 and provides links to other relevant material.

- Ambient Networks **Phase 2** (2006-2007) focuses on validation aspects. One key result of phase 2 is an integrated prototype that will be used to study the feasibility of the Ambient Networks concept for a number of typical network scenarios. The ACS prototype will be used to iteratively test the components developed by the project in a real implementation. In parallel, the top-down work is being continued which will lead to a refined System Specification. This document, referred to as the System Description, is available on the Ambient Networks website. Furthermore, standardization of the composition concept is addressed in 3GPP.

Interfaces and their use

The ACS (Ambient Control Space) is the internal of an Ambient Network. It has the functions that can be accessed and it is in full control of the resources of the network. The Ambient Networks infrastructure does not deal with nodes, instead it deals with networks, though at the beginning, all the "networks" might only consist of just one node: these "networks" need to merge to form a network in the original sense of the word. A composition establishment consists of the negotiation and then the realization of a Composition Agreement. This merging can happen be fully automatic. The decision to merge or not is decided using pre-configured policies.

There are three interfaces present to communicate with an ACS. These are:

- ANI: Ambient Network Interface. If a network wants to join in, it has to do so through this interface.
- ASI: Ambient Service Interface. If a function needs to be accessed inside the ACS, this Interface is used.
- ARI: Ambient Resource Interface. If a resource inside a network needs to be accessed (e.g. the volume of the traffic), this interface is used.

Interfaces are used to hide the internal structures of the underlying network.

If two networks meet, and decide to merge, a new ACS will be formed of the two (though the two networks will have their own ACS along with the interfaces inside this global, new ACS). The newly composed ACS will of course have its own ANI, ASI and ARI, and will use these interfaces to merge with other Ambient Networks. Other options for composition are to not merge the two Ambient Networks (Network Interworking) or to establish a new virtual ACS that exercises joint control over a given set of shared resources (Control Sharing).

ACS Functional Entities

Functions are divided into Functional Entities (FEs). The ACS provides a flexible and extensible framework to run these FEs as a distributed system. Examples are

- Composition Functional Entity: Controlling composition of ANs

- Bearer Management FEs
- Overlay Management FEs

More information on FEs is contained in the Ambient Networks Framework Architecture and the latest version of the System Description.

Example Situation

Alice has a PAN, a Personal Area Network on her body: she has a Bluetooth enabled PDA, mobile phone and laptop that she is carrying, and are all currently turned on, and forming a network. Her laptop also has the ability to connect using an available WLAN, and her mobile phone has the ability to connect through GPRS, though GPRS is slower and much more costly for Alice to use. She is now on the move, and her laptop is downloading her emails using the GPRS connection on the mobile:

Laptop -> (Bluetooth) -> Mobile -> (GPRS) -> Mobile phone network

While walking, she passes into an area covered by a free WLAN hotspot: Her PAN now immediately starts to initiate a connection with the hotspot. This is called "merging" of the networks (that of the hotspot and that of her PAN). Once this merging is complete, the downloading of her email continues totally unaffected, but instead of using the expensive and slow GPRS connection, it is now using the newly established WLAN connection. If she now wants to browse the web with her PDA, the PDA will also use the WLAN connection of the laptop:

PDA-> (bluetooth) -> Laptop-> (WLAN) -> Hotspot

Bus network

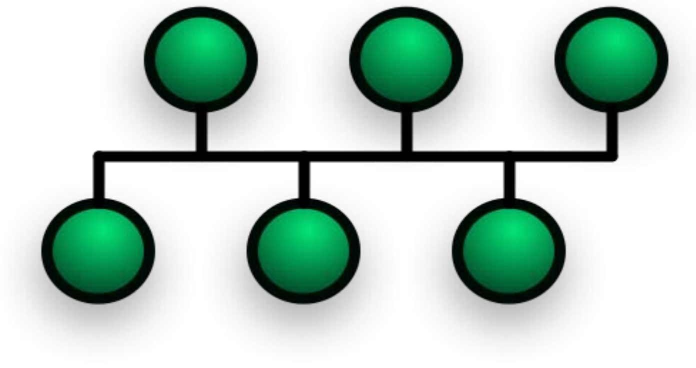

Bus network layout

A **bus network topology** is a network architecture in which a set of clients are connected via a shared communications line, called a bus. There are several common instances of the bus architecture, including one in the motherboard of most computers, and those in some versions of Ethernet networks.

How it works

Bus networks are the simplest way to connect multiple clients, but may have problems when two clients want to transmit at the same time on the same bus. Thus systems which use bus network architectures normally have some scheme of collision handling or collision avoidance for communication on the bus, quite often using Carrier Sense Multiple Access or the presence of a bus master which controls access to the shared bus resource.

A true bus network is passive – the computers on the bus simply listen for a signal; they are not responsible for moving the signal along. However, many active architectures can also be described as a "bus", as they provide the same logical functions as a passive bus; for example, switched Ethernet can still be regarded as a logical network, if not a physical one. Indeed, the hardware may be abstracted away completely in the case of a software bus.

With the dominance of switched Ethernet over passive Ethernet, passive bus networks are uncommon in wired networks. However, almost all current wireless networks can be viewed as examples of passive bus networks, with radio propagation serving as the shared passive medium.

The bus topology makes the addition of new devices straightforward. The term used to describe clients is station or workstation in this type of network. Bus network topology uses a broadcast channel which means that all attached stations can hear every transmission and all stations have equal priority in using the network to transmit data.

The Ethernet bus topology works like a big telephone party line — before any device can send a packet, devices on the bus must first determine that no other device is sending a packet on the cable. When a device sends its packet out over the bus, every other network card on the bus sees and reads the packet. Ethernet’s scheme of having devices communicate like they were in chat room is called Carrier Sense Multiple Access/ Collision Detection (CSMA/CD). Sometimes two cards talk (send packets) at the same time. This creates a collision, and the cards themselves arbitrate to decide which one will resend its packet first. All PCs on a bus network share a common wire, which also means they share the data transfer capacity of that wire – or, in tech terms, they share its bandwidth.

This creates an interesting effect. Ten PCs chatting on a bus each get to use a much higher proportion of its total bandwidth than, for instance, 100 PCs on the same bus (in this case, one – tenth compared to one – hundredth). The more PCs on a bus, the more likely you’ll have a communication traffic jam.

Advantages and disadvantages of a bus network

Advantages

- Easy to implement and extend.
- Easy to install.
- Well-suited for temporary or small networks not requiring high speeds (quick setup), resulting in faster networks.
- Cheaper than other topologies (But in recent years has became less important due to devices like a switch)
- Cost effective; only a single cable is used.
- Easy identification of cable faults.
- Reduced weight due to fewer wires.

Disadvantages

- Limited cable length and number of stations.
- If there is a problem with the cable, the entire network breaks down.
- Maintenance costs may be higher in the long run.
- Performance degrades as additional computers are added or on heavy traffic (shared bandwidth).
- Proper termination is required (loop must be in closed path).
- Significant Capacitive Load (each bus transaction must be able to stretch to most distant link).
- It works best with limited number of nodes.
- Commonly has a slower data transfer rate than other topologies.
- Only one packet can remain on the bus during one clock pulse.

Chapter- 11

Delay-tolerant Networking

Delay-tolerant networking (DTN) is an approach to computer network architecture that seeks to address the technical issues in heterogeneous networks that may lack continuous network connectivity. Examples of such networks are those operating in mobile or extreme terrestrial environments, or planned networks in space.

Recently, the term **disruption-tolerant networking** has gained currency in the United States due to support from DARPA, which has funded many DTN projects. Disruption may occur because of the limits of wireless radio range, sparsity of mobile nodes, energy resources, attack, and noise.

History

In the 1970s, spurred by the micronization of computing, researchers began developing technology for routing between non-fixed locations of computers. While the field of ad-hoc routing was inactive throughout the 1980s, the widespread use of wireless protocols reinvigorated the field in the 1990s as mobile ad-hoc networking (MANET) and vehicular ad-hoc networking became areas of increasing interest.

Concurrently with (but separate from) the MANET activities, DARPA had funded NASA, MITRE and others to develop a proposal for the Interplanetary Internet (IPN). Internet pioneer Vint Cerf and others developed the initial IPN architecture, relating to the necessity of networking technologies that can cope with the significant delays and packet corruption of deep-space communications. In 2002, Kevin Fall started to adapt some of the ideas in the IPN design to terrestrial networks and coined the term delay-tolerant networking and the DTN acronym. A paper published in 2003 SIGCOMM conference gives the motivation for DTNs. The mid-2000s brought about increased interest in DTNs, including a growing number of academic conferences on delay and disruption-tolerant networking, and growing interest in combining work from sensor networks and MANETs with the work on DTN. This field saw many optimizations on classic ad-hoc and delay-tolerant networking algorithms and began to examine factors such as security, reliability, verifiability, and other areas of research that are well understood in traditional computer networking.

Routing

The ability to transport, or route, data from a source to a destination is a fundamental ability all communication networks must have. Delay and disruption-tolerant networks (DTNs), are characterized by their lack of connectivity, resulting in a lack of instantaneous end-to-end paths. In these challenging environments, popular ad hoc routing protocols such as AODV and DSR fail to establish routes. This is due to these protocols trying to first establish a complete route and then, after the route has been established, forward the actual data. However, when instantaneous end-to-end paths are difficult or impossible to establish, routing protocols must take to a "store and forward" approach, where data is incrementally moved and stored throughout the network in hopes that it will eventually reach its destination. A common technique used to maximize the probability of a message being successfully transferred is to replicate many copies of the message in the hope that one will succeed in reaching its destination. This is feasible only on networks with large amounts of local storage and internode bandwidth relative to the expected traffic. In many common problem spaces, this inefficiency is outweighed by the increased efficiency and shortened delivery times made possible by taking maximum advantage of available unscheduled forwarding opportunities. In others, where available storage and internode throughput opportunities are more tightly constrained, a more discriminate algorithm is required.

Other concerns

Bundle protocols

In efforts to provide a shared framework for algorithm and application development in DTNs, RFC 4838 and RFC 5050 were published in 2007 to define a common abstraction to software running on disrupted networks. Commonly known as the Bundle Protocol, this protocol defines a series of contiguous data blocks as a bundle—where each bundle contains enough semantic information to allow the application to make progress where an individual block may not. Bundles are routed in a store and forward manner between participating nodes over varied network transport technologies (including both IP and non-IP based transports). The transport layers carrying the bundles across their local networks are called bundle convergence layers. The bundle architecture therefore operates as an overlay network, providing a new naming architecture based on Endpoint Identifiers (EIDs) and coarse-grained class of service offerings.

Protocols using bundling must leverage application-level preferences for sending bundles across a network. Due to the store and forward nature of delay-tolerant protocols, routing solutions for delay-tolerant networks can benefit from exposure to application-layer information. For example, network scheduling can be influenced if application data must be received in its entirety, quickly, or without variation in packet delay. Bundle protocols collect application data into bundles that can be sent across heterogeneous network configurations with high-level service guarantees. The service guarantees are generally set by the application level, and the RFC 5050 Bundle Protocol specification includes 'bulk', 'normal', and 'expedited' markings.

Security

Addressing security issues has been a major focus of the bundle protocol.

Security concerns for delay-tolerant networks vary depending on the environment and application, though authentication and privacy are often critical. These security guarantees are difficult to establish in a network without persistent connectivity because the network hinders complicated cryptographic protocols, hinders key exchange, and each device must identify other intermittently-visible devices. Solutions have typically been modified from mobile ad hoc network and distributed security research, such as the use of distributed certificate authorities and PKI schemes. Original solutions from the delay-tolerant research community include: 1) the use of identity-based encryption, which allows nodes to receive information encrypted with their public identifier, and 2) the use of tamper-evident tables with a gossiping protocol;

Research efforts

Various research efforts are currently investigating the issues involved with DTN:

- The The Delay-Tolerant Networking Research Group.
- The Technology and Infrastructure for Developing Regions project at UC Berkeley
- The KioskNet research project at the University of Waterloo.
- The DieselNet research project at the University of Massachusetts, Amherst.
- The ResiliNets Research Initiative at the University of Kansas and Lancaster University.
- The Haggle EU research project.
- The N4C EU/FP7 research project.
- The WNaN DARPA project.
- The EMMA project at TU Braunschweig
- The DTN networking at Helsinki University of Technology.
- The SARAH project, funded by the French National Research Agency (ANR).
- The development of the DoDWAN platform at the University of South Brittany.
- The CROWD project, funded by the French National Research Agency (ANR).
- The PodNet project at KTH Stockholm and ETH Zurich.

Some research efforts look at DTN for the Interplanetary Internet by examining use of the Bundle Protocol in space:

- The Saratoga project at the University of Surrey, which was the first to test the bundle protocol in space on the UK-DMC Disaster Monitoring Constellation satellite in 2008.
- NASA JPL's Deep Impact Networking (DINET) Experiment on board the Deep Impact/EPOXI spacecraft.
- BioServe Space Technologies, one of the first payload developers to adopt the DTN technology, has utilized their CGBA (Commercial Generic Bioprocessing

Apparatus) payloads onboard the ISS, which provide computational/communications platforms, to implement the DTN protocol.

Chapter- 12

Fiber to the x

Fiber to the x (**FTTx**) is a generic term for any broadband network architecture that uses optical fiber to replace all or part of the usual metal local loop used for last mile telecommunications. The generic term originated as a generalization of several configurations of fiber deployment (FTTN, FTTC, FTTB, FTTH...), all starting by FTT but differentiated by the last letter, which is substituted by an x in the generalization.

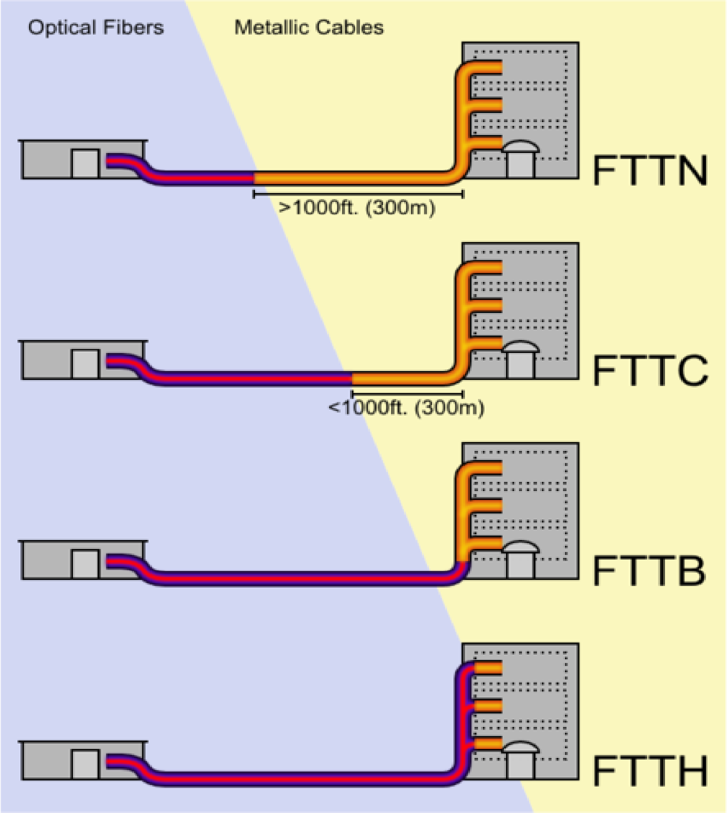

A schematic illustrating how FTTx architectures vary — with regard to the distance between the optical fiber and the end-user. The building on the left is the central office; that on the right is one of the buildings served by the central office. Dotted rectangles represent separate living or office spaces within the same building.

Definition of terms

The telecommunications industry differentiates between several distinct configurations. The terms in most widespread use today are:

- FTTN - Fiber-to-the-node - fiber is terminated in a street cabinet up to several kilometers away from the customer premises, with the final connection being copper.
- FTTC - Fiber-to-the-cabinet or fiber-to-the-curb - this is very similar to FTTN, but the street cabinet is closer to the user's premises; typically within 300 m.
- FTTB - Fiber-to-the-building or Fiber-to-the-basement - fiber reaches the boundary of the building, such as the basement in a multi-dwelling unit, with the final connection to the individual living space being made via alternative means.
- FTTH - Fiber-to-the-home - fiber reaches the boundary of the living space, such as a box on the outside wall of a home.
- FTTP - Fiber-to-the premises - this term is used in several contexts: as a blanket term for both FTTH and FTTB, or where the fiber network includes both homes and small businesses.

To promote consistency, especially when comparing FTTH penetration rates between countries, the three FTTH Councils of Europe, North America and Asia-Pacific have agreed upon definitions for FTTH and FTTB. The FTTH Councils do not have formal definitions for FTTC and FTTN.

It is worth pointing out that fiber to the telecom enclosure (FTTE) is not considered to be part of the FTTx group of technologies, despite the similarity in name. FTTE is a form of structured cabling typically used in the enterprise local area network, where fiber is used to link the main computer equipment room to an enclosure close to the desk or workstation. Similarly, in fiber-to-the-desk a fiber connection is installed from the main computer room to a terminal at the desk.

Benefits of fiber in the access network

The speeds of fiber optic and copper cables are both limited by length, but copper is much more sharply limited in this respect. For example, gigabit Ethernet runs over relatively economical category 5e, category 6, or augmented category 6 unshielded twisted pair copper cabling but only to 100 meters. However, over the right kind of fiber, gigabit ethernet can easily reach distances of tens of kilometers.

Even in the commercial world, most computers have copper communication cables. But these cables are short, typically tens of meters. Most metropolitan network links (e.g., those based on telephone or cable television services) are several kilometers long, in the range where fiber significantly outperforms copper. Replacing at least part of these links with fiber shortens the remaining copper segments and allows them to run much faster.

Fiber configurations that bring fiber right into the building can offer the highest speeds since the remaining segments can use standard Ethernet or coaxial cable. Fiber configurations that transition to copper in a street cabinet are generally too far from the users for standard Ethernet configurations over existing copper cabling. They generally use VDSL at (downstream) speeds of several tens of megabits per second.

Fiber is often said to be 'future proof' because the speed of the broadband connection is usually limited by the terminal equipment rather than the fiber itself, permitting at least some speed improvements by equipment upgrades before the fiber itself must be upgraded. Still, the type and length of employed fibers chosen, e.g. multimode vs single mode, are critical for applicability for future high gigabit connections.

Fiber to the node

Fiber to the node (FTTN), also called fiber to the neighborhood or fiber to the cabinet (FTTCab), is a telecommunication architecture based on fiber-optic cables run to a cabinet serving a neighborhood. Customers typically connect to this cabinet using traditional coaxial cable or twisted pair wiring. The area served by the cabinet is usually less than 1,500 m in radius and can contain several hundred customers. (If the cabinet serves an area of less than 300 m in radius then the architecture is typically called fiber to the curb.)

Fiber to the node allows delivery of broadband services such as high speed Internet. High speed communications protocols such as broadband cable access (typically DOCSIS) or some form of DSL are used between the cabinet and the customers. The data rates vary according to the exact protocol used and according to how close the customer is to the cabinet.

Unlike the competing fiber to the premises technology, fiber to the node often uses the existing coaxial or twisted pair infrastructure to provide last mile service. For this reason, fiber to the node is less costly to deploy. In the long-term, however, its bandwidth potential is limited relative to implementations which bring the fiber still closer to the subscriber.

Fiber to the last amplifier

FTTLA are the initials of Fiber To The Last Amplifier. The network cables being able to use several amplifiers, the FTTLA aims at replacing the coaxial cable to the last amplifier (towards the subscriber) by optical fiber. It acts as a new technology aiming at re-using the network cables existing in particular on the final part while installing of optical fiber more closely to the subscriber while using the coaxial cable of the networks cables for the "last mile" or "last meters" connected with the subscriber.

Fiber to the last amplifier (FttLA) node is an efficient tool to deploy fiber deeper into the CATV network architecture and add most desirable aspects of scalability (performance and reliability) which are necessary when new services (i.e. "triple play", video on demand, gaming) are introduced.

FTTLA is a technology which assists hybrid fiber-coaxial CATV networks to provide to their customers more bandwidth. Using a replacement of all coaxial active equipments by nodes (optical receiver) with high power output (up to 117 dBuV). The coaxial is maintained from the node to the customer without any active equipment in between.

From the optical sender to the node, it uses fiber which is split by 4 or by 8 depending on the distance and on the output power of the optical sender (from 6 to 16 dBm).

Also, IM2, IM3 and C/N are modified for a better network and it also has other benefits such as power saving in the network, as the power consumption is lower than a normal HFC network (up to 40%).

Fiber to the curb

Fiber to the curb (FTTC) is a telecommunications system based on fiber-optic cables run to a platform that serves several customers. Each of these customers has a connection to this platform via coaxial cable or twisted pair.

Fiber to the curb allows delivery of broadband services such as high speed internet. High speed communications protocols such as broadband cable access (typically DOCSIS) or some form of DSL are used between the cabinet and the customers. The data rates vary according to the exact protocol used and according to how close the customer is to the cabinet.

FTTC is subtly distinct from FTTN or FTTP (all are versions of Fiber in the Loop). The chief difference is the placement of the cabinet. FTTC will be placed near the "curb" which differs from FTTN which is placed far from the customer and FTTP which is placed right at the serving location.

Unlike the competing fiber to the premises (FTTP) technology, fiber to the curb can use the existing coaxial or twisted pair infrastructure to provide last mile service. For this reason, fiber to the curb costs less to deploy. However, it also has lower bandwidth potential than fiber to the premises.

In the United States of America and Canada, the largest deployment of FTTC was carried out by BellSouth Telecommunications. With the acquisition of BellSouth by AT&T, deployment of FTTC will end. Future deployments will be based on either FTTN or FTTP. Existing FTTC plant may be removed and replaced with FTTP.

Fiber to the premises

Fiber to the premises is a form of fiber-optic communication delivery in which an optical fiber is run from the central office all the way to the premises occupied by the subscriber.

Fiber to the premises is often abbreviated with the acronym FTTP. However, this acronym has become ambiguous and may instead refer to a form of fiber to the curb where the fiber terminates at a utility pole without reaching the premises.

FTTH vs. FTTB

Fiber to the premises can be categorized according to where the optical fiber ends:

- FTTH (fiber to the home) is a form of fiber optic communication delivery in which the fiber extends from the central office to the subscriber's living or working space. Once at the subscriber's living or working space, the signal may be conveyed throughout the space using any means, including twisted pair, coaxial cable, wireless, power line communication, or optical fiber.

- FTTB (fiber to the building, also called fiber to the basement) is a form of fiber optic communication delivery in which the optical fiber terminates before actually reaching the subscribers living or working space itself, but does extend to the property containing that living or working space. The signal is conveyed the final distance using any non-optical means, including twisted pair, coaxial cable, wireless, or power line communication. By definition, FTTB necessarily applies only to those properties which contain multiple living or working spaces.

An apartment building may provide an example of the distinction between FTTH and FTTB. If a fiber is run to a panel at each subscriber's apartment, this is FTTH. If instead the fiber goes only as far as the apartment building's shared electrical room, then this is FTTB.

Direct fiber

The simplest optical distribution network can be called direct fiber. In this architecture, each fiber leaving the central office goes to exactly one customer. Such networks can provide excellent bandwidth since each customer gets their own dedicated fiber extending all the way to the central office. However, this approach is about 10% more costly due to the amount of fiber and central office machinery required. The approach is generally favored by new entrants and competitive operators. A benefit of this approach is that it doesn't exclude any layer 2 networking technologies, be they Passive optical network, Active Optical Network, etc. From a regulatory point of view it leads to least implications as any form of regulatory remedy is still possible using this topology.

Shared fiber

More commonly each fiber leaving the central office is actually shared by many customers. It is not until such a fiber gets relatively close to the customers that it is split into individual customer-specific fibers. There are two competing optical distribution network architectures which achieve this split: active optical networks (AONs) and passive optical networks (PONs).

Active optical network

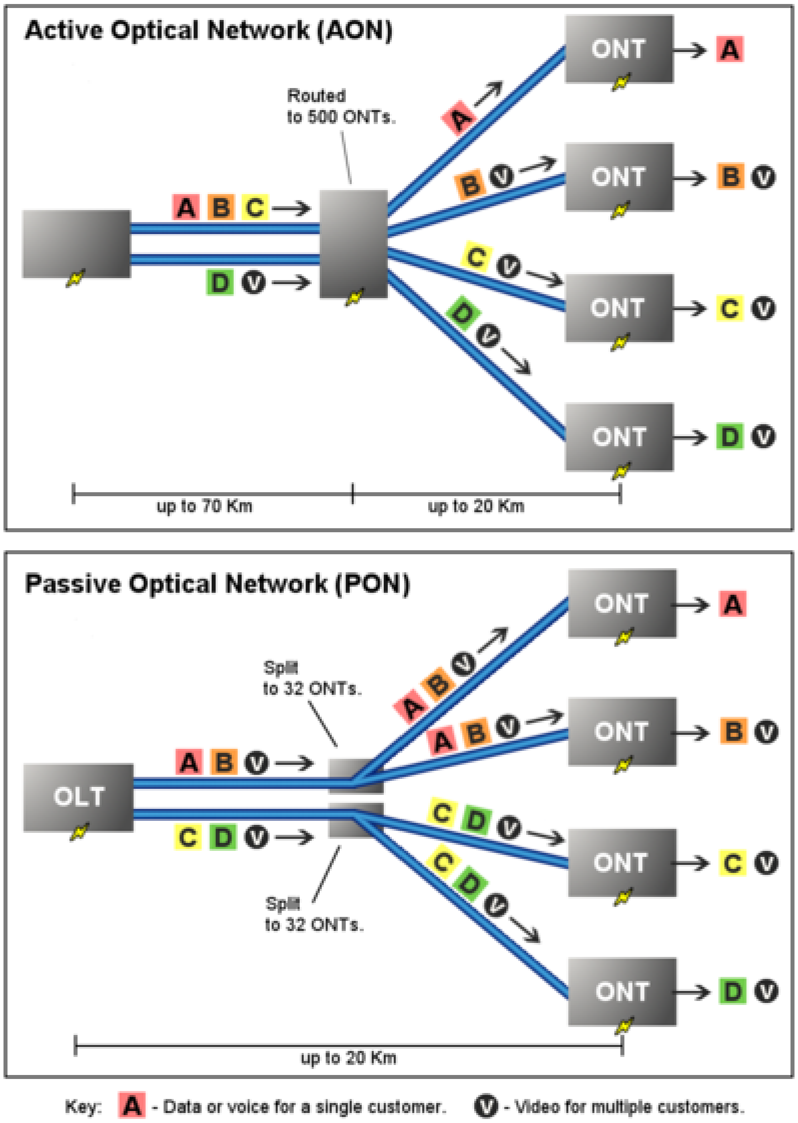

Comparison showing how a typical active optical network handles downstream traffic differently than a typical passive optical network. The type of active optical network shown is a star network capable of multicasting. The type of passive optical network shown is a star network having multiple splitters housed in the same cabinet.

Active optical networks rely on some sort of electrically powered equipment in Optical Distribution Network(ODN) to distribute the signal, such as a switch or router. Normally, optical signals need O-E-O transformation in ODN. Each signal leaving the central office is directed only to the customer for which it is intended. Incoming signals from the customers avoid colliding at the intersection because the powered equipment there provides buffering.

As of 2007, the most common type of active optical networks are called active Ethernet, a type of Ethernet in the first mile (EFM). Active Ethernet uses optical Ethernet switches to distribute the signal, thus incorporating the customers' premises and the central office into one giant switched Ethernet network. Such networks are identical to the Ethernet computer networks used in businesses and academic institutions, except that their purpose is to connect homes and buildings to a central office rather than to connect computers and printers within a campus. Each switching cabinet can handle up to 1,000 customers, although 400-500 is more typical. This neighborhood equipment performs layer 2/layer 3 switching and routing, offloading full layer 3 routing to the carrier's central office. The IEEE 802.3ah standard enables service providers to deliver up to 100 Mbit/s full-duplex over one single-mode optical fiber to the premises depending on the provider. Speeds of 1Gbit/s are becoming commercially available.

Passive optical network

A passive optical network (PON) is a point-to-multipoint, fiber to the premises network architecture in which unpowered optical splitters are used to enable a single optical fiber to serve multiple premises, typically 32-128. A PON configuration reduces the amount of fiber and central office equipment required compared with point to point architectures.

Downstream signal coming from the central office is broadcast to each customer premises sharing a fiber. Encryption is used to prevent eavesdropping.

Upstream signals are combined using a multiple access protocol, invariably time division multiple access (TDMA). The OLTs "range" the ONUs in order to provide time slot assignments for upstream communication.

Electrical portion

Once on private property, the signal typically travels the final distance to the end user's equipment using an electrical format.

A device called an Optical Network Terminal (ONT), also called an Optical Network Unit (ONU), converts the optical signal into an electrical signal. (ONT is an ITU-T term, whereas ONU is an IEEE term, but the two terms mean exactly the same thing.) Optical network terminals require electrical power for their operation, so some providers connect them to back-up batteries in case of power outages. Optical network units use thin film filter technology to convert between optical and electrical signals.

For fiber to the home and for some forms of fiber to the building, it is common for the building's existing phone systems, local area networks, and cable TV systems to connect directly to the ONT.

If all three systems cannot directly reach the ONT, it is possible to combine signals and transport them over a common medium. Once closer to the end-user, equipment such as a router, modem, and/or network interface module can separate the signals and convert them into the appropriate protocol. For example, one solution for apartment buildings uses VDSL to combine data (and / or video) with voice. With this approach, the combined signal travels through the building over the existing telephone wiring until it reaches the end-user's living space. Once there, a VDSL modem copies the data and video signals and converts them into Ethernet protocol. These are then sent over the end user's category 5 cable. A network interface module can then separate out the video signal and convert it into an RF signal that is sent over the end-user's coaxial cable. The voice signal continues to travel over the phone wiring and is sent through DSL filters to remove the video and data signals. An alternative strategy allows data and / or voice to be transmitted over coaxial cable. In yet another strategy, some office buildings dispense with the telephone wiring altogether, instead using voice over Internet Protocol phones that can plug directly into the local area network.

Chapter- 13

GINA: Global Information Network Architecture

The concept for the **Global Information Network Architecture (GINA)** evolved from a realization that the current technologies provided an unprecedented opportunity to create a useful Global Information Grid (GIG) that could transform the possibilities for Net-Centric Operations.

The Global Information Network Architecture (GINA) Team was created in 2004 to address this possibility. Originally developed under Cooperative Research and Development Agreement (CRADA) with The US Naval Postgraduate School (NPS) in Monterey CA, the projects initial title was Network Aware Business Data Management System (NABDMS).

In late 2008, the United States Army Corps of Engineers (USACE) Engineer Research and Development Center (ERDC) began the second phase of GINA's development. Currently focus is on purposing GINA as a High Level Architecture (HLA) for System Fusion Networks (SFN), GINA is being evolved and deployed globally to facilitate interoperability and a new form of computational design.

The Global Information Grid

“The globally interconnected, end-to-end set of information capabilities, associated processes, and personnel for collecting, processing, storing, disseminating and managing information on demand to warfighters, policy makers, and support personnel. The Global Information Grid (GIG) includes all owned and leased communications and computing systems and services, software (including applications), data, security services, and other associated services necessary to achieve Information Superiority. It also includes National Security Systems as defined in section 5142 of the Clinger-Cohen Act of 1996. “

The United States Department of Defense (DoD) recognized back in 1996 the need to have a GIG. We have struggled, largely unsuccessfully to bring a reasonable facsimile of a GIG forward. There have been numerous research projects directed at creating a GIG, but actually turning DoD networks into a functioning GIG has proved elusive. The GIG is a very hard problem that requires a rethinking of current approaches to interoperability.

Vector Relational Data Modeling (VRDM)... Modeling Models.

The dictionary defines modeling as, “The representation, often mathematical, of a process, concept, or operation of a system.”

While software modeling languages such as Unified Modeling Language (UML) or Object Role Modeling (ORM) attempt to provide an iconic representation to articulate the structure or architecture of an application, they are not executable and have limited applicability beyond the software design environment.

The GINA Team created GINA to enable model-based software engineering, but we did so in such a way that the model, once-defined, represented a working application. By properly bounding the problem space to "Information Applications", i.e., non-algorithmically-intense applications with linear relationships representing the vast majority of software applications, the GINA team was able to create a configurable Component Based Object Model (CBOM) for information management where the configuration represented the instructions for assembling a working implementation of the model. Moreover, the GINA team made the decision early in its development to make GINA a GINA model. By doing so the configuration itself could be controlled by configuration. As will be illustrated later, that turned out to be an important decision. Enabling GINA’s deep configurability required the development and implementation of multiple models.

The **Control Model** assembles the components of the component model, according to the assembly instructions in the Application Model into the structures defined in the Implementation Model to create GINA information objects.

The **Application Model** describes actual GINA applications. Both the description of applications and the GINA and Application Models themselves. These applications are described in terms of components in the Component Model. It represents the set of components that are assembled in order to create GINA information objects as specified in the application model.

The Control Model assembles the components of the component model, according to the assembly instructions in the Application Model into the structures defined in the **Implementation Model** to create GINA information objects.

Ultimately, we have to define GINA applications using a development model that is appropriate for developing GINA applications. And again, the GINA **Development Model** is itself described as a GINA application.

GINA, at a high level, is a model for modeling. Designed to facilitate a principle of development where relationships between objects can be objectified and wielded as objects in their own right, GINA itself is an executable model configured using VRDM.

VRDM is a core concept that is embodied by GINA. GINA could be looked at as an environment that turns collected data into a multi-dimensional object environment with each object being connected to other objects through vectors. This environment makes many of the information-centric tasks that a user might want to perform far easier than any other approach.

A key concept of VRDM is that relationships among information objects should themselves be defined as information objects, and be fully configurable.

Taking relationships and implementing them as GINA objects enables GINA to take configurations and assemble models that can perform most, if not all of the work done by typical hard-coded information applications, such as most enterprise systems, integrations, and information sharing systems.

As a result of VRDM, it is possible to specify an application through describing the required components as a series of objects and their relationships, or vectors. VRDM enables disparate data, from disparate sources to be invoked and configured to relate in a “System of Systems” model called a specification. The behavior of a specification can be the much same as a contemporary application. The difference, however, is that with VRDM, there is no programming. GINA is a true CBOM for Object Modeling, where the models are themselves executable.

Critical to this approach is the concept of **Reflexivity**, i.e., the GINA model is described as a GINA model, that permits deep configurability. When one is using the interactive development environment used to create GINA applications, one is using a GINA application. More importantly, GINA assembles GINA applications according to a GINA model for GINA applications. Deeper still, the GINA model is itself an example of a GINA model. This deep configurability is the key to GINA’s power as interoperability and multi-level security ("MLS") engine.

As a combination of hardware and software-based components, which meets the requirements of a true GIG, GINA is configured as a universal virtual network of data of any type from any source in any location on collected physical networks. GINA is a product of several key concepts which collectively enable it to represent a complete, configurable interoperability environment. These key concepts are embodied in a series of layers and components which collectively allow GINA to perform the types of functions and provide the type of services it needs to be a fully functional interoperability environment.

Vector Relational Data Modeling Core Concepts

Descriptive Programming

Just as Assembler made machine language programming faster and more accurate, and descriptive languages for specifying procedures (“4GL”s like SQL) made procedural programming faster and more accurate, GINA's VRDM brings a new level of speed and accuracy to object-oriented programming.

In the long run GINA's encapsulation of the management of network-available data may be more important that even its speed and accuracy for large organizations.

Component Objects

With VRDM, Data Agnostic Objects can be created to represent common relationships called Mechanisms. These Mechanisms can be reused and combined with others, new and existing, to create systems and subsystems. This facilitates rapid deployment and non-programmatic implementation.

Complexity

Just as everything in the world is assembled from remarkably few elements, arbitrarily complex systems can be assembled from relatively few objects. The key in both cases is that objects have to be designed for interaction and assembly. GINA is designed in that way. Fundamentally, at the bottom level are very few objects, e.g., objects, or XTypes in VRDM, and relationships between XTypes, or Vectors in VRDM. These objects are the primitives on which the VRDM object management model is based. In turn, instances of these primitives are assembled into the basic building blocks of VRDM: fully defined objects representing XTypes and Vectors, as well as constraints and simple entities. These can then be assembled to fully describe the GINA environment, and to allow the administrator to create the data objects to support a Task Oriented User Interface (TOUI), or a specific application.

WorldSpace

A central concept in GINA is that objects can be referenced in multiple WorldSpaces, depending on how a user gets to that object. A WorldSpace determines the applicability of an object's vectors, e.g. attributes and relationships, when assembled for a particular event or usage. WorldSpaces are inherently hierarchical: as one more tightly defines the WorldSpace associated with an event or usage, the more tightly one must define, and the more granularly one needs to specify associated behaviors.

HyperPlanes

If we look back at the concepts associated with GINA, we could say that an object exists in a 3 dimensional data object space. Its location in that space is defined by its order of complexity, its usage and related components, and the user and WorldSpace in which it is being accessed. At any given time the behavior of a system is dictated by all of its objects locations in this 3-dimensional object- space. However, this behavior is not the same for every user, and is influenced by the characteristics of that user that effectively define hyperplanes in this object space. Thus, the appropriate object model is only summarized in three dimensions, with multiple user-based dimensions of behavior effectively being summarized on this graphic. In effect, at any given time an object exists as a point in 7+ dimensional object behavior space. Remarkably, GINA not only models this space effectively, but does so in an environment where most specifications are created through configuration, not programming.

Directory Sub System (DSS)

GINA is implemented through a software-based, multi-layer, configurable data object management environment. Just as the entirety of GINA can be viewed as a series of well-structured layers, the data object management environment is also structured and layered, with multiple layers of the object management environment corresponding to each of the top three layers in the overall GINA. GINA’s “DSS” layer is actually composed of two separate implementation layers: a content server layer that consists of a collection of configurable objects that know how to navigate the network, acquire data, and present it in a consistent way; and an aggregation layer that homogenizes all incoming data, in both format and name, and presents itself as a universal object repository that insulates the information consumer from the complexities of managing the underlying data stores.

Data Access Layer (DAL)

GINA collects data from aggregated systems using a collection of adaptors called Content Servers which structure the protocols, formats, and syntax of collected data into a common representation that then becomes the base data that can then be managed through the GINA model. Just as the providers of data to GINA operate on multiple protocols, formats, and syntaxes, the prospective consumers of GINA data may require information using their own protocols, formats, and syntax. GINA exposes itself using a standard "Data Access Layer" ("DAL") that can be—and has been—used to provide data to standardized or customized DALs such as SOAP Web Services, ODBC interfaces, etc.

Task Oriented User Interface (TOUI)

Another model that has been built in GINA is called the Task-Oriented User Interface, or "TOUI". The current mainstream approach to user interfaces ("UIs") involves a process where a developer "paints", or in some other way creates a mark-up of the UI, and then defines the binding of components of that the UI to the underlying application using some standardized approach. The TOUI model takes a different approach: a UI is assembled at the time of request from components according to a set of vectors that take into the account model states and the user during the assembly process. As a result, the UI no longer represents the application that uses information, but rather becomes the external expression of the information model that represents the application. Moreover, because the definition of the UI is done as a set of metadata-defined GINA components, the expression of those components can be done in any environment that has sufficiently strong semantics for representing applications, whether that is Java, .NET, Python, or even a 3-D visualization environment.interfaces, etc.

Chapter- 14

IBM Systems Network Architecture

Systems Network Architecture (**SNA**) is IBM's proprietary networking architecture created in 1974. It is a complete protocol stack for interconnecting computers and their resources. SNA describes the protocol and is, in itself, not actually a program. The implementation of SNA takes the form of various communications packages, most notably Virtual telecommunications access method (VTAM) which is the mainframe package for SNA communications. SNA is still used extensively in banks and other financial transaction networks, as well as in many government agencies. While IBM is still providing support for SNA, one of the primary pieces of hardware, the 3745/3746 communications controller has been withdrawn from marketing by the IBM Corporation. However, there are an estimated 20,000 of these controllers installed and IBM continues to provide hardware maintenance service and micro code features to support users. A robust market of smaller companies continues to provide the 3745/3746, features, parts and service. VTAM is also supported by IBM, as is the IBM Network Control Program (NCP) required by the 3745/3746 controllers.

Objectives of SNA

IBM in the mid-1970s saw itself mainly as a hardware vendor and hence all its innovations in that period aimed to increase hardware sales. SNA's objective was to reduce the costs of operating large numbers of terminals and thus induce customers to develop or expand interactive terminal based-systems as opposed to batch systems. An expansion of interactive terminal based-systems would increase sales of terminals and more importantly of mainframe computers and peripherals - partly because of the simple increase in the volume of work done by the systems and partly because interactive processing requires more computing power per transaction than batch processing.

Hence SNA aimed to reduce the main non-computer costs and other difficulties in operating large networks using earlier communications protocols. The difficulties included:

- A communications line could not be shared by terminals whose users wished to use different types of application, for example one which ran under the control of CICS and another which ran under TSO.
- Often a communications line could not be shared by terminals of different types, as they used different "dialects" of the existing communications protocols. Up to

the early 1970s, computer components were so expensive and bulky that it was not feasible to include all-purpose communications interface cards in terminals. Every type of terminal had a hard-wired communications card which supported only the operation of one type of terminal without compatibility with other types of terminals on the same line.

- The protocols which the primitive communications cards could handle were not efficient. Each communications line used more time transmitting data than modern lines do.
- Telecommunications lines at the time were of much lower quality. For example, it was almost impossible to run a dial-up line at more than 300 bits per second because of the overwhelming error rate, as comparing with 56,000 bits per second today on dial-up lines; and in the early 1970s few leased lines were run at more than 2400 bits per second (these low speeds are a consequence of Shannon's Law in a relatively low-technology environment). Telecommunications companies had little incentive to improve line quality or reduce costs, because at the time they were mostly monopolies and sometimes state-owned.

As a result running a large number of terminals required a lot more communications lines than the number required today, especially if different types of terminals needed to be supported, or the users wanted to use different types of applications (.e.g. under CICS or TSO) from the same location. In purely financial terms SNA's objectives were to increase customers' spending on terminal-based systems and at the same time to increase IBM's share of that spending, mainly at the expense of the telecommunications companies.

SNA also aimed to overcome a limitation of the architecture which IBM's System/370 mainframes inherited from System/360. Each CPU could connect to at most 16 "channels" (devices which acted as controllers for peripherals such as tape and disk drives, printers, card-readers) and each channel could handle up to 16 peripherals - i.e. there was maximum of 256 peripherals per CPU. At the time when SNA was designed, each communications line counted as a peripheral. Thus the number of terminals with which powerful mainframe could otherwise communicate is severely limited.

Principal components and technologies

Improvements in computer component technology made it feasible to build terminals that included more powerful communications cards which could operate a single standard communications protocol rather than a very stripped-down protocol which suited only a specific type of terminal. As a result several multi-layer communications protocols were proposed in the 1970s, of which IBM's SNA and ITU-T's X.25 became dominant later.

The most important elements of SNA include:

- IBM Network Control Program (NCP) is a primitive switching protocol, implemented in 3705 communications processors. The protocol performed two main functions:

    - It is a packet forwarding protocol, acting like modern switch - forwarding data packages to the next node, which might be a mainframe, a terminal or another 3705. The communications processors supported only hierarchical networks with a mainframe at the center, unlike modern routers which support peer-to-peer networks in which a machine at the end of the line can be both a client and a server at the same time.
    - It is a multiplexer that connected multiple terminals into one communication line to the CPU, thus relieved the constraints on the maximum number of communication lines per CPU. A 3705 could support a larger number of lines (352 initially) but only counted as one peripheral by the CPUs and channels. Since the launch of SNA IBM has introduced improved communications processors, of which the latest is the 3745.
- Synchronous Data Link Control (SDLC), a protocol which greatly improved the efficiency of data transfer over a single link:
    - SDLC included much more powerful error detection and correction codes than earlier protocols. These codes often enabled the communications cards to correct minor transmission errors without requesting re-transmission, and therefore made it possible to pump data down a line much faster.
    - It enabled terminals and 3705 communications processors to send "frames" of data one after the other without waiting for an acknowledgement of the previous frame - the communications cards had sufficient memory and processing capacity to "remember" the last 7 frames sent or received, request re-transmission of only those frames which contained errors that the error detection and correction codes could not repair, and slot the re-transmitted frames into the right place in the sequence before forwarding them to the next stage.
    - These frames all had the same type of "envelope" (frame header and trailer) which contained enough information for data packages from different types of terminal to be send along the same communications line, leaving the mainframe to deal with any differences in the formatting of the content or in the rules governing dialogs with different types of terminal.

  Remote terminals (i.e. those connected to the mainframe by telephone lines) and 3705 communications processors would have SDLC-capable communications cards.
  This is the precursor of the so called "packet communication" that eventually evolved into today's IP technology, and SDLC itself evolved into HDLC that is one of the base technology for dedicated telecommunication circuit.

- VTAM, a software package to provide log-in, session keeping and routing services within the mainframe. A terminal user would log-in via VTAM to a specific application or application environment (e.g. CICS or TSO). A VTAM device would then route data from that terminal to the appropriate application or application environment until the user logged out and possibly logged in to another application. The original versions of IBM hardware could only keep one session per terminal. In the 1980s further software (mainly from third-party vendors) made it possible for a terminal to have simultaneous sessions with different applications or application environments.

Advantages and disadvantages

SNA removed link control from the application program and placed it in the NCP. This had the following advantages and disadvantages:

Advantages

- Localization of problems in the telecommunications network was easier because a relatively small amount of software actually dealt with communication links. There was a single error reporting system.
- Adding communication capability to an application program was much easier because the formidable area of link control software that typically requires interrupt processors and software timers was relegated to system software and NCP.
- With the advent of APPN, routing functionality was the responsibility of the computer as opposed to the router (as with TCP/IP networks). Each computer maintained a list of Nodes that defined the forwarding mechanisms. A centralized node type known as a Network Node maintained Global tables of all other node types. APPN stopped the need to maintain APPC routing tables that explicitly defined endpoint to endpoint connectivity. APPN sessions would route to endpoints through other allowed node types until it found the destination. This was similar to the way that TCP/IP routers function today.

Disadvantages

- Connection to non-SNA networks was difficult. An application which needed access to some communication scheme, which was not supported in the current version of SNA, faced obstacles. Before IBM included X.25 support (NPSI) in SNA, connecting to an X.25 network would have been awkward. Conversion between X.25 and SNA protocols could have been provided either by NCP software modifications or by an external protocol converter.

- A sheaf of alternate pathways between every pair of nodes in a network had to be predesigned and stored centrally. Choice among these pathways by SNA was rigid and did not take advantage of current link loads for optimum speed.

- SNA network installation and maintenance are complicated and SNA network products are (or were) expensive. Attempts to reduce SNA network complexity by adding IBM Advanced Peer-to-Peer Networking functionality were not really successful, if only because the migration from traditional SNA to SNA/APPN was very complex, without providing much additional value, at least initially. SNA software licences (VTAM) cost as much as $10000 a month for high-end systems.

And SNA IBM 3745 Communications Controllers typically cost over $100K. TCP/IP was still seen as unfit for commercial applications e.g. in the finance industry until the late 1980s, but rapidly took over in the 1990s due to its peer-to-peer networking and packet communication technology it deployed.

- The design of SNA was in the era when the concept of layered communication was not fully adopted by the computer industry. Applications, databases and communication functions were mingled into the same protocol or product, to make it difficult to maintain or manage. That was very common for the products created in that time. Even after TCP/IP was fully developed, X window system was designed with the same model where communication protocols were embedded into graphic display application.

- SNA's connection based architecture invoked huge state machine logic to "keep track" of everything. APPN added a new dimension to state logic with its concept of differing node types. While it was solid when everything was running correctly, there was still a need for manual intervention. Simple things like watching the Control Point sessions had to be done manually. APPN wasn't without issues; in the early days many shops abandoned it due to issues found in APPN support. Over time, however, many of the issues were worked out but not before the advent of the Web Browser which was the beginning of the end for SNA.

Logical unit types

Network Addressable Units in an SNA network are any components that can be assigned an address and can send and receive information. They are distinguished further as follows:

- System Service Control Points, provide services to manage a network or subnetwork (typically in the mainframe),
- Physical Units, a physical device or communications link (relating to boxes),
- Logical Units, an access point to the network (relating to applications or subsystems such as CICS and TSO) or terminals.

SNA essentially offers transparent communication: equipment specifics don't impose any constraints onto LU-LU communication. But eventually it serves a purpose to make a distinction between LU types, as the application must take the functionality of the terminal equipment into account (e.g. screen sizes and layout).

SNA defines several kinds of devices, called Logical Unit types:

- LU0 provides for undefined devices, or build your own protocol.
- LU1 devices are printers.
- LU2 devices are dumb IBM 3270 display terminals.
- LU3 devices are printers using 3270 protocols.

- LU4 devices are batch terminals.
- LU5 has never been defined.
- LU6 provides for protocols between two applications.
- LU7 provides for sessions with IBM 5250 terminals.

The primary ones in use are LU1, LU2, and LU6.2 (an advanced protocol for application to application conversations).

Within SNA there are two types of data stream to connect local terminals and printers; there is the 3270 data stream mainly used by mainframes (zSeries family) and the 5250 data stream mainly used by minicomputers/servers such as the S/36, S/38, and AS/400 (now System i).

Starting from version 5.2 of OS/400, SNA for client-access is no longer supported.

The term 37xx refers to IBM's family of SNA communications controllers. The 3745 supports up to eight high-speed T1 circuits, the 3725 is a large-scale node and front-end processor for a host, and the 3720 is a remote node that functions as a concentrator and router.

Implementation and publication

SNA was made public as part of IBM's "Advanced Function for Communications" announcement in September, 1974, which included the implementation of the SNA/SDLC (Synchronous Data Link Control) protocols on new communications products:

- IBM 3767 communication terminal (printer)
- IBM 3770 data communication system

They were supported by IBM 3704/3705 communication controllers and their Network Control Program, and by System/360 and System/370 and their VTAM and other software such as CICS and IMS. This announcement was followed by another announcement in July, 1975, which introduced the IBM 3760 data entry station, the IBM 3790 communication system, and the new models of the IBM 3270 display system.

SNA was mainly designed by the IBM Systems Development Division laboratory in Research Triangle Park, North Carolina, USA, helped by other laboratories that implemented SNA/SDLC. The details were later made public by IBM's System Reference Library manuals and IBM Systems Journal.

Competitors

The proprietary networking architecture for Honeywell Bull mainframes is Distributed Systems Architecture (DSA). Communications package for DSA is VIP. Like SNA, DSA is also no longer supported for client access. Bull mainframes are fitted with Mainway for translating DSA to TCP/IP and VIP devices are replaced by TNVIP Terminal Emulations (GLink, Winsurf). GCOS 8 supports TNVIP SE over TCP/IP.

Chapter- 15

Next Generation Network and Open Access Network

Next generation network

A **Next generation network** (**NGN**) is a broad term to describe key architectural evolutions in telecommunication core and access networks. The general idea behind the NGN is that one network transports all information and services (voice, data, and all sorts of media such as video) by encapsulating these into packets, like it is on the Internet. NGNs are commonly built around the Internet Protocol, and therefore the term "all-IP" is also sometimes used to describe the transformation toward NGN.

Description

Challenges
for New Generation Network
Seoul, Korea
Hanyang University, December 9th, 2009
NiCT
Hiroaki Harai (harai@nict.go.jp)
NICT (National Institute of Information and Communications Technology)
http://{akari-project, nag}.nict.go.jp/

NGN Seminar in Fusion Technology Center by NICT(Japan) researcher

According to ITU-T, the definition is:

> A Next generation network (NGN) is a packet-based network which can provide services including Telecommunication Services and able to make use of multiple broadband, Quality of Service-enabled transport technologies and in which service-related functions are independent from underlying transport-related technologies. It offers unrestricted access by users to different service providers. It supports generalized mobility which will allow consistent and ubiquitous provision of services to users..

From a practical perspective, NGN involves three main architectural changes that need to be looked at separately:

- In the core network, NGN implies a consolidation of several (dedicated or overlay) transport networks each historically built for a different service into one core transport network (often based on IP and Ethernet). It implies amongst others the migration of voice from a circuit-switched architecture (PSTN) to VoIP, and also migration of legacy services such as X.25, Frame Relay (either commercial migration of the customer to a new service like IP VPN, or technical emigration by emulation of the "legacy service" on the NGN).
- In the wired access network, NGN implies the migration from the dual system of legacy voice next to xDSL setup in local exchanges to a converged setup in which the DSLAMs integrate voice ports or VoIP, making it possible to remove the voice switching infrastructure from the exchange.
- In the cable access network, NGN convergence implies migration of constant bit rate voice to CableLabs PacketCable standards that provide VoIP and SIP services. Both services ride over DOCSIS as the cable data layer standard.

In an NGN, there is a more defined separation between the transport (connectivity) portion of the network and the services that run on top of that transport. This means that whenever a provider wants to enable a new service, they can do so by defining it directly at the service layer without considering the transport layer - i.e. services are independent of transport details. Increasingly applications, including voice, tend to be independent of the access network (de-layering of network and applications) and will reside more on end-user devices (phone, PC, set-top box).

Underlying technology components

Next Generation Networks are based on Internet technologies including Internet Protocol (IP) and Multiprotocol Label Switching (MPLS). At the application level, Session Initiation Protocol (SIP) seems to be taking over from ITU-T H.323.

Initially H.323 was the most popular protocol, though its popularity decreased in the "local loop" due to its original poor traversal of Network address translation (NAT) and firewalls. For this reason as domestic VoIP services have been developed, SIP has been more widely adopted. However in voice networks where everything is under the control of the network operator or telco, many of the largest carriers use H.323 as the protocol of choice in their core backbones. So really SIP is a useful tool for the "local loop" and H.323 is like the "fiber backbone". With the most recent changes introduced for H.323, it is now possible for H.323 devices to easily and consistently traverse NAT and firewall devices, opening up the possibility that H.323 may again be looked upon more favorably in cases where such devices encumbered its use previously. Nonetheless, most of the telcos are extensively researching and supporting IP Multimedia Subsystem (IMS), which gives SIP a major chance of being the most widely adopted protocol.

For voice applications one of the most important devices in NGN is a Softswitch - a programmable device that controls Voice over IP (VoIP) calls. It enables correct integration of different protocols within NGN. The most important function of the Softswitch is creating the interface to the existing telephone network, PSTN, through Signalling Gateways and Media Gateways. However, the Softswitch as a term may be defined differently by the different equipment manufacturers and have somewhat different functions.

One may quite often find the term Gatekeeper in NGN literature. This was originally a VoIP device, which converted (using gateways) voice and data from their analog or digital switched-circuit form (PSTN, SS7) to the packet-based one (IP). It controlled one or more gateways. As soon as this kind of device started using the Media Gateway Control Protocol, the name was changed to Media Gateway Controller (MGC).

A Call Agent is a general name for devices/systems controlling calls.

The IP Multimedia Subsystem (IMS) is a standardised NGN architecture for an Internet media-services capability defined by the European Telecommunications Standards Institute (ETSI) and the 3rd Generation Partnership Project (3GPP).

Implementations

In the UK another popular acronym was introduced by BT (British Telecom) as 21CN (21st Century Networks, sometimes mistakenly quoted as C21N) — this is another loose term for NGN and denotes BT's initiative to deploy and operate NGN switches and networks in the period 2006-2008 (the aim being by 2008 BT to have only all-IP switches in their network)

The first company in the UK to roll out a NGN was THUS plc which started deployment back in 1999. THUS' NGN contains 10,600 km of fibre optic cable with more than 190 points of presence throughout the UK. The core optical network uses Dense Wave Division Multiplexing (DWDM) technology to provide scalability to many hundreds of gigabits per second of bandwidth, in line with growth demand. On top of this, the THUS backbone network uses MPLS technology to deliver the highest possible performance. IP/MPLS-based services carry voice, video and data traffic across a converged infrastructure, potentially allowing organisations to enjoy lower infrastructure costs, as well as added flexibility and functionality. Traffic can be prioritised with Classes of

Service, coupled with Service Level Agreements (SLAs) that underpin quality of service performance guarantees. The THUS NGN accommodates seven Classes of Service, four of which are currently offered on MPLS IP VPN.

In the Netherlands, KPN is developing a NGN network in a network transformation program called all-IP — this is another loose term for NGN that is increasingly used. Next Generation Networks also extends into the messaging domain and in Ireland, Openmind Networks has designed, built and deployed Traffic Control to handle the demands and requirements of all IP networks.

In Bulgaria, BTC (Bulgarian Telecommunications Company) has implemented the NGN as underlying network of its telco services on a large scale project in 2004. The inherent flexibility and scalability of the new core network approach resulted in an unprecedented rise of classical services deployment as POTS/ISDN, Centrex, ADSL, VPN, as well as implementation of higher bandwidths for the Metro and Long-distance Ethernet / VPN services, cross-national transits and WebTV/IPTV application.

In Israel, Bezeq announced in a June 2009 press release the move to NGN in selected areas. The service will allow enhanced services to phone subscribers as well as upgraded speed capabilities for ADSL users (up to 50Mbps DL, 1000Kbps UL).

In Canada, upstart Wind Mobile owned by Globalive is deploying an all-ip wireless backbone for its mobile phone service.

Open Access Network

In telecommunications, **Open Access Network** (**OAN**) refers to horizontally layered network architecture and business model that separates physical access to the network from service provisioning. The same OAN will be used by a number of different providers that share the investments and maintenance cost.

The OAN concept is appropriate for both fiber and WiFi Access Networks, especially where exclusivity can not be allowed. The shared maintenance costs make it very appropriate for distant rural areas, where traditional ISP are reluctant to offer their services. An open access network uses a different business model than traditional telecom networks. In the twentieth century, analog telephone and cable TV networks were designed around the limitations of the technology; copper-based twisted pair networks were not able to carry TV programming, and copper-based coaxial networks were not able to carry voice telephony. Near the end of the twentieth century, with the rise of packet-based switching (i.e. the Internet) and IP-based fiber and wireless technologies, it became possible to design, build, and operate a single high performance network capable of delivering dozens or even hundreds of services from multiple, competing providers.

Open access networks are also viewed as a feasible way of deploying next-generation broadband networks in low population density areas where service providers cannot obtain a sufficient return on investment to cover the high costs associated with trenching, right-of-way encroachment permits, and the requisite network infrastructure. In contrast to traditional municipal networks where the municipality owns the network and there is only one service provider, the open access model allows multiple service providers to compete over the same network at wholesale prices. This allows service providers to make money in the short-term and the municipality or cooperative to recoup its costs over the long-term. The build-out and infrastructure is typically financed through low-cost bonds. Open access networks have been very successful in the U.S., Europe, and Asia. One of the best known and most mature open access networks is Vasteras, city of about 40,000 homes in Sweden. The Vasteras open access network has dozens of providers and more than a hundred services available on the network.

In the United States, open access networks like The Wired Road have been able to attract both local and regional service providers quickly, and the cost of Internet access and telephone service for business users in The Wired Road service area have declined by 50% to 70% because of the increased competition between providers. The Wired Road is a municipal network owned by the counties of Carroll and Grayson and the City of Galax, and is operated as a regional authority. The Wired Road Broadband Authority sells no services, so it does not compete with private sector providers. The Wired Road provides open access transport to any service provider that meets minimum technical and financial qualifications, and incumbent providers are also able to use the system to deliver enhanced services to existing and new customers.

Chapter- 16

Optical Mesh Network

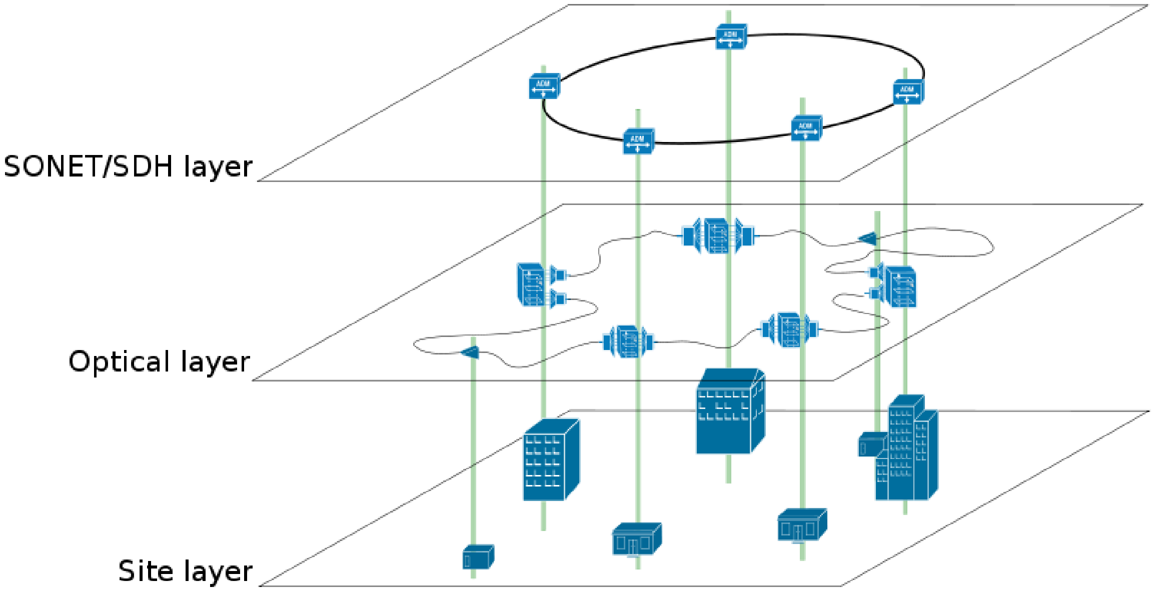

Transport network based on SONET/SDH ring architecture

Optical mesh networks are a type of telecommunications network.

Transport networks, the underlying optical fiber-based layer of telecommunications networks, have evolved from DCS (Digital Cross-connect Systems)-based mesh architectures in the 1980s, to SONET/SDH (Synchronous Optical Networking/Synchronous Digital Hierarchy) ring architectures in the 1990s. Technological advancements in optical transport equipment in the first decade of the 21st century, along with continuous deployment of DWDM systems, have led telecommunications service providers to replace their SONET ring architectures by mesh-based architectures. The new optical mesh networks support the same fast recovery previously available in ring networks while achieving better capacity efficiency and resulting in lower capital cost.

Optical mesh networks today not only provide trunking capacity to higher-layer networks, such as inter-router or inter-switch connectivity in an IP, MPLS, or Ethernet-centric infrastructure, but also support efficient routing and fast failure recovery of high-bandwidth services. This was made possible by the emergence of optical network elements that have the intelligence required to automatically control certain network functions, such as fault recovery.

Optical mesh networks enable a variety of dynamic services such as bandwidth-on-demand, Just-In-Time bandwidth, bandwidth scheduling, bandwidth brokering, and optical virtual private networks that open up new opportunities for service providers and their customers alike.

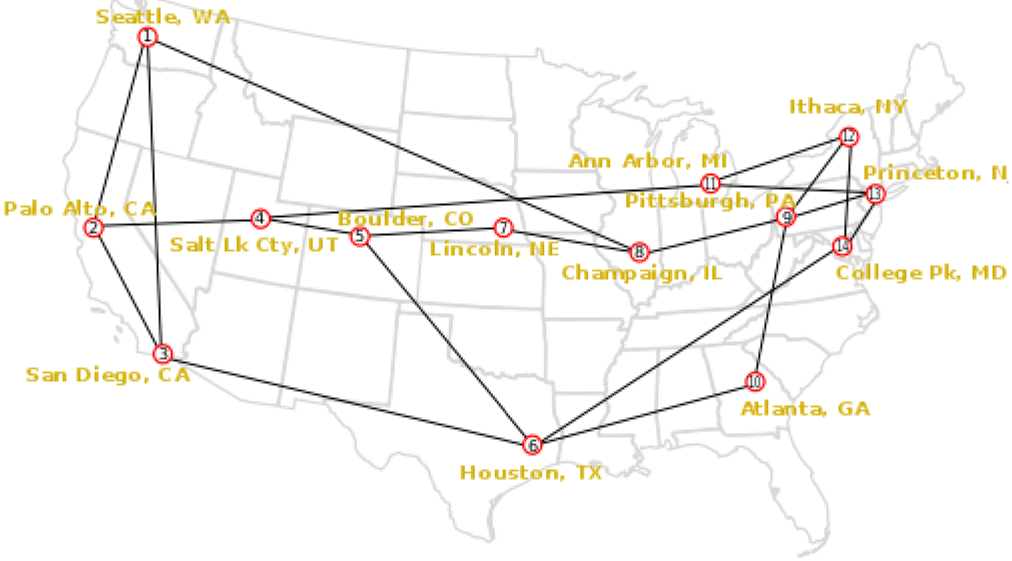

Example of mesh network: NSFNET 14nodes

History of transport networks

Transport networks, the underlying optical fiber-based layer of telecommunications networks, have evolved from Digital cross connect system (DCS)-based mesh architectures in the 1980s, to SONET/SDH (Synchronous Optical Networking/Synchronous Digital Hierarchy) ring architectures in the 1990s. In DCS-based mesh architectures, telecommunications carriers deployed restoration systems for DS3 circuits such as at&t FASTAR (FAST Automatic Restoration) and MCI Real Time Restoration (RTR), restoring circuits in minutes after a network failure. In SONET/SDH rings, carriers implemented ring protection such as SONET Universal Path Switched Ring (UPSR) (also called Sub-Network Connection Protection (SCNP) in SDH networks) or SONET Bidirectional Line Switched Ring (BLSR) (also called Multiplex Section - Shared Protection Ring (MS-SPRing) in SDH networks), protecting against and recovering from a network failure in 50 msecs or less, a significant improvement over the recovery time supported in DCS-based mesh restoration, and a key driver for the deployment of SONET/SDH ring-based protection.

There have been attempts at improving and/or evolving traditional ring architectures to overcome some of its limitations, with trans-oceanic ring architecture (ITU-T Rec.

G.841), “P-cycles” protection, next-generation SONET/SDH equipment that can handle multiple rings, or have the ability to not close the working or protection ring side, or to share protection capacity among rings (e.g., with Virtual Line Switched Ring (VLSR).

Technological advancements in optical transport switches in the first decade of the 21st century, along with continuous deployment of dense wavelength-division multiplexing (DWDM) systems, have led telecommunications service providers to replace their SONET ring architectures by mesh-based architectures for new traffic. The new optical mesh networks support the same fast recovery previously available in ring networks while achieving better capacity efficiency and resulting in lower capital cost. Such fast recovery (in the 10's to 100's of msecs) in case of failures (e.g., network link or node failure) is achieved through the intelligence embedded in these new optical transport equipment, which allows recovery to be automatic and handled within the network itself as part of the network control plane, without relying on an external network management system.

Optical mesh networks

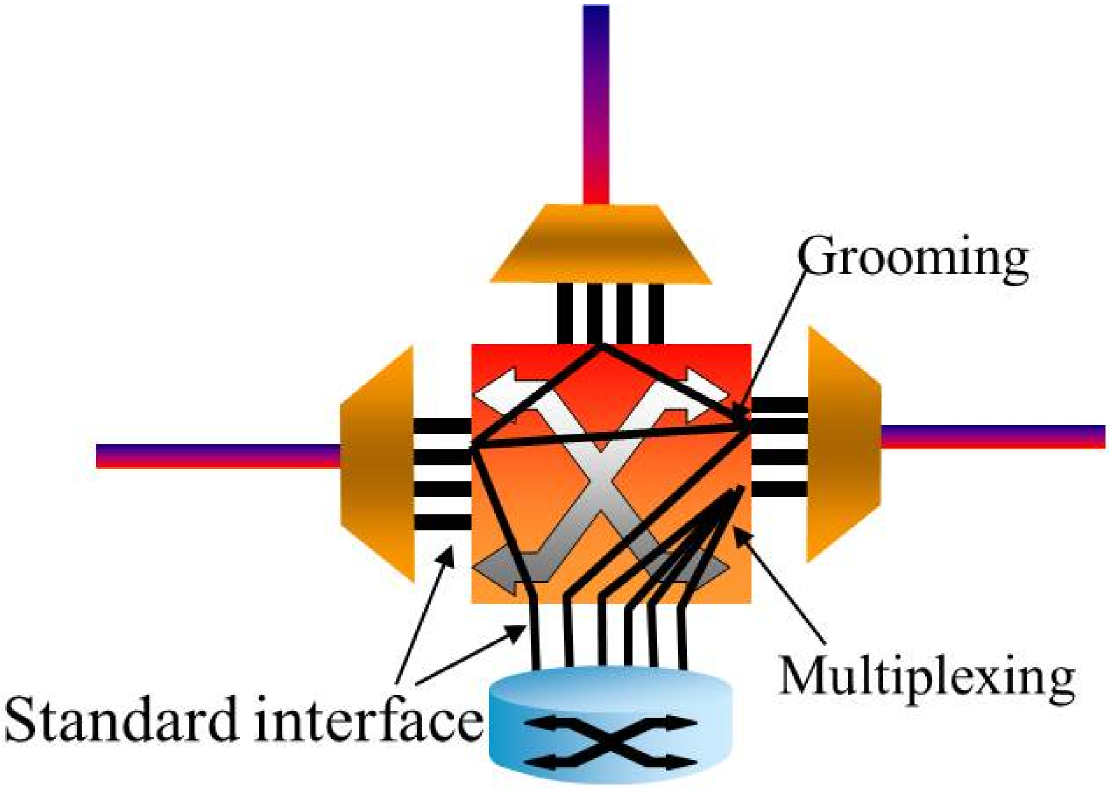

Switching, multiplexing, and grooming of traffic in an OEO device

Optical mesh networks refer to transport networks that are built directly off the mesh-like fiber infrastructure deployed in metropolitan, regional, national, or international (e.g., trans-oceanic) areas by deploying optical transport equipment that are capable of switching traffic (at the wavelength or sub-wavelength level) from an incoming fiber to an outgoing fiber. In addition to switching wavelengths, the equipment is typically also able to multiplex lower speed traffic into wavelengths for transport, and to groom traffic. Finally, these equipment also provide for the recovery of traffic in case of a network failure. As most of the transport networks evolve toward mesh topologies utilizing intelligent network elements (optical cross-connects or optical switches) for provisioning and recovery of services, new approaches have been developed for the design, deployment, operations and management of mesh optical networks.

Optical mesh networks today not only provide trunking capacity to higher-layer networks, such as inter-router or inter-switch connectivity in an IP, MPLS, or Ethernet-centric packet infrastructure, but also support efficient routing and fast failure recovery of high-bandwidth point-to-point Ethernet and SONET/SDH services.

Recovery in optical mesh networks

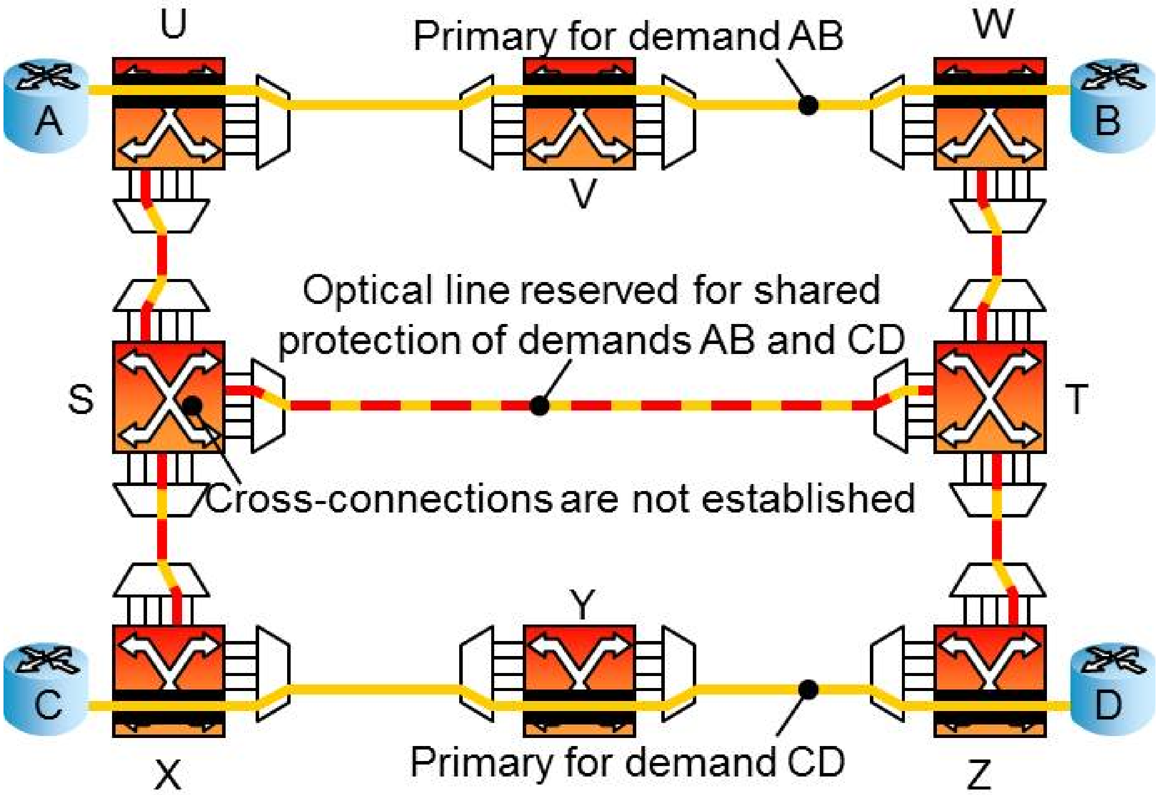
U
Primary for demand AB
W
A
B
V
Optical line reserved for shared protection of demands AB and CD
S
T
Cross-connections are not established
Y
C
D
X
Primary for demand CD
Z

Shared backup path protection - before failure

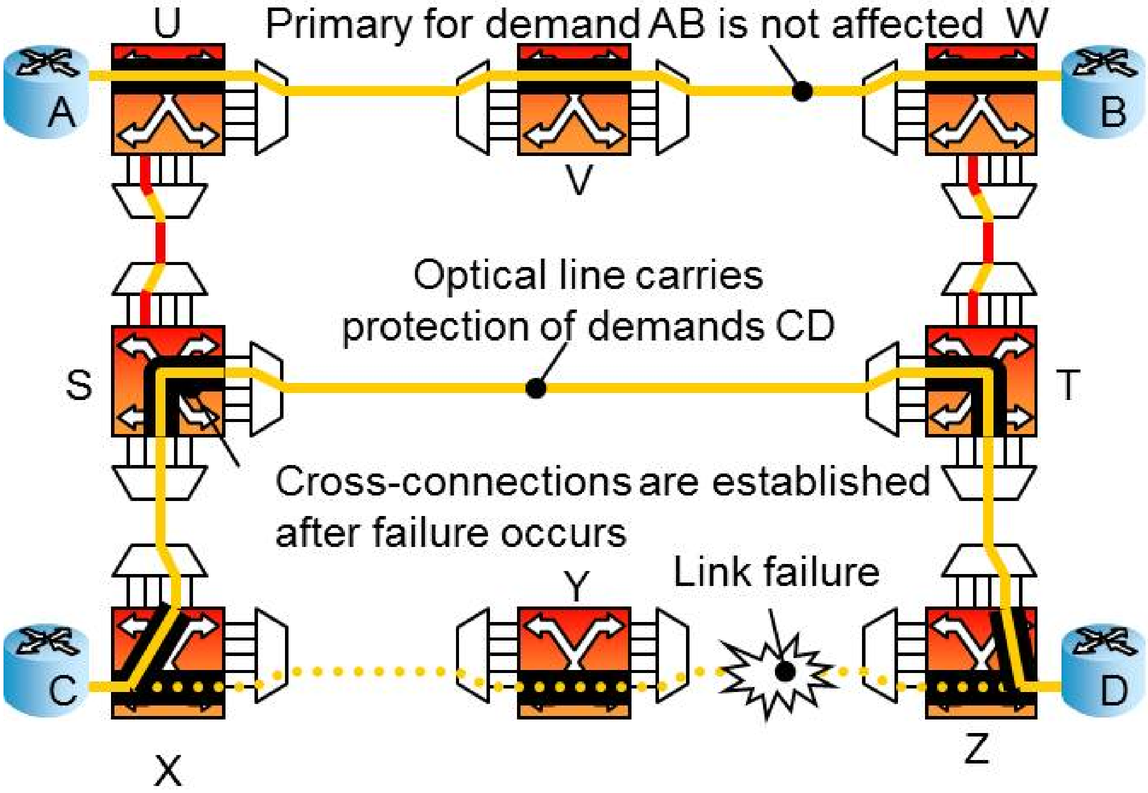

Shared backup path protection - after failure and recovery

Optical mesh networks support the establishment of circuit-mode connection-oriented services. Multiple recovery mechanisms that provide different levels of protection or restoration against different failure modes are available in mesh networks. Channel, link, segment and path protection are the most common protection schemes. P-cycles is another type of protection that leverages and extends ring-based protection. Restoration is another recovery method that can work on its own or complement faster protection schemes in case of multiple failures.

In path-protected mesh networks, some connections can be unprotected; others can be protected against single or multiple failures in various ways. A connection can be protected against a single failure by defining a backup path, diverse from the primary path taken by the connection over the mesh network. The backup path and associated resources can be dedicated to the connection (aka Dedicated Backup Path Protection), or shared among multiple connections (aka Shared Backup Path Protection), typically ones whose primary paths are not likely to fail at the same time, thereby avoiding contention for the shared resources in case of a single link or node failure. A number of other protection schemes such as the use of pre-emptible paths, or only partially diverse backup paths, can be implemented. Finally, multiple diverse routes can be designed so that a connection has multiple recovery routes and can recover even after multiple failures (examples of mesh networks across the Atlantic and Pacific oceans).

Transparency

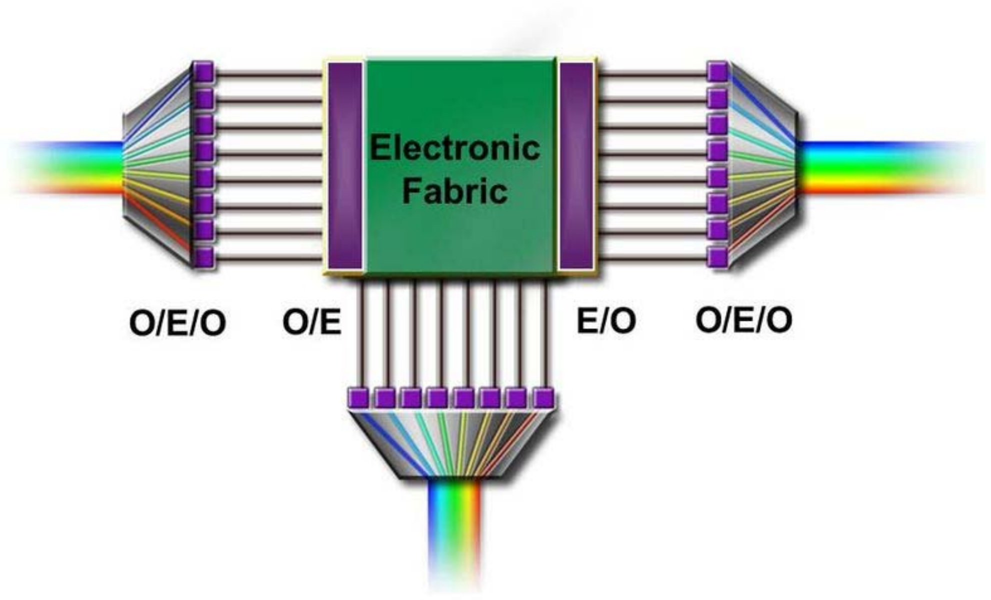

Opaque switching of traffic between fiber links

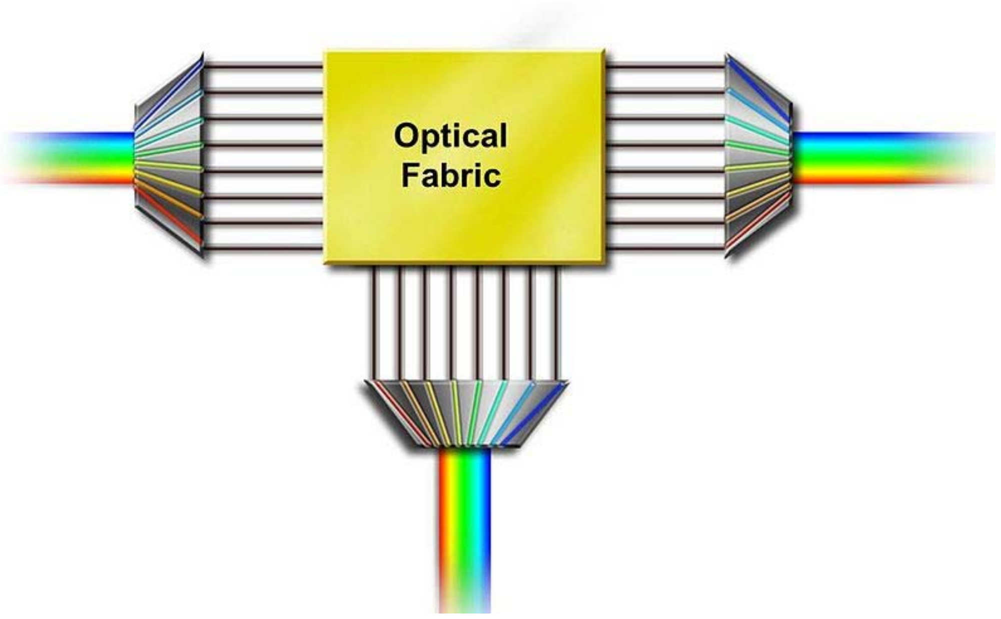

Transparent switching of traffic between fiber links

Traditional transport networks are made of optical fiber-based links between telecommunications offices, where multiple wavelengths are multiplexed to increase the capacity of the fiber. The wavelengths are terminated on electronic devices called transponders, undergoing an optical-to-electrical conversion for signal Reamplification, Reshaping, and Retiming (3R). Inside a telecommunications office, the signals are then handled to and switched by a transport switch (aka optical cross-connect or optical switch) and either are dropped at that office, or directed to an outgoing fiber link where they are again carried as wavelengths multiplexed into that fiber link towards the next telecommunications office. The act of going through Optical-Electrical-Optical (O-E-O) conversion through a telecommunications office causes the network to be considered opaque. When the incoming wavelengths do not undergo an optical-to-electrical conversion and are switched through a telecommunications office in the optical domain using all-optical switches (also called photonic cross-connect, optical add-drop multiplexer, or Reconfigurable Optical Add-Drop Multiplexer (ROADM) systems), the network is considered to be transparent. Hybrid schemes can provide limited O-E-O conversions at key locations across the network.

Transparent optical mesh networks have been deployed in metropolitan and regional networks. In 2010, operational long distance networks still tend to remain opaque.

Routing in optical mesh networks

Routing is a key control and operational aspect of optical mesh networks. In transparent or all-optical networks, routing of connections is tightly linked to the wavelength selection and assignment process (so-called routing and wavelength assignment, or "RWA"). This is due to the fact that the connection remains on the same wavelength from end-to-end throughout the network (sometimes referred to as wavelength continuity constraint, in the absence of devices that can translate between wavelengths in the optical domain). In an opaque network, the routing problem is one of finding a primary path for a connection and if protection is needed, a backup path diverse from the primary path. Wavelengths are used on each link independently of each other's. Several algorithms can be used to determine a primary path and a diverse backup path (with or without sharing of resource along the backup path) for a connection or service, such as shortest path, including Dijkstra's algorithm, k-shortest path, edge and node-diverse or disjoint routing, including Suurballe's algorithm, and numerous heuristics.

Applications

The deployment of optical mesh networks is enabling new services and applications for service providers to offer their customers, such as

- Dynamic services such as Bandwidth-on-Demand (BoD), Just-In-Time (JIT) bandwidth, bandwidth scheduling, and bandwidth brokering
- Optical virtual private networks

It also supports new network paradigms such as

- IP-over-optical network architectures

Chapter- 17

Internet Protocol Suite

The **Internet Protocol Suite** is the set of communications protocols used for the Internet and other similar networks. It is commonly also known as **TCP/IP**, named from two of the most important protocols in it: the Transmission Control Protocol (TCP) and the Internet Protocol (IP), which were the first two networking protocols defined in this standard. Modern IP networking represents a synthesis of several developments that began to evolve in the 1960s and 1970s, namely the Internet and local area networks, which emerged during the 1980s, together with the advent of the World Wide Web in the early 1990s.

The Internet Protocol Suite, like many protocol suites, is constructed as a set of layers. Each layer solves a set of problems involving the transmission of data.

Every layer provides a well-defined service to the upper layer protocols and relies on using services from the lower layers. The upper layers provide a higher level of abstraction, being closer to the application and the end user. The lower layer protocols are more concerned providing specific solutions to deal with the actual physical transmission of the data.

The TCP/IP model consists of 4 layers (RFC 1122). From lowest to highest, these are the Link Layer, the Internet Layer, the Transport Layer, and the Application Layer.

History

The Internet Protocol Suite resulted from research and development conducted by the Defense Advanced Research Projects Agency (DARPA) in the early 1970s. After initiating the pioneering ARPANET in 1969, DARPA started work on a number of other data transmission technologies. In 1972, Robert E. Kahn joined the DARPA Information Processing Technology Office, where he worked on both satellite packet networks and ground-based radio packet networks, and recognized the value of being able to communicate across both. In the spring of 1973, Vinton Cerf, the developer of the existing ARPANET Network Control Program (NCP) protocol, joined Kahn to work on open-architecture interconnection models with the goal of designing the next protocol generation for the ARPANET.

By the summer of 1973, Kahn and Cerf had worked out a fundamental reformulation, where the differences between network protocols were hidden by using a common internetwork protocol, and, instead of the network being responsible for reliability, as in the ARPANET, the hosts became responsible. Cerf credits Hubert Zimmerman and Louis Pouzin, designer of the CYCLADES network, with important influences on this design.

The design of the network included the recognition that it should provide only the functions of efficiently transmitting and routing traffic between end nodes and that all other intelligence should be located at the edge of the network, in the end nodes. Using a simple design, it became possible to connect almost any network to the ARPANET, irrespective of their local characteristics, thereby solving Kahn's initial problem. One popular expression is that TCP/IP, the eventual product of Cerf and Kahn's work, will run over "two tin cans and a string."

A computer called a router (a name changed from gateway to avoid confusion with other types of gateways) is provided with an interface to each network, and forwards packets back and forth between them. Requirements for routers are defined in (Request for Comments 1812).

The idea was worked out in more detailed form by Cerf's networking research group at Stanford in the 1973–74 period, resulting in the first TCP specification.(Request for Comments 675) (The early networking work at Xerox PARC, which produced the PARC Universal Packet protocol suite, much of which existed around the same period of time, was also a significant technical influence; people moved between the two.)

DARPA then contracted with BBN Technologies, Stanford University, and the University College London to develop operational versions of the protocol on different hardware platforms. Four versions were developed: TCP v1, TCP v2, a split into TCP v3 and IP v3 in the spring of 1978, and then stability with TCP/IP v4 — the standard protocol still in use on the Internet today.

In 1975, a two-network TCP/IP communications test was performed between Stanford and University College London (UCL). In November, 1977, a three-network TCP/IP test was conducted between sites in the US, UK, and Norway. Several other TCP/IP prototypes were developed at multiple research centres between 1978 and 1983. The migration of the ARPANET to TCP/IP was officially completed on flag day January 1, 1983, when the new protocols were permanently activated.

In March 1982, the US Department of Defense declared TCP/IP as the standard for all military computer networking. In 1985, the Internet Architecture Board held a three day workshop on TCP/IP for the computer industry, attended by 250 vendor representatives, promoting the protocol and leading to its increasing commercial use.

Layers in the Internet Protocol Suite

The concept of layers

The TCP/IP suite uses encapsulation to provide abstraction of protocols and services. Such encapsulation usually is aligned with the division of the protocol suite into layers of general functionality. In general, an application (the highest level of the model) uses a set of protocols to send its data down the layers, being further encapsulated at each level.

This may be illustrated by an example network scenario, in which two Internet host computers communicate across local network boundaries constituted by their internetworking gateways (routers).

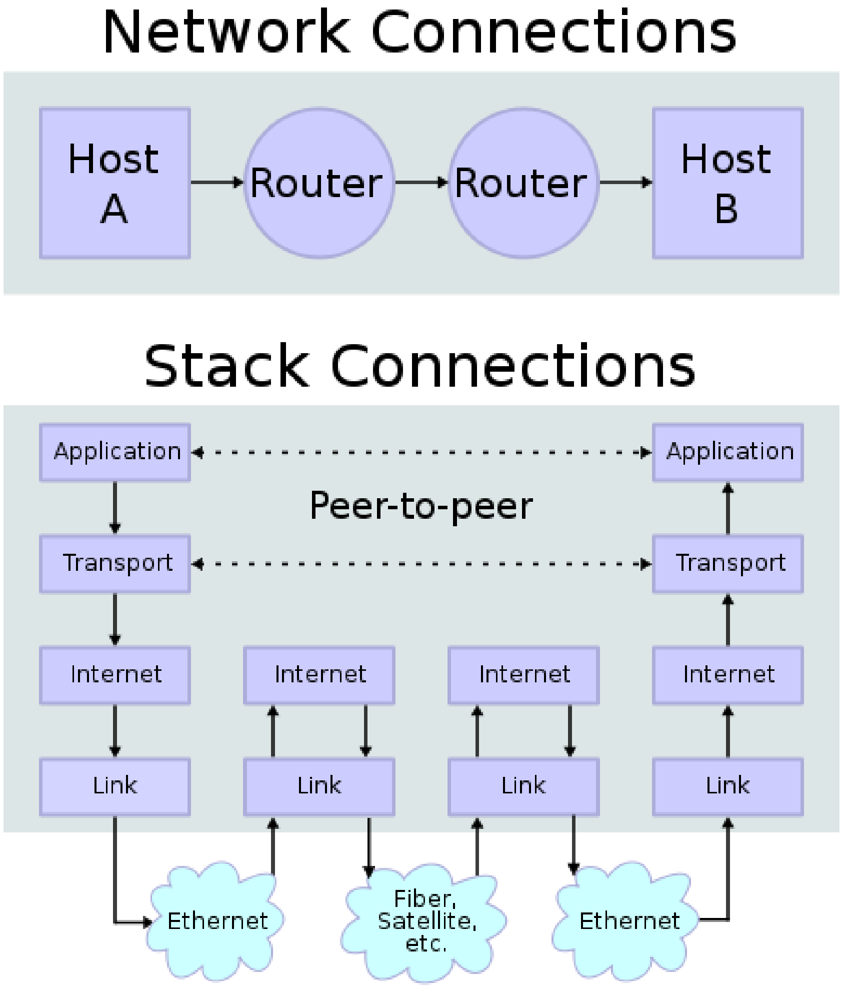

TCP/IP stack operating on two hosts connected via two routers and the corresponding layers used at each hop

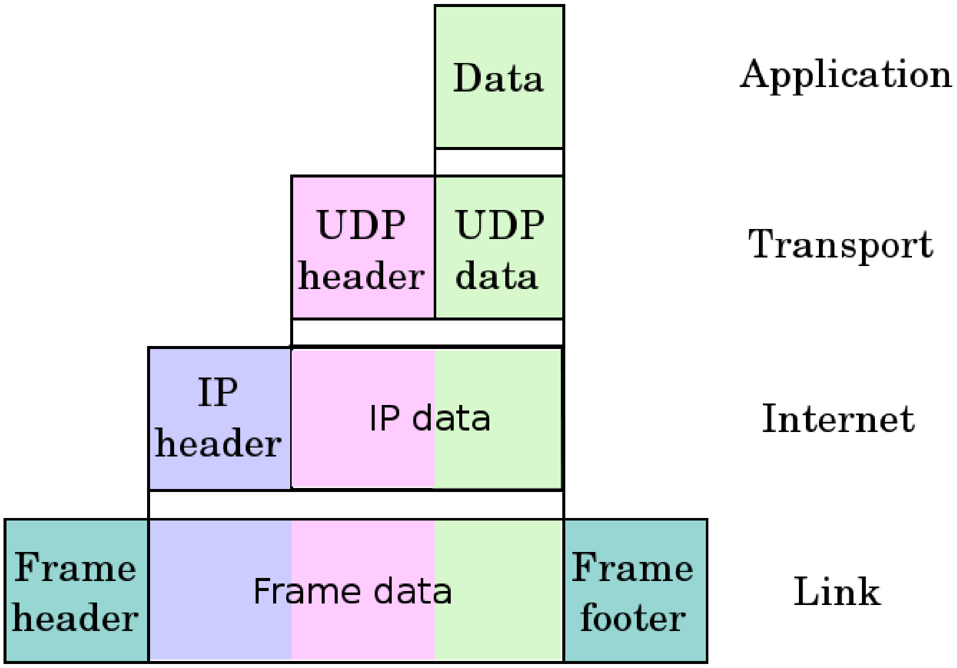

Encapsulation of application data descending through the protocol stack

The functional groups of protocols and methods are the Application Layer, the Transport Layer, the Internet Layer, and the Link Layer (RFC 1122). This model was not intended to be a rigid reference model into which new protocols have to fit in order to be accepted as a standard.

The following table provides some examples of the protocols grouped in their respective layers.

Application	DNS, TFTP, TLS/SSL, FTP, Gopher, HTTP, IMAP, IRC, NNTP, POP3, SIP, SMTP, SMPP, SNMP, SSH, Telnet, Echo, RTP, PNRP, rlogin, ENRP Routing protocols like BGP and RIP which run over TCP/UDP, may also be considered part of the Internet Layer.
Transport	TCP, UDP, DCCP, SCTP, IL, RUDP, RSVP
Internet	IP (IPv4, IPv6), ICMP, IGMP, and ICMPv6 OSPF for IPv4 was initially considered IP layer protocol since it runs per IP-subnet, but has been placed on the Link since RFC 2740.
Link	ARP, RARP, OSPF (IPv4/IPv6), IS-IS, NDP

Layer names and number of layers in the literature

The following table shows the layer names and the number of layers of networking models presented in RFCs and textbooks in widespread use in today's university computer networking courses.

RFC 1122	Tanenbaum	Cisco Academy	Kurose Forouzan	Comer Kozierok	Stallings	Arpanet Reference Model 1982 (RFC 871)
Four layers	Four layers	Four layers	Five layers	Four+one layers	Five layers	Three layers
"Internet model"	"TCP/IP reference model"	"Internet model"	"Five-layer Internet model" or "TCP/IP protocol suite"	"TCP/IP 5-layer reference model"	"TCP/IP model"	"Arpanet reference model"
Application	Application	Application	Application	Application	Application	Application/Process
Transport	Transport	Transport	Transport	Transport	Host-to-host or transport	Host-to-host
Internet	Internet	Internetwork	Network	Internet	Internet	
Link	Host-to-network	Network interface	Data link	Data link (Network interface)	Network access	Network interface
			Physical	(Hardware)	Physical	

These textbooks are secondary sources that may contravene the intent of RFC 1122 and other IETF primary sources.

Different authors have interpreted the RFCs differently regarding the question whether the Link Layer (and the TCP/IP model) covers Physical Layer issues, or if a hardware layer is assumed below the Link Layer. Some authors have tried to use other names for the Link Layer, such as network interface layer, in view to avoid confusion with the Data Link Layer of the seven layer OSI model. Others have attempted to map the Internet Protocol model onto the OSI Model. The mapping often results in a model with five layers where the Link Layer is split into a Data Link Layer on top of a Physical Layer. In literature with a bottom-up approach to Internet communication, in which hardware issues are emphasized, those are often discussed in terms of Physical Layer and Data Link Layer.

The Internet Layer is usually directly mapped into the OSI Model's Network Layer, a more general concept of network functionality. The Transport Layer of the TCP/IP model, sometimes also described as the host-to-host layer, is mapped to OSI Layer 4 (Transport Layer), sometimes also including aspects of OSI Layer 5 (Session Layer) functionality. OSI's Application Layer, Presentation Layer, and the remaining functionality of the Session Layer are collapsed into TCP/IP's Application Layer. The argument is that these OSI layers do usually not exist as separate processes and protocols in Internet applications.

However, the Internet protocol stack has never been altered by the Internet Engineering Task Force from the four layers defined in RFC 1122. The IETF makes no effort to follow the OSI model although RFCs sometimes refer to it. The IETF has repeatedly stated that Internet protocol and architecture development is not intended to be OSI-compliant.

RFC 3439, addressing Internet architecture, contains a section entitled: "Layering Considered Harmful".

Implementations

Most computer operating systems in use today, including all consumer-targeted systems, include a TCP/IP implementation.

Minimally acceptable implementation includes implementation for (from most essential to the less essential) IP, ARP, ICMP, UDP, TCP and sometime IGMP. It is in principle possible to support only one of transport protocols (i.e. simple UDP), but it is rarely done, as it limits usage of the whole implementation. IPv6, beyond own version of ARP (NBP), and ICMP (ICMPv6), and IGMP (IGMPv6) have some additional required functionalities, and often is accompanied with integrated IPSec security layer. Other protocols could be easily added later (often they can be implemented entirely in the userspace), for example DNS for resolving domain names to IP addresses or DHCP client for automatic configuration of network interfaces.

Most of the IP implementations are accessible to the programmers using socket abstraction (usable also with other protocols) and proper API for most of the operations. This interface is known as BSD sockets and was used initially in C.

Unique implementations include Lightweight TCP/IP, an open source stack designed for embedded systems and KA9Q NOS, a stack and associated protocols for amateur packet radio systems and personal computers connected via serial lines.

Chapter- 18

Network Planning & Design and TCP/IP Model

Network planning and design

Network planning and design is an iterative process, encompassing topological design, network-synthesis, and network-realization, and is aimed at ensuring that a new network or service meets the needs of the subscriber and operator. The process can be tailored according to each new network or service.

This is an extremely important process which must be performed before the establishment of a new telecommunications network or service.

A network planning methodology

A traditional network planning methodology involves five layers of planning, namely:

- business planning
- long-term and medium-term network planning
- short-term network planning
- IT asset sourcing
- operations and maintenance.

Each of these layers incorporates plans for different time horizons, i.e. the business planning layer determines the planning that the operator must perform to ensure that the network will perform as required for its intended life-span. The Operations and Maintenance layer, however, examines how the network will run on a day-to-day basis.

The network planning process begins with the acquisition of external information. This includes:

- forecasts of how the new network/service will operate;
- the economic information concerning costs; and
- the technical details of the network's capabilities.

It should be borne in mind that planning a new network/service involves implementing the new system across the first four layers of the OSI Reference Model. This means that even before the network planning process begins, choices must be made, involving protocols and transmission technologies.

Once the initial decisions have been made, the network planning process involves three main steps:

- **Topological design**: This stage involves determining where to place the components and how to connect them. The (topological) optimisation methods that can be used in this stage come from an area of mathematics called Graph Theory. These methods involve determining the costs of transmission and the cost of switching, and thereby determining the optimum connection matrix and location of switches and concentrators.

- **Network-synthesis**: This stage involves determining the size of the components used, subject to performance criteria such as the Grade of Service (GoS). The method used is known as "Nonlinear Optimisation", and involves determining the topology, required GoS, cost of transmission, etc., and using this information to calculate a routing plan, and the size of the components.

- **Network realization**: This stage involves determining how to meet capacity requirements, and ensure reliability within the network. The method used is known as "Multicommodity Flow Optimisation", and involves determining all information relating to demand, costs and reliability, and then using this information to calculate an actual physical circuit plan.

These steps are interrelated and are therefore performed iteratively, and in parallel with one another. The planning process is highly complex, meaning that at each iteration, an analyst must increase his planning horizons, and in so doing, he must generate plans for the various layers outlined above.

The role of forecasting

During the process of Network Planning and Design, it is necessary to estimate the expected traffic intensity and thus the traffic load that the network must support. If a network of a similar nature already exists, then it may be possible to take traffic measurements of such a network and use that data to calculate the exact traffic load. However, as is more likely in most instances, if there are no similar networks to be found, then the network planner must use telecommunications forecasting methods to estimate the expected traffic intensity.

The forecasting process involves several steps as follows:

- Definition of problem;
- Data acquisition;

- Choice of forecasting method;
- Analysis/Forecasting;
- Documentation and analysis of results.

Dimensioning

The purpose of dimensioning a new network/service is to determine the minimum capacity requirements that will still allow the Teletraffic Grade of Service (GoS) requirements to be met. To do this, dimensioning involves planning for peak-hour traffic, i.e. that hour during the day during which traffic intensity is at its peak.

The dimensioning process involves determining the network's topology, routing plan, traffic matrix, and GoS requirements, and using this information to determine the maximum call handling capacity of the switches, and the maximum number of channels required between the switches.. This process requires a complex model that simulates the behavior of the network equipment and routing protocols.

A dimensioning rule is that the planner must ensure that the traffic load should never approach a load of 100 percent. To calculate the correct dimensioning to comply with the above rule, the planner must take on-going measurements of the network's traffic, and continuously maintain and upgrade resources to meet the changing requirements.. Another reason for "overprovisioning" is to make sure that traffic can be rerouted in case a failure occurs in the network.

Because of the complexity of network dimensioning, this is typically done using specialized software tools. Whereas researchers typically develop custom software to study a particular problem, network operators typically make use of commercial network planning software (e.g. OPNET Technologies, SevOne, WANDL, VPISystems, Cariden, Aria Networks). However, there is one notable open source network planning software available by the name of TOTEM named after TOolbox for Traffic Engineering Methods.

Traffic engineering

Comparing to network engineering, which adds resources such as links, routers and switches into the network, traffic engineering targets to change traffic paths on the existing network to alleviate traffic congestion or accommodate more traffic demand.

This technology is critical when the cost of network expansion is prohibitively high and network load is not optimally balanced. The first part provides financial motivation for traffic engineering while the second part grants the possibility of deploying this technology.

The available technologies for traffic engineering include MPLS and ATM for current Internet backbone. For example, MPLS allows carriers to provision LSPs with dynamic or explicit routes. The dynamic routes is controlled by CSPF while the explicit routes are optimized in an offline tool or through a path computation element which is under study by IETF. Fast reroute has been implemented by major vendors, such as Cisco and Juniper Networks, to provide localized resilient capability for MPLS networks. End-to-end protection is an alternative resilient approach. It provisions a backup route for each primary route. Pre-planning enough bandwidth for these backup routes is one of the active topic for survivable network design.

Provisioning a large number of LSPs also brought up a scalability problem. Various solutions have been proposed and it is still an active topic under study.

Survivability

Network survivability enables the network to maintain maximum network connectivity and quality of service under failure conditions. It has been one of the critical requirements in network planning and design. It involves design requirements on topology, protocol, bandwidth allocation, etc.. Topology requirement can be maintaining a minimum two-connected network against any failure of a single link or node. Protocol requirements include using dynamic routing protocol to reroute traffic against network dynamics during the transition of network dimensioning or equipment failures. Bandwidth allocation requirements pro-actively allocate extra bandwidth to avoid traffic loss under failure conditions. This topic has been actively studied in conferences, such as the International Workshop on Design of Reliable Communication Networks DRCN.

TCP/IP model

The **TCP/IP model** is a description framework for computer network protocols created in the 1970s by DARPA, an agency of the United States Department of Defense. It evolved from ARPANET, which was the world's first wide area network and a predecessor of the Internet. The TCP/IP Model is sometimes called the Internet Model or the DoD Model.

The TCP/IP model, or Internet Protocol Suite, describes a set of general design guidelines and implementations of specific networking protocols to enable computers to communicate over a network. TCP/IP provides end-to-end connectivity specifying how data should be formatted, addressed, transmitted, routed and received at the destination. Protocols exist for a variety of different types of communication services between computers.

TCP/IP, sometimes referred to as the Internet model, has four abstraction layers as defined in RFC 1122. This layer architecture is often compared with the seven-layer OSI Reference Model; using terms such as Internet reference model, incorrectly, however, because it is descriptive while the OSI Reference Model was intended to be prescriptive, hence being a reference model.

The TCP/IP model and related protocols are maintained by the Internet Engineering Task Force (IETF).

Key architectural principles

An early architectural document, RFC 1122, emphasizes architectural principles over layering.

- End-to-End Principle: This principle has evolved over time. Its original expression put the maintenance of state and overall intelligence at the edges, and assumed the Internet that connected the edges retained no state and concentrated on speed and simplicity. Real-world needs for firewalls, network address translators, web content caches and the like have forced changes in this principle.
- Robustness Principle: "In general, an implementation must be conservative in its sending behavior, and liberal in its receiving behavior. That is, it must be careful to send well-formed datagrams, but must accept any datagram that it can interpret (e.g., not object to technical errors where the meaning is still clear)." "The second part of the principle is almost as important: software on other hosts may contain deficiencies that make it unwise to exploit legal but obscure protocol features."

Even when the layers are examined, the assorted architectural documents—there is no single architectural model such as ISO 7498, the OSI reference model—have fewer and less rigidly-defined layers than the OSI model, and thus provide an easier fit for real-world protocols. In point of fact, one frequently referenced document, RFC 1958, does not contain a stack of layers. The lack of emphasis on layering is a strong difference between the IETF and OSI approaches. It only refers to the existence of the "internetworking layer" and generally to "upper layers"; this document was intended as a 1996 "snapshot" of the architecture: "The Internet and its architecture have grown in evolutionary fashion from modest beginnings, rather than from a Grand Plan. While this process of evolution is one of the main reasons for the technology's success, it nevertheless seems useful to record a snapshot of the current principles of the Internet architecture."

RFC 1122, entitled Host Requirements, is structured in paragraphs referring to layers, but the document refers to many other architectural principles not emphasizing layering. It loosely defines a four-layer model, with the layers having names, not numbers, as follows:

- Application Layer (process-to-process): This is the scope within which applications create user data and communicate this data to other processes or applications on another or the same host. The communications partners are often called peers. This is where the "higher level" protocols such as SMTP, FTP, SSH, HTTP, etc. operate.
- Transport Layer (host-to-host): The Transport Layer constitutes the networking regime between two network hosts, either on the local network or on remote networks separated by routers. The Transport Layer provides a uniform

networking interface that hides the actual topology (layout) of the underlying network connections. This is where flow-control, error-correction, and connection protocols exist, such as TCP. This layer deals with opening and maintaining connections between Internet hosts.

- Internet Layer (internetworking): The Internet Layer has the task of exchanging datagrams across network boundaries. It is therefore also referred to as the layer that establishes internetworking, indeed, it defines and establishes the Internet. This layer defines the addressing and routing structures used for the TCP/IP protocol suite. The primary protocol in this scope is the Internet Protocol, which defines IP addresses. Its function in routing is to transport datagrams to the next IP router that has the connectivity to a network closer to the final data destination.
- Link Layer: This layer defines the networking methods with the scope of the local network link on which hosts communicate without intervening routers. This layer describes the protocols used to describe the local network topology and the interfaces needed to affect transmission of Internet Layer datagrams to next-neighbor hosts. (cf. the OSI Data Link Layer).

The Internet Protocol Suite and the layered protocol stack design were in use before the OSI model was established. Since then, the TCP/IP model has been compared with the OSI model in books and classrooms, which often results in confusion because the two models use different assumptions, including about the relative importance of strict layering.

Layers in the TCP/IP model

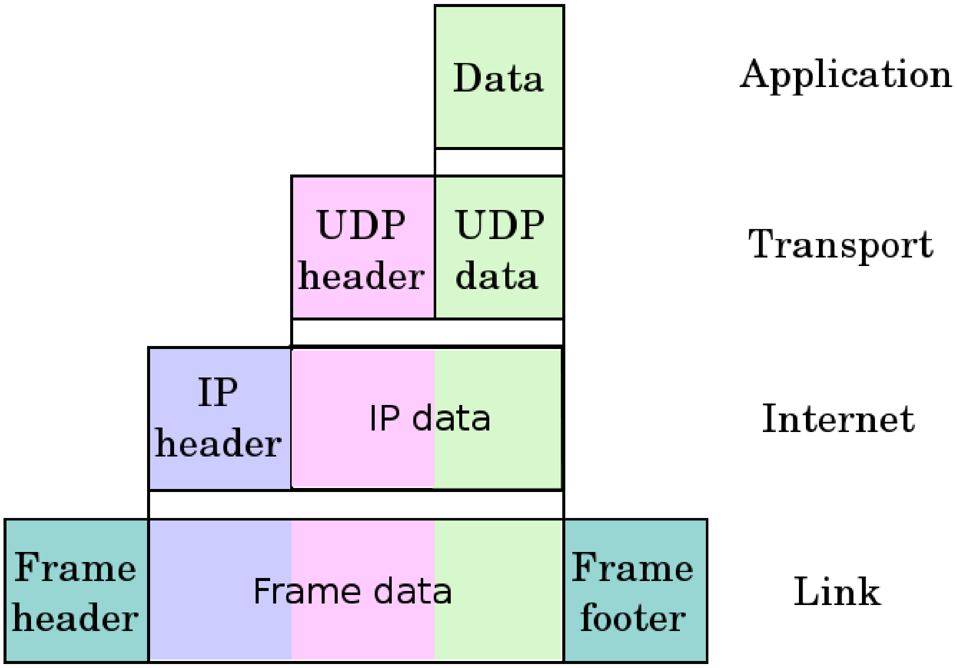

Encapsulation of application data descending through the TCP/IP layers

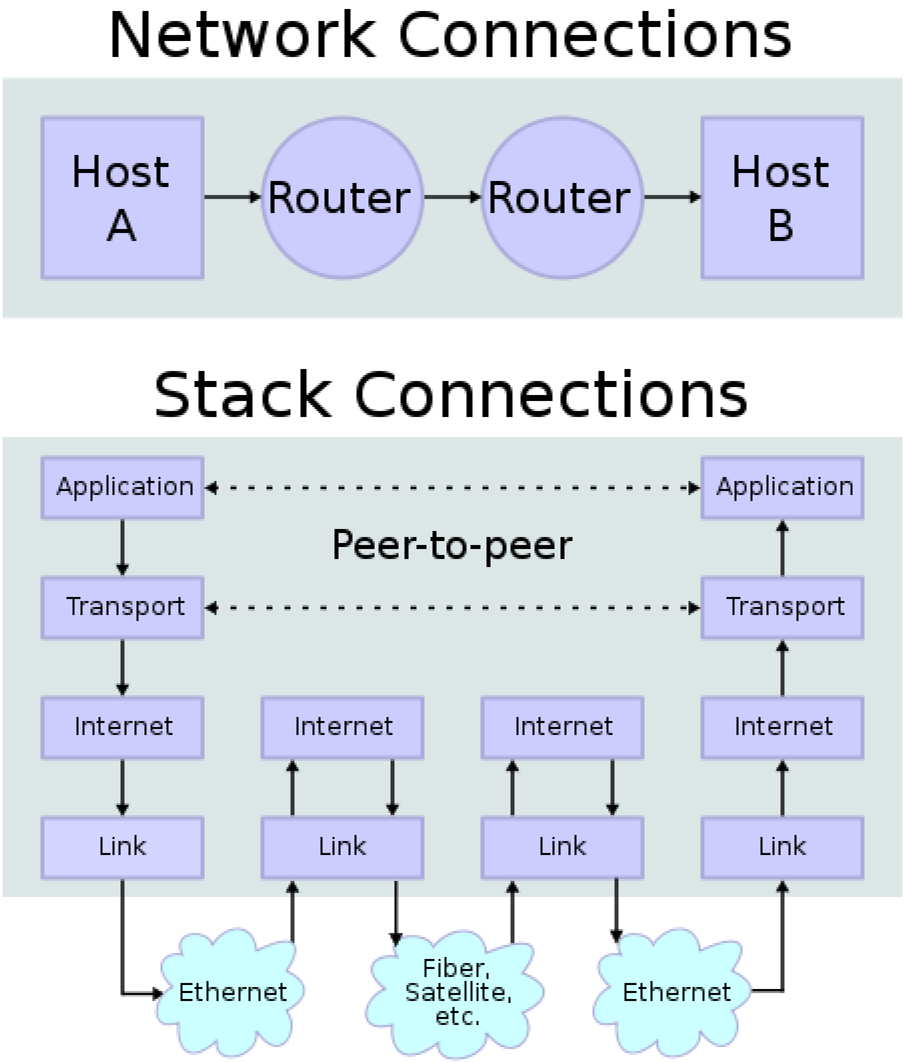

Two Internet hosts connected via two routers and the corresponding layers used at each hop

The layers near the top are logically closer to the user application, while those near the bottom are logically closer to the physical transmission of the data. Viewing layers as providing or consuming a service is a method of abstraction to isolate upper layer protocols from the nitty-gritty detail of transmitting bits over, for example, Ethernet and collision detection, while the lower layers avoid having to know the details of each and every application and its protocol.

This abstraction also allows upper layers to provide services that the lower layers cannot, or choose not to, provide. Again, the original OSI Reference Model was extended to include connectionless services (OSIRM CL). For example, IP is not designed to be reliable and is a best effort delivery protocol. This means that all transport layer implementations must choose whether or not to provide reliability and to what degree. UDP provides data integrity (via a checksum) but does not guarantee delivery; TCP provides both data integrity and delivery guarantee (by retransmitting until the receiver acknowledges the reception of the packet).

This model lacks the formalism of the OSI reference model and associated documents, but the IETF does not use a formal model and does not consider this a limitation, as in the comment by David D. Clark, "We reject: kings, presidents and voting. We believe in: rough consensus and running code." Criticisms of this model, which have been made with respect to the OSI Reference Model, often do not consider ISO's later extensions to that model.

1. For multiaccess links with their own addressing systems (e.g. Ethernet) an address mapping protocol is needed. Such protocols can be considered to be below IP but above the existing link system. While the IETF does not use the terminology, this is a subnetwork dependent convergence facility according to an extension to the OSI model, the Internal Organization of the Network Layer (IONL).
2. ICMP & IGMP operate on top of IP but do not transport data like UDP or TCP. Again, this functionality exists as layer management extensions to the OSI model, in its Management Framework (OSIRM MF)
3. The SSL/TLS library operates above the transport layer (uses TCP) but below application protocols. Again, there was no intention, on the part of the designers of these protocols, to comply with OSI architecture.
4. The link is treated like a black box here. This is fine for discussing IP (since the whole point of IP is it will run over virtually anything). The IETF explicitly does not intend to discuss transmission systems, which is a less academic but practical alternative to the OSI Reference Model.

The following is a description of each layer in the TCP/IP networking model starting from the lowest level.

Link Layer

The Link Layer is the networking scope of the local network connection to which a host is attached. This regime is called the link in Internet literature. This is the lowest component layer of the Internet protocols, as TCP/IP is designed to be hardware independent. As a result TCP/IP has been implemented on top of virtually any hardware networking technology in existence.

The Link Layer is used to move packets between the Internet Layer interfaces of two different hosts on the same link. The processes of transmitting and receiving packets on a given link can be controlled both in the software device driver for the network card, as well as on firmware or specialized chipsets. These will perform data link functions such as adding a packet header to prepare it for transmission, then actually transmit the frame over a physical medium. The TCP/IP model includes specifications of translating the network addressing methods used in the Internet Protocol to data link addressing, such as

Media Access Control (MAC), however all other aspects below that level are implicitly assumed to exist in the Link Layer, but are not explicitly defined.

The Link Layer is also the layer where packets may be selected to be sent over a virtual private network or other networking tunnel. In this scenario, the Link Layer data may be considered application data which traverses another instantiation of the IP stack for transmission or reception over another IP connection. Such a connection, or virtual link, may be established with a transport protocol or even an application scope protocol that serves as a tunnel in the Link Layer of the protocol stack. Thus, the TCP/IP model does not dictate a strict hierarchical encapsulation sequence.

Internet Layer

The Internet Layer solves the problem of sending packets across one or more networks. Internetworking requires sending data from the source network to the destination network. This process is called routing.

In the Internet Protocol Suite, the Internet Protocol performs two basic functions:

- Host addressing and identification: This is accomplished with a hierarchical addressing system.
- Packet routing: This is the basic task of getting packets of data (datagrams) from source to destination by sending them to the next network node (router) closer to the final destination.

IP can carry data for a number of different upper layer protocols. These protocols are each identified by a unique protocol number: for example, Internet Control Message Protocol (ICMP) and Internet Group Management Protocol (IGMP) are protocols 1 and 2, respectively.

Some of the protocols carried by IP, such as ICMP (used to transmit diagnostic information about IP transmission) and IGMP (used to manage IP Multicast data) are layered on top of IP but perform internetworking functions. This illustrates the differences in the architecture of the TCP/IP stack of the Internet and the OSI model.

Transport Layer

The Transport Layer's responsibilities include end-to-end message transfer capabilities independent of the underlying network, along with error control, segmentation, flow control, congestion control, and application addressing (port numbers). End to end message transmission or connecting applications at the transport layer can be categorized as either connection-oriented, implemented in Transmission Control Protocol (TCP), or connectionless, implemented in User Datagram Protocol (UDP).

The Transport Layer can be thought of as a transport mechanism, e.g., a vehicle with the responsibility to make sure that its contents (passengers/goods) reach their destination safely and soundly, unless another protocol layer is responsible for safe delivery.

The Transport Layer provides this service of connecting applications through the use of service ports. Since IP provides only a best effort delivery, the Transport Layer is the first layer of the TCP/IP stack to offer reliability. IP can run over a reliable data link protocol such as the High-Level Data Link Control (HDLC). Protocols above transport, such as RPC, also can provide reliability.

For example, the Transmission Control Protocol (TCP) is a connection-oriented protocol that addresses numerous reliability issues to provide a reliable byte stream:

- data arrives in-order
- data has minimal error (i.e. correctness)
- duplicate data is discarded
- lost/discarded packets are resent
- includes traffic congestion control

The newer Stream Control Transmission Protocol (SCTP) is also a reliable, connection-oriented transport mechanism. It is Message-stream-oriented — not byte-stream-oriented like TCP — and provides multiple streams multiplexed over a single connection. It also provides multi-homing support, in which a connection end can be represented by multiple IP addresses (representing multiple physical interfaces), such that if one fails, the connection is not interrupted. It was developed initially for telephony applications (to transport SS7 over IP), but can also be used for other applications.

User Datagram Protocol is a connectionless datagram protocol. Like IP, it is a best effort, "unreliable" protocol. Reliability is addressed through error detection using a weak checksum algorithm. UDP is typically used for applications such as streaming media (audio, video, Voice over IP etc) where on-time arrival is more important than reliability, or for simple query/response applications like DNS lookups, where the overhead of setting up a reliable connection is disproportionately large. Real-time Transport Protocol (RTP) is a datagram protocol that is designed for real-time data such as streaming audio and video.

TCP and UDP are used to carry an assortment of higher-level applications. The appropriate transport protocol is chosen based on the higher-layer protocol application. For example, the File Transfer Protocol expects a reliable connection, but the Network File System (NFS) assumes that the subordinate Remote Procedure Call protocol, not transport, will guarantee reliable transfer. Other applications, such as VoIP, can tolerate some loss of packets, but not the reordering or delay that could be caused by retransmission.

The applications at any given network address are distinguished by their TCP or UDP port. By convention certain well known ports are associated with specific applications.

Application Layer

The Application Layer refers to the higher-level protocols used by most applications for network communication. Examples of application layer protocols include the File Transfer Protocol (FTP) and the Simple Mail Transfer Protocol (SMTP). Data coded according to application layer protocols are then encapsulated into one or (occasionally) more transport layer protocols (such as the Transmission Control Protocol (TCP) or User Datagram Protocol (UDP)), which in turn use lower layer protocols to effect actual data transfer.

Since the IP stack defines no layers between the application and transport layers, the application layer must include any protocols that act like the OSI's presentation and session layer protocols. This is usually done through libraries.

Application Layer protocols generally treat the transport layer (and lower) protocols as "black boxes" which provide a stable network connection across which to communicate, although the applications are usually aware of key qualities of the transport layer connection such as the end point IP addresses and port numbers. As noted above, layers are not necessarily clearly defined in the Internet protocol suite. Application layer protocols are most often associated with client–server applications, and the commoner servers have specific ports assigned to them by the IANA: HTTP has port 80; Telnet has port 23; etc. Clients, on the other hand, tend to use ephemeral ports, i.e. port numbers assigned at random from a range set aside for the purpose.

Transport and lower level layers are largely unconcerned with the specifics of application layer protocols. Routers and switches do not typically "look inside" the encapsulated traffic to see what kind of application protocol it represents, rather they just provide a conduit for it. However, some firewall and bandwidth throttling applications do try to determine what's inside, as with the Resource Reservation Protocol (RSVP). It's also sometimes necessary for Network Address Translation (NAT) facilities to take account of the needs of particular application layer protocols. (NAT allows hosts on private networks to communicate with the outside world via a single visible IP address using port forwarding, and is an almost ubiquitous feature of modern domestic broadband routers).

Hardware and software implementation

Normally, application programmers are concerned only with interfaces in the Application Layer and often also in the Transport Layer, while the layers below are services provided by the TCP/IP stack in the operating system. Microcontroller firmware in the network adapter typically handles link issues, supported by driver software in the operational system. Non-programmable analog and digital electronics are normally in charge of the physical components in the Link Layer, typically using an application-specific integrated circuit (ASIC) chipset for each network interface or other physical standard.

However, hardware or software implementation is not stated in the protocols or the layered reference model. High-performance routers are to a large extent based on fast non-programmable digital electronics, carrying out link level switching.

OSI and TCP/IP layering differences

The three top layers in the OSI model—the Application Layer, the Presentation Layer and the Session Layer—are not distinguished separately in the TCP/IP model where it is just the Application Layer. While some pure OSI protocol applications, such as X.400, also combined them, there is no requirement that a TCP/IP protocol stack needs to impose monolithic architecture above the Transport Layer. For example, the Network File System (NFS) application protocol runs over the eXternal Data Representation (XDR) presentation protocol, which, in turn, runs over a protocol with Session Layer functionality, Remote Procedure Call (RPC). RPC provides reliable record transmission, so it can run safely over the best-effort User Datagram Protocol (UDP) transport.

The Session Layer roughly corresponds to the Telnet virtual terminal functionality, which is part of text based protocols such as the HTTP and SMTP TCP/IP model Application Layer protocols. It also corresponds to TCP and UDP port numbering, which is considered as part of the transport layer in the TCP/IP model. Some functions that would have been performed by an OSI presentation layer are realized at the Internet application layer using the MIME standard, which is used in application layer protocols such as HTTP and SMTP.

Since the IETF protocol development effort is not concerned with strict layering, some of its protocols may not appear to fit cleanly into the OSI model. These conflicts, however, are more frequent when one only looks at the original OSI model, ISO 7498, without looking at the annexes to this model (e.g., ISO 7498/4 Management Framework), or the ISO 8648 Internal Organization of the Network Layer (IONL). When the IONL and Management Framework documents are considered, the ICMP and IGMP are neatly defined as layer management protocols for the network layer. In like manner, the IONL provides a structure for "subnetwork dependent convergence facilities" such as ARP and RARP.

IETF protocols can be encapsulated recursively, as demonstrated by tunneling protocols such as Generic Routing Encapsulation (GRE). While basic OSI documents do not consider tunneling, there is some concept of tunneling in yet another extension to the OSI architecture, specifically the transport layer gateways within the International Standardized Profile framework. The associated OSI development effort, however, has been abandoned given the overwhelming adoption of TCP/IP protocols.

Layer names and number of layers in the literature

The following table shows the layer names and the number of layers of networking models presented in RFCs and textbooks in widespread use in today's university computer networking courses.

Kurose, Forouzan	**Comer, Kozierok**	**Stallings**	**Tanenbaum**	**RFC 1122, Internet STD 3 (1989)**	**Cisco Academy**	**Mike Padlipsky's 1982 "Arpanet Reference Model" (RFC 871)**
Five layers	Four+one layers	Five layers	Four layers	Four layers	Four layers	Three layers
"Five-layer Internet model" or "TCP/IP protocol suite"	"TCP/IP 5-layer reference model"	"TCP/IP model"	"TCP/IP reference model"	"Internet model"	"Internet model"	"Arpanet reference model"
Application	Application	Application	Application	Application	Application	Application/Process
Transport	Transport	Host-to-host or transport	Transport	Transport	Transport	Host-to-host
Network	Internet	Internet	Internet	Internet	Internetwork	
Data link	Data link (Network interface)	Network access	Host-to-network	Link	Network interface	Network interface
Physical	(Hardware)	Physical				

These textbooks are secondary sources that may contravene the intent of RFC 1122 and other IETF primary sources such as RFC 3439.

Different authors have interpreted the RFCs differently regarding the question whether the Link Layer (and the TCP/IP model) covers Physical Layer issues, or if a hardware layer is assumed below the Link Layer. Some authors have tried to use other names for the Link Layer, such as network interface layer, in view to avoid confusion with the Data Link Layer of the seven layer OSI model. Others have attempted to map the Internet Protocol model onto the OSI Model. The mapping often results in a model with five layers where the Link Layer is split into a Data Link Layer on top of a Physical Layer. In literature with a bottom-up approach to Internet communication, in which hardware issues are emphasized, those are often discussed in terms of physical layer and data link layer.

The Internet Layer is usually directly mapped into the OSI Model's Network Layer, a more general concept of network functionality. The Transport Layer of the TCP/IP model, sometimes also described as the host-to-host layer, is mapped to OSI Layer 4 (Transport Layer), sometimes also including aspects of OSI Layer 5 (Session Layer) functionality. OSI's Application Layer, Presentation Layer, and the remaining functionality of the Session Layer are collapsed into TCP/IP's Application Layer. The argument is that these OSI layers do usually not exist as separate processes and protocols in Internet applications.

However, the Internet protocol stack has never been altered by the Internet Engineering Task Force from the four layers defined in RFC 1122. The IETF makes no effort to follow the OSI model although RFCs sometimes refer to it and often use the old OSI layer numbers. The IETF has repeatedly stated that Internet protocol and architecture development is not intended to be OSI-compliant. RFC 3439, addressing Internet architecture, contains a section entitled: "Layering Considered Harmful".