TCP/IP

Andrei.doncescu@laas.fr

Networks and Protocols

- A **network** is a collection of computers or computer-like devices that can communicate across a common transmission medium.
  Often the transmission medium is an insulated metal wire that carries electrical pulses between the computers, but the transmission medium could also be a phone line, or even no line at all in the case of a wireless network.

- Regardless of how the computers are connected, the communication process requires that data from one computer pass across the transmission medium to another computer.

- Computer A must be able to send a message or request to computer B. Computer B must be able to understand computer A's message and respond to it by sending a message back to computer A.

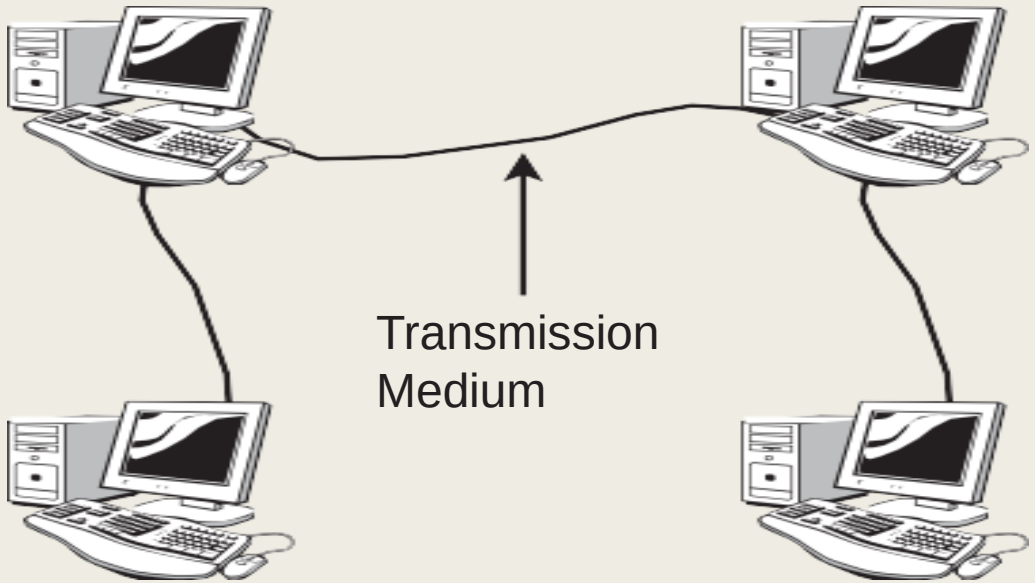

- A computer interacts with the world through one or more applications that perform specific tasks and manage the communication process. On modern systems, this network communication is so effortless that the user hardly even notices it.
- For instance, when you surf to a website, your web browser is communicating with the web server specified in the URL. When you view a list of neighboring computers in Windows Explorer or the Mac OS Finder, the computers on your local network are communicating to announce their presence.
- In every case, if your computer is part of a network, an application on the computer must be capable of communicating with applications on other network computers.

Network Protocol

- A **network protocol** is a system of common rules that helps define the complex process of network communication.
  Protocols guide the process of sending data from an application on one computer, through the networking components of the operating system, to the network hardware, across the transmission medium, and up through the destination computer's network hardware and operating system to a receiving application

Network Protocol

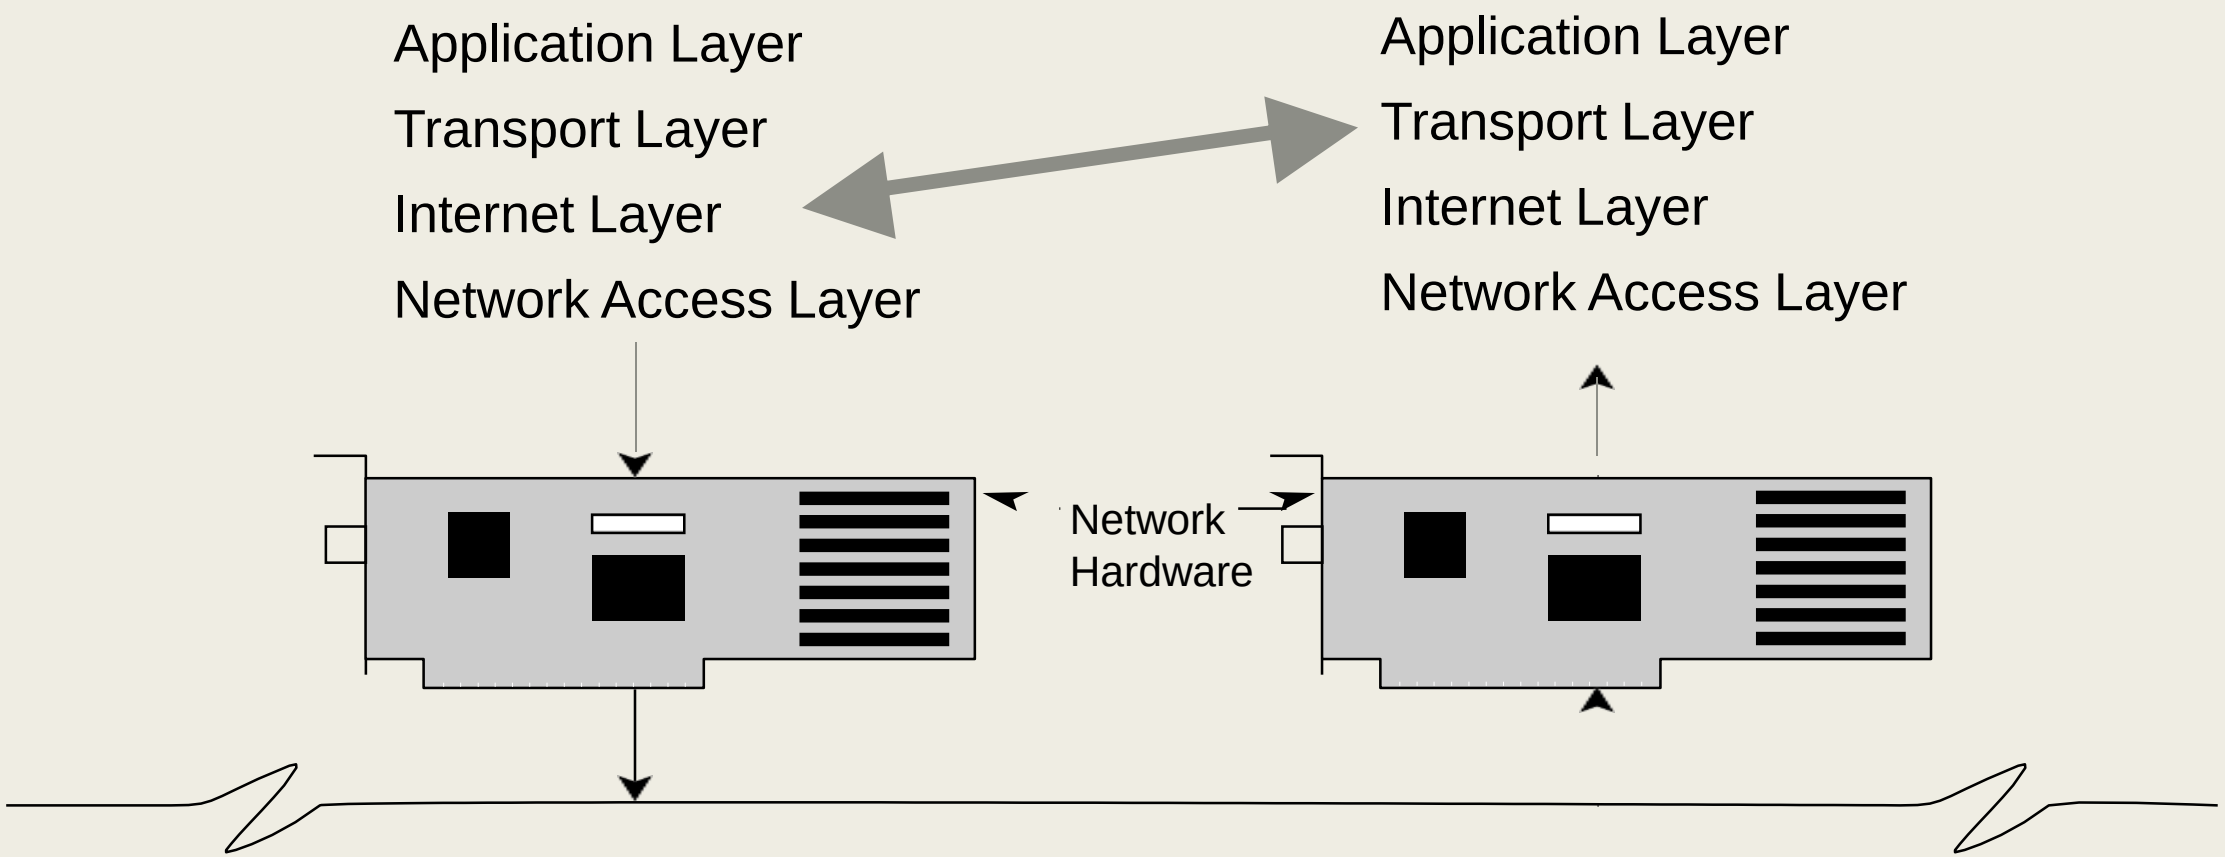

TCP/IP Network Protocol

- The protocols of TCP/IP define the network communication process and, more importantly, define how a unit of data should look and what information it should contain so that a receiving computer can interpret the message correctly.
- TCP/IP and its related protocols form a complete system defining how data should be processed, transmitted, and received on a TCP/IP network.
- A system of related protocols, such as the TCP/IP protocols, is called a **protocol suite**.

TCP/IP

- The actual act of formatting and processing TCP/IP transmissions is performed by a software component known as the vendor's **implementation** of TCP/IP.
- For instance, a TCP/IP software component in Microsoft Windows enables Windows computers to process TCP/IP-formatted data and thus to participate in a TCP/IP network.

  - *A **TCP/IP standard** is a system of rules defining communication on TCP/IP networks.*
  - *A **TCP/IP implementation** is a software component that performs the functions that enable a computer to participate in a TCP/IP network.*

The Development of TCP/IP

- TCP/IP's design is a result of its historical role as the protocol system for what was to become the Internet. The Internet, like so many other high-tech developments, grew from research originally performed by the United States Department of Defense.
- In the late 1960s, Defense Department officials began to notice that the military was accumulating a large and diverse collection of computers. Some of those computers weren't networked, and others were grouped in small, closed networks with incompatible proprietary protocols.
- **Proprietary**, in this case, means that the technology is controlled by a private entity (such as a corporation). That entity might not have any interest in divulging enough information about the protocol so that users can use it to connect to other (rival) network protocols.

The Development of TCP/IP

- Defense officials began to wonder whether it would be possible for these disparate computers to share information. These visionary soldiers created a network that became known as **ARPAnet**, named for the Defense Department's Advanced Research Projects Agency (ARPA).
- As this network began to take shape, a group of computer scientists, led by Robert E. Kahn and Vinton Cerf, started to work on a versatile protocol system that would support a wide range of hardware and provide a resilient, redundant, and decentralized system for delivering data on a massive, global scale.
- The result of this research was the beginning of the TCP/IP protocol suite.
  The National Science Foundation wanted to build a network to connect research institutions, it adopted ARPAnet's protocol system and began to build what we know as the Internet. University College of London and other European research institutes contributed to the early development of TCP/IP, and the first transatlantic communications tests began around 1975.

The Development of TCP/IP

- Two important features of TCP/IP that provide for this decentralized environment are as follows:
- **End-node verification:** The two computers that are actually communicating called the end nodes because they are at each end of the chain passing the message are responsible for acknowledging and verifying the transmission.
  All computers basically operate as equals, and there is no central scheme for overseeing communications.
- **Dynamic routing:** Nodes are connected through multiple paths, and the routers choose a path for the data based on present conditions.

The Personal Computing Revolution

- Around the time the Internet was catching on, most computers were multiuser systems. Several users in a single office (or campus) connected to a single computer through a text-screen interface device known as a terminal. Users worked independently, but in fact, they were all accessing the same computer, which required only one Internet connection to serve a large group of users. The proliferation of personal computers in the 1980s and 1990s began to change this scenario.
- In the early days of personal computers, most users didn't even bother with networking. But as the Internet began to reach beyond its original academic roots, users with personal computers started looking for ways to connect.
- One solution was a dial-up connection through a modem, which offered network connectivity through a phone line.
- But users also wanted to connect to other nearby computers in their own office to share files and access peripheral devices. To address this need, another network concept, the **local area network (LAN)** began to take form.

The Personal Computing Revolution

- Early LAN protocols did not provide Internet access and were designed around proprietary protocol systems. Many did not support routing of any kind. Computers in a single workgroup would talk to each other using one of these proprietary protocols, and users would either do without the Internet, or they would connect separately using a dial-up line.
- Internet service providers grew more numerous, and Internet access became more affordable, companies began to ask for a fast, permanent, always-on Internet connection. A variety of solutions began to emerge for getting LAN users connected to the TCP/IP-based Internet.
- Specialized **gateways** offered the protocol translation necessary for these local networks to reach the Internet.
- Gradually, however, the growth of the World Wide Web, and the accompanying need for end-user Internet connectivity, made TCP/IP essential, leaving little purpose for proprietary LAN protocols such as AppleTalk, NetBEUI, and Novell's IPX/SPX.

TCP/IP Features

- TCP/IP includes many important features
- In particular, pay close attention to the way the TCP/IP protocol suite addresses the following problems:
  - *Logical addressing*
  - *Routing*
  - *Name resolution*
  - *Error control and flow control*
  - *Application support*

These issues are at the heart of TCP/IP.

Logical Addressing

- A network adapter has a unique physical address. In the case of ethernet, the **physical address** (which is sometimes called a Media Access Control [MAC] address) is typically assigned to the adapter at the factory, although some contemporary devices now provide a means for changing the physical address.
- On a LAN, low-lying hardware-conscious protocols deliver data across the physical network using the adapter's physical address.
- There are many network types, and each has a different way of delivering data. On a basic ethernet network, for example, a computer sends messages directly onto the transmission medium.
- The network adapter of each computer listens to every transmission on the local network to determine whether a message is addressed to its own physical address.

Logical Addressing

- On large networks, of course, every network adapter can't listen to every message. (Imagine your computer listening to *every* piece of data sent over the Internet.) As the transmission medium becomes more populated with computers, a physical addressing scheme cannot function efficiently.
- Network administrators often segment networks using devices such as routers to reduce network traffic.
- On routed networks, administrators need a way to subdivide the network into smaller subnetworks (called **subnets**) and impose a hierarchical design so that a message can travel efficiently to its destination.
- TCP/IP provides this subnetting capability through logical addressing.
- A **logical address** is an address configured through the network software.

Logical Addressing

An IP address can include

- *A network ID number identifying a network*
- *A subnet ID number identifying a subnet on the network*
- *A host ID number identifying the computer on the subnet*

- The IP addressing system also lets the network administrator impose a sensible numbering scheme on the network so that the progression of addresses reflects the internal organization of the network.

Internet-Ready Addresses

- If your network is isolated from the Internet, you are free to use any IP addresses you want (as long as your network follows the basic rules for IP addressing).
- If your network will be part of the Internet, however, Internet Corporation for Assigned Names and Numbers (ICANN), which was formed in 1998, will assign a network ID to your network, and that network ID will form the first part of the IP address.
- One interesting development is a system called Network Address Translation (NAT), which lets you use a private, nonroutable IP address on the local network that the router will translate into an official Internet ready address for Internet communications.

Routing

- A **router** is a special device that can read logical addressing information and direct data across the network to its destination. At the simplest level, a router divides a local subnet from the larger network

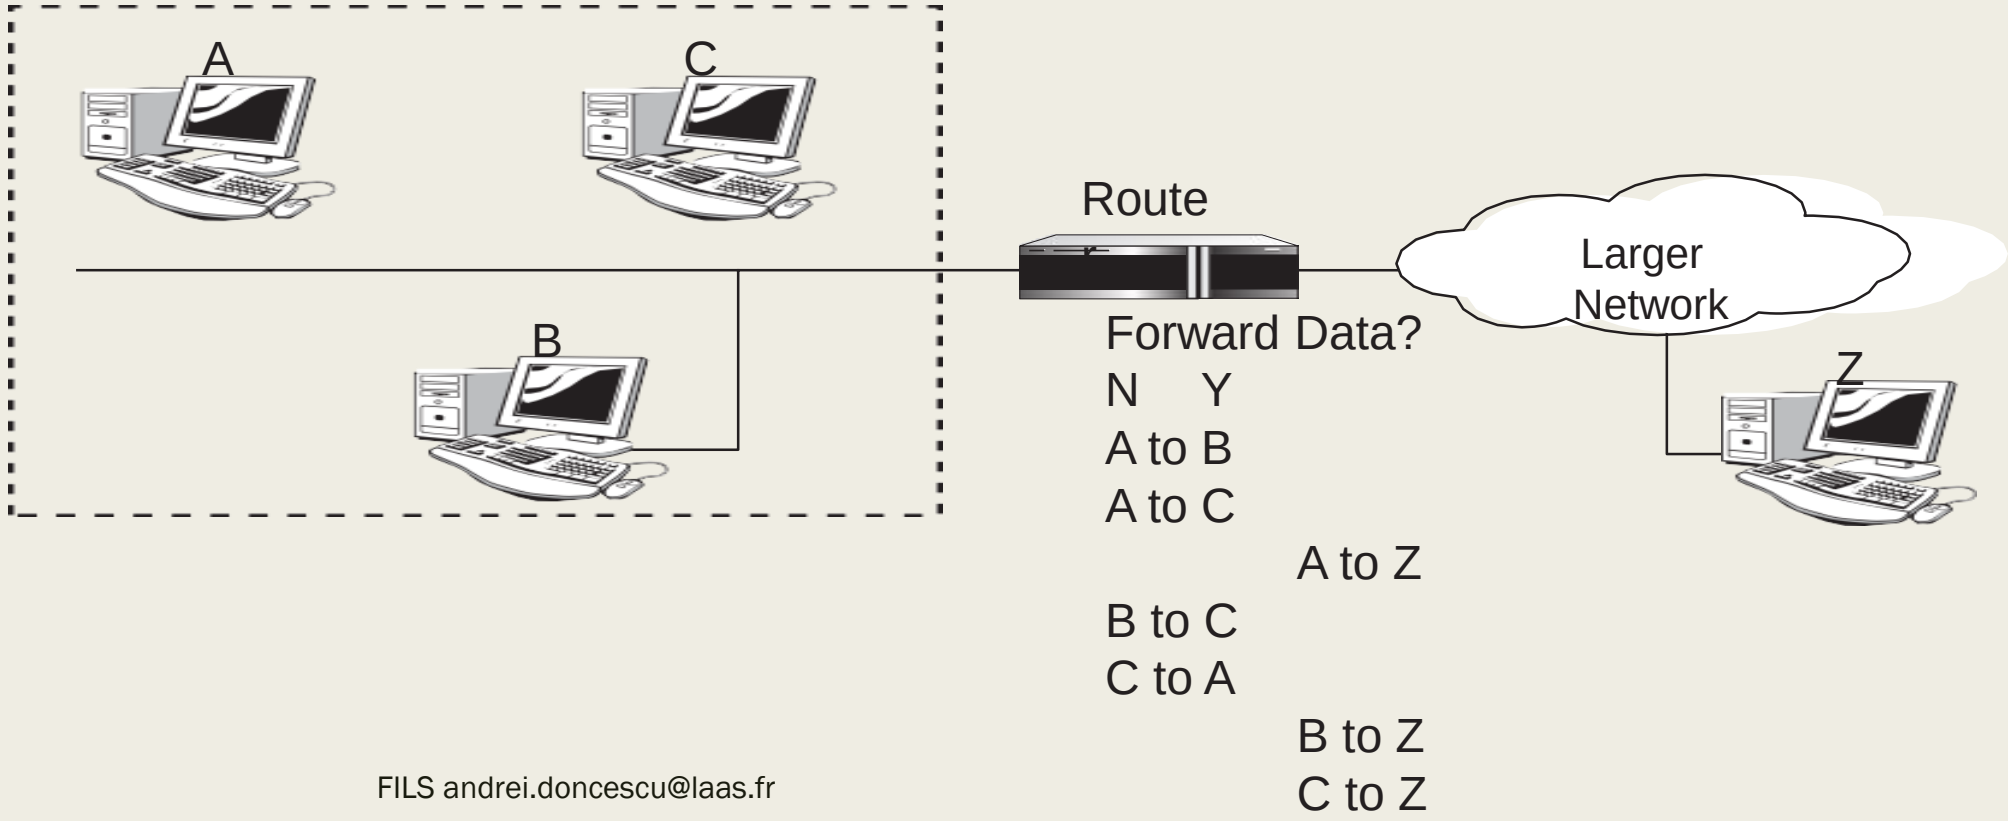

Routing

- Data addressed to another computer or device on the local subnet does not cross the router and, therefore, doesn't clutter up the transmission lines of the greater network. If data is addressed to a computer outside the subnet, the router forwards the data accordingly. As previously mentioned in this hour, large networks such as the Internet include many routers and provide multiple paths from the source to the destination.
- TCP/IP includes protocols that define how the routers find a path through the network.

Other Filtering Devices

- Network devices such as bridges, switches, and intelligent hubs can also filter traffic and reduce network traffic. Because these devices work with physical addresses rather than logical addresses, they cannot perform the complex routing functions

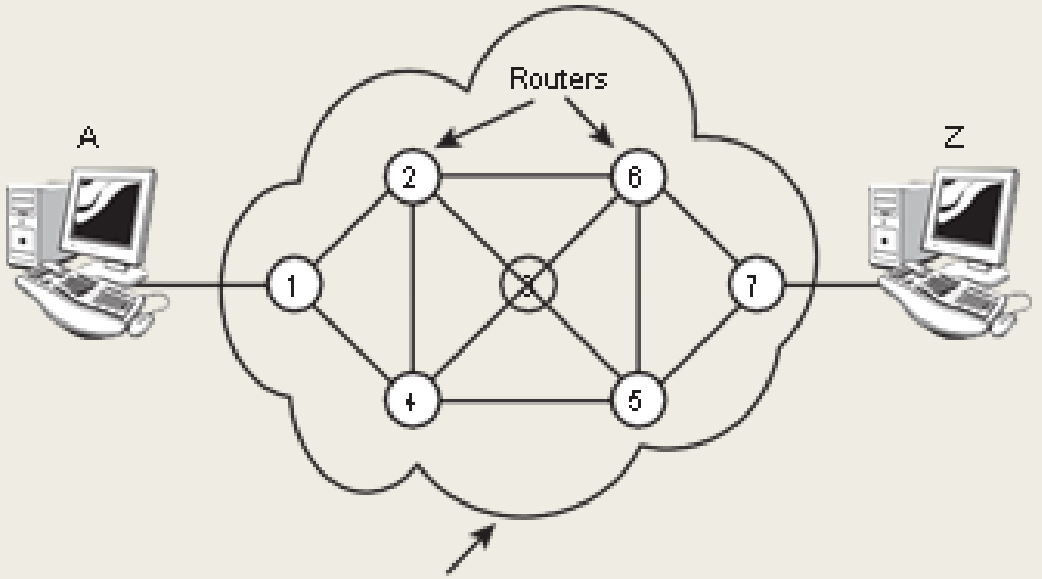

Name Resolution

- Although the numeric IP address is probably more user friendly than the network adapter's prefabricated physical address, the IP address is still designed for the convenience of the computer rather than the convenience of the user.
- People might have trouble remembering whether a computer's address is 111.121.131.146 or 111.121.131.156. TCP/IP, therefore, provides for a parallel structure of user-oriented alphanumeric names, called **domain names** or Domain Name System (DNS) names.
- This mapping of domain names to an IP address is called **name resolution**. Special computers called **name servers** store tables showing how to translate these domain names to and from IP addresses.

Name Resolution

- The computer addresses commonly associated with email or the World Wide Web are expressed as DNS names (for example, www.microsoft.com, falcon.ukans.edu, and idir.net). TCP/IP's **name service** system provides for a hierarchy of name servers that supply domain name/IP address mappings for DNS-registered computers on the network. This means that the everyday user rarely has to enter or decipher an actual IP address.
- DNS is the name resolution system for the Internet and is the most common name resolution method.
- However, other techniques also exist for resolving alphanumeric names to IP addresses. These alternative systems have gradually faded in importance in recent years, but name resolution services such as the Windows Internet Name Services (WINS), which resolves NetBIOS names to IP addresses, are still in operation around the world.

Error Control and Flow Control

- The TCP/IP protocol suite provides features that ensure the reliable delivery of data across the network.
- These features include checking data for transmission errors (to ensure that the data that arrives is exactly what was sent) and acknowledging successful receipt of a network message.
- TCP/IP's Transport defines many of these error-control, flow-control, and acknowledgment functions through the TCP protocol.
- Lower-level protocols at TCP/IP's Network Access layer also play a part in the overall system of error control.

Application Support

- Several network applications might be running on the same computer.
- The protocol software must provide some means for determining which incoming packet belongs with each application.
- In TCP/IP, this interface from the network to the applications is accomplished through a system of logical channels called **ports**.
- Each port has a number that is used to identify the port.
- Think of these ports as logical pipelines within the computer through which data can flow from the application to (and from) the protocol software

Application Support

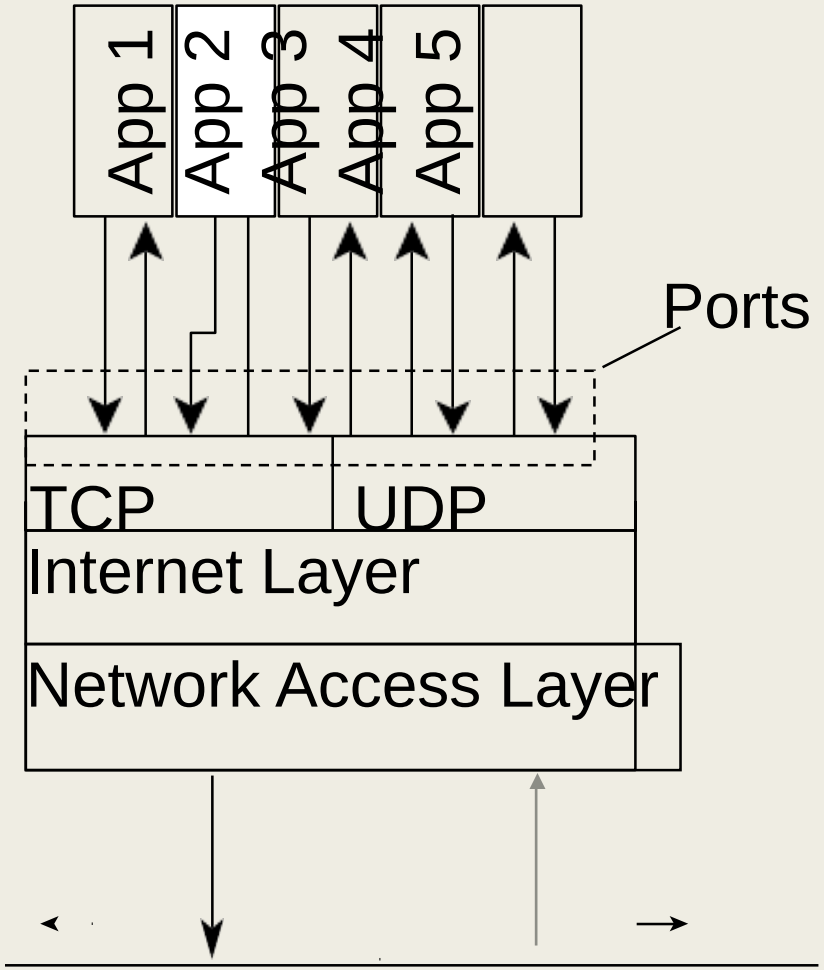

Application Support

The TCP/IP suite also includes a number of ready-made applications designed to assist with various network tasks.

`ftp`	File transfer
`Lpr`	Printing
`Ping`	Configuration/troubleshooting
`Route`	Configuration/troubleshooting
`telnet`	Remote terminal access
`Traceroute`	Configuration/troubleshooting

Standards Organizations and RFCs

Several organizations have been instrumental in the development of TCP/IP and the Internet. Another way in which TCP/IP reveals its military roots is in the quantity and obscurity of its acronyms. Still, a few organizations in the past and present of TCP/IP deserve mention, as follows:

- **Internet Architecture Board (IAB):** The governing board that sets policy for the Internet and sees to the further development of TCP/IP standards.
- **Internet Engineering Task Force (IETF):** An organization that studies and rules on engineering issues. The IETF is divided into workgroups that study particular aspects of TCP/IP and the Internet, such as applications, routing, network management, and so forth.
- **Internet Research Task Force (IRTF):** The branch of the IAB that sponsors long-range research.

- **Internet Corporation for Assigned Names and Numbers (ICANN):** An organization established in 1998 that coordinates the assignment of Internet domain names, IP addresses, and globally unique protocol parameters such as port numbers (www.icann.com).

Standards Organizations and RFCs

- Because TCP/IP is a system of open standards that are not owned by any company or individual, the Internet community needs a comprehensive, independent, vendor-neutral process for proposing, discussing, and releasing additions and changes.
- Most of the official documentation on TCP/IP is available through a series of **Requests for Comment (RFCs)**. The library of RFCs includes Internet standards and reports from workgroups. IETF official specifications are published as RFCs.
- Many RFCs are intended to illuminate some aspect of TCP/IP or the Internet
- A majority of the RFCs were created by industry workgroups and research institutions, anyone can submit an RFC for review.

Standards Organizations and RFCs

- The RFCs provide essential technical background for anyone wanting a deeper understanding of TCP/IP.

Number	Title
791	Internet Protocol (IP)
792	Internet Control Message Protocol (ICMP)
793	Transmission Control Protocol
959	File Transfer Protocol
968	Twas the Night Before Start-up
1180	TCP/IP Tutorial
1188	Proposed Standard for Transmission of Datagrams over FDDI Networks
2097	The PPP NetBIOS Frames Control Protocol

Q&A

- Q. *What is the difference between a protocol standard and a protocol implementation?*
  - *A protocol standard is a system of rules. A protocol implementation is a soft- ware component that applies those rules to provide networking capability to a computer.*

- Q. *Why was end-node verification an important feature of ARPAnet?*
  - *By design, the network was not supposed to be controlled from any central point. The sending and receiving computers, therefore, had to take charge of verifying their own communication.*

- Q. *Why do larger networks employ name resolution?*
  - *IP addresses are difficult to remember and easy to get wrong. DNS-style domain names are easier to remember because they let you associate a word or name with the IP address.*

Quiz

- *What is a network protocol?*
- *What are two features of TCP/IP that allow it to operate in a decentralized manner?*
- *What system is responsible for mapping domain names to IP addresses?*
- *What are RFCs?*
- *What is a port?*

Exercises

- Visit www.rfc-editor.org and browse some of the RFCs.
- Visit the IETF and explore the various active working groups at datatracker.ietf.org/wg/.
- Visit the IRTF at www.irtf.org and explore some of the ongoing research.
- Visit the ICANN About page at www.icann.org/en/about/ and learn about the ICANN mission.
- Read RFC 1160 for an early history (up to 1990) of the IAB and IETF.

Key Terms

- **ARPAnet:** An experimental network that was the birthplace of TCP/IP.
- **Domain name:** An alphanumeric name associated with an IP address through TCP/IP's DNS name service system.
- **Gateway:** A router that connects a LAN to a larger network. In the days of proprietary LAN protocols, the term *gateway* sometimes applied to a router that performed some kind of protocol conversion.
- **IP address:** A logical address used to locate a computer or other networked device (such as a printer) on a TCP/IP network.
- **Local Area Network (LAN):** A small network belonging to a single office, organization, or home, usually occupying a single geographical location.
- **Logical address:** A network address configured through the protocol software.

Key Terms

- **Name service:** A service that associates human-friendly alphanumeric names with network addresses. A computer that provides this service is known as a **name server**, and the act of resolving a name to an address is called **name resolution**.
- **Network Protocol:** A set of common rules defining a specific aspect of the communication process.
- **Physical address:** An address associated with the network hardware. In the case of an ethernet adapter, the physical address is typically assigned at the factory.
- **Port:** An internal channel or address that provides an interface between an application and TCP/IP's Transport layer.
- **Proprietary:** A technology controlled by a private entity, such as a corporation.
- **Protocol implementation:** A software component that implements the communication rules defined in a protocol standard.

Key Terms

Protocol system or protocol suite: A system of interconnected standards and procedures (protocols) that enables computers to communicate over a network.

RFC (Request for Comment): An official technical paper providing relevant information on TCP/IP or the Internet. You can find the RFCs at sever- al places on the Internet; try www.rfc-editor.org.

Router: A network device that forwards data by logical address and can also be used to segment large networks into smaller subnetworks.

Transport Control Protocol/Internet Protocol (TCP/IP): A network protocol suite used on the Internet and also on many other networks around the world.