

# CGS 3175: Internet Applications Fall 2009

## Introduction to the Internet

Instructor : Dr. Mark Llewellyn  
markl@cs.ucf.edu  
HEC 236, 407-823-2790  
<http://www.cs.ucf.edu/courses/cgs3175/fall2009>

School of Electrical Engineering and Computer Science  
University of Central Florida



# What is the Internet?

- The Internet traces its roots to a project of the U.S. Dept. of Defense's then named Advanced Research Projects Agency (ARPA) in the 1960s.
- The **ARPANET** project was intended to support DoD research on computer networking.
  - When this project began, there had only been a few small experimental networks providing communication between geographically dispersed computers from different manufacturers running different operating systems.
- The purpose of ARPANET was to create a much larger network, both in order to electronically connect DoD-sponsored researchers and to experiment with and develop tools for heterogeneous computer networking.



# What is the Internet?

- ARPANET was launched in 1969 and by the end of that year consisted of 4 computers at 4 sites running 4 different operating systems.
  - The first site in ARPANET was at UCLA.
- ARPANET grew steadily, but because it was restricted to DoD-funded organizations and was a research project, it was never large.
- When many ARPANET nodes were split off to form a separate network called MILNET in 1983, there were only 113 nodes in the entire ARPANET network (mostly universities).



# What is the Internet?

- Despite the relatively small number of machines actually on the ARPANET, the benefits of networking were becoming known to a wide audience.
- For example, e-mail was available on ARPANET beginning in 1972, and it soon became an extremely popular application for those who had ARPANET access.
- It wasn't long before other networks were being built, both internationally and regionally within the US.
  - The regional US networks were often cooperative efforts between universities.
  - SURAnet (Southeastern University Research Association Network) was organized by the University of Maryland 1982 and eventually included all of the major universities in the southeastern US). CSNET (Computer Science Network) was partially funded by NSF to aid scientists at universities without ARPANET access.

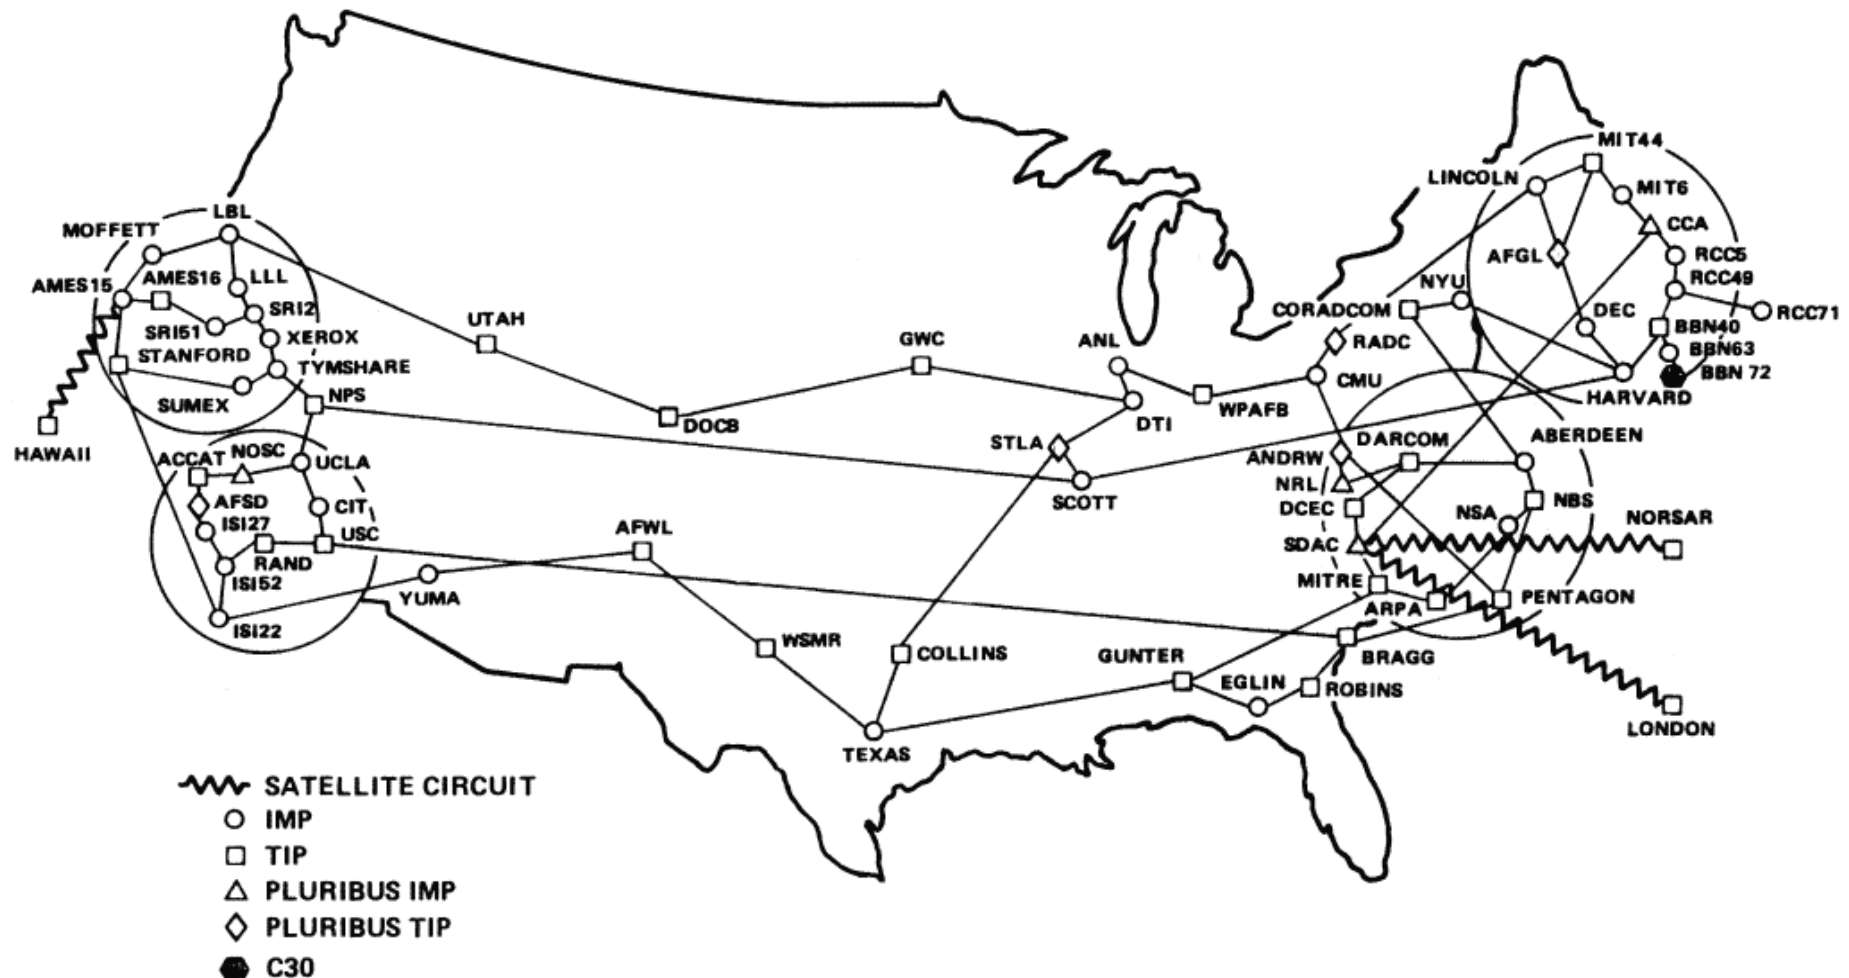


# What is the Internet?

- While other networks were being constructed during the 1970s and 1980s, the ARPANET project continued to fund research on networking. (A map of ARPANET in 1980 is shown on the next page.)
- Several of the most widely used Internet protocols - including the **File Transfer Protocol (FTP)** and **Simple Mail Transfer Protocol (SMTP)** which underlie many of the Internet's file transfer and e-mail operations, respectively – were initially developed under ARPANET.
- However, the most important development of ARPANET, as far as the Internet was concerned, was the development of the **TCP/IP (Transmission Control Protocol/Internet Protocol)** communication protocol.
  - TCP/IP was designed to be used for host-to-host communication both within *local area networks* and between networks.



# ARPANET GEOGRAPHIC MAP, OCTOBER 1980



(NOTE: THIS MAP DOES NOT SHOW ARPA'S EXPERIMENTAL SATELLITE CONNECTIONS)  
 NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES



# What is the Internet?

- ARPANET switched from using an earlier protocol to TCP/IP during 1982.
- At around the same time an ARPA Internet was created, allowing computers on some outside networks such as CSNET to communicate via TCP/IP with computers on the ARPANET.
- A “connection” from CSNET to the APRA Internet often meant that a modem connection was made from one computer to another for the purposes of sending along an e-mail message.
  - This form of communication was **asynchronous** – meaning that the e-mail might be delayed some time before it was actually delivered, which precluded interactive communication.



# What is the Internet?

- Each institution connecting to CSNET was largely on its own in determining how it was going to connect to the network.
- At first, many institutions connected through the so-called PhoneNet (modems) approach from passing e-mail messages. This generally, involved long distance calls, and the expense of these calls was a mounting problem for most organizations. Other options, which included leasing phone lines for dedicated use, were typically even more expensive.
- It quickly became obvious to everyone that the CSNET institutions were not enjoying all of the potential benefits of the ARPA Internet.





# What is the Internet?

- Beginning in 1985, the NSF began work on a new network based on the TCP/IP, called NSFNET.
  - This network connected supercomputer centers in Ithaca, Pittsburgh, Princeton, Champaign, Boulder, and San Diego.
- One of the primary goals of this network was to connect the NSF's new regional supercomputing centers.
- It was also decided that regional networks should be able to connect to NSFNET, so that NSFNET would provide a **backbone** through which other networks could interconnect synchronously.
  - Regional networks that connected to NSFNET included SURANET and NYSERNet (through Ithaca), JvNCnet (through Princeton), SDSCnet (through San Diego).



# What is the Internet?

- The original NSFNET backbone operated at 56 kbps, the maximum speed of a home dial-up line today. The primary network traffic was textual, so this was a reasonable starting point at the time.
- Once operational, the number of machines connected to NSFNET grew quickly, in part due to NSF encouragement through grants and funding.
  - The backbone rate was upgraded to 1.5 Mbps (T1 level) in 1988 and then to 45 Mbps (T2 level) in 1991. Around this time, several additional research networks were added to the NSFNET.
  - In 1988, networks in Canada and France were connected to NSFNET.
  - For the remaining seven years of its life, networks from 10 or so new countries were added to NSFNET each year.



# What is the Internet?

- NSFNET quickly supplanted ARPANET, which was officially decommissioned in 1990.
- At this point, NSFNET was at the center of the Internet, that is, the collection of computer networks connected via the public backbone and communicating across networks using TCP/IP.
- This same year (1990), commercial Internet dial-up access was first offered. However, the NSFNET terms of usage stipulated that purely commercial traffic was not to be carried over the backbone – the purpose of the Internet was still, in the eyes of the NSF, research and education.



# What is the Internet?

- It became increasingly clear that there could be significant benefits to allowing commercial traffic on the Internet.
- One of the arguments for allowing commercial traffic was economic – commercial traffic would increase network usage, leading to reduced unit costs through economics of scale. This would in turn provide a less expensive network for research and education.
- Whatever the motivation, the restriction on commercial traffic was rescinded in 1991 (it lasted just about a year), which spurred further growth of the Internet and laid the groundwork for the metamorphosis of the Internet from a tool used primarily by scientists and research institutions to the conduit for information, entertainment, and commerce that we know today.



# What is the Internet?

- This rapid growth and commercialization of the Internet lead fairly quickly to the NSF being able to leave its role as the operator of the Internet backbone in the United States.
- Those responsibilities were assumed by private telecommunication firms in 1995.
  - These firms are paid by other firms, such as some of the larger **Internet service providers (ISPs)**, who connect directly with the Internet backbone.
  - These ISPs, in turn, are paid by their users, which may include smaller ISPs as well as end users.
- In summary, the Internet is a collection of computers that can communicate with one another using TCP/IP over an open, global communication network.



# What is the Internet?

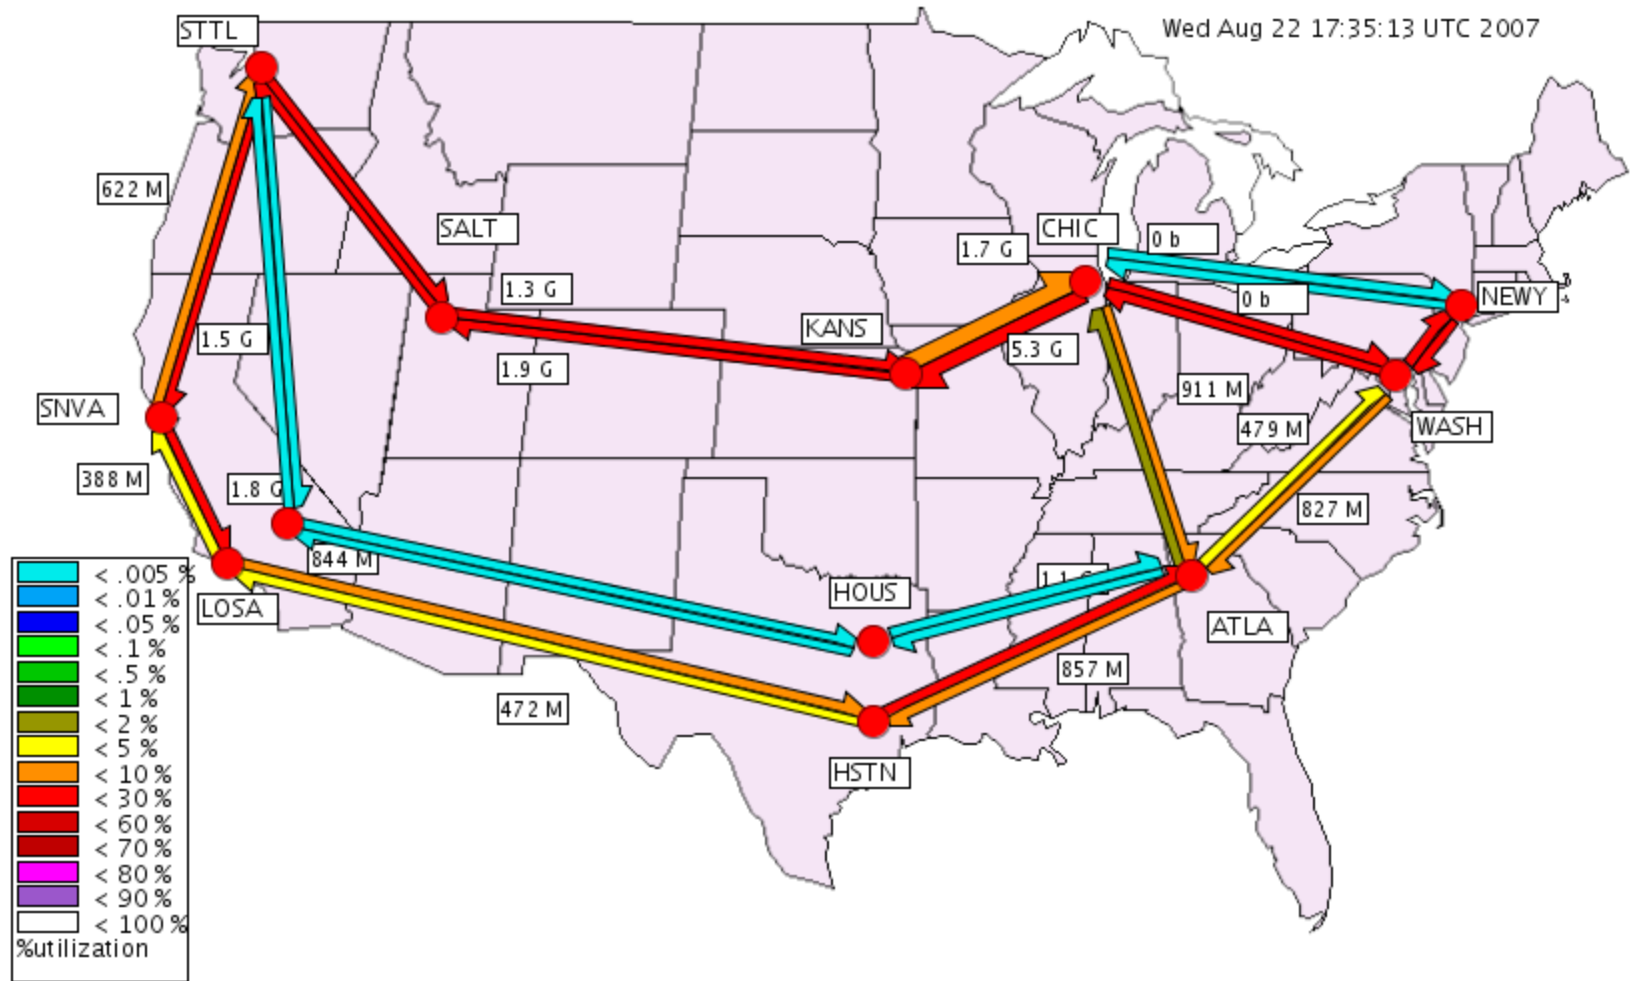


# A Note on Internet2

- There is a new research oriented Internet known as Internet 2, which has been built and is maintained by a consortium of about 60 different organizations, primarily, IU, Cisco, and Quest Communications.
- The backbone speed is currently 10 Gbps, which is about 1000 times faster than the Internet.
  - This means that using Internet2 you could download a DVD-quality copy of the movie The Matrix in about 30 seconds!
- The picture on the following page shows the current Internet2 configuration.



# Internet2



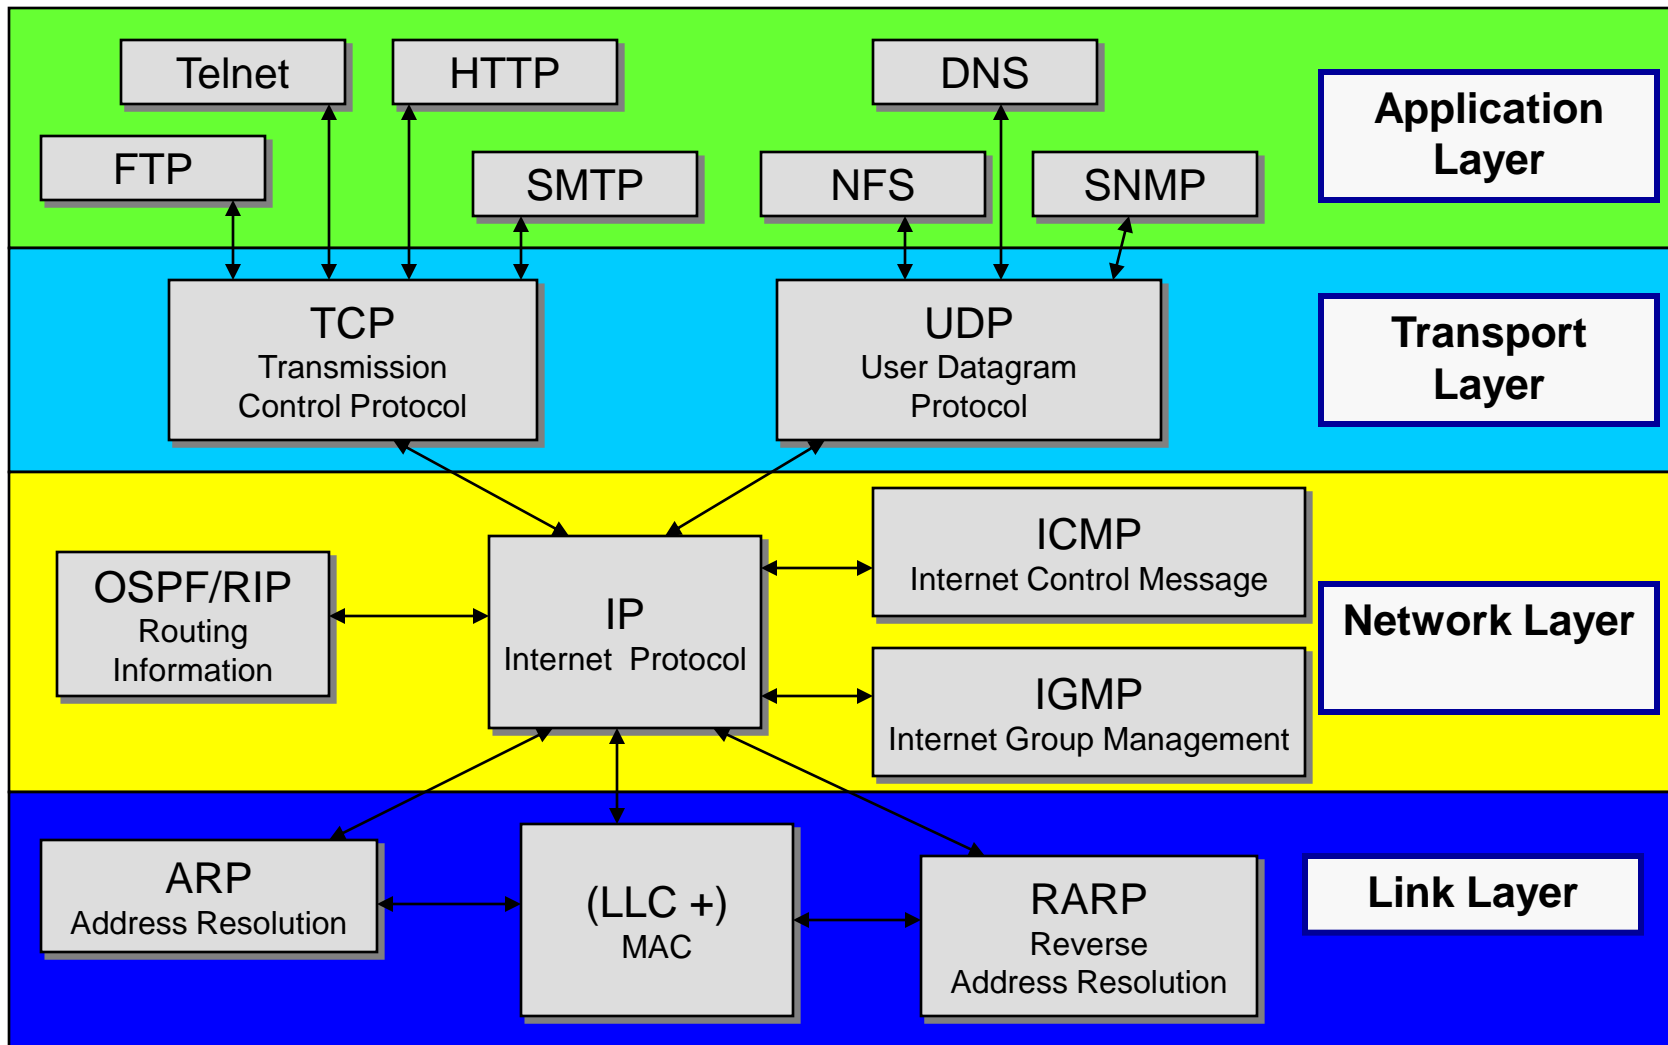


# Basic Internet Protocols

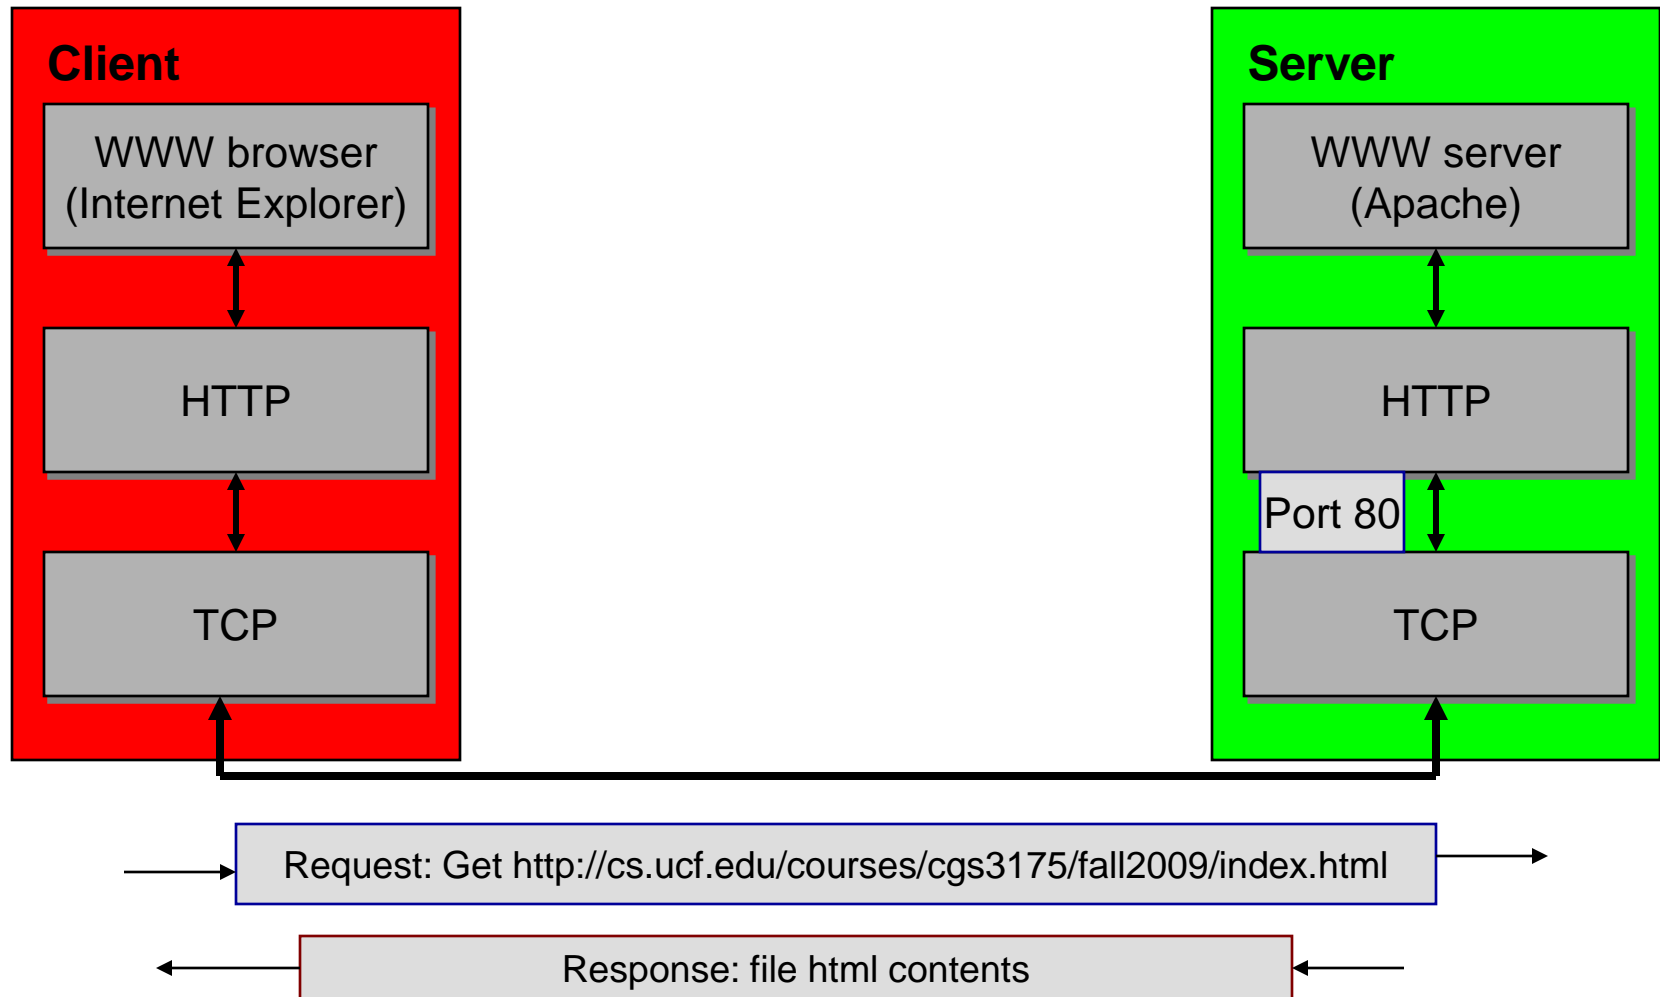
- A computer communication protocol is a detailed specification of how communication between two computers will be carried out in order to serve some purpose.
  - For example, as we will see later, the Internet Protocol (the IP in TCP/IP) specifies both the high-level behavior of the software implementing the protocol and the low-level details such as the specific fields of information that will be contained in a communication message, the order in which these fields will appear, the number of bits in each field, and how these bits should be interpreted.
- In this course, we will be primarily interested in a high-level view of general-purpose Internet protocols.



# Internet Reference Model



# The Architecture of a Web Service



# TCP/IP

- Since TCP/IP is fundamental to the definition of the Internet, it's natural to look at these protocols first.
- Note: TCP/IP are actually two protocols, although they are often treated and referred to as a single protocol.
  - The reason they are often treated as one is that the bulk of the services we associate with the Internet – e-mail, Web browsing, file downloads, accessing remote databases – are built on top of both the TCP and IP protocols.
  - In reality, only the IP, the Internet Protocol, is fundamental to the definition of the Internet.
- So we'll look at IP first.



# Internet Protocol – IP

- A key element of IP is the **IP address**, which is simply a 32-bit number.
- At any given moment, each device on the Internet has one or more IP addresses associated with it (although the device associated with a given address may change over time).
- IP addresses are normally written as a sequence of four decimal numbers separated by periods (called “dots”), as in 132.170.7.155 (which is the IP address of the computer in my office). Each decimal number represents one byte of the IP address.
- To find the IP address of your computer, go to [www.ip-address.com](http://www.ip-address.com) or [www.ipaddress.com](http://www.ipaddress.com) (See next two pages for examples.)



http://www.ip-address.com/

File Edit View Favorites Tools Help

Google Search

Check AutoFill Sign In

Favorites Suggested Sites Free Hotmail Web Slice Gallery KeepVid- Download and s...

My IP address? Free IP Address tracer, IP address l...

Home RSS Mail Print Page Safety Tools



## IP-address.com - My IP address and DNS tools

What is my IP? IP address finder, Speedtest and more.

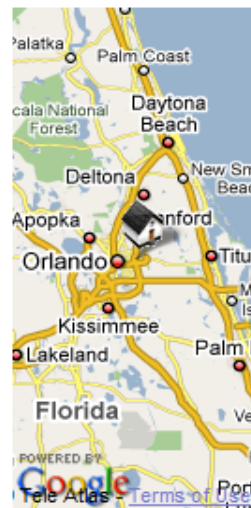


An IP address (Internet Protocol Address) is a logical address of a network adapter. The IP address is unique and identifies computers on a network.

# My IP address: 132.170.10.195

My IP Address Location: Orlando in United States

ISP of my IP: University of Central Florida



### [Reverse Email Lookup](#)

1) Type in Email Address. 2) Get Owner Name & Info.

EmailFinder.com

Ads by Google

Done



Internet | Protected Mode: Off



100%



ip address my lookup find at ipaddress.com - Windows Internet Explorer

http://www.ipaddress.com/online/address/lookup/find/my\_ip\_address.htm?yt=qs%3d06oENya4Z

File Edit View Favorites Tools Help

Google Search

ip address my lookup find at ipaddress.com

Your IP Address is? 132.170.10.195

## Welcome to IPAddress.com

Sponsored listings for My IP Address:

- Hide Your IP Address**  
Protect Your Data & Your Privacy With GoTrusted Security. Free Trial.  
[GoTrusted.com](#)
- IP Address Management**  
Gain Control of Your IP Addresses With Ease. Freeware Available.  
[www.infoblox.com](#)
- Free Address Lookup**  
Verify addresses nationwide. City, state, zip code & map. Easy to use.  
[www.MelissaData.com](#)
- Hide Your IP Address NOW**  
Your IP Is A Tracking Bug ! Surf The Web In Total Privacy & Invisible  
[www.Hide-My-IP-Address.com](#)
- Limited Time Only - \$9.99**  
Save On New Domain Registrations. Includes 24/7 Support And More.  
[networksolutions.com](#)

### Related Links

- [My IP Address](#)
- [Static IP Address](#)
- [ISP Provider](#)
- [Address IP Search](#)
- [Internet Address](#)
- [IP Address Lookup](#)
- [ISP](#)
- [Hide IP Address](#)
- [Address](#)
- [Find Out My IP Address](#)
- [Address IP Track](#)
- [Dynamic IP](#)
- [IP Address To Computer](#)
- [Router IP Address](#)
- [IP](#)

Internet | Protected Mode: Off 100%



ARIN: WHOIS Database Search - Windows Internet Explorer

http://ws.arin.net/whois/?queryinput=132.170.10.195

File Edit View Favorites Tools Help

Google Search

ARIN: WHOIS Database Search

## ARIN WHOIS Database Search

Relevant Links: [ARIN Home Page](#) [ARIN Site Map](#) Training: [Querying ARIN's WHOIS](#)

Search ARIN WHOIS for: 132.170.10.195

OrgName: University of Central Florida  
OrgID: **UCF**  
Address: CS Building 54, Room 301  
Address: 4000 Central Florida Blvd.  
City: Orlando  
StateProv: FL  
PostalCode: 32816-2500  
Country: US

NetRange: **132.170.0.0 - 132.170.255.255**  
CIDR: 132.170.0.0/16  
NetName: **UCF**  
NetHandle: **NET-132-170-0-0-1**  
Parent: **NET-132-0-0-0-0**  
NetType: Direct Assignment  
NameServer: AQUARIUS.CC.UCF.EDU  
NameServer: NAMESRV2.CC.UCF.EDU  
Comment:  
RegDate: 1989-02-27

Internet | Protected Mode: Off

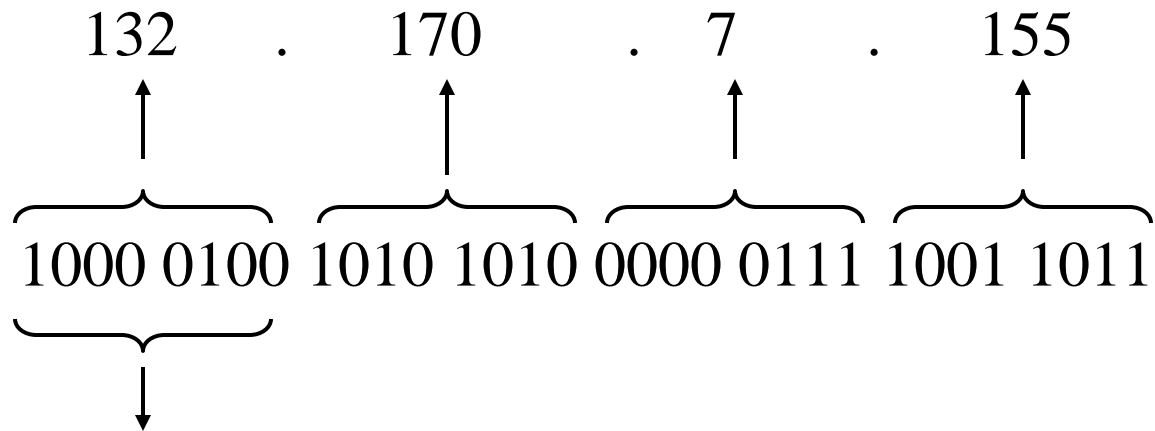
100%





# IP Example

- Using the IP address of my office computer as an example, the binary form of 132.170.7.155 would be:



$$\begin{aligned} 132 &= (1 \times 2^7) + (0 \times 2^6) + (0 \times 2^5) + (0 \times 2^4) + (0 \times 2^3) + \\ &\quad (0 \times 2^2) + (0 \times 2^1) + (0 \times 2^0) \\ &= 128 + 4 \end{aligned}$$



# Internet Protocol – IP

- Organizations are assigned blocks of IP addresses, which they in turn assign to their machines that need Internet access – this includes most computers today!
  - Small organizations may be assigned 256 IP addresses, such as 192.57.126.0 to 199.57.126.255. Very large organizations, such as the Department of Defense, may be assigned as many as 16 million IP addresses, which would include all IP addresses with one particular first 8-bit number, such as 12.0.0.0 to 12.255.255.255.
- In late 1998, a new IP standard, IPv6, was approved, although it has not yet been widely implemented (although it is implemented in Internet2).
  - The most significant change was to expand the address size from 32 bits to 128 bits. This change will soon become essential as the number of remaining unused IP addresses is diminishing rapidly.



# Internet Protocol – IP

- The function of IP software is to transfer data from one computer (the **source**) to another computer (the **destination**).
- When an application (software) on the source computer wants to send information to a destination, the application calls IP software on the source machine and provides it with data to be transferred along with an IP address for each of the source and destination computers.
- The IP software running on the source creates a **packet**, which is a sequence of bits representing the data to be transferred along with the source and destination IP addresses and some other header information, such as the length of the data.



# Internet Protocol – IP

- If the destination computer is on the same local network (LAN) as the source computer, then the IP software will send the packet to the destination directly via this network.
- If the destination computer is on another network, the IP software will send the packet to a **gateway**, which is a device that is connected to the source computer's network as well as to at least one other network.
- The gateway will select a computer on one of the other networks to which it is attached and send the packet on to that computer.



# Internet Protocol – IP

- This process will continue, with the packet going through perhaps a dozen or more **hops**, until the packet reaches the destination computer.
- Finally, the IP software on the destination computer will receive the packet and pass the data in the packet up to the application that is waiting for the data.
- The sequence of computers that a packet travels through from source to destination is known as its **route**.
  - A separate protocol (the current standard is BGP-4, the Border Gateway Protocol) is used to pass network connectivity information between gateways so that each can choose a good next hop for each packet it receives.



# Internet Protocol – IP

- IP software also adds some error detection information, called a **checksum**, to each packet it creates, so that if a packet is corrupted during transmission, this can usually be detected by the recipient.
- The IP standard calls for IP software to simply discard any corrupted packets. This means that IP-based communication is unreliable: packets can be lost.
- Obviously, this means that IP *alone* is not a particularly good form of communication for many Internet applications.



# Transmission Control Protocol - TCP

- TCP, Transmission Control Protocol, is a higher-level protocol that extends IP to provide additional functionality, including reliable communication based on the concept of a **connection**.
- A connection is established between TCP software running on two machines by one of the machines (let's call it machine A) sending a connection-request message (message 1) via IP to the other machine (call it B).
  - That is, the IP message contains a message conforming to the TCP protocol and representing a TCP connection request.
- If the connection is accepted by B, then B returns a message to A (message 2) requesting a connection in the other direction. If A responds in the affirmative (message 3), then the connection is established.



# Transmission Control Protocol - TCP

- Once the connection is established, A and B can both send messages to one another at the same time. This is known as **full-duplex communication**.
- When A and B are both done sending messages to one another, a similar set of three messages is used to close the connection.
- Once a connection has been established, the TCP provides reliable data transmission by demanding an **acknowledgement** for each packet it sends via IP. Essentially, the software sets a timer after sending each packet. The TCP software on the receiving side sends a packet containing an acknowledgement for every TCP-based packet it receives that passes the checksum test.





# Transmission Control Protocol - TCP

- If the TCP software sending the packet does not receive an acknowledgement before its timer expires, then it resends the packet and restarts the timer.
- Another important feature that TCP adds to IP is the concept of a **port**.
- The Internet Assigned Numbers Authority (IANA) sets port numbers for common applications in the range 0-1023. Typically only used by applications by the system at boot-up or by a user with administrative permission on the system. All other possible port numbers, from 1024 to 65535, can be used by the first application on a system that requests the port.

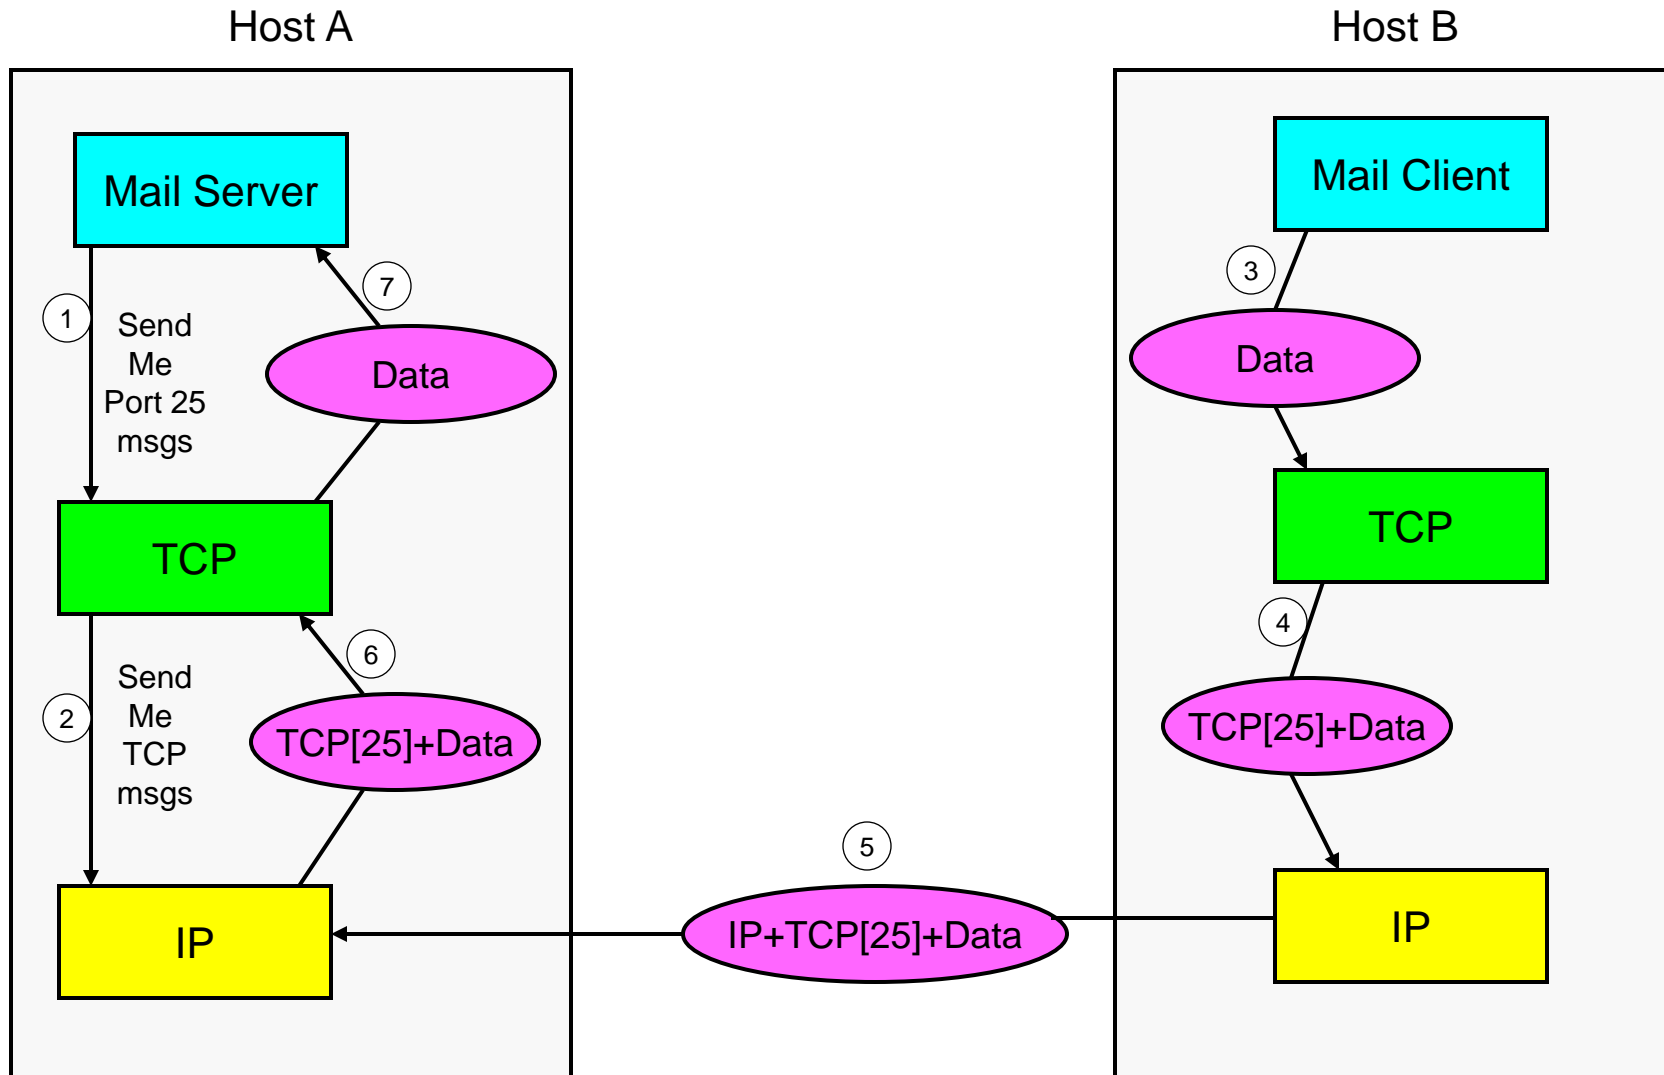


# Transmission Control Protocol - TCP

- The port concept allows TCP to be used to communicate with many different applications on a single machine.
  - For example, a machine connected to the Internet may run a mail server for users on its local network, a file download server, and also a server that allows users to log in to the machine and execute commands from remote locations.
- Consider the example shown on the next page.



# Simplified View of Communication Using TCP/IP



# TCP/IP Example Explanation

- The illustration on the previous page presents a simplified view (without the connections and acknowledgements) of how TCP/IP works.
1. The mail server (application) running on Host A, makes a call to the TCP software on its system to request that any incoming TCP connection requests that specify port 25 as part of the TCP/IP message be sent to the mail application. (Typically, mail servers running SMTP ask TCP to listen for requests for port 25).
  2. TCP software requests all IP messages conforming to TCP protocol. IP just “forwards” them to TCP.
  3. A mail client running on Host B sends an email message to the server on Host A – specifying port 25.



# TCP/IP Example Explanation

4. The TCP software on Host B creates the packet with the email message appending port number 25 to the packet and forwards the packet to the IP software on Host B.
5. The IP software on Host B transmits the packet to Host A.
6. IP software on Host A, receives the packet, discovers it is a TCP protocol and forwards it to the TCP software on Host A.
7. TCP software on Host A sends the email data to the mail server on Host A completing the connection.

