INTERNETWORKING CONCEPTS

Why Internetworking!!

- Each network technology is designed to fit a specific set of constraints
- No single networking technology best for all needs
  - LANs
    - Low cost
    - Provide high speed communication
    - Cover a limited distance
  - WANs
    - High cost
    - Cover an unlimited distance
- If an organization chooses the network type that is best for each task, the organization will have several types of networks

- Keeping all these networks unconnected to each others will leads to:
  - A huge loss of resources, due to hardware duplications; an employee may be given access to multiple monitors and keyboards to access multiple networks
  - Losing data integrity, due to data duplications
  - Less users satisfaction/productivity, for example, if a user wants to send a message to another user working in another network, he/she is forced to move from one computer to another to send the message across the appropriate network

Internetworking

- Internetworking is a scheme which uses both hardware and software to provide universal connection (service) among networks, even heterogeneous networks
- The hardware is used to interconnect a set of physical networks
  - the basic hardware component used is called a router, a.k.a. a gateway
    - A special-purpose computer dedicated to the task of interconnecting networks
    - Has a conventional processor and memory as well as a separate I/O interface for each network to which it connects
- The software (a.k.a. a set of protocols) is used to make the resulting system appears homogeneous
- the resulting system is known as *internetwork* or *internet*

Internetworkings

- When two, or more, networks are connected together, they become an *internetwork*, or simply *internet*

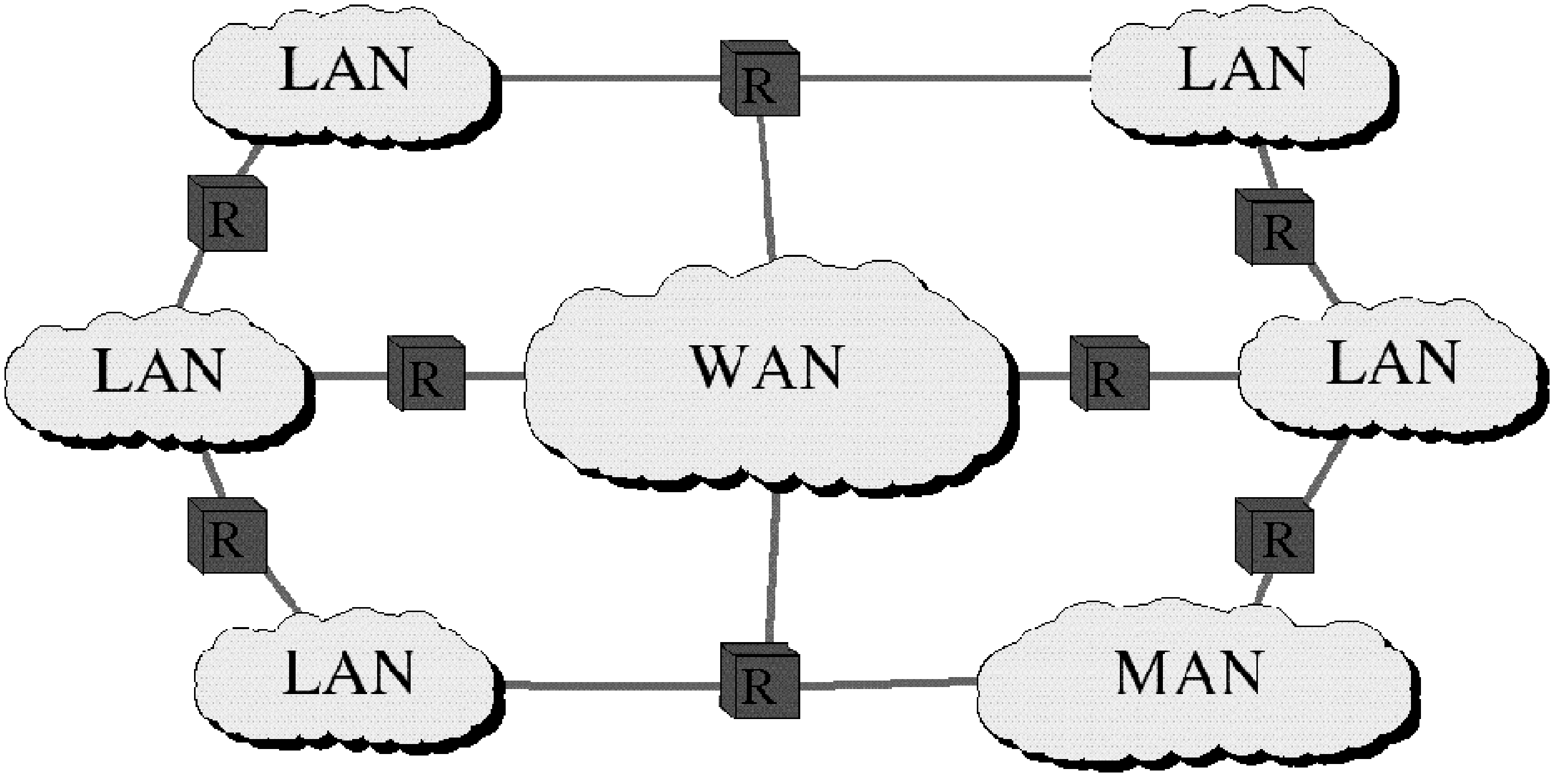

- An organization seldom uses a single router to connect all of its network because:
  - The CPU power and memory in one router is insufficient to handle the traffic passing among an arbitrary number of networks
  - Redundancy improves internet reliability; when a router fails, an alternative paths may be used
- When planning an internet, an organization must choose a design that meets the organization's need for
  - Reliability (multiple routers)
  - Capacity (expected traffic/physical network/bandwidth)
  - Cost

Virtual Network

- Internet software provides the appearance of a single communication system to which many computers are attached
- The system offers universal service: each computer is assigned an address and any computer can send a packet to any other computer
- Internet protocol software hides the details of
  - physical network connections
  - physical addresses
  - routing information
- Although a combination of hardware and software provides the illusion of a uniform network system, no such network exist, hence the name *virtual network*

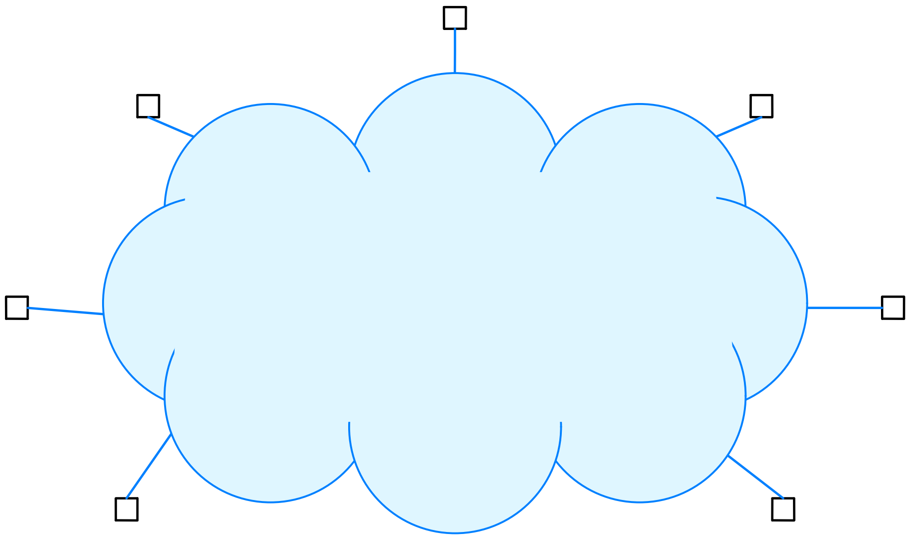

(a)

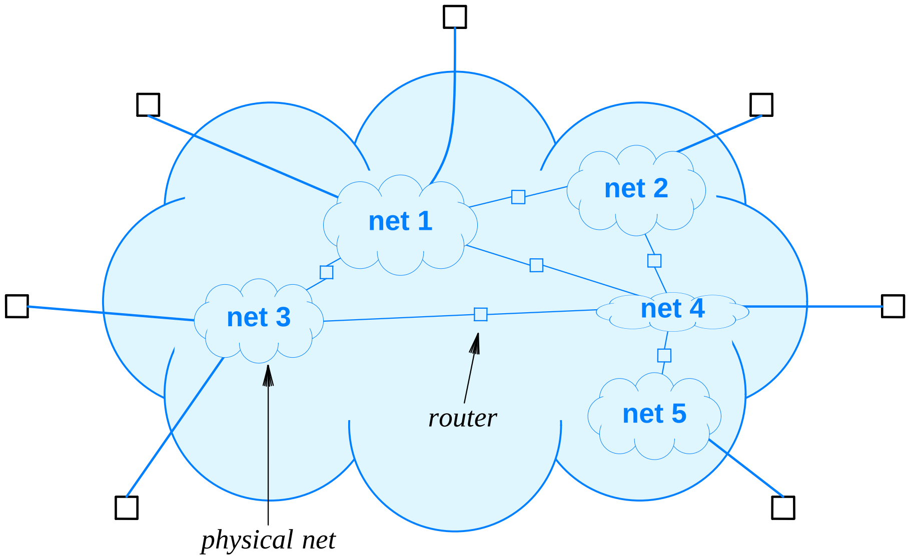

(b)

The internet concept

INTERNET PROTOCOL ADDRESSES

Addresses For the Virtual Internet

- The goal of internetworking is to provide a seamless communication system
- To achieve this goal, the internet protocol software must hide the details of physical networks and offer the facilities of a large virtual network
- The virtual internet operates much like any physical network
- The difference between an internet and a physical network is that the internet is
  - Simply an abstraction imagined by its designers
  - Created entirely by software
- The designers, independent of the details of the physical hardware, are free to choose:
  - Addresses
  - Packet formats

Uniform Addressing Scheme

- Each network technology may define its own address format
- Thus, the addresses used by two technologies may be incompatible because
  - They are different in size
  - They have different formats
- To give the appearance of a single, uniform system,
  - All host computers must use a uniform addressing scheme
  - Each address must be unique
- To guarantee uniform addressing for all hosts, protocol software defines an addressing scheme that is independent of the underlying physical addresses
- Uniform addressing Scheme
  - Helps creating the illusion of a large network
  - Hides the details of underlying physical network addresses;
    two application programs can communicate without knowing either hardware address

TCP/IP Addressing Scheme

- In the TCP/IP protocol stack, addressing is specified by the internet protocol
- Each host is assigned a unique 32-bit number known as
  - The host's internet protocol address,
  - IP address, or
  - Internet address
- To transmit information across a TCP/IP internet, a computer must know the IP address of the remote computer to which the information is being sent
- IP addresses are only understood by software

The IP Address Hierarchy

- Conceptually, each 32-bit IP address is divided into two parts
  - A prefix
    - A network number
    - Identifies a network
    - Global authority assigns unique prefix to each network
  - A suffix
    - An address within a network
    - Identifies a connection between a computer and a network
    - Local administrator assigns unique suffix to each host
- Note that, a computer with multiple network connections, e.g., router, must be assigned one IP address for each connection
- This two-level hierarchy is designed to make routing efficient
- *How many bits should be placed in each part?*

Classes of IP Addresses

- A compromise to accommodate a combination of large and small networks
- The scheme divides the IP address space into:
  - Three primary classes, class A, B, and C, where each class has different prefix/suffix sizes
  - One multicast class, class D
  - One reserved class, class E, for future use
- Initial bits determine class

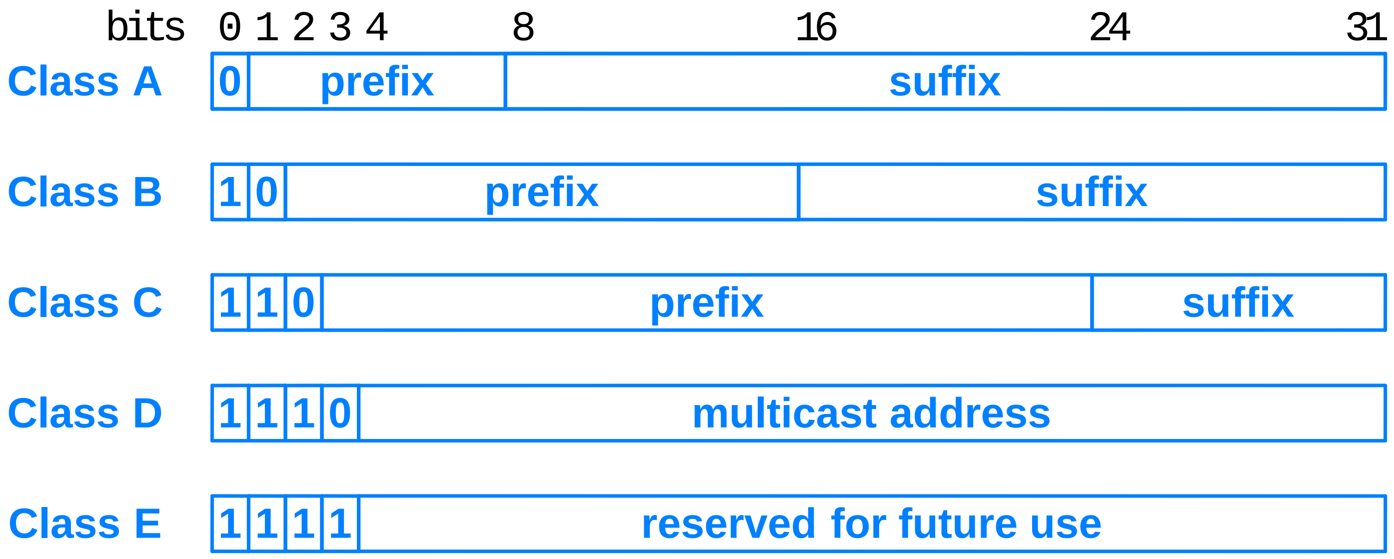

- Note that
  - The IP addressing scheme does not divide the 32-bit address space into equal size classes
  - The classes do not contain the same number of networks

Address class	Bits in Prefix	Maximum Number of Network	Bits in Suffix	Maximum number of Hosts Per Network
A	8: *the MSB is fixed*	128	24	16777216
B	16: *the two MSBs are fixed*	16384	16	65536
C	24: *the three MSBs are fixed*	2097152	8	256

MSB means *Most Significant Bit*

Note that, there is a different between this table and the table in Figure 18.5 at page 287 in the text book.

Please follow the table in the lecture note, i.e., the table above.

Dotted Decimal Notation

- Although IP addresses are 32-bit numbers, users seldom enter or read the values in binary, instead, the dotted decimal notation is used
- Dotted decimal notation
  - Is a shorthand for IP address
  - Allows human to avoid binary
  - Treats each byte as an unsigned integer
  - Represents these 4 integer numbers in decimal and uses a dot to separate them
- In a class A address, the last three bytes correspond to a host suffix
- In a class B address, the last two bytes correspond to a host suffix
- In a class C address, the last byte corresponds to a host suffix
- Not the same as names like www.somewhere.com

32-bit binary number	Equivalent dotted decimal
10000001 00110100 00000110 00000000	129.52.6.0
11000000 00000101 00110000 00000011	192.5.48.3
00001010 00000010 00000000 00100101	10.2.0.37
10000000 00001010 00000010 00000011	128.10.2.3
10000000 10000000 11111111 00000000	128.128.255.0

- Because dotted decimal notation does not make individual bits of an address visible, the class must be recognized from the decimal value of the first byte

Class	Range of Values
A	0 through 127
B	128 through 191
C	192 through 223
D	224 through 239
E	240 through 255

An Addressing Example

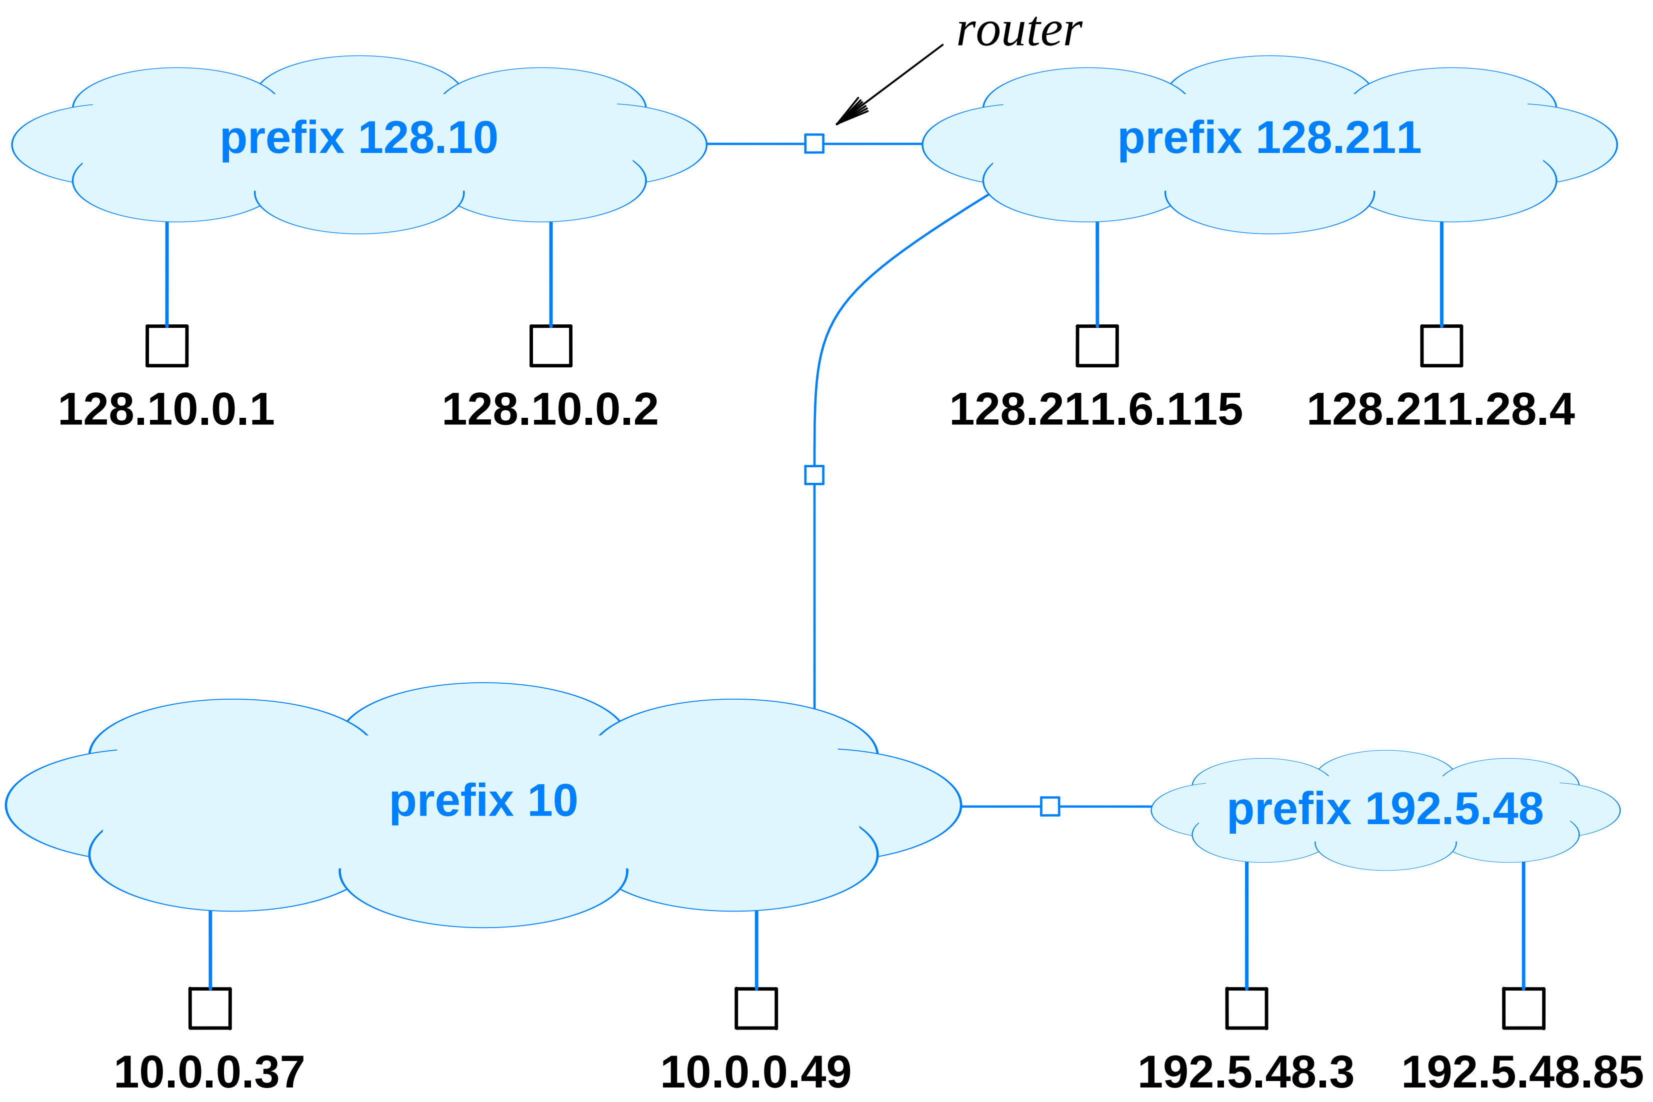

A Router Addressing Example

- A router has connections to multiple physical networks
- A one address is needed for each router connection
- IP does not require that the same suffix be assigned to all interfaces, but as a practical matter, using the same suffix may help people who manage the internet to remember it

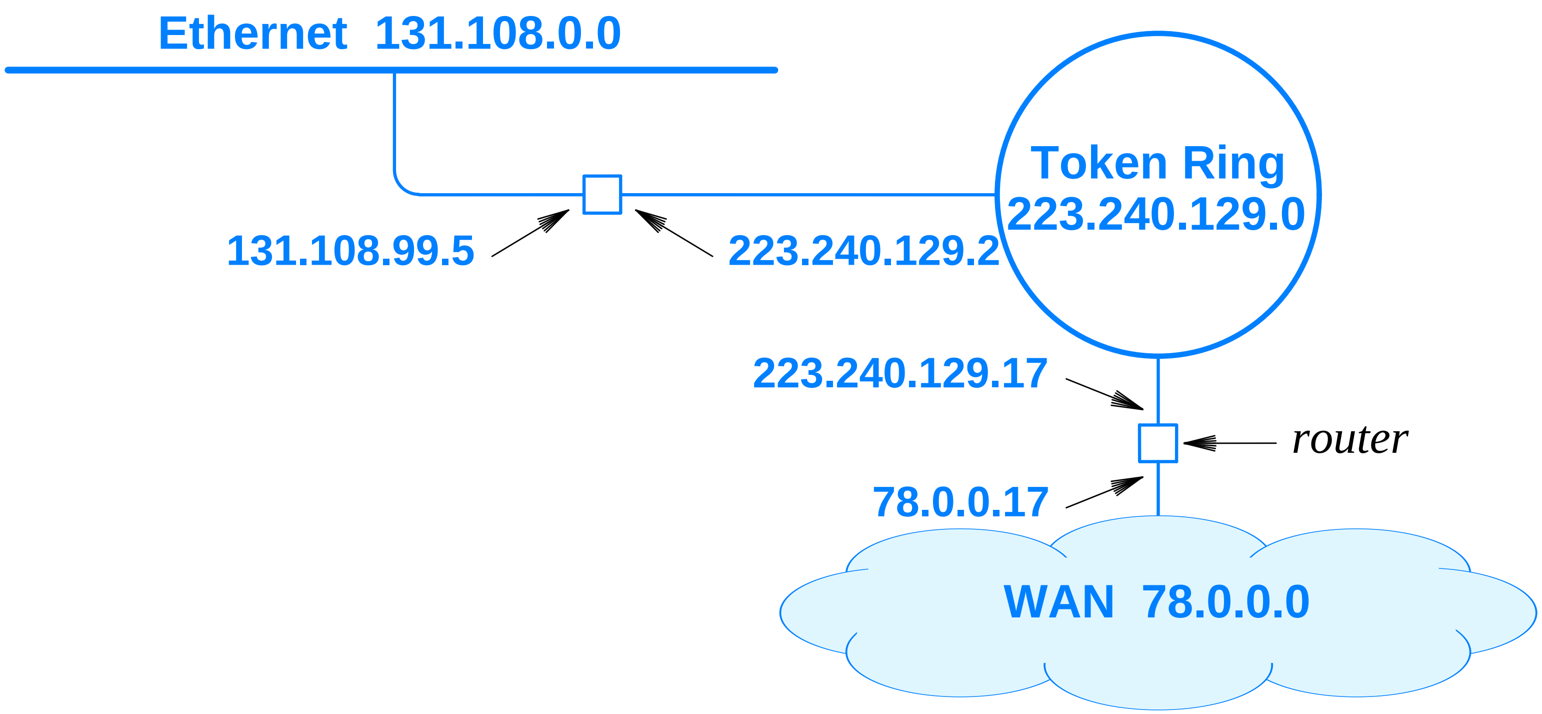

Multi-Homed Hosts

- A host which has multiple network connections is called *multi-homed*
- Multi-homing is sometime used to increase the reliability and improve the performance
- Like a router, a multi-homed host has multiple IP addresses, one for each network connection

Special IP Addresses

- Special IP Addresses never assigned to a host
- Network address
  - It is convenient to have an address that can be used to denote the prefix assigned to a given network
  - The network address is formed by adding a suffix that consists of all 0s bits to the network prefix
  - For example, 128.211.0.0 is not a host IP address, but it is a network address with a network prefix equals 128.211
- Directed broadcast address
  - To make broadcast easy, IP defines a directed broadcast address for each physical network
  - When a packet is sent to a network's directed broadcast address, a single copy of the packet travels across the internet until it reaches the specified network, the packet is then delivered to all hosts on the network

- When a directed broadcast is sent to a network which does not have hardware support for broadcast, software must send a separate copy of the packet to each host on the network
- The directed broadcast address for a network is formed by adding a suffix that consists of all 1s bits to the network prefix
- For example, 128.211.255.255 is not a host IP address, but it is the directed broadcast address for a network with a prefix equals 128.211

• Limited broadcast address

- Refers to a broadcast on a local physical network
- Usually is used during system startup by a computer which does not yet know the network number
- IP reserves the 255.255.255.255 address (i.e., all 1s bits) to refer to the limited broadcast address

- Loopback address
  - To test network applications
    - A programmer may have two application programs that are intended to communicate across the network
    - Instead of executing each program on a separate computer, the programmer runs both programs on a single computer and instructs them to use a loopback IP address when communicating
    - In this case, data travels down the protocol stack and then forwards back, through the same protocol stack, to the second program
  - During loopback testing, no packets ever leave a computer, i.e., a loopback address never appears in a packet traveling across a network
  - IP reserves the class A network prefix 127, with any suffix, for use with loopback
  - Although any suffix can be used, by convention, programmers often used host number 1, i.e, 127.0.0.1, as a loopback address

- This computer address
  - A computer needs to know its IP address to send/receive internet packets
  - When a computer boots, there are some TCP/IP protocols the computer can use them to obtain its IP address automatically
  - However, when using such startup protocols, a just booted computer can not supply a correct IP source address
  - IP handle such case by reserving the 0.0.0.0 address (i.e., all 0s bits) to mean *this computer*
  - The 0.0.0.0 address must not be used after a host completes the startup procedure and has obtained an IP address

Prefix	Suffix	Type of Address	Purpose
all 0s	all 0s	this computer	used during bootstrap
network	all 0s	network	identifies a network
network	all 1s	directed broadcast	broadcast on specified net
all 1s	all 1s	limited broadcast	broadcast on local net
127	any	loopback	testing

MORE ON INTERNET PROTOCOL ADDRESSES (SUBNETTING)

Subnets

- All hosts in a network must have the same network number
- This property of IP addressing can cause problems as networks grow, e.g.,
  - Consider a company that starts out with one class C LAN
  - As time goes on, it may acquire more than 254 machines; hence, it needs:
    - A second class C address
    - A new IP address for the second network
    - A new router for the second network
  - As the number of distinct local networks grows, managing them can become a serious headache, for example,
    - Moving a machine from one LAN to another requires
      - Changing its IP address
      - Modifying its configuration files
      - Announcing the new IP address to the world
    - If some other machine is given the newly-released IP address, that machine will get emails and other data intended for the original machine until the address has propagated all over the world

- These problems can be solved by starting with a larger network and at the same time allowing this network to be split into several parts for internal use, but still acting like a one single network to the outside world
- In the internet literature, these parts are called *subnets*
- The word *subnet* is also means the set of all routers and communication lines in a network
- It will be clear from the context which meaning is intended

- If a growing company started up with a class B address, for example, 128.211, instead of a class C address
  - It could decide to split the 16-bit host number into a 6-bit subnet number and a 10-bit host number
  - This split allows 62 networks *Why?*, each with up to 1022 hosts *Why?*
    - Network 128.211.4.0, i.e., <u>10000000 11010011 000001</u>00 00000000:
      Hosts within the subnetwork:
      from 128.211.4.1, i.e., <u>10000000 11010011 000001</u>00 00000001
      to 128.211.7.254, i.e., <u>10000000 11010011 000001</u>11 11111110

    - Network 128.211.8.0, i.e., <u>10000000 11010011 000010</u>00 00000000:
      Hosts within the subnetwork:
      from 128.211.8.1, i.e., <u>10000000 11010011 000010</u>00 00000001
      to 128.211.11.254, i.e., <u>10000000 11010011 000010</u>11 11111110

    - Network 128.211.12.0, i.e., <u>10000000 11010011 000011</u>00 00000000:
      Hosts within the subnetwork:
      from 128.211.12.1, i.e., <u>10000000 11010011 000011</u>00 00000001
      to 128.211.15.254, i.e., <u>10000000 11010011 000011</u>11 11111110

* $\cdots\cdots\cdots$
* Network 128.211.244.0, i.e., 10000000 11010011 11110100 00000000:
  Hosts within the subnetwork:
  from 128.211.244.1, i.e., 10000000 11010011 11110100 00000001
  to 128.211.247.254, i.e., 10000000 11010011 11110111 11111110
* Network 128.211.248.0, i.e., 10000000 11010011 11111000 00000000:
  Hosts within the subnetwork:
  from 128.211.248.1, i.e., 10000000 11010011 11111000 00000001
  to 128.211.251.254, i.e., 10000000 11010011 11111011 11111110

– Note that the following are special addresses
  * 128.211.0.0, 128.211.4.0, 128.211.8.0, 128.211.12.0, $\cdots$, 128.211.244.0, and 128.211.248.0 are network and subnetwork addresses
  * 128.211.255.255, 128.211.7.255, 128.211.11.255, 128.211.15.255, $\cdots$, 128.211.247.255, and 128.211.251.255 are directed broadcast addresses

- Outside the network, the subnetting is not visible, so allocating a new subnet does not require applying for a new IP address, or updating any external databases
- At routers, IP packets are processed as follow:
  - Each router has a table contains: *Destination*, *Mask*, and *Next_hop* entries
  - When an IP packet arrives, its destination address is masked using *Mask* to get rid of the host number and the resulting address is looked up to the *Destination*
    - If the packet is for a remote network, it is forwarded to the *next_hop* router
    - If it is a local host, it is sent directly to this host
    - If the network number is not exist, the packet is forwarded to the default *next_hop* router which has a more extensive table
- When subnetting is introduced, only the local routing tables are changed by adding entries as: *Local_subnet*, *Mask*, and *Next_hop*
- In other words, each router knows, in addition, how to go to
  - All other subnets in its network
- It does not have to know the details about hosts on other subnets
- Subnetting reduces router tables space, by creating a three-level hierarchy

Subnet Mask

- For a class B network address, the subnet mask can be

255.255.192.0	255.255.224.0	255.255.240.0	255.255.248.0
255.255.252.0	255.255.254.0	255.255.255.0	255.255.255.128
255.255.255.192	255.255.255.224	255.255.255.240	255.255.255.248
255.255.255.252			

- *Why are these subnet masks, 255.255.128.0, 255.255.255.254, and 255.255.255.255, not used?*
- Many network administrators prefer byte-oriented masking, i.e., 255.255.255.0, because it is easy to read and understand when addresses are written in dotted decimal notation, i.e.,
  - The first two bytes define the original network
  - The third byte defines the subnet address
  - The fourth byte defines the host on the subnet
- However, byte-oriented masking does not take advantage of the bit-oriented power in subnetting

Subnets and Organizations

- Organizations usually decide to do subnet in order to overcome
  - Topological problems
    - Distance limitation
    - Hardware differences
  - Organizational problems
    - With the standard addressing scheme, a central administrator is responsible for managing host addresses for the entire network
    - By subnetting, the local administrator can delegate address assignment his/her own department within the overall organization, without returning to the central administrator
      - Technical reasons
      - Political reasons

Subnetting and Class C addresses

- Same subnetting concept can be applied to class C
- For example, the mask 255.255.255.192 divides a class C address into 4 subnets of 64 host addresses, however
  - Host-number 0 and host-number 63 will not be used (special address)
  - Subnet-number 0 and subnet-number 3 will not be used (special address)
- Therefore, the address space of this class C network number is reduced from 254 host to 124 hosts

IP IS RUNNING OUT OF ADDRESSES!!

IP Is Running Out of Addresses

- IP has been in heavily use for over a decade
- Unfortunately, IP is rapidly becoming a victim of its own popularity
- In 1987, a few visionaries predicted that some day the internet might grow to 100,000 networks
- But, most experts saw this as being decades in the future, if ever
- In 1996, the 100,000th network was connected
- In the January 2002 *Internet Domain Survey* (done by the *Internet Software Consortium*), there were 147,344,723 hosts counted, i.e., approximately 150 millions hosts
- In principle, over 4 billion addresses are exist, i.e., 2^{32}
- The practice of organizing the address space by classes wastes millions of them

- For most organizations,
  - A class A network, with $16M - 2$ addresses is too much
  - A class C network, with $256 - 2$ addresses is too little
  - A class B network, with $65,536 - 2$ addresses is just right
- In reality, a class B address is far too much for most organizations
- Studies have shown that more than half of all class B networks have fewer than 50 hosts, i.e., a class C network would have done the job
- However, every organization that asked for a class B address thought that one day, it would outgrow the 8-bit host field

What If

- What if class C IP addresses use 10 bits, instead of 8 bits, to describe host number?
  - It will allow 1022 hosts per network
  - Most organizations would probably settle for a class C network
  - We would have about half a million of such class C network addresses
  - Currently, there are about 2 millions class C networks, but most of them are unused

The Size of Routing Tables Is Another Problem

- Most of the existing routers software/firmware were designed at a time when the Internet had 1000 connected networks and 10,000 networks seemed decades away
- Having more and more network addresses might let routing tables explode;
  - Memory problem
  - Complexity problem

IP by Country

- The routing table problem could have been solved by going to a deeper hierarchy
- For example, having each IP address to contain a country, a state/province, a city, a network, and a host fields might work
- Then each router would only need to know how to get to
  - Each country
  - States/provinces in a country
  - Cities in a state/province
  - Networks in a city
- Unfortunately, this solution would
  - Require considerably more than 32 bits to represent an IP address
  - Use addresses inefficiently, a very small country would have as many bits as the united States

Classless Inter-Domain Routing (CIDR)

- The main purpose of CIDR is to allocate the remaining class C network addresses in consecutive variable-sized blocks with
  - A single entry indicates the first class C address
  - A subnet mask indicates the range of addresses in the CIDR block
- With this CIDR addressing scheme, a range of class C addresses that would otherwise each require a separate routing table entry can be represented by a single "CIDRized" entry
- In other words, CIDR addressing compresses routing table entries

CIDR Examples

- If a site needs, for example, 2000 address
  - It is given a block of 2048 addresses, e.g., 194.24.0.0 through 194.24.7.255, i.e., 8 contiguous class C networks

11000010	00011000	00000000	00000000
11000010	00011000	00000001	00000000
11000010	00011000	00000010	00000000
11000010	00011000	00000011	00000000
11000010	00011000	00000100	00000000
11000010	00011000	00000101	00000000
11000010	00011000	00000110	00000000
11000010	00011000	00000111	00000000

  - Using these 8 contiguous class C networks, is a way more efficient than a full class B address
  - The mask for this CIDRized entry is 255.255.248.0, i.e,

    11111111 11111111 11111000 00000000
  - All these 2048 addresses have same first 21 bits

- If another site asks for 4096 addresses
  - Since a block of 4096 addresses must lie on a 4096 boundary, they cannot be given address starting at 194.24.8.0
  - Instead, they get 194.24.16.0 through 194.24.31.255,i.e.,

16 contiguous class C networks

11000010	00011000	00010000	00000000
11000010	00011000	00010001	00000000
11000010	00011000	00010010	00000000
11000010	00011000	00010011	00000000
11000010	00011000	00010100	00000000
11000010	00011000	00010101	00000000
11000010	00011000	00010110	00000000
11000010	00011000	00010111	00000000
11000010	00011000	00011000	00000000
11000010	00011000	00011001	00000000
11000010	00011000	00011010	00000000
11000010	00011000	00011011	00000000
11000010	00011000	00011100	00000000
11000010	00011000	00011101	00000000
11000010	00011000	00011110	00000000
11000010	00011000	00011111	00000000

  - The mask for this CIDRized entry is 255.255.240.0, i.e.,

    11111111 11111111 11110000 00000000

  - All these 4096 addresses have same first 20 bits

- If another site asks for 1000 addresses
  - It is given a block of 1024 addresses, e.g., 194.24.8.0 through 194.24.11.255, i.e., 4 contiguous class C networks

    11000010 00011000 00001000 00000000
    11000010 00011000 00001001 00000000
    11000010 00011000 00001010 00000000
    11000010 00011000 00001011 00000000
  - The mask for this CIDRized entry is 255.255.252.0

    11111111 11111111 11111100 00000000
  - All these 1024 addresses have same first 22 bits
- If another site asks for 1000 addresses
  - It is given a block of 1024 addresses, e.g., 194.24.12.0 through 194.24.15.255, i.e., 4 contiguous class C networks

    11000010 00011000 00001100 00000000
    11000010 00011000 00001101 00000000
    11000010 00011000 00001110 00000000
    11000010 00011000 00001111 00000000
  - The mask for this CIDRized entry is 255.255.252.0

    11111111 11111111 11111100 00000000
  - All these 1024 addresses have same first 22 bits

- To summarize the example:

Network Address	Network Mask
194.24.0.0	255.255.248.0
194.24.16.0	255.255.240.0
194.24.8.0	255.255.252.0
194.24.12.0	255.255.252.0

Network Address				**Network Mask**			
11000010	00011000	00000000	00000000	11111111	11111111	11111000	00000000
11000010	00011000	00010000	00000000	11111111	11111111	11110000	00000000
11000010	00011000	00001000	00000000	11111111	11111111	11111100	00000000
11000010	00011000	00001100	00000000	11111111	11111111	11111100	00000000

- Now, consider what happens when a packet comes which addressed to
  - 194.24.17.4, i.e., 11000010 00011000 00010001 00000100
  - 194.24.15.4, i.e., 11000010 00011000 00001111 00000100

Geographical IP Addresses

- In addition of using CIDR to allocate contiguous class C networks as units,
  - The world was partitioned into four zones
  - Each zone is given a portion of the class C address space
- The allocation was as follows:
  - Address 194.0.0.0 to 195.255.255.255 are for Europe, i.e.,
    $\underline{1100001}0$ 00000000 00000000 00000000 to
    $\underline{1100001}1$ 11111111 11111111 11111111
  - Address 198.0.0.0 to 199.255.255.255 are for North America
    $\underline{1100011}0$ 00000000 00000000 00000000 to
    $\underline{1100011}1$ 11111111 11111111 11111111
  - Address 200.0.0.0 to 201.255.255.255 are for Central and South America
    $\underline{1100100}0$ 00000000 00000000 00000000 to
    $\underline{1100100}1$ 11111111 11111111 11111111
  - Address 202.0.0.0 to 203.255.255.255 are for Asia and Pacific
    $\underline{1100101}0$ 00000000 00000000 00000000 to
    $\underline{1100101}1$ 11111111 11111111 11111111

- In this way, each region was given about 32M, i.e., 2^{25}, addresses to allocate, i.e., 128K, i.e., 2^{17}, class C network address
- Another 320M class C addresses,
  from 204.0.0.0 through 223.255.255.255 are held in reserve for the future, i.e.,
  11001100 00000000 00000000 00000000 through
  11011111 11111111 11111111 11111111

Some Comments About CIDR

- Network Masks can be represented by only the number of ones, i.e., in the previous example, it should be 194.24.0.0/21, 194.24.16.0/20, 194.24.8.0/22, and 194.24.12.0/22, i.e., less than a byte to represent a mask, instead of 4 bytes
- Indexing tricks are used to speed up the search;
  In other words, the router entries are not tried sequentially
- The same idea can be applied to all addresses, not just the new class C addresses
- With CIDR, the old class A, B, and C networks are no longer used for routing; that is why CIDR is called *Classless Inter-Domain Routing*
- CIDR provides address and routing relief for many more years to come
- The long-term solution is to replace the current IPv4 with a new version, which is IPv6

