SE 4C03 Winter 2006

03 Internet Addressing

Instructor: W. M. Farmer

Revised: 10 January 2006

IP Addresses

- There are two Internet naming systems:
  1. The primary system is the **internet address system** which uses binary **IP addresses**
  2. The secondary system is the **domain name system (DNS)** which uses natural language **DNS names**

- IP addresses are 32-bit integers
  - Composed of four 8-bit octets
  - Represented as four integers, usually in base 2 or base 10, separated by dots

    base 2: 11000111.00010001.00101000.11010010

    base 10: 199.17.40.210

IP Address Assignment

- IP addresses are assigned to **network interfaces**, **not hosts**
  - A host with one network interface is assigned an IP address by default
- A network interface is normally assigned a unique IP address
  - In practice, an interface may be assigned more than one address
  - In some rare cases, an interface may be assigned no address at all
  - In certain cases, different interfaces may have the same IP address

Class Networks

- IP addresses are organized into **class networks** to facilitate address assignment and packet routing
  - Class A: 0nnnnnnn.iiiiiiii.iiiiiiii.iiiiiiii
  - Class B: 10nnnnnn.nnnnnnnn.iiiiiiii.iiiiiiii
  - Class C: 110nnnnn.nnnnnnnn.nnnnnnnn.iiiiiiii
  - Class D: 1110bbbb.bbbbbbbb.bbbbbbbb.bbbbbbbb
- A class network is really a set of IP addresses and not a network
- The Internet Corporation for Assigned Names and Numbers (ICANN) is responsible for assigning class A, B, and C networks to organizations
- Each address in a class A, B, or C network is a pair (N, I) where N is its **network identification** and I is its a **interface (or host) identification**. n and i denote bits in the network and interface identifications, respectively.

Class Networks (cont.)

- By convention, the address in a class A, B, or C network whose interface bits are all 0 (e.g., 199.17.40.0) is the **network address** for the class

- By convention, the address in a class A, B, or C network whose interface bits are all 1 (e.g. 199.17.40.255) is the **(direct) broadcast address** for the class

- The **limited broadcast address** is 255.255.255.255

- The addresses in a class D network are for multicasting

- Some addresses in the above classes and the addresses of the form 1111bbbb.bbbbbbbb.bbbbbbbb.bbbbbbbb are reserved

The Loopback

- Each host running TCP/IP has a virtual interface called the **loopback interface** which is the only interface on a virtual network called the **loopback network**

- The network and interface addresses of the loopback are 127.0.0.0 and 127.0.0.1, respectively

Weaknesses of IP Address System

- Some hosts (e.g., multi-homed host) have more than one IP address
- The class networks are too rigid
- There are not enough IP addresses for future expansion

Subnets

- IP addresses are also organized into **subnets** to facilitate address assignment, network organization, and routing

- Each subnet is a set of addresses determined by:
  1. A **subnet address** (e.g., 199.17.35.96)
  2. A **subnet mask** (e.g., 255.255.255.240)

- Each address in a subnet is pair (S, I) where S is its **subnet identification** and I is its **interface (or host) identification**

- Special cases:
  - Set of all IP addresses
  - Class A, B, and C networks
  - Individual interface (or host) IP address

Subnet Conventions

- Usually, but not necessarily, the subnet identification of a class A, B, or C address is an extension of the network identification of the address

- Usually, but not necessarily, the subnet mask consists of a block of 1s followed by a block of 0s

- By convention, there is one subnet corresponding to each SPN
  - Each interface on the SPN is assigned the same subnet address and subnet mask

Address Resolution Problem

- High-level IP addresses are used for communication across an internet and are assigned independently of physical hardware addresses

- Low-level physical addresses are needed for physically delivering a packet to an interface on a network

- How are IP addresses mapped to physical addresses?
  - A solution is a function f that maps each IP address i to a physical address $f(i)$
  - The function must be changed as the internet changes
  - The function must be represented efficiently

Address Resolution Solutions

1. Physical addresses are encoded in IP addresses
   - Possible for proNET networks
   - Not viable for Ethernet

2. Each machine contains a table that represents the local part of an address resolution function
   - Awkward for Ethernet because physical addresses change when a host or interface is replaced

3. IP addresses are bound to physical addresses dynamically

Address Resolution Protocol (ARP)

- Used for dynamically binding an IP address to a physical address (especially on Ethernet networks)

- ARP process:

  1. A host h_A broadcasts a request for the physical address which resolves an IP address i
  2. The host h_B with the network interface having the address i sends a reply to h_A containing the physical address of the interface

- The results of ARP queries are kept in a cache on each host

- When a sender requests a physical address, it can include its physical address in the message

Complications

- Several packets may simultaneously need to know the same physical address
- The host of the requested physical address may be down
- The cache may contain out-of-date bindings
  - When a host boots it can send a broadcast message informing the other computers on the network of its physical address
- Hosts may provide bogus address bindings
- At boot-time a diskless host knows its physical address but not its IP address
  - The host must get its IP address from a server on another computer

Reverse Address Resolution Protocol (RARP)

- Used for obtaining the IP address that is bound to a physical address

- RARP process:

  1. A host h_A broadcasts a request for the IP address which reversely resolves a physical address p
  2. The RARP servers which receive the request send replies back to the h_A containing the requested IP address

- Some scheme is needed to keep all the RARP servers from sending replies at the same time (and causing collisions on an Ethernet network)

ARP and RARP Messages

- ARP and RARP messages are **encapsulated** in a physical frame
  - ARP and RARP share the same message format
  - Type field says the data is an ARP or RARP message
  - The message itself is held in the data portion of the frame

- Each message has the following address fields:
  - Sender IP address
  - Sender physical address
  - Target IP address
  - Target physical address