CS 640 Introduction to Computer Networks

Lecture 8

Today's lecture

- IP
  - Addressing and forwarding
  - ARP
  - DHCP

The Internet Protocol

- IP implements best effort end to end datagram delivery service
- All computers in the Internet use IP (version 4)
- Store and forward handling of packets
- Forwarding: routers decide which way to send a packet based on its destination IP address
  - Uses local database of networks called forwarding table
  - Forwarding tables configured statically or built dynamically by routing protocols

Internet Structure

Recent Past

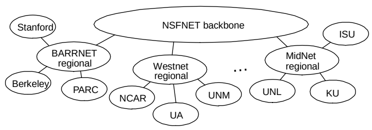

Internet Structure

Today

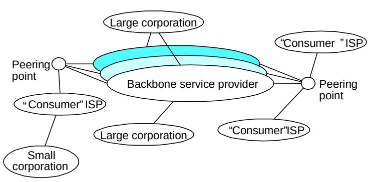

Forwarding Tables

- Suppose there are *n* possible destinations, how many bits are needed to represent addresses in a forwarding table?
  - $\log_2 n$
- So, we need to store and search $n * \log_2 n$ bits in forwarding tables?
  - We're smarter than that!

Addressing

- IP Address: 4byte-string that identifies a node
  - usually unique (some exceptions)
  - dotted decimal notation: 128.92.54.32
- Types of addresses
  - unicast: node-specific
  - broadcast: all nodes on the network
  - multicast: some subset of nodes on the network

Global Addresses

- Properties
  - globally unique
  - hierarchical: network + host
- Dotted Decimal Notation
  - 12.3.218.4
  - 138.96.33.81
  - 195.12.69.77
- Address classes
  - A, B, C (shown)
- Network respresented as Network Part / Num. Bits
  - E.g. 120.0.0.0/8 or 128.96.0.0/16

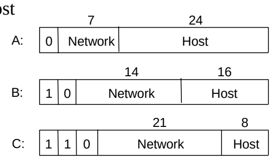

Other Addresses

- Private address (RFC 1918):
  - 10.0.0.0 to 10.255.255.255 (10.0.0.0/8)
  - 172.16.0.0 to 172.16.255.255 (172.16.0.0/12)
  - 192.168.0.0 to 192.168.255.255 (192.168.0.0/16)
- Class D: multicast addresses: 224.0.0.0 to 224.255.255.255

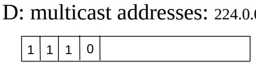

- Host part all 1's: broadcast in local network
- Host part all 0's: unspecified (not allowed)

Forwarding

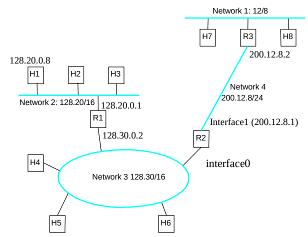

CS 640

Datagram Forwarding

- Strategy
  - every datagram contains destination's address
  - if directly connected to dest. network, forward to host
  - if not directly connected to destination network, then forward to some router
  - forwarding table maps network number to next hop
  - each router has forwarding table
  - each host has a default router
- Example
  for router R2
  in previous figure

Network	Next Hop
1	R3
2	R1
3	interface 0
4	interface 1
default	R3

CS 640

Subnetting and Supernetting

- Fixed network sizes are wasteful
  - What happens if a site asks for 300 IP addresses?
  - Subnetting
- Too many entries at a router can be combined
  - Keep routing tables small
  - Supernetting
- Classless Inter-Domain Routing (CIDR)

CS 640

Subnetting

- Add another level to address/routing hierarchy: *subnet*
- *Subnet masks* define variable partition of host part
- Subnets visible only within site

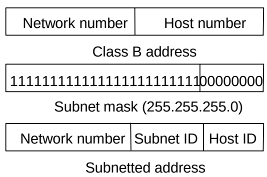

CS 640

Subnet Example

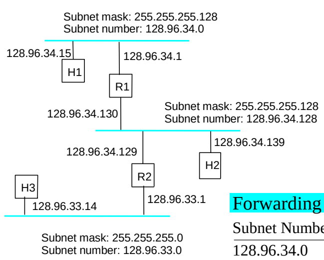

Forwarding table at router R1

Subnet Number	Subnet Mask	Next Hop
128.96.34.0	255.255.255.128	interface 0
128.96.34.128	255.255.255.128	interface 1
128.96.33.0	255.255.255.0	R2

CS 640

Forwarding Algorithm

```
D = destination IP address
for each entry(SubnetNum,SubnetMask,NextHop)
   D1 = SubnetMask & D
   if D1 = SubnetNum
      if NextHop is an interface
         deliver datagram directly to D
      else
         deliver datagram to NextHop
```

- Use a default router if nothing matches
- Can put multiple subnets on one physical network
- Subnets not visible from the rest of the Internet

CS 640

Supernetting

- Assign block of contiguous network numbers to nearby networks
- Restrict block sizes to powers of 2
- Use a bit mask to identify block size
- CIDR: Classless Inter-Domain Routing
  - Routers work with prefixes (subnets and supernets)
- All routers must understand CIDR addressing

Forwarding Table Lookup

- What if more than one prefix matches?
- Longest prefix match
  - Each entry in the forwarding table is:

    < Network Number / Num. Bits> | interface-id

  Suppose we have:

  192.20.0.0/16 | i0

  192.20.12.0/24 | i1

  And destination address is: 192.20.12.7, choose i1

Today’s lecture

- IP
  - Addressing and forwarding
  - ARP
  - DHCP

Address Translation

- Map IP addresses into physical addresses
  - destination host
  - next hop router
- Techniques
  - encode MAC address in host part of IP address
  - table-based
- ARP
  - table of IP to MAC address bindings
  - broadcast request if IP address not in table
  - target machine responds with its MAC address
  - table entries are discarded if not refreshed

ARP Details

- Request Format
  - HardwareType: type of physical network (e.g., Ethernet)
  - ProtocolType: type of higher layer protocol (e.g., IP)
  - HLEN & PLEN: length of physical and protocol addresses
  - Operation: request or response
  - Source/Target-Physical/Protocol addresses
- Notes
  - table entries timeout in about 10 minutes
  - update table with source when you are the target
  - update table if already have an entry
  - do not refresh table entries upon reference

ARP Packet Format

0 8 16 31

Hardware type = 1		ProtocolType = 0x0800
HLen = 48	PLen = 32	Operation
SourceHardwareAddr (bytes 0 – 3)		
SourceHardwareAddr (bytes 4 – 5)		SourceProtocolAddr (bytes 0 – 1)
SourceProtocolAddr (bytes 2 – 3)		TargetHardwareAddr (bytes 0 – 1)
TargetHardwareAddr (bytes 2 – 5)		
TargetProtocolAddr (bytes 0 – 3)		

Reverse Address Resolution Protocol

- RARP is part of the TCP/IP specification
- RARP operates much like ARP
  - A requestor broadcasts is RARP request
  - Servers respond by sending response directly to requestor
  - Requestor keeps IP delivered by first responder
  - Requestor keeps sending requests until it gets an IP
- Need redundant RARP servers for reliability
  - Timeouts can be used to activate backup RARP servers
    - Backup servers reply to a RARP request if they don't hear the RARP response from the primary server after some time

Today's lecture

- IP
  - Addressing and forwarding
  - ARP
  - DHCP

Alternatives to RARP

- RARP has shortcomings
  - Most are subtle, all deal with fact that RARP operates at data link level
- BOOTstrap Protocol (BOOTP) was developed as an alternative to RARP – moves process to network level
  - Uses UDP/IP packets to carry messages
    - Hosts are still identified by MAC address
  - How can UDP running over IP be used by a computer to discover its IP address?
    - Use special case IP address 255.255.255.255 – limited broadcast – not forwarded by routers

Dynamic Configuration

- BOOTP was designed for a static environment where each host has a permanent network addr.
  - Manager creates a BOOTP config file with parameters for each host – file stable for long time
- Wireless networking enables much more dynamic environments
  - BOOTP does not provide for dynamic address assignment
- Dynamic configuration is the primary method for IP address allocation used today
  - Not only facilitates mobility but also efficient use of IPs

Dynamic Host Configuration Protocol

- DHCP extends BOOTP
  - Still supports static allocation
  - Supports automatic configuration where addresses are permanent but assigned by DHCP
  - Supports temporary allocation
- Relies on existence of a DHCP server
  - Repository for host configuration information
  - Maintains a pool of available IP's for use on demand
  - Considerably reduces administration overhead
  - Uses UDP to send messages

DHCP Implementation

- State machine (6 states) determines DHCP operation
  - Host boots into *INITIALIZE* state
- To contact the DHCP server(s) a client sends DHCPDISCOVER message to IP broadcast address and moves to *SELECT* state
  - Unique header format with variable length options field
  - UDP packet sent to well known BOOTP port 67
- Server(s) respond with DHCPOFFER message
  - Client can receive 0 or more responses and responds to one

DHCP Implementation contd.

- Client moves to *REQUEST* state to negotiate IP lease with 1 server
  - Sends DHCPREQUEST message to server which responds with DHCPACK
- Client is then in *BOUND* (normal) state
- From *BOUND*, client can issue DHCPRELEASE and return to *INITIALIZE* state
  - This is simply client deciding it no longer needs the IP

DHCP implementation contd.

- When lease reaches 50% of lease expiration time, it issues DHCPREQUEST to extend lease of current IP with server and moves to *RENEW* state
  - Receipt of DHCPACK moves client back to *BOUND* state
  - Receipt of DHCPNACK moves client back to *INITIALIZE* state
- If no response is received by 87.5% of lease expiration time, the client resends the DHCPREQUEST and moves to *REBIND* state
  - Receipt of DHCPACK moves client back to *BOUND* state
  - Receipt of DHCPNACK or timeout moves client back to *INITIALIZE* state