CS 640 Introduction to Computer Networks

Lecture 8

Today's lecture

- IPv4
  - Packet format, fragmentation
  - Addressing and forwarding
  - NAT

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

Packet format

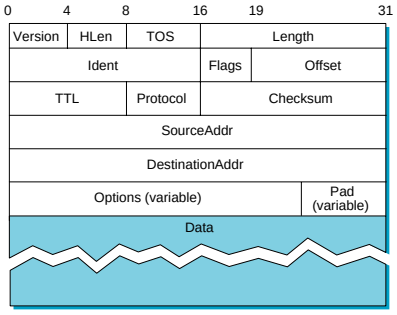

Fragmentation

- Different data link layer technologies impose different limits on the size of packets
- IP fragmentation/reassembly: break packet into pieces and reconstruct it at receiver

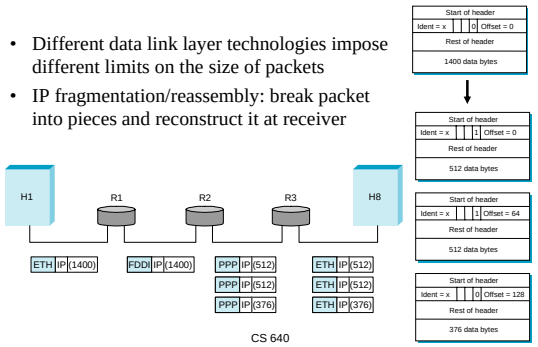

Today's lecture

- IPv4
  - Packet format, fragmentation
  - Addressing and forwarding
  - NAT

Internet Structure

Recent Past

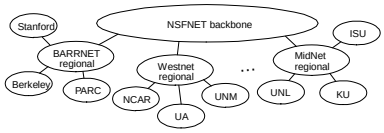

Internet Structure

Today

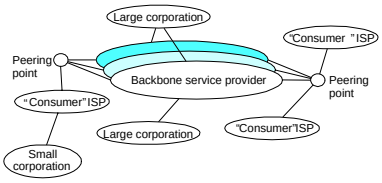

Forwarding Tables

- 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

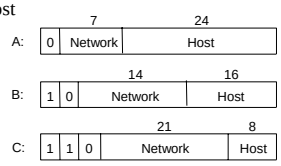

Other Addresses

- 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

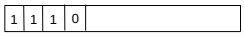

- Host part all 1's: broadcast in local network
- Host part all 0's: unspecified (not allowed)

Forwarding

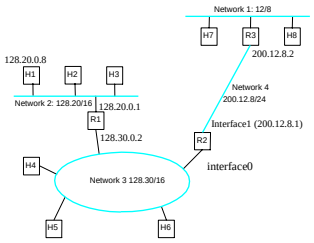

Datagram Forwarding

- 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

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)

Subnetting

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

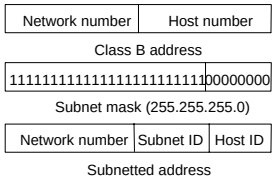

Subnet Example

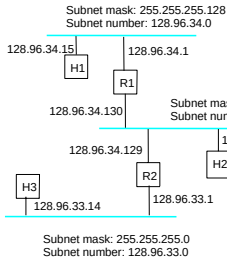

Forwarding table at router R1

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

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

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

- IPv4
  - Packet format, fragmentation
  - Addressing and forwarding
  - NAT

Network Address Translation

- Multiple clients access network through NAT box using single IP address
- **Advantages**:
  - Fewer IP addresses used (clients have private RFC 1918 addresses)
  - Outside hosts cannot connect to NATed clients
- **Disadvantages**:
  - Breaks end to end reachability
  - Outside hosts cannot connect to NATed clients
- Works by rewriting IP, TCP, and UDP headers

Network Address Translation

- Rewrites <hostaddr, port> to <extaddr, port> and back
- Source address, port of outgoing packet changed
- Destination address, port of incoming packet changed

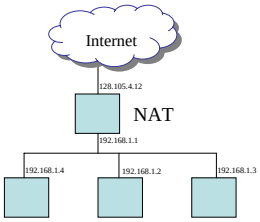

NAT box maintains table with translations

Internal	External
192.168.1.4, 1336	128.105.4.12, 1330
192.168.1.3, 1455	128.105.4.12, 1331
192.168.1.2, 1336	128.105.4.12, 1332
192.168.1.2, 1771	128.105.4.12, 1333