CLASSFUL IPv4 ADDRESSES + DATAGRAM FORWARDING

Internet Protocols

CSC / ECE 573

Fall, 2005

N. C. State University

Today's Lecture

I. IPv4 Addresses

II. Address Classes

III. "Special Case" Addresses

IV. Forwarding Basics

V. Forwarding Decisions

VI. Next-Hop vs. Destination Addresses

IPv4 ADDRESSES

How Do Addresses Get Assigned?

1. ICANN (Internet Corp. for Assigned Numbers and Names)
   – establishes policy for address and name allocation
   – Allocates top-level address space to regional registries
2. Regional registries allocate address space to ISPs, companies, and other organizations
   – APNIC (Asia-Pacific)
   – ARIN (North America)
   – RIPE (Europe)
   – LACNIC (Latin America and Caribbean)
3. Sys admins assign individual host addresses

IP Allocation Goals (RFC 2050)

1. Conservation: fair distribution of globally unique Internet address space, no stockpiling
2. Routability: distribution in a hierarchical manner, makes routing easier
   – good? bad?
3. Public registries document address space allocation and assignment

How Do I Get to www.ietf.org?

User specifies destination of …

www.ietf.org

DNS translates this to…

132.151.6.21

Router forwarding tables determine the path is…

(…some hops omitted…)
24.93.64.53
66.15.132.33
66.185.152.29
66.185.139.129
66.185.145.6
152.63.43.178
152.63.41.138
152.63.39.254
152.63.39.97
157.130.44.142
132.151.6.21

IPv4 Addresses

- 32-bits long, globally unique
- Each interface has an IP address

Example: a router

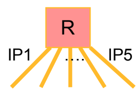

Example: a multi-homed host

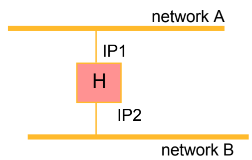

Dotted Decimal Notation

- A convenient way to describe (and remember) IPv4 addresses
- Example

32-bit address	10011000		00000001		00110110		00110000
Dotted decimal representation	152	.	1	.	54	.	48

IPv4 ADDRESS CLASSES

Classful Addresses

- Addresses are organized in a two-level hierarchy
  1. the *network part* (leftmost, most significant)
  2. the *host part* (rightmost, least significant)

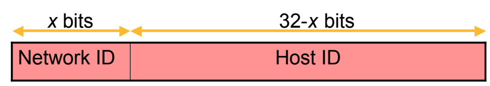

- More networks (= larger network part) means fewer hosts per network (= smaller host part), and vice versa

Classful Address Formats

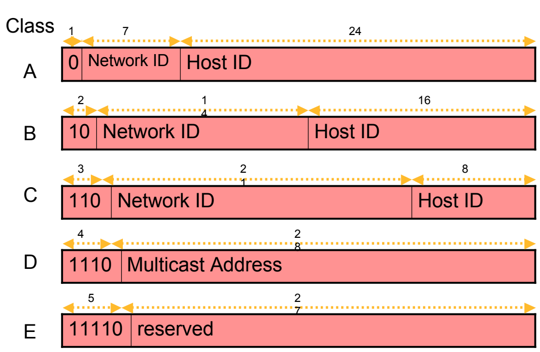

Classful Address Ranges

- The size (number of bits) in the network part is not fixed
  - the first few bits of the address indicate this size
- Classes
  - **A** = addresses 0.0.0.0—127.255.255.255
  - **B** = addresses 128.0.0.0—191.255.255.255
  - **C** = addresses 192.0.0.0—223.255.225.225
  - **D** = addresses 224.0.0.0—239.255.255.255
  - **E** = addresses 240.0.0.0—255.255.255

Classful Network Sizes

Class	Potential Number of Networks	Potential Number of Hosts Per Network
A	2^7 (128)	2^{24} (16M)
B	2^{14} (16K)	2^{16} (64K)
C	2^{21} (2M)	2^8 (256)

Example

- Why is class B address range 128.0.0.0—191.255.255.255?

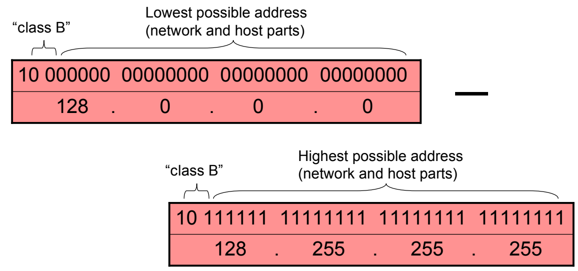

Good or Bad?

1. Good: simple, easy to understand
2. Bad: limited address space
   - 2^{32} = 4G addresses not enough?
3. Bad: limited network size choices (3)
   - ex.: what if a class C net needs to grow beyond 255 hosts?
3. Bad: moving to a new network requires changing IP addresses
   - and may require updating DNS records

How Much of the Address Space is in Use?

SPECIAL-CASE IP ADDRESSES

Network Addresses

- An IP address with host ID part == 0 and network ID part != (all 0's or all 1's) refers to the entire network

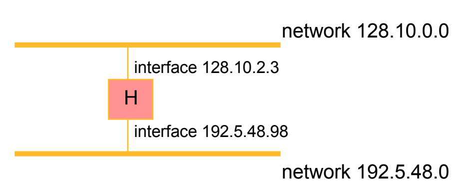

Directed Broadcast Addresses

- An IP <u>destination address</u> with
  Host ID part = all 1's
  means "all hosts attached to the specified network"
- Ex.: Packet sent to address 128.10.255.255 from host H5 will reach H1…H4

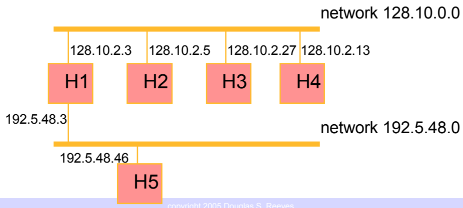

Limited Broadcast Addresses

- An IP <u>destination address</u>
  == all 1's
  means "all hosts part of the same network as me"
- Ex.: Packet sent to 255.255.255.255 from host H3 reaches H1—H4

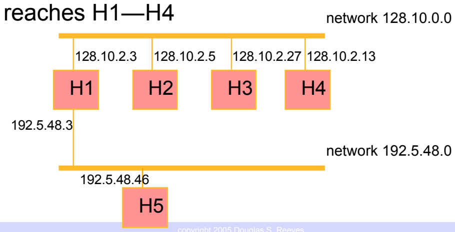

Another Special Case

- An IP <u>source address</u> with
  network ID part = all 0's
  means "from this network"
- Only allowed at startup (during bootstrapping)
  - allows a machine to communicate temporarily before it learns its own IP address
  - thereafter it must not use network 0

The Loopback Address

- An IP <u>destination address</u> with
  network ID part = all 1's
  means "this computer" (i.e., the one sending the packet)
- Used in testing network applications without sending data over a network
  - ex.: "ping 127.0.0.1" should always get a reply!
  - a datagram with destination address 127.x.x.x should never appear on any network

Summary of Special Addresses

For Address of Type…	If Network part is…	And Host part is …	Then this means …
--	Anything other than all 0's or all 1's	All 0's	The address of the whole network
Destination	Anything other than all 0's or all 1's	All 1's	Broadcast address for the specified network
Destination	All 1's	All 1's	Broadcast address for same network as originating host
Source	All 0's	Anything other than all 0's or all 1's	(host which doesn't yet know what network it is attached to)
Destination	127 (Class A, all 1's)	Anything	"This computer" (source of the packet)

RFC 3330: Special-Use IPv4 Addresses

- 0.0.0.0—0.255.255.255 "This" Network [RFC1700]
- 10.0.0.0—10.255.255.255 Private-Use Networks [RFC1918]
- 24.0.0.0—24.255.255.255 Cable Television Networks
- 169.254.0.0—169.254.255.255 Link Local
- 172.16.0.0—172.23.255.255 Private-Use Networks [RFC1918]
- 192.168.0.0--192.168.255.255 Private-Use Networks [RFC1918]
- 224.0.0.0—239.255.255.255 Multicast [RFC3171]
- 240.0.0.0—255.255.255.255 Reserved for Future Use [RFC1700]

FORWARDING BASICS

Routers and Neighbors

- *Routers* (also called *Gateways*)
  - receive packets on one network, send out on another
- *Neighbors* (or directly-connected computers)
  - are attached to the same physical network
  - can communicate directly with each other (i.e., no router needed)

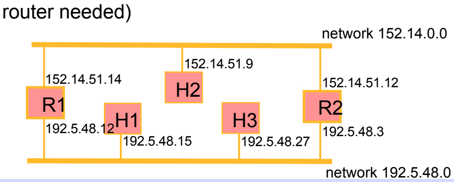

Packet Forwarding

- Deciding what neighbor to send a packet to is a forwarding decision
- Ex.: for H1 to send a packet to H2, should it forward the packet to…
  - 192.5.48.12 (router R1)
  - or 192.5.48.3 (router R2)?

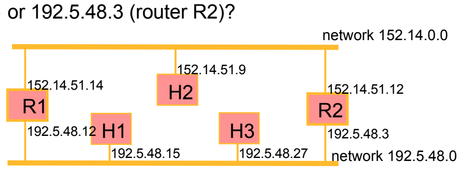

Direct Packet Delivery

- Host H_x wishes to send packet to a **neighboring** host H_y
  - how does H_x know they are on the same network?
- H_x frames (encapsulates) the datagram according to the requirements of the network connecting H_x and H_y
- H_x sends this frame directly to H_y
  - there are **no** intervening routers involved

Indirect Datagram Delivery

- Needed if hosts H_x and H_y are not neighbors
  - Q: how does H_x figure this out?
- H_x picks a neighboring router R1 to forward the datagram to
- H_x frames the packet, sends directly to R1

Indirect Datagram Delivery (cont'd)

- R1 extracts the packet, picks a neighboring router R2 to forward to, frames the packet, sends to R2
- …
- Rn extracts packet, determines Hy is a neighbor (how does Rn know this?), frames the packet, sends directly to Hy

Forwarding (or routing) Tables

- Forwarding decisions are based on information computed by routing protocols
  - this information is stored in a forwarding table
- For router R, each entry in its table consists roughly of
  1. Key
  2. IP address of "next hop" router
  3. Which interface to use

Routing Tables (cont'd)

- Notes!
  - the forwarding table does not specify the complete path to the destination
  - the next router must be directly connected to R

Example

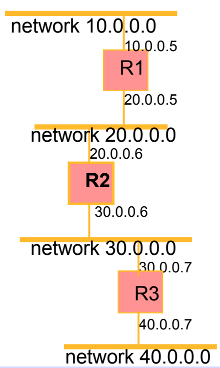

R2's Forwarding Table

To reach hosts on network…	Forward to address …	Which interface to use
20.0.0.0	(direct)	20.0.0.6
30.0.0.0	(direct)	30.0.0.6
10.0.0.0	20.0.0.5	20.0.0.6
40.0.0.0	30.0.0.7	30.0.0.6

FOWARDING DECISIONS

Forwarding Decisions

- The *key* on which forwarding decisions are based is (usually) the destination network ID
- Note that path from H_x to H_y may not be the same as the path from H_y to H_x
- Traffic for destination network N not split across multiple paths
- Why this approach?
  - extremely simple, fast lookup decision
  - drawbacks / limitations?

Forwarding Decisions (cont'd)

- Another benefit: routing tables can be (relatively) small
  - many fewer network addresses than there are host addresses
  - ex.: to deliver packets to one class A network having 16M hosts, only one routing table entry needed!

Other Consequences

- Forwarding (generally) does not consider…
  - application type
  - quality of service requirements
  - bandwidth available
  - congestion
  - reliability
  - … !

"Default" Routes

- Frequently, a single router R is used for most outgoing traffic
  - may need to specify a few destination-specific network routes
  - "everything else" goes through R

"Default" Routes (cont'd)

- In the forwarding table, there will be an entry with *key* = "all other (non-specified) destination networks"
  - normal meaning: "the rest of the Internet"
  - simplifies forwarding tables

"Host-Specific" Routes

- The *key* may be a single destination host address
  - allows specifying a route to a single computer
- Useful for
  - testing and debugging purposes
  - security purposes
  - what else?

The "Datagram Forwarding" Algorithm

```
/* M is a machine (router or host) making  */
/* a forwarding decision about a packet    */
Extract destination address Hd,
compute network part N
if (N matches any directly connected networks)
      deliver to Hd directly
else if (there is a host-specific route for Hd)
      forward datagram to specified next hop
else if (there is a route for network N)
      forward datagram to specified next hop
…
```


The "Datagram Forwarding" Algorithm (cont'd)

```
…
else if (there is a default route)
      forward datagram to default router
else /* Hd is not directly connected and we */
     /* don't know how to get to it…        */
      discard the datagram and declare routing error
```

Host Forwarding Tables

- Hosts also need forwarding tables to pick the appropriate "first hop" router

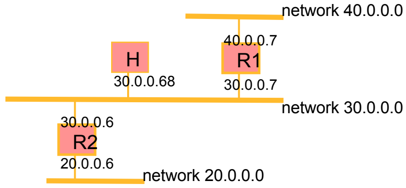

- Frequently there is only one directly-connected router, and the only entry in the table is the default route

Example

```
engr01ras-linux.eos.ncsu.edu% /sbin/ifconfig
eth0      Link encap:Ethernet  HWaddr 00:90:27:71:99:08
          inet addr:152.1.68.208  Bcast:152.1.68.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:2851908 errors:0 dropped:0 overruns:0 frame:0
          TX packets:2506310 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:737519380 (703.3 Mb)  TX bytes:1770287804 (1688.2 Mb)
          Interrupt:19 Base address:0xfce0 Memory:fe800000-fe800038

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:1599968 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1599968 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:1745849721 (1664.9 Mb)  TX bytes:1745849721 (1664.9 Mb)

engr01ras-linux.eos.ncsu.edu% netstat -r
Kernel IP routing table
Destination     Gateway         Genmask         Flags   MSS Window  irtt Iface
152.1.68.0      *               255.255.255.0   U         0 0          0 eth0
169.254.0.0     *               255.255.0.0     U         0 0          0 eth0
127.0.0.0       *               255.0.0.0       U         0 0          0 lo
default         poehub-6509msfc 0.0.0.0         UG        0 0          0 eth0
```


The “Datagram Receiving” Algorithm

```
if (Hd is one of M’s IP addresses)
      receive the datagram
else if (Hd is a limited or directed broadcast
      address for the network on which it was sent)
      receive the datagram
else if (M is a router)
      forward the datagram if possible
else /* M is a host and this packet is not intended for it */
      discard the datagram
```


Should Multi-Homed Hosts Forward?

- Since they don’t participate in routing protocols… probably not!
  - inefficient routes
  - can create loops
  - leads to broadcast "storms“
  - etc.

DESTINATION vs. NEXT-HOP ADDRESSES

Destination vs. Next Hop IP Addresses

- The destination IP address in a IP datagram never changes
- At router R, the datagram is framed and a physical address is added to get it to the "next hop router"

Example

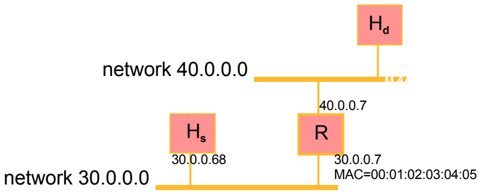

1. H_s wants to send IP packet to destination IP address 40.0.0.42 (H_d)
2. H_s sends on network 30.0.0.0 the encapsulated IP packet with <MAC address 00:01:02:03:04:05 of R, IP address 40.0.0.42>

Example (cont'd)

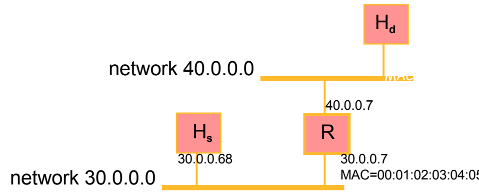

3. R sends on network 40.0.0.0 the encapsulated IP packet with <MAC address 66:77:88:99:AA:BB of H_d, IP address 40.0.0.42>
4. Demultiplexed IP packet received by H_d

Example (cont'd)

- The next hop router IP address (from the routing table) is never stored in the packet
  - must be translated into a physical address instead
- So… why not just store MAC addresses in routing tables?
  - routing is IP-layer function (i.e., should be independent of the link layer)

Summary

1. IP Addresses use two levels of hierarchy
2. First few bits of address specify what class it is
3. Special addresses reserved for particular uses
4. Both hosts and routers have to make forwarding decisions
5. Forwarding tables contain the information needed to make these decisions
6. Forwarding decisions are based on the destination only

Next Lecture

- ICMP