Lecture 17: Network Layer Addressing, Control Plane, and Routing COMP 411, Fall 2022 Victoria Manfredi





Acknowledgements: materials adapted from Computer Networking: A Top Down Approach 7th edition: ©1996-2016, J.F Kurose and K.W. Ross, All Rights Reserved as well as from slides by Abraham Matta at Boston University, and some material from Computer Networks by Tannenbaum and Wetherall.

Today

- 1. Internet protocol
- 2. Network programming
 - raw sockets and byte packing

3. Addressing

- IPV4 addresses
- usage in routing
- how to get an IP address

Network Layer INTERNET PROTOCOL

IP packet format

Q: Why is version 1st? So you know how to parse pkt



Bits transmitted left to right, top to bottom

Wireshark

Look at IP headers and ping/traceroute

IP fragmentation and reassembly

Network links have MTU

- largest possible link-level frame
- different link types have different MTUs

Fragment when pkt > MTU

- 1 pkt becomes several pkts
 - IP header bits used to identify and order related fragments
- reassembled only at final dst
- re-fragmentation possible
- don't recover from lost fragments
- (IPv6 does not support)



DoS attack: send fragmented pkts but leave one out

IP fragmentation and reassembly

4000 byte packet

- 3980 bytes payload
- IP hdr >=20 bytes •



segment

Network Programming RAW SOCKETS

Raw sockets

Take bytes put into socket and push out of network interface – no IP or transport layer headers added by operating system!

Lets you create your own transport and network layer headers

- set field values as you choose
 - e.g., time-to-live fields

Raw sockets

Create send and receive sockets
send_sock = socket.socket(
 socket.AF_INET, socket.SOCK_RAW, socket.IPPROTO_RAW)
recv_sock = socket.socket(
 socket.AF_INET, socket.SOCK_RAW, socket.IPPROTO_ICMP)

Set IP_HDRINCL flag so kernel does not rewrite header fields send_sock.setsockopt(socket.IPPROTO_IP, socket.IP_HDRINCL, 1)

Set receive socket timeout to 2 seconds
recv_sock.settimeout(2.0)

https://docs.python.org/3/library/socket.html

Byte packing and structs

```
def create_icmp_header(self):
```

```
ECHO_REQUEST_TYPE = 8
ECHO_CODE = 0
```

```
# ICMP header info from https://tools.ietf.org/html/rfc792
icmp_type = ECHO_REQUEST_TYPE  # 8 bits
icmp_code = ECHO_CODE  # 8 bits
icmp_checksum = 0  # 16 bits
icmp_identification = self.icmp_id # 16 bits
icmp_seq_number = self.icmp_seqno  # 16 bits
```

return icmp_header

https://docs.python.org/3/library/struct.html

Addressing IPV4 ADDRESSES

IPv4 addresses

Globally unique 32-bit identifier

- associated with host or router interface
- interface: connection between host/router and physical link
 - host: usually 1 or 2 interfaces
 - router: usually many interfaces

Address format is hierarchical

- CIDR: Classless InterDomain Routing
- split into subnet part and host part
 - a.b.c.d/x, where x is # bits in subnet part



What's a subnet?

Subnet

- set of interfaces with same subnet part of IP addr
- devices reachable without intervening routers

Subnet mask

- divides IP addr into subnet addr + host addr
- included in routing info given to routers

Recipe to find subnets

- detach each interface from its host or router
- create islands of isolated networks, i.e., subnets

Network comprising 3 subnets



wireless WiFi interfaces connected by WiFi base station



Subnet mask example



Subnet mask

- zeroes out host part
- e.g., 200.23.16.0/23
 - 11111111 1111111 1111110 0000000
- take logical "and" of subnet mask with address to get subnet part
 - 1 AND $1 \rightarrow 1$
 - 1 AND $0 \rightarrow 0$
 - 0 AND $1 \rightarrow 0$
 - 0 AND 0 \rightarrow 0

Ifconfig example

> ifconfig lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384 options=1203<RXCSUM,TXCSUM,TXSTATUS,SW_TIMESTAMP> inet 127.0.0.1 netmask 0xff000000 inet6 ::1 prefixlen 128 inet6 fe80::1%lo0 prefixlen 64 scopeid 0x1 nd6 options=201<PERFORMNUD,DAD> gif0: flags=8010<POINTOPOINT,MULTICAST> mtu 1280 stf0: flags=0<> mtu 1280 en0: flags=8863<UP, BROADCAST, SMART, RUNNING, SIMPLEX, MULTICAST> mtu 1500 ether 78:4f:43:73:43:26 inet6 fe80::1c8d:4bcb:b52d:9d1d%en0 prefixlen 64 secured scopeid 0x5 nd6 options=201<PERFORMNUD,DAD> media: autoselect status: active

Hex is [0:15] where A=10, B=11, C=12, D=13, E=14, F=15

1111	1111	1111	1111	1111	1100	0000	0000	Q: Why is broadcast
f	f	f	f	f	С	0	0	addr 10.66.107.255?

Subnet masks and address blocks

Suppose

- we must have 223.1.1 as network prefix
- we need block of 90 addresses

What should subnet mask be?

- how many bits for 90 addresses?

 2^7 2^6 2^5 2^4 2^3 2^2 2^1 2^0 128 64 32 16 8 4 2 1

> 223.1.1.0/24 gives 256 addresses [0-255]

223.1.1.0/25 gives 128 addresses [0-127]

223.1.1.128/25 gives a different set of 128 addresses [128-255]

н

IP addresses are hierarchical

Pros

- scalable: routers don't need to look at host part
- all pkts on same network forwarded in same direction
 - · only when pkt reaches network does host matter

Cons

- every IP addr belongs to specific network
- what if host moves networks and wants to keep same addr?
 - mobile IP
 - contrast with fixed Ethernet link layer addr

Special addresses

Private subnet (used in NAT), do not appear on Internet

- 172.16-31.*.*
- 10.*.*.*
- 192.168.*.*

Loopback address:

- 127.*.*.*

Addresses you can't assign to devices

- *.*.*.255: broadcast addr
- *.*.*.0: used for subnet name

Broadcast address

- 255.255.255.255: broadcast to all hosts on network indicated
 - if no mask: local network
 - if mask: broadcast on that network

Address when device booting up

- 0.0.0.0

Addressing USAGE IN ROUTING

Routers forward traffic to networks not hosts

Forwarding table

- does not contain row for every dest IP address
- instead computes routes between subnets (blocks of addresses)

D	Link Interface			
11001000 through	00010111	00010000	0000000	0
11001000	00010111	00010111	11111111	
11001000 through	00010111	00011000	0000000	1
11001000	00010111	00011000	11111111	•
11001000 through	00010111	00011001	0000000	2
11001000	00010111	00011111	11111111	
otherwise				3

What if address ranges don't divide up nicely?

Longest prefix matching

use longest address prefix that matches destination address

Destination Address Range	Link interface
11001000 00010111 00010*** ********	0
11001000 00010111 00011000 ********	1
11001000 00010111 00011*** ********	2
otherwise	3

Question

DA: 11001000 00010111 00010110 10100001 which interface? DA: 11001000 00010111 00011000 10101010 which interface?

How big is a routing table for a core router?

800000

700000

600000

From http://www.cidr-report.org/as2.0/

Table History



Q: If a core router processes 1million pkts+ per second, how fast does it need to be able to search table?

Hierarchical addressing

Route aggregation

- combine multiple small prefixes into a single larger prefix
- allows efficient advertisement of routing information



Longest prefix matching

More specific routes

- ISPs-R-Us has a more specific route to Organization 1



Addressing HOW TO GET AN IP ADDRESS?

How does ISP get block of addresses?

ICANN

- Internet Corporation for Assigned Names and Numbers
- <u>http://www.icann.org/</u>

ICANN functions

- allocates addresses
- manages DNS
- assigns domain names, resolves disputes

— ...

How does network get net part of IP address?

Allocated portion of its provider ISP's address space

ISP's block	11001000	00010111	<u>0001</u> 0000	00000000	200.23.16.0/20
Organization 0	<u>11001000</u>	00010111	<u>0001000</u> 0	00000000	200.23.16.0/23
Organization 1	<u>11001000</u>	00010111	<u>0001001</u> 0	00000000	200.23.18.0/23
Organization 2	<u>11001000</u>	00010111	<u>0001010</u> 0	00000000	200.23.20.0/23
Organization 7	11001000	00010111	<u>0001111</u> 0	00000000	200.23.30.0/23

How does host get an IP address?

Option 1

- hard-coded by system admin in a file on your host

Option 2:

- dynamically get address from a server
 - DHCP: Dynamic Host Configuration Protocol

We're running out of IPv4 addresses

Why?

- inefficient use of address space
 - from pre-CIDR use of address classes (A: /8, B: /16, C: /24)
- too many networks (and devices)
 - Internet comprises 100,000+ networks
 - routing tables and route propagation protocols do not scale

Q: how many IPv4 addresses are there?

- 2³²

Solutions

- IPv6 addresses
- DHCP: Dynamic Host Configuration Protocol
- NAT: Network Address Translation