15-441 Computer Networking

Lecture 4 –Internet design and IP Addressing

Peter Steenkiste

Fall 2016

www.cs.cmu.edu/~prs/15-441-F16

Outline

- IP design goals
- Traditional IP addressing
  - Addressing approaches
  - Class-based addressing
  - Subnetting
  - CIDR
- IP protocol and friends
- Routing

Logical Structure of an Internet

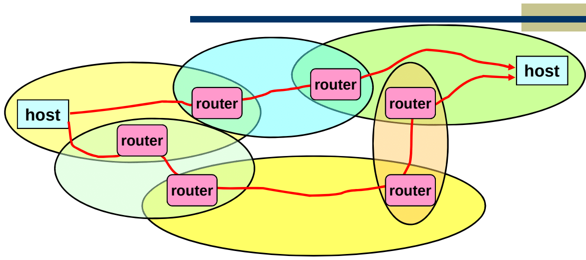

- Interconnection of separately managed networks using routers
  - Topology has emerged over time – not designed
  - Individual networks can use different (layer 1-2) technologies
  - The public Internet is a special (highly successful) example
- Send packets from source to destination by hopping through networks
  - "Network" layer responsiblity

Internet Protocol (IP)

- Inter-network connectivity provided by the Internet protocol
- Implemented on *routers:* forward IP packets between networks
- IP creates abstraction layer that hides underlying technology from network application software
  - Allows range of current & future technologies
  - WiFi, traditional and switched Ethernet, personal area networks, …

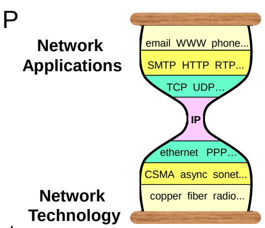

- But is a router not just a switch?

Frame Forwarding

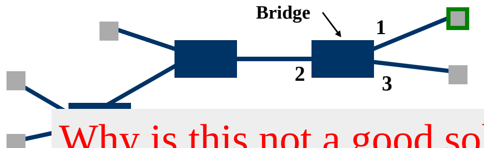

Why is this not a good solution for the Internet?

MAC Addr	Port	Age
A21032C9A591	1	36
99A323C90842	2	01
8711C98900AA	2	15
301B2369011C	2	16
695519001190	3	11

- … which port to use for forwarding
- For every packet, the bridge "looks up" the entry for the packets destination MAC address and forwards the packet on that port.
  - Other packets are broadcast – why?
- Timer is used to flush old entries

5

What are the Goals?

- LANs: "Connect hosts" → switching:
  - "Wire" abstraction: behaves like Ethernet
  - Only has to scale up a "LAN size"
  - Availability
- Internet: "Connect networks" → routing:
  - Scalability
  - Manageability of individual networks
  - Availability
- Affects addressing, protocols, routing

6

Outline

- IP design goals
- Traditional IP addressing
  - Addressing approaches
  - Class-based addressing
  - Subnetting
  - CIDR
- IP protocol and friends
- Routing

7

Getting to a Network Destination

- How do you get driving directions?
- Intersections→routers
- Roads→links/networks
- Roads change slowly

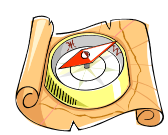

8

Forwarding Packets

- (Table of virtual circuits ids)
  - More on this later
- Table of global destination addresses (IP)
  - Routers keep next hop for destination
  - Packets carry destination address
- Source routing - no forwarding table!
  - Packet carries a path

Source Routing

- List entire path in packet
  - Driving directions (north 3 hops, east, etc..)
- Router processing
  - Strip first step from packet
  - Examine next step in directions and forward
- Rarely used
  - End points need to know a lot about network
  - Economic and security concerns
  - Variable header size

Global Address Example

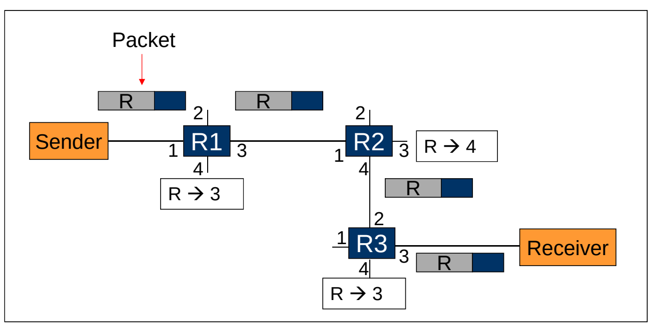

Forwarding based on Global Addresses

- Advantages
  - Conceptually simple
  - Lines up with roles of actors (ISPs, endpoints)
  - Stateless (soft state) – simple error recovery
- Disadvantages - challenges
  - Every switch knows about every destination
    - Potentially large tables – today's topic
  - All packets to destination take same route
    - Potentially inefficient - "Traffic engineering" lecture
  - Need routing protocol to fill table
    - Next couple of lectures

Addressing in IP

- IP addresses identify interfaces
  - E.g., 128.2.1.1
  - Multiple interfaces -> multiple IP addresses
- Domain Name System (DNS) names are names of hosts
  - E.g., www.cmu.edu
- DNS binds host names to interfaces
  - More on DNS later in the course
- Routing binds interface addresses to paths

13

Addressing Considerations

- Flat addresses – one address for every host
  - Peter Steenkiste: 123-45-6789
  - Does not scale – router table size explodes
  - 630M (1/09) entries,doubling every 2.5 years
  - Why does it work for Ethernet?
- Hierarchical – add structure
  - Pennsylvania / Pittsburgh / Oakland / CMU / Steenkiste or Peter Steenkiste: (412)268-0000
  - Common "trick" to simplify forwarding, reduce forwarding table
- What type of Hierarchy?
  - How many levels?
  - Same hierarchy depth for everyone?
  - Who controls the hierarchy?

14

IP Address Structure

Challenge: Accommodate networks of different very sizes

Initially: classful structure (1981) (not relevant now!!!)

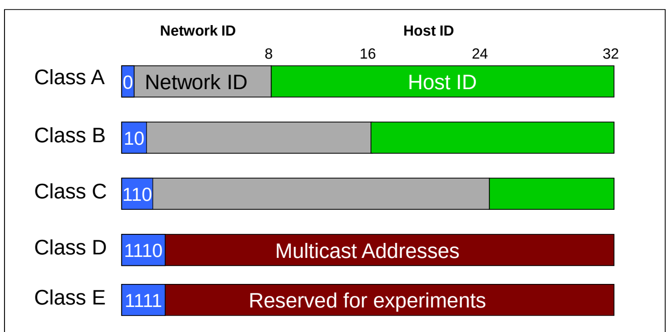

15

Original IP Route Lookup

- Address specifies prefix for forwarding table
  - Extract address type and network ID
- Forwarding table contains
  - List of class+network entries
  - A few fixed prefix lengths (8/16/24)
  - Prefix – part of address that really matters for routing
- www.cmu.edu address 128.2.11.43
  - Class B address – class + network is 128.2
  - Lookup 128.2 in forwarding table for class B
- Tables are still large!
  - 2 Million class C networks

16

Subnet Addressing RFC917 (1984)

- Some "LANs" are very big, class A & B networks
  - Large companies, universities, …
  - Internet became popular quickly
- Cannot manage this as a single LAN
  - Hard to manage, becomes inefficient
- Need simple way to partition large networks
  - Partition into multiple IP networks that share the same prefix – called a "subnet", part of a network
- CMU case study in RFC
  - Chose not to adopt – concern that it would not be widely supported ☺

Subnetting

- Add another layer to hierarchy
- Variable length subnet masks
  - Could subnet a class B into several chunks
- Subnetting is done internally in the organization
  - It is not visible outside – important for management

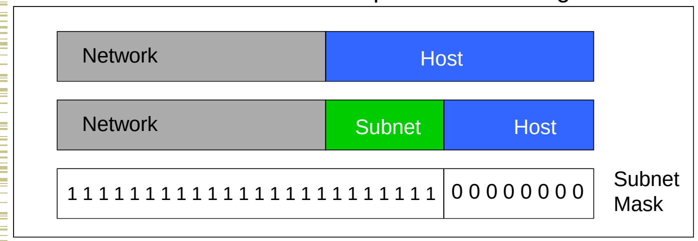

Subnetting Example

- Assume an organization was assigned address 150.100
- Assume < 100 hosts per subnet
- How many host bits do we need?
  - Seven
- What is the network mask?
  - 11111111 11111111 11111111 10000000
  - 255.255.255.128

Forwarding Example

- Assume a packet arrives with address 150.100.12.176
- Step 1: AND address with class + subnet mask
  - Subnet masks stored on router

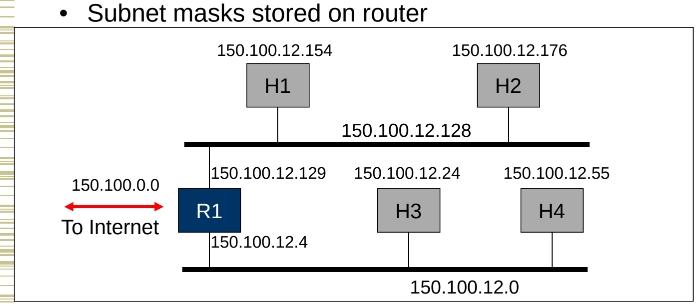

Outline

- IP design goals
- Traditional IP addressing
  - Addressing approaches
  - Class-based addressing
  - Subnetting
  - CIDR
- IP protocol and friends
- Routing

IP Address Problem (1991)

- Address space depletion
  - Suppose you need 2^{16} + 1 addresses?
  - Class A too big for all but a few domains
  - Class C too small for many domains but they don't need a class B address
  - Class B address pool allocated at high rate
  - Many allocated address block are sparsely used
- Developed a strategy based on a three solutions
  - Switch to a "classless" addressing model
  - Network address translation
  - Definition of IPv6 with larger IP addresses

Classless Inter-Domain Routing (CIDR) – RFC1338

- Arbitrary split between network & host part of address → more efficient use of address space
  - Do not use classes to determine network ID
  - Use "prefix" that is propagated by routing protocol
  - E.g., addresses 192.4.16 - 192.4.31 have the first 20 bits in common. Thus, we use these 20 bits as the prefix (network number) → 192.4.16/20
- Merge forwarding entries → smaller tables
  - Use single entry for range in forwarding tables even if they belong to different destination networks
    - "Adjacent" in address space and same egress

CIDR Example

- Network is allocated 8 class C chunks, 200.10.0.0 to 200.10.7.255
  - Move 3 bits of class C address to host address
  - Network address is 21 bits: 201.10.0.0/21
- Replaces 8 class C routing entries with 1 entry
- But how do routers know size of network address?
  - Routing protocols must carry prefix length with address

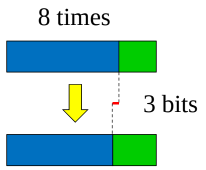

IP Addresses: How to Get One?

Network (network portion):

- Get allocated portion of ISP's address space:

ISP's block	11001000 00010111 00010000	00000000	200.23.16.0/20
Organization 0	11001000 00010111 00010000	00000000	200.23.16.0/23
Organization 1	11001000 00010111 00010010	00000000	200.23.18.0/23
Organization 2	11001000 00010111 00010100	00000000	200.23.20.0/23
...	.....	....	....
Organization 7	11001000 00010111 00011110	00000000	200.23.30.0/23

25

IP Addresses: How to Get One?

- How does an ISP get block of addresses?
  - From **Regional Internet Registries** (RIRs)
    - ARIN (North America, Southern Africa), APNIC (Asia-Pacific), RIPE (Europe, Northern Africa), LACNIC (South America)
- How about a single host?
  - Assigned by sys admin (static or dynamic)
  - DHCP: Dynamic Host Configuration Protocol: dynamically get address: "plug-and-play"
    - Host broadcasts "DHCP discover" msg
    - DHCP server responds with "DHCP offer" msg
    - Host requests IP address: "DHCP request" msg
    - DHCP server sends address: "DHCP ack" msg

26

IP Address Availability Remains a Major Challenge

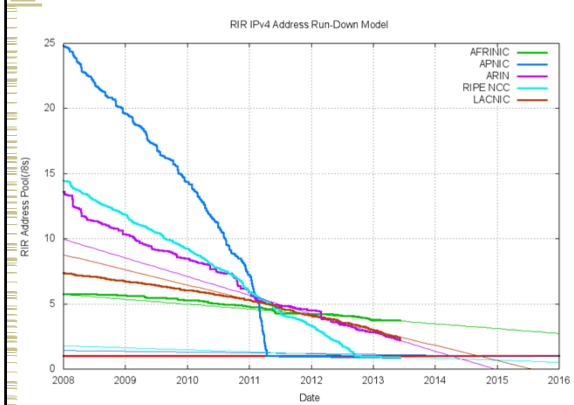

- Some are in big trouble!
- APNIC:
  - Asia
- AFRINIC:
  - Africa
- ARIN:
  - North America
- LACNIC:
  - Latin America
- RIPE NCC:
  - Europe, Middle East, parts of central Asia

27

CIDR Illustration

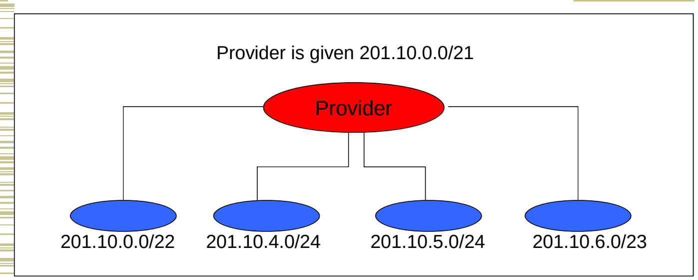

28

CIDR Implication: Longest Prefix Match

- How to deal with multi-homing, legacy addresses, …

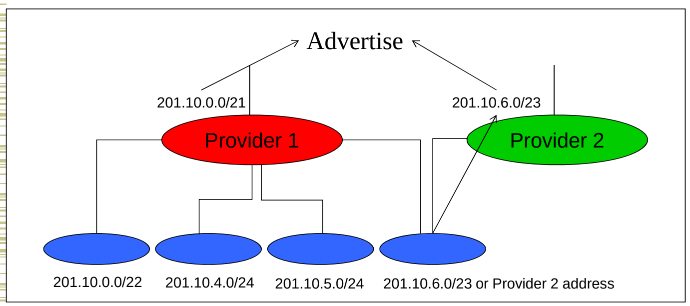

Host Routing Table Example

```
Destination     Gateway         Genmask         Iface
128.2.209.100   0.0.0.0         255.255.255.255 eth0
128.2.0.0       0.0.0.0         255.255.0.0     eth0
127.0.0.0       0.0.0.0         255.0.0.0       lo
0.0.0.0         128.2.254.36    0.0.0.0         eth0
```

- From "netstat –rn"
- Host 128.2.209.100 when plugged into CS ethernet
- Dest 128.2.209.100 → routing to same machine
- Dest 128.2.0.0 → other hosts on same ethernet
- Dest 127.0.0.0 → special loopback address
- Dest 0.0.0.0 → default route to rest of Internet
  - Main CS router: gigrouter.net.cs.cmu.edu (128.2.254.36)

Routing to the Network

- Packet to 10.1.1.3 arrives
- Path is R2 – R1 – H1 – H2
- H1 serves as a router for the 10.1.1.2/31 network

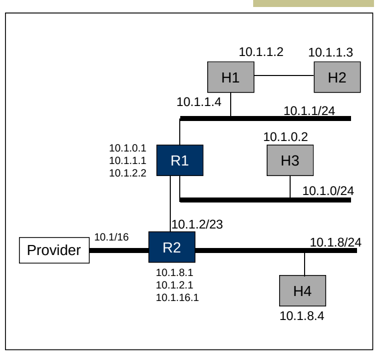

Routing Within the Subnet

- Packet to 10.1.1.3
- Matches 10.1.0.0/23

Routing table at R2

Destination	Next Hop	Interface
127.0.0.1	-	lo0
Default or 0/0	provider	10.1.16.1
10.1.8.0/24	-	10.1.8.1
10.1.2.0/23	-	10.1.2.1
10.1.0.0/23	10.1.2.2	10.1.2.1

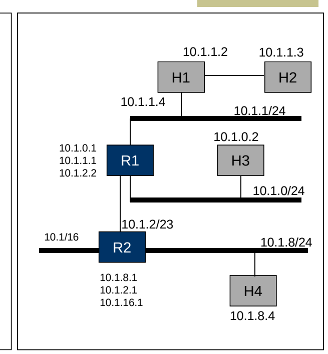

Routing Within the Subnet

- Packet to 10.1.1.3
- Matches 10.1.1.2/31
  - Longest prefix match

Routing table at R1

Destination	Next Hop	Interface
127.0.0.1	-	lo0
Default or 0/0	10.1.2.1	10.1.2.2
10.1.2.0/23	10.1.2.1	10.1.2.2
10.1.0.0/24	-	10.1.0.1
10.1.1.0/24	-	10.1.1.1
10.1.1.2/31	10.1.1.4	10.1.1.1

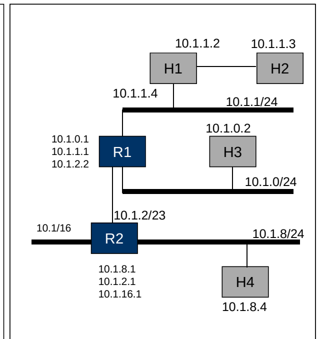

33

Routing Within the Subnet

- Packet to 10.1.1.3
- Direct route
  - Longest prefix match

Routing table at H1

Destination	Next Hop	Interface
127.0.0.1	-	lo0
Default or 0/0	10.1.1.1	10.1.1.4
10.1.1.0/24	-	10.1.1.1
10.1.1.2/31	-	10.1.1.2

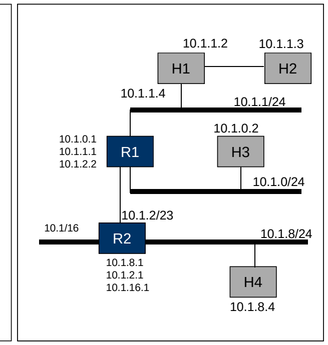

34

Important Concepts

- Hierarchical addressing critical for scalable system
  - Don't require everyone to know everyone else
  - Reduces number of updates when something changes
- Classless inter-domain routing supports more efficient use of address space
  - Adds complexity to routing, forwarding, …
  - Not a problem today

35