Simple Internetworking: Global Addressing and Datagram Forwarding

Hui Chen

Department of Computer & Information Science

CUNY Brooklyn College

Outline

- Topic: internetworking
  - Case study: Internet Protocol (IP) Suite
- Simple interworking
  - Overview of internet and the Internet
  - Global addressing scheme
  - **Best effort service model and datagram forwarding**
  - Packet fragmentation and assembly
  - Address translation
  - Host configuration
  - Error reporting

IP Service Model

- Packet Delivery Model
  - Connectionless model for data delivery
  - Best-effort delivery (unreliable service)
    - packets may be lost
    - packets may be delivered out of order
    - duplicate copies of a packet may be delivered
    - packets may be delayed for a long time
  - Why motivated this design requirement (policy)?
- Mechanism (solution) to support this design requirement (policy)
  - Datagram forwarding between networks (application of datagram switching)
  - Two problems:
    - Identify hosts
    - Identify networks

IP Address Space

- Global unicast addresses
- Multicast addresses
- Broadcast addresses
- Private addresses
- Link-local addresses
- And more (such as for testing and documentation)

IPv4 Address Space

- 32-bit long
- 2^{32} Addresses
  - Global unicast addresses
  - Multicast addresses
  - Broadcast addresses
  - Private addresses (private unicast addresses)
  - Link-local addresses
  - …

Global Unicast IPv4 Addresses

- Binary representation begins with 0, 10, 110
- Globally unique
- Human-readable form: IPv4 numbers-and-dots notation
  - 146.245.201.50
- hierarchical: network + host
  - An unicast IPv4 address specifies both network and the host on the network

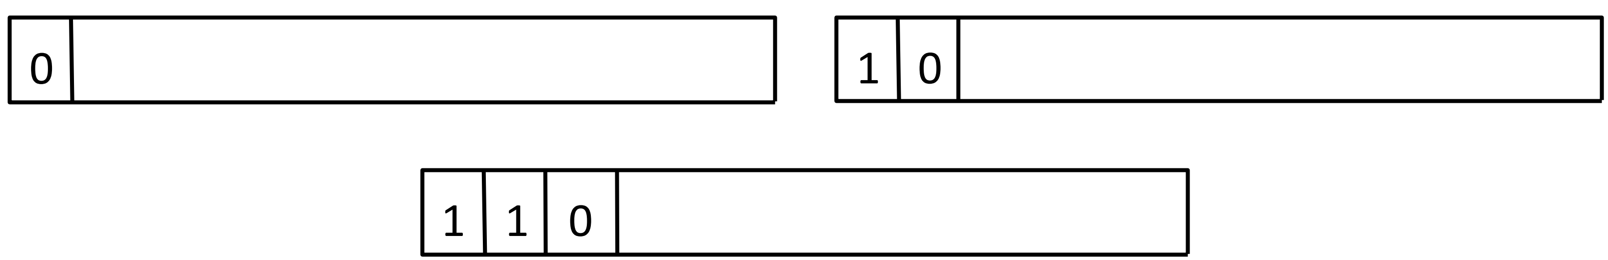

IPv4 Multicast Address

- Addresses starting with 1110

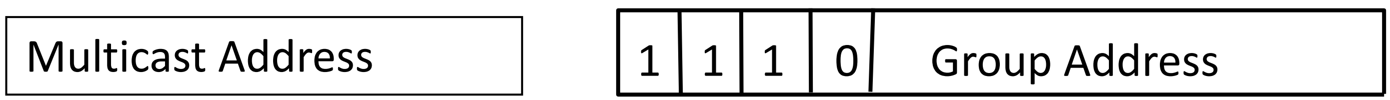

IPv4 Broadcast Address

- IP unicast address: network number + host number
- Setting all the host bits to 1

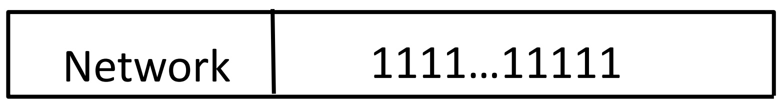

- We shall discuss methods to divide the bits into network number and host number

IPv4 Address Spaces for Private Networks

- See RFC 1918
- Private networks
  - 24-bit block 10.0.0.0–10.255.255.255
    - First 8 bits: 0000 1010
  - 20-bit block 172.16.0.0–172.31.255.255
    - First 12 bits: 1010 1100 0001
  - 16-bit block 192.168.0.0–192.168.255.255
    - First 16bits: 1100 0000 1010 1000

Link Local IPv4 Address

- See RFC 3927
- Link-Local IPv4 Address
  - 16-bit block 169.254.0.0–169.254.255.255
  - First 16 bits: 1010 1001 1111 1110

Exercise 1

- Find out IPv4 addresses of following hosts (with the given domain names) and indicate the class to which the IP addresses belong
  - www.brooklyn.cuny.edu
  - www.sme.sk
  - www.google.com
- Convert the first number (from left) to a binary number, then take a look at the 1st, and/or 2nd, and/or 3rd bit
- Are the addresses global unicast addresses?
- Remark
  - There are many ways to find out the IP address of a host given a domain name
    - Example: nslookup (which works on most platforms including Windows, Unix/Linux, and Mac OS X) and dig (works on Linux)

Summary and Questions?

- IP service model
- IPv4 addresses
  - Unicast addresses
  - Multicast addresses
  - Broadcast addresses
  - Private addresses
  - Link-local addresses
- Look-up global unicast address assignment for some hosts with domain names

IP Datagram Forwarding

- Set-up:
  - Every datagram contains source and destination's addresses
    - Each address: network number + host number
  - Each network has network gateways/routers
  - Each host and router maintains a forwarding table
    - each host/router can have a default router
  - forwarding table maps network number into next hop
- Forwarding algorithm
  - if source and destination are on the same network, then forward to the destination host
  - If source and destination are not connected to the same network, then forward to some router (representing a network)

Forwarding Algorithm

```
if (NetworkNum of destination = NetworkNum of one
  of my interfaces) then

  deliver packet to destination over that interface

else

  if (NetworkNum of destination is in my forwarding
  table) then

       deliver packet to NextHop router

  else

       deliver packet to default router
```

Forwarding Algorithm Simplified

- For a host with only one interface and only a default router in its forwarding table, this simplifies to

```
if (NetworkNum of destination = my NetworkNum) then
  deliver packet to destination directly
else
  deliver packet to default router
```

Forwarding Table: Example

- Forwarding table at router R2 that has two interfaces 0 and 1

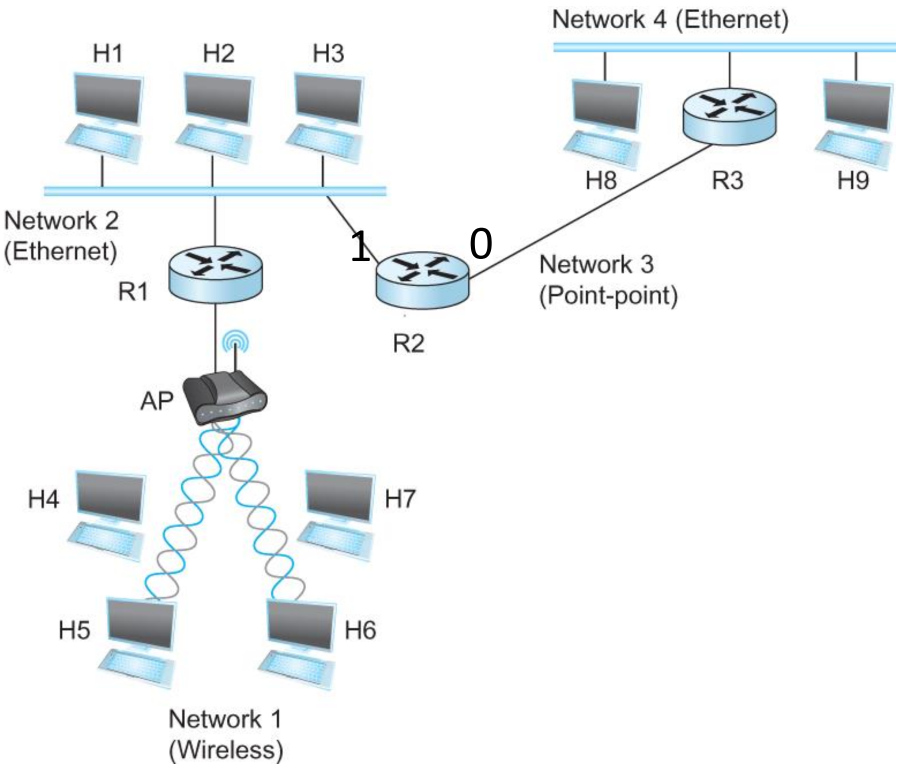

NetworkNum	NextHop
1	R1
2	Interface 1
3	Interface 0
4	R3

Exercise 2

- Construct forwarding tables for routers R1 and R3. Interfaces of routers are marked

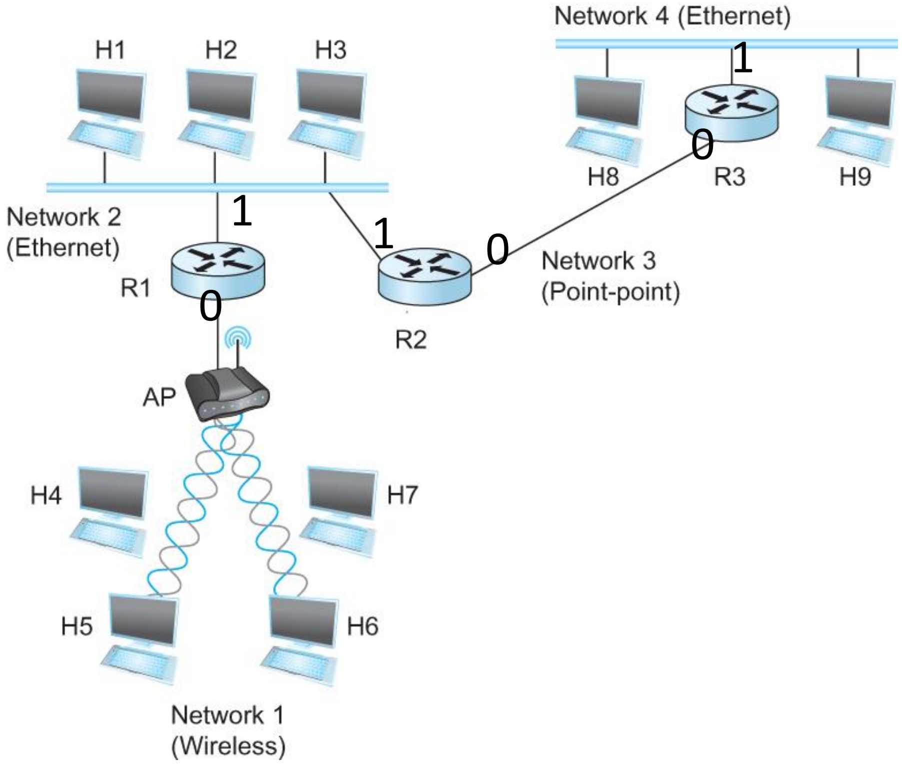

Summary and Questions?

- IP Datagram Forwarding
- Forwarding table
- Forwarding algorithm
- But how do we divide an address into a network number and a host number?

Network + host numbers?

- We will need to divide a unicast address into two parts: network number + host numbers
  - (legacy) IPv4 address classes
  - (legacy) subnetting/supernetting
  - CIDR addressing
- Forwarding algorithms will need to be adjusted accordingly

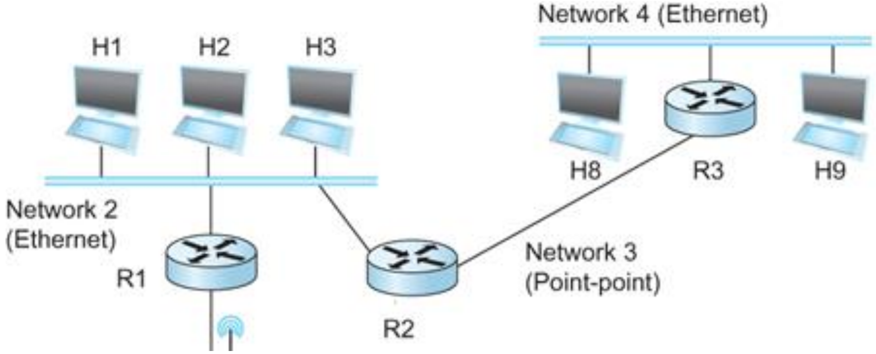

NetworkNum	NextHop
1	R1
2	Interface 1
3	Interface 0
4	R3

IPv4 Unicast Address Classes

- To express network and host
  - Divide into classes (legacy)

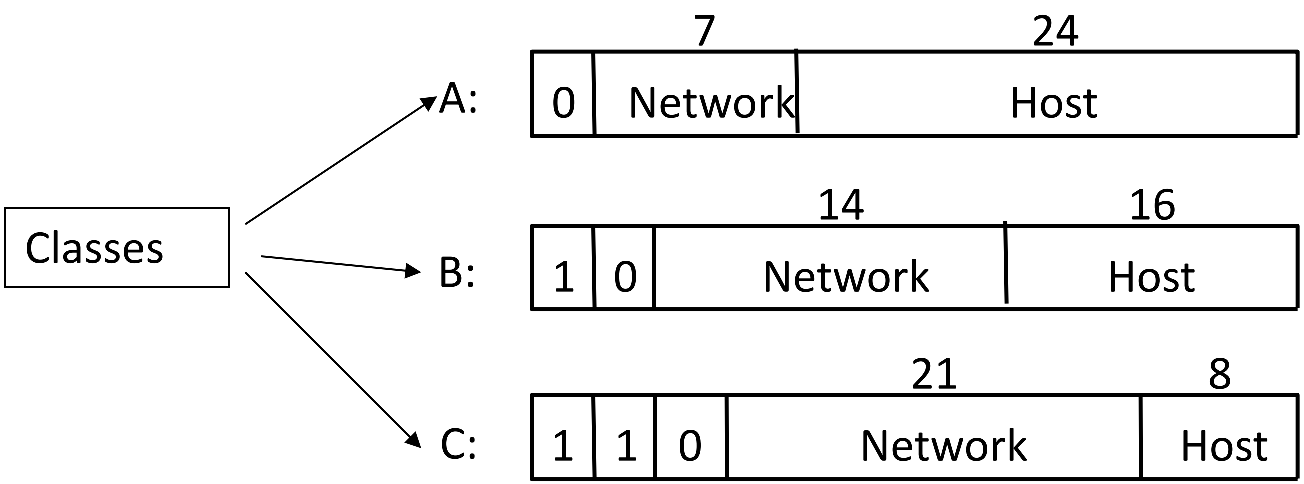

Subnetting

- Subnetting:
  - divide large network into smaller networks using a network mask
  - merge smaller networks into a large network using a network mask

Subnetting: Example

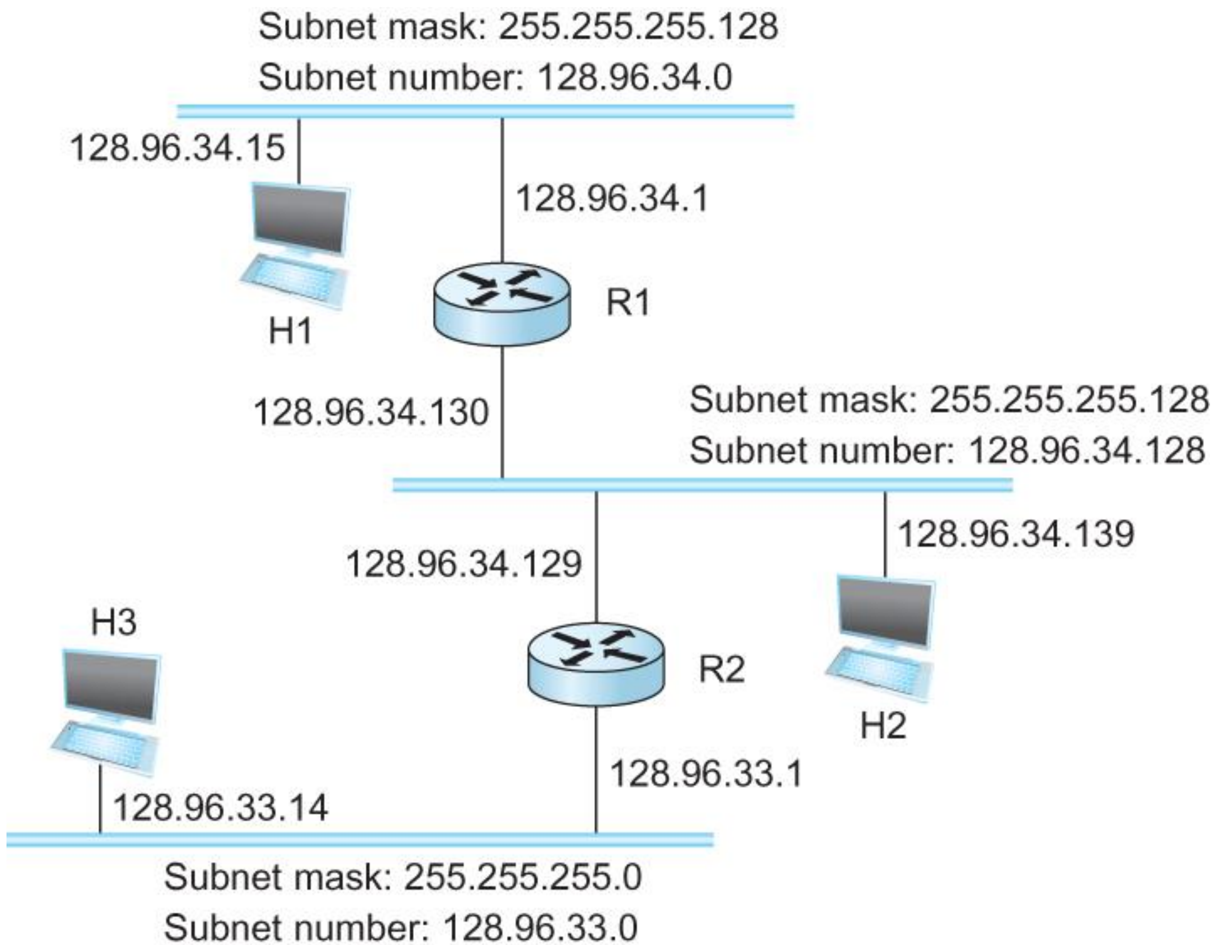

Subnetting: Example: Forwarding Table

- Forwarding Table at Router R1

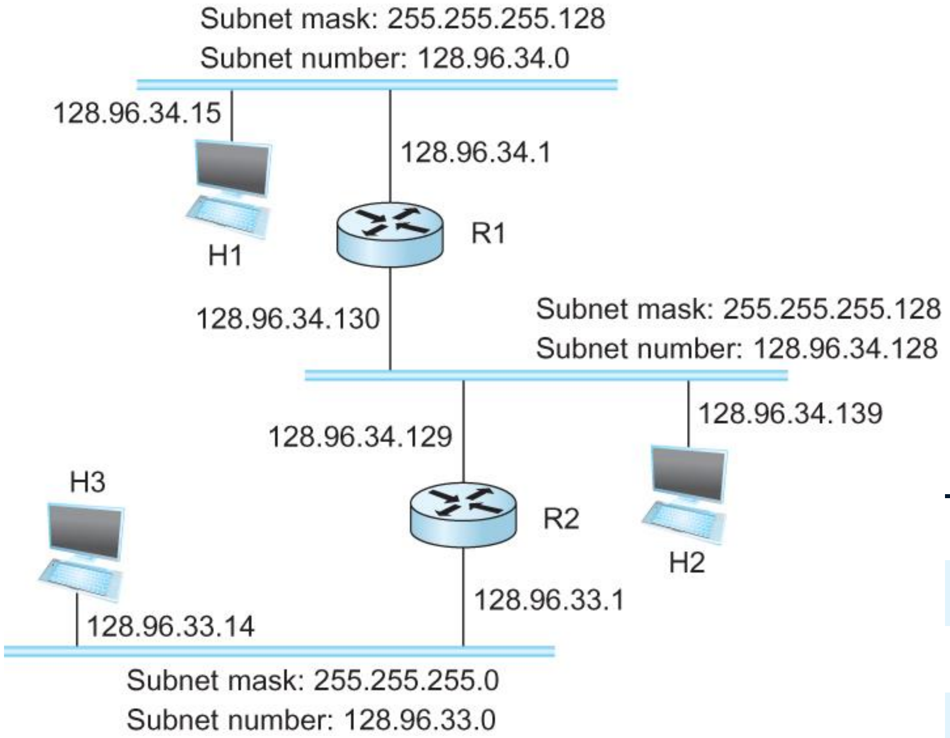

SubnetNumber	SubnetMask	NextHop
128.96.34.0	255.255.255.128	Interface 0
128.96.34.128	255.255.255.128	Interface 1
128.96.33.0	255.255.255.0	R2

Subnetting: Discussion

- Would use a default router if nothing matches
- Subnet masks do not have to align with a byte boundary
- Subnet masks need **not** to be contiguous 1's
  - 255.255.1.0 is OK
    - 11111111 11111111 00000001 00000000
    - What is subnet number of IP address 128.96.34.1?

      10000000 01100000 00100010 00000000 &
      11111111 11111111 00000001 00000000 →
      10000000 01100000 00000000 00000000 →
      128.96.0.0 → can not directly tell from the IP address
  - In practice, use contiguous 1's
- Multiple subnets can be on a single physical network
- Subnets not visible from the rest of the Internet

Subnetting: Discussion

- How do you tell whether an IP address is on a given subnet?

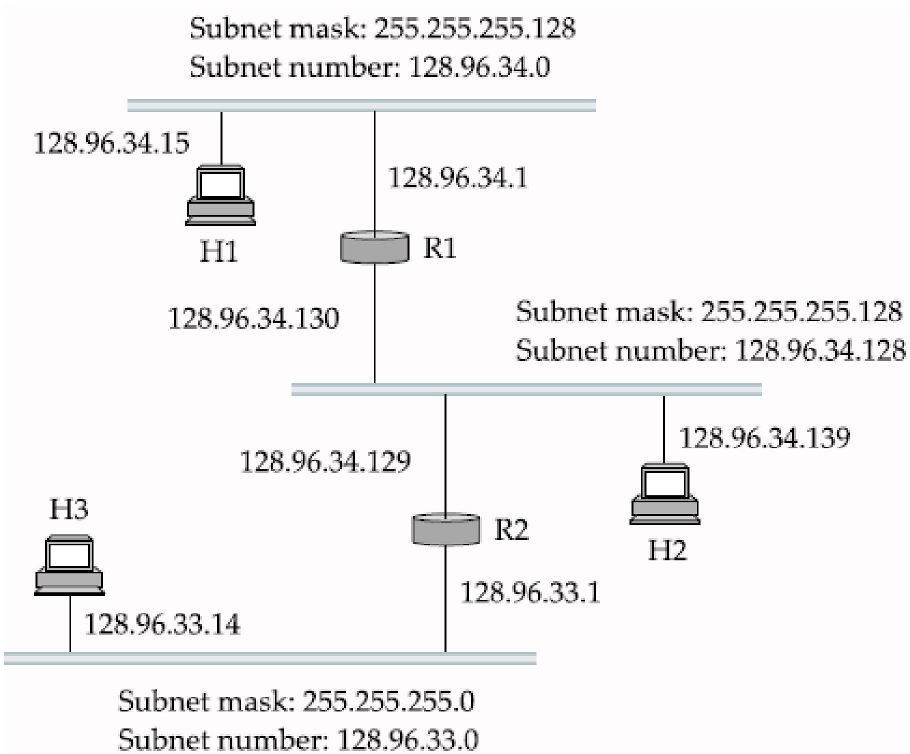

SubnetNumber	SubnetMask	NextHop
128.96.34.0	255.255.255.128	Interface 0
128.96.34.128	255.255.255.128	Interface 1
128.96.33.0	255.255.255.0	R2

Forwarding Algorithm

```
D = destination IP address
for each entry < SubnetNum, SubnetMask, NextHop>
  in the forwarding table
   D1 = SubnetMask & D
   if D1 = SubnetNum
      if NextHop is an interface
         deliver datagram directly to destination
      else
         deliver datagram to NextHop (a router)
```

Exercise 3

- The router in an IPv4 network uses subnetting. State to what next hop the IP packets addressed to each of the following destinations will be delivered

(a) 128.96.171.92
(b) 128.96.167.151
(c) 128.96.163.151
(d) 128.96.169.192
(e) 128.96.165.121

Table 3.19 Routing Table for Exercise 56

SubnetNumber	SubnetMask	NextHop
128.96.170.0	255.255.254.0	Interface 0
128.96.168.0	255.255.254.0	Interface 1
128.96.166.0	255.255.254.0	R2
128.96.164.0	255.255.252.0	R3
⟨default⟩		R4

Classless Addressing

- Represent network number with a single pair

```
<length, value>
```

- All routers must understand CIDR addressing

Classless Addressing: Notation

- Also called CIDR addressing (**C**lassless **I**ntra**d**omain **R**outing addressing)
- Convention
  - Place a /X after the prefix where X is the prefix length in bits
- Example
  - 20-bit prefix for all the networks 192.4.16 through 192.4.31: 192.4.16/20
  - A single class C network number, which is 24 bits long: 192.4.16/24

Handling Classless Addresses

- How do the routing protocols handle this classless addresses
  - It must understand that the network number may be of any length
  - Requires to hand out blocks of addresses that share a common prefix
  - Revising IP forwarding algorithm

IP Forwarding Revisited

- Original *assumptions* in IP forwarding mechanism
  - It can find the network number from destination IP address in a packet
  - Then look up that number in the forwarding table
- Need to *change* this assumption in case of CIDR
  - How? (next slide)

IP Forwarding in CIDR

- Prefixes may be of any length
  - Prefixes in the forwarding tables may overlap
- Some addresses may match more than one prefix
  - Example
    - Both 171.69/16 and 171.69.10/24 may coexist in the forwarding table of a single router
    - A packet destined to 171.69.10.5 clearly matches both prefixes.
  - Principle of “longest match”
    - A packet destined to 171.69.10.5 matches prefix 171.69.10/24
    - A packet destined to 171.69.20.5 matches 171.69/16

Exercise 4

- State to what next hop the IP packets addressed to each of the following destinations will be delivered

(a) C4.4B.31.2E

(b) C4.5E.05.09

(c) C4.4D.31.2E

(d) C4.5E.03.87

(e) C4.5E.7F.12

(f) C4.5E.D1.02

Table 3.21 Routing Table for Exercise 73

Net/MaskLength	Nexthop
C4.5E.2.0/23	A
C4.5E.4.0/22	B
C4.5E.C0.0/19	C
C4.5E.40.0/18	D
C4.4C.0.0/14	E
C0.0.0.0/2	F
80.0.0.0/1	G

Benefits of Classless Addressing

- Addressing two scaling challenges
  - Potential exhaustion of the 32-bit address space
  - The growth of backbone routing table as more and more network numbers need to be stored in them (route aggregation)
- Address assignment efficiency
  - Arises because of the IP address structure with class A, B, and C addresses
  - Forces us to hand out network address space in fixed-size chunks of three very different sizes
    - A network with two hosts needs a class C address:
      - Address assignment efficiency = 2/255 = 0.78
    - A network with 256 hosts needs a class B address
      - Address assignment efficiency = 256/65535 = 0.39

Classless Addressing: Address Assignment: Example

- Consider an organization has 16 class C network numbers.
- Instead of handing out 16 addresses at random, hand out a block of contiguous class C addresses
- Suppose we assign the class C network numbers from 192.4.16 through 192.4.31
- Observe that top 20 bits of all the addresses in this range are the same (11000000 00000100 0001)
  - We have created a 20-bit network number (which is in between class B network number and class C number)
- Requires to hand out blocks of class C addresses that share a common prefix (sometimes, called supnetting)

Route Aggregation with CIDR: Example

- Route aggregation with CIDR

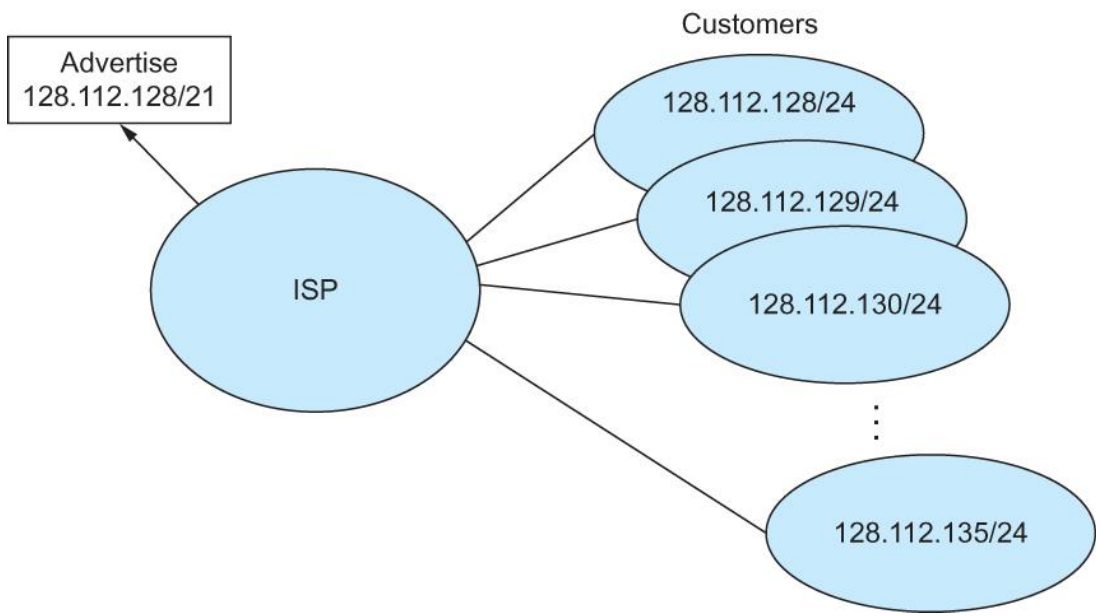

Summary

- Global addressing scheme
- Best effort service model and datagram forwarding
  - Defining a network
    - Legacy IPv4 address classes
    - Subnetting (supernetting)
    - Classless addressing (CIDR addressing)