SUBNETS, CIDR, AND CLASSLESS ADDRESSING

Internet Protocols

CSC / ECE 573

Fall, 2005

N. C. State University

Announcements

No office hours tomorrow (Wednesday) – out of town

No class on Thursday – Fall break!

Midterm exam next Tuesday! Study guide is online, as well as old exams

HW3 Part 1 solutions will be online today

Today's Lecture

I. Problems with Classful Addresses

II. Subnetting

III. Supernetting

IV. CIDR

V. IP Routing example

CLASSFUL ADDRESS PROBLEMS

Classful Addresses (Review)

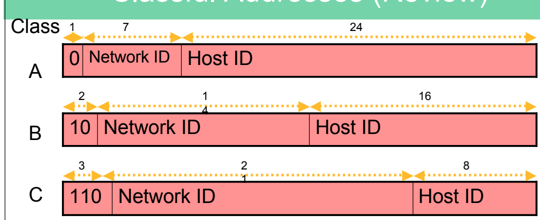

Class	Potential Number of Networks	Potential Number of Hosts Per Network
A	2^7 (128)	2^{24} (16M)
B	2^{14} (16K)	2^{16} (64K)
C	2^{21} (2M)	2^8 (256)

Classful Addressing "Issues"

1. "Three Bears" problem
   - class A (16M hosts) too large for most organizations, and there aren't very many of them (128)
   - class C (254 hosts) too small for most
   - class B (65,534 hosts) is "just right", but there aren't enough of them (16,384)

"Flat" host ID part of the address means very large routing tables within an organization

- e.g., one network with 50,000 hosts → each router has to know 50,000 routes?!

A Hypothetical Solution: Redefine Class C

- E.g., use 19 bits for network ID, 10 bits for host ID
  - result: 512K such networks, each with 1022 host addresses
- Problem: not backwards compatible with original addressing scheme
  - i.e., old class B and C address space allocations would have to be reclaimed and reassigned

An Actual Solution: Assign Multiple Networks

- Assign one organization multiple class C networks
  - E.g., an organization having 1000 hosts would get 4 class C networks (4 * 254 host addresses = 1016 > 1000)
- Problem: no longer have one network address for one organization
  - increases routing table sizes for routing between organizations
  - is that serious?
- Doesn't fix massive size of class A networks

Requirements for Address Modifications

- Any solution should be backwards compatible with the original addressing scheme
- Each organization should be free to interpret the "host ID" part of its addresses any way it wishes
  - only the "network ID" part of addresses is used by other organizations
- If there is a new interpretation of the "host ID"…
  - all hosts and routers in the organization's network must use the new addressing scheme
  - the internal routing protocol must support this new scheme

SUBNETTING

Subnet Addressing (RFCs 950, 1122)

- Particularly useful for class B sized networks
- Example

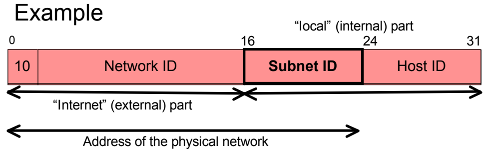

- Result: three-level hierarchical addressing
  - good for routing inside an organization (reduces the size of internal routing tables)
  - network structure not revealed to outsiders, since structure of "local" part only locally known

Subnet Addressing (cont'd)

- Result: large organizations may have many separate LANs, sharing a single network address
  - network administrators must decide how many bits to allocate to the subnet ID and host ID
- Hosts must support this capability
- Does not change classful addressing
  - only affects interpretation of local (internal) part of address

Example: Without Subnetting

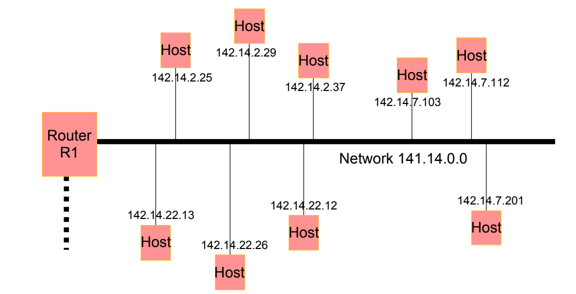

Example: With Subnetting

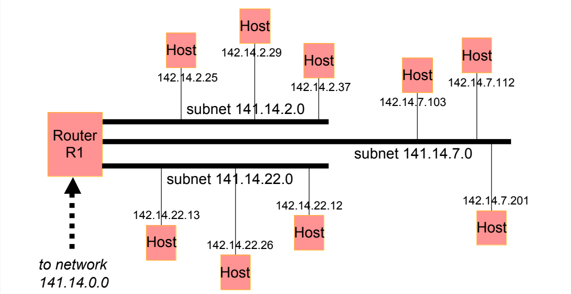

How Long is the Subnet ID?

- *Masking*: the process of extracting a portion of the IP address
  - *subnet masking*: extracting the subnet ID
  - masks are 32 bits long for IPv4 addresses
- If a bit in a subnet mask = …
  - 1 → this bit position is part of the "physical network ID" (network ID + subnet ID)
  - 0 → this bit position is part of the host ID
- How split "physical network ID" into network and subnet parts?

Masking (Bitwise-AND Operation)

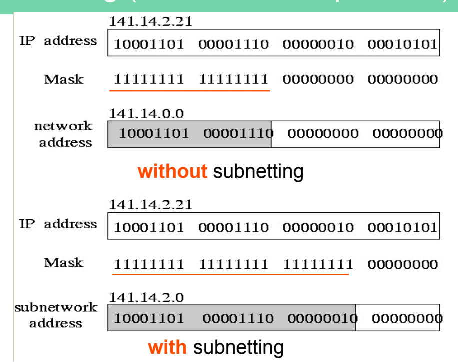

Example

- Step 1: determine network ID

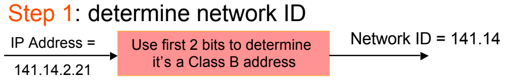

- Step 2: determine physical network address

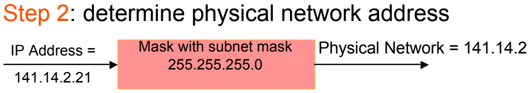

- Step 3: determine subnet ID

Physical Network = 141.14.2, Class B address = 141.14 → Subnet ID = 2

Questions

- How can you tell if a destination IP address is…
  1. On the same subnet as you???
  2. On the same network as you???

Special Addresses (Subnetting)

Net ID	Subnet ID	Host ID	Description
(normal)	(normal)	All 1's	Broadcast to a specific subnet

Possible Subnet Sizes for a Class B Network

# of bits for subnet ID	# of subnets possible	# of hosts / subnet
0 (i.e., subnetting not used)	1	$2^{16} - 2$
1	2	$2^{15} - 2$
2	2^2	$2^{14} - 2$
3	2^3	$2^{13} - 2$
. . .	. . .	. . .
13	2^{13}	$2^3 – 2 = 4$
14	2^{14}	$2^2 – 2 = 2$

What happened to **15** bits for subnet ID???

What happened to **16** bits for subnet ID???

Subnet Routing Recommendations

1. Use simple subnet masks (sequence of all 1's followed by string of all 0's)
   – e.g., 1111111111111111111110000111110000 is not OK
2. Use the same mask for all subnets of the network
3. All subnets should be physically contiguous (connected)
   – drawbacks?

Forwarding Decisions with Subnetting

- Reminder: a forwarding table consists of <*key*, *next IP address*> entries
- **key** in forwarding table entries needs to consist of
  1. a mask
  2. a network identifier
- Examples
  – for network 152.14.0.0…
    mask = 255.255.0.0, identifier = 152.14.0.0
  – for subnet 152.14.22.0…
    mask = 255.255.255.0, identifier = 152.14.22.0

Forwarding Decisions (cont'd)

- Steps in determining a forwarding table match (in order of the table entries)
  1. AND the packet's destination IP address with the mask
  2. compare the result to the identifier to see if there is a match
  3. use the first match found

Forwarding Table Example

Mask	Destination Identifier	Next hop IP address
255.255.255.255	152.14.51.15	…
255.255.0.0	152.14.0.0	…
0.0.0.0	0.0.0.0	…

– Which entry does 152.14.51.15 match?
– Which entry does 152.14.51.129 match?
– Which entry does 152.14.1.36 match?
– Which entry does 152.12.1.230 match?
– Why choose this ordering of the routing table entries?

Remember This Algorithm?

```
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
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
```

The New Forwarding Algorithm

```
extract destination address Hd
for forwarding table entries i=1..n {
    if (Hd AND mask[i]) == destination[i] {
        forward datagram to specified next_hop[i]
        exit
    }
}
/* we don't know how to get to Hd */
discard the datagram and declare routing error
```


Variable-Length Subnetting

- Site with a Class C address 192.15.34.0, needs five subnets with 60, 60, 60, 30, and 30 hosts
- Problems
  - cannot use subnet mask 255.255.255.192, since it allows only 4 subnets, each with (2^8/ 4) - 2 = 62 hosts
  - cannot use subnet mask 255.255.255.224, since it allows 8 subnets, each with only (2^8 / 8) - 2 = 30 hosts
- Solution: use two masks of different lengths
  - router requirements RFC requires support for this
  - however, not all routing protocols provide the subnet mask along with the destination network ID

SUPERNETTING

Supernet Addressing

- Possible solution to "exhaustion" of Class B addresses → use multiple Class C addresses instead
- Assign organizations blocks of 2^n contiguous Class C addresses
  - ex.: if 2000 addresses (hosts) are needed, allocate a block of 2048 addresses (= 8 contiguous Class C networks)
- Problem: no longer have one network address for each organization
  - inter-organization routing tables will be larger

Supernetting Example

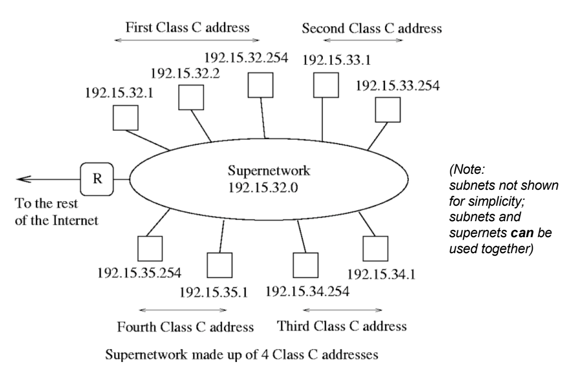

CLASSLESS INTERDOMAIN ROUTING (CIDR)

Classless InterDomain Routing (CIDR) (RFC1519)

- Idea: allocate IP addresses in a block of size 2^n, for n=2,3,….
  - "In a block" means IP addresses must be allocated contiguously (i.e., share the same high-order bits)
- A block is *summarized* by a single routing table key of the form <*network_address*, *CIDR_mask*>
  - *network_address*: the smallest (first) address in the block
  - *CIDR_mask*: the size of the address block
- No longer need classful addresses; the network size is now explicit
  - but, backwards compatible with classful addresses

Classless InterDomain Routing (CIDR) (RFC1519) (cont'd)

- CIDR_mask: number of 0's indicates value of *n*
  - `11111111 11111111 11111000 00000000`: *n* = 11
  - `11111111 00000000 00000000 00000000`: *n* = 24 (i.e., class A address)
- More conveniently for people: *"/32-n"* to indicate number of 1's in the mask
  - `11111111 11111111 11111000 00000000`: /21
  - `11111111 00000000 00000000 00000000`: /8

CIDR Masking + Subnet Masking

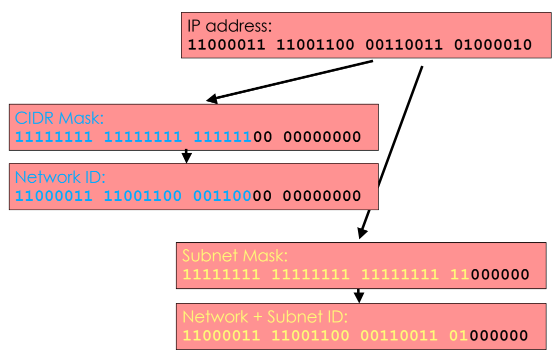

Address Allocation and Masks

- Ex.: allocate a block of 2^{23} (= 8M) addresses, starting at address 194.0.0.0

Address	Dotted Decimal	Binary
Lowest	194.0.0.0	11000010 00000000 00000000 00000000
Highest	194.127.255.255	11000010 01111111 11111111 11111111
Mask	255.128.0.0 (/9)	11111111 10000000 00000000 00000000

- How can you tell what the mask should be
  - knowing the number of addresses?
  - knowing the lowest and highest addresses?

Regional Aggregation of Class C Addresses

- All sites in a region (e.g., Europe) have a common prefix

Lowest	Highest	Region
194.0.0.0	195.255.255.255	Europe
198.0.0.0	199.255.255.255	North America
200.0.0.0	201.255.255.255	Central and South America
202.0.0.0	203.255.255.255	Asia and the Pacific

Regional Allocation of Class C Addresses (cont'd)

- CIDR summarizes these 32M addresses into one entry
  - i.e., only a single entry needed in most U.S. routers for all European destinations
- Once the packet gets to Europe, more detailed routing tables are needed to reach the specific network / host

Example: Allocating Classless Addresses

- Three sites X, Y, and Z in Europe ask for 2048, 1024, and 4096 addresses, in that order, starting from 194.24.0.0.
  Result:

Site	# of Addresses	Lowest	Highest	Mask
X	2048 = 2^{11}	194.24.0.0	194.24.7.255	255.255.248.0 (/21)
Y	1024 = 2^{10}	194.24.8.0	194.24.11.255	255.255.252.0 (/22)
Z	4096 = 2^{12}	194.24.16.0	194.24.31.255	255.255.240.0 (/20)

- Why is there a gap between site Y's allocation and site Z's allocation???

Allocating Classless Addresses (cont'd)

- Result: routers all over Europe have the following 3 entries added to their routing tables

Network Address	11000010 00011000 00000000 00000000 (194.24.0.0)
Mask	11111111 11111111 11111000 00000000 (255.255.248.0, or /21)
Network Address	11000010 00011000 00001000 00000000 (194.24.8.0)
Mask	11111111 11111111 11111100 00000000 (255.255.252.0, or /22)
Network Address	11000010 00011000 00010000 00000000 (194.24.16.0)
Mask	11111111 11111111 11110000 00000000 (255.255.240.0, or /20)

Example: CIDR Forwarding

- A packet arrives at a router with destination IP address = 194.24.17.4; where should it be forwarded?
  `11000010 00011000 00010001 00000100`
  - ANDed with Site 1 mask =
    `11000010 00011000 00010000 00000000`
    match site 1 address?
  - ANDed with Site 2 mask =
    `11000010 00011000 00010000 00000000`
    match site 2 address?
  - ANDed with Site 3 mask =
    `11000010 00011000 00010000 00000000`
    match site 3 address?

Longest Prefix Matching

- A forwarding table entry with a longer mask (i.e, more 1's) is called *more specific* than another entry with a shorter mask
- There is a dilemma in classless network masking: address matching ambiguity
  - X = `10110000 00000000 00000000 00000000 (/4)`
    Y = `10110100 00000000 00000000 00000000 (/6)`
  - destination IP address
    `10110100 00000000 00000000 00110000`
    will match both of these entries!
- For forwarding purposes, the longest prefix (most specific match) wins! Ordering doesn't matter

The CIDR Forwarding Algorithm

```
Extract destination address Hd
Use mask to find longest matching prefix
if (match found)
        forward datagram to specified next hop
else /* we don't know how to get to Hd */
        discard the datagram and declare routing error
```

- Unified routing table lookup!
- Questions
  - How deal with "default" addresses?
  - How deal with "host-specific" addresses?

CIDR Blocks Reserved for "Private" Networks

- Useful for intranets, testing, etc.
- Should never see one of these addresses in IP packets in the public Internet

Prefix	Lowest Address	Highest Address
10 /8	10.0.0.0	10.255.255.255
169.254 /16	169.254.0.0	169.254.255.255
172.16 /12	172.16.0.0	172.31.255.255
192.168 /16	192.168.0.0	192.168.255.255

FORWARDING EXAMPLE

Example Network Configuration

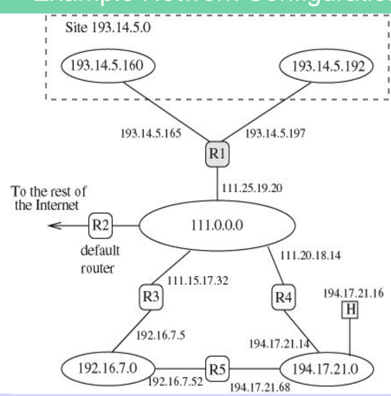

Routing Table for Router R1

Mask	Destination	Next Hop	Flags
255.0.0.0	111.0.0.0	-----	U
255.255.255.224	193.14.5.160	-----	U
255.255.255.224	193.14.5.192	-----	U
....	....	....	....
255.255.255.255	194.17.21.16	111.20.18.14	UGH
255.255.255.0	192.16.7.0	111.15.17.32	UG
255.255.255.0	194.17.21.0	111.20.18.14	UG
0.0.0.0	0.0.0.0	111.30.31.18	UG

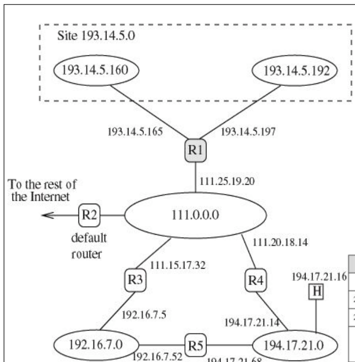

- What happens when R1 receives a packet with destination:
  - 192.16.7.14?
  - 193.14.5.176?
  - 200.34.12.34?

Mask	Destination	Next Hop	Flags
255.0.0.0	111.0.0.0	-----	U
255.255.255.224	193.14.5.160	-----	U
255.255.255.224	193.14.5.192	-----	U
....	....	....	....
255.255.255.255	194.17.21.16	111.20.18.14	UGH
255.255.255.0	192.16.7.0	111.15.17.32	UG
255.255.255.0	194.17.21.0	111.20.18.14	UG
0.0.0.0	0.0.0.0	111.30.31.18	UG

Summary

1. Classful addressing is simple, but too restrictive
   - Supernetting is a cheap (but not very clean) fix
2. Subnetting creates a third level of address hierarchy
   - useful for routing inside an organization
   - requires the use of a mask
3. CIDR decreases routing table sizes, and increases flexibility in allocating addresses
4. Both subnetting and CIDR require the use of masks

Next Lecture

- Routing, Part I