

CSc 450/550

Computer Networks

Internet Addressing and Routing

Jianping Pan
Summer 2006

7/10/06

CSc 450/550

1

Review

- IP
 - IP header
 - addressing
 - class, classless
 - routing
 - link state, distance vector
- What really happens on the Internet?!

7/10/06

CSc 450/550

2

Internet addressing

- Currently IPv4
 - IPv6 designed and implemented
 - 128-bit address space
 - 6Bone: experimental IPv6 networks over IPv4
- Address allocation
 - was class-based allocation
 - class A, B, and C
 - now hierarchical allocation with CIDR
 - classless

7/10/06

CSc 450/550

3

Hierarchical addressing

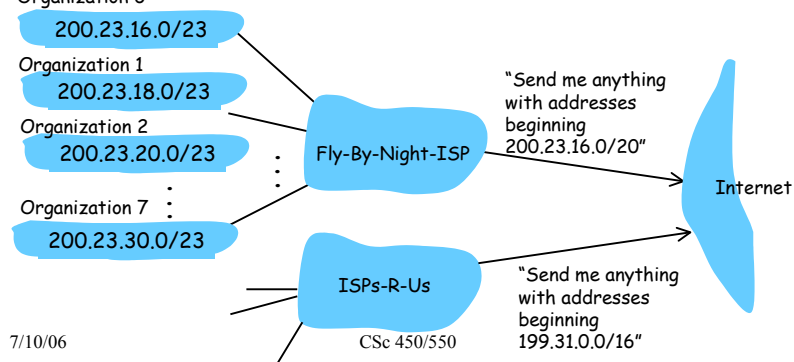
ISP's block 11001000 00010111 00010000 00000000 200.23.16.0/20

Organization 0 11001000 00010111 00010000 00000000 200.23.16.0/23

Organization 1 11001000 00010111 00010010 00000000 200.23.18.0/23

... ..

Organization 7 11001000 00010111 00011110 00000000 200.23.30.0/23



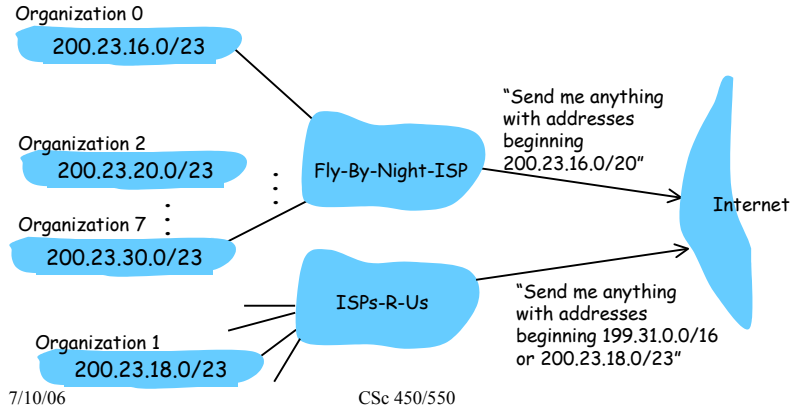
7/10/06

CSc 450/550

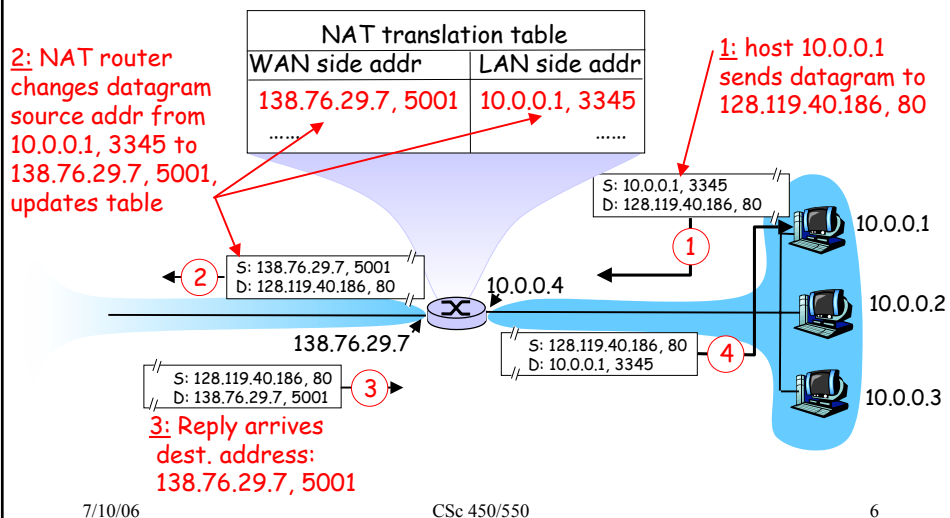
4

Hierarchical addressing: more

- Longest-prefix match



Network address translation



NAT: more

- NAT: mapping created by outgoing packets
 - (srcIP, srcPt, dstIP, dstPt) => (natIP, natPt, dstIP, dstPt)
 - full cone, IP restricted, port-restricted NAT
 - symmetric NAT
- Issues with NAT
 - session initiated by incoming packets
- NAT traversal
 - see <http://www.cs.uvic.ca/~pan/seng490>

7/10/06

CSc 450/550

7

Internet control message protocol

- Ping
 - ICMP echo request
 - ICMP echo reply
- Traceroute

Message type	Description
Destination unreachable	Packet could not be delivered
Time exceeded	Time to live field hit 0
Parameter problem	Invalid header field
Source quench	Choke packet
Redirect	Teach a router about geography
Echo request	Ask a machine if it is alive
Echo reply	Yes, I am alive
Timestamp request	Same as Echo request, but with timestamp
Timestamp reply	Same as Echo reply, but with timestamp

- UDP with small, increasing TTL
 - ICMP TTL expires (type 11, code 0)
- UDP with unusual port number
 - ICMP destination port unreachable (type 3, code 3)

7/10/06

CSc 450/550

8

Routing information protocol

- RIP: using distance vector routing algorithm
 - included in BSD Unix in 1982; max hops: 15
- Distance vector
 - exchanged between neighbors every 30s
 - up to 25 destinations within an AS (UDP 520)
 - if no advertisement for 180s: neighbor is dead
 - invalidate routes going through the neighbor
 - poisoned reverse to speed up “bad news”
 - infinite: 16 hops

7/10/06

CSc 450/550

9

Open shortest path first

- OSPF: using link state routing algorithm
 - link state dissemination
 - flooding, directly over IP
 - topology map at each node
 - Dijkstra’s algorithm at each node
- Hierarchical OSPF
 - intra-domain areas: backbone and areas
 - flooding in an area
 - area border routers

7/10/06

CSc 450/550

10

Border gateway protocol

- BGPv4: based on distance vector
 - the de facto inter-domain routing standard
 - heavily policy-influenced, over TCP
- Reachability information
 - “you can reach X through me”
- AS path
 - “with a path of AS numbers”
 - AS: autonomous system (e.g., ISP domains)

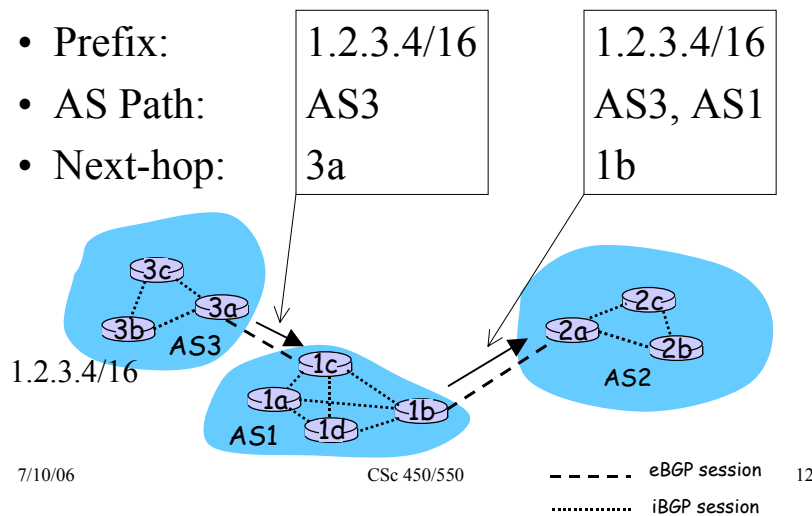
7/10/06

CSc 450/550

11

BGP advertisement

- Prefix: 1.2.3.4/16
 - AS Path: AS3
 - Next-hop: 3a
- Prefix: 1.2.3.4/16
 - AS Path: AS3, AS1
 - Next-hop: 1b



7/10/06

CSc 450/550

12

BGP: more

- BGP routing
 - longest-prefix match
 - e.g., 1.2.3.4/16 vs 1.2.3.4/24
 - multiple AS path
 - e.g., shortest AS path
 - multiple next-hop router
 - e.g., nearest next-hop router; “hot potato”
- BGP security
 - prefix authenticity; AS path authenticity

7/10/06

CSc 450/550

13

This lecture

- Internet addressing and routing
 - addressing: hierarchical, NAT
 - control: ICMP
 - routing: RIP, OSPF, BGP
- Explore further
 - advanced computer networks (topics course)
 - NAT traversal
 - routing security

7/10/06

CSc 450/550

14

Next lecture

- The link layer
 - read CN chapter 4