CS 356: Computer Network Architectures

Lecture 16: IPv6, IP tunnels, and (brief) Midterm Review

[PD] chapter 4.1.3, 3.2.9

Xiaowei Yang
xwy@cs.duke.edu

History

- In early 90s, IPv4 is running out of addresses
- Changing to a larger address space requires many changes
- IETF solicited other desired features
- Chose one for IPv6 (**RFC 2460**) in 1998

IPv6 features

- Large address space (128-bit)
- Hierarchical addressing and routing
- Autoconfiguration
- Built-in security
- Better support for QoS
- New protocols for neighboring node interactions
- Extensibility

Addressing

- 128-bit addresses
  - 2^{128}
- "if the earth were made entirely out of 1 cubic millimetre grains of sand, then you could give a unique [IPv6] address to each grain in 300 million planets the size of the earth"
- http://en.wikipedia.org/wiki/IP_address
- Or, using a more earthly analogy:
- "The optimistic estimate would allow for 3,911,873,538,269,506,102 addresses per square meter of the surface of the planet Earth." "IP Next Generation Overview"
- R. Hinden, Communications of the ACM, Vol. 39, No. 6 (June 1996) pp 61 - 71, ISSN:0001-0782 http://portal.acm.org/citation.cfm?coll=GUIDE&dl=GUIDE&id=228517

IPv6 Addresses

- Classless addressing/routing (similar to CIDR)
- Notation: x:x:x:x:x:x:x:x (x = 16-bit hex number)
  - contiguous 0s are compressed: 47CD::A456:0124
  - IPv6 compatible IPv4 address:
    ::FFFF:128.42.1.87

IPv6 addressing architecture

- RFC 4291
- All addresses are assigned to interfaces, not nodes

Types of IPv6 addresses

Address type	Binary prefix	IPv6 notation
Unspecified	00...0 (128 bits)	::/128
Loopback	00...1 (128 bits)	::1/128
Multicast	11111111	FF00::/8
Link-local unicast	1111111010	FE80::/10
Global unicast	Everything else	
Anycast	Allocated from unicast space	

Global Unicast Addresses

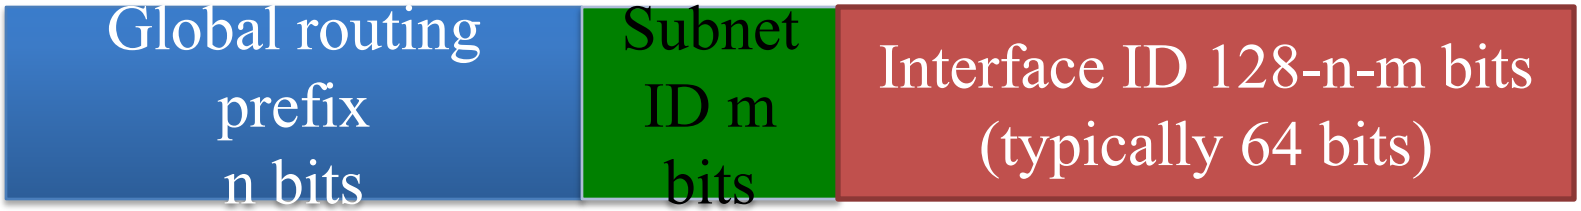

- For all unicast addresses, except those that start with the binary value 000, Interface IDs are required to be 64 bits long
  - Can be derived from 48-bit Ethernet address

IPv6 Header

- 40-byte “base” header
- Extension headers (fixed order, mostly fixed length)
  - fragmentation
  - source routing
  - authentication and security
  - other options

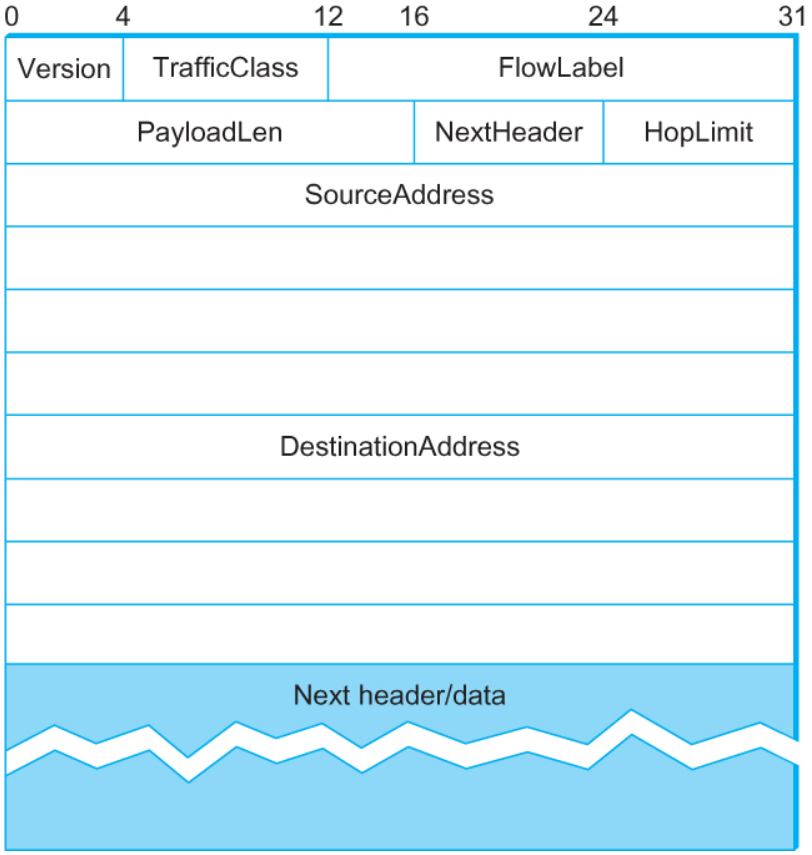

Autoconfiguration

- Link-local prefix + interface ID
- Routers advertise global prefixes

IPv6 Anycast Addresses

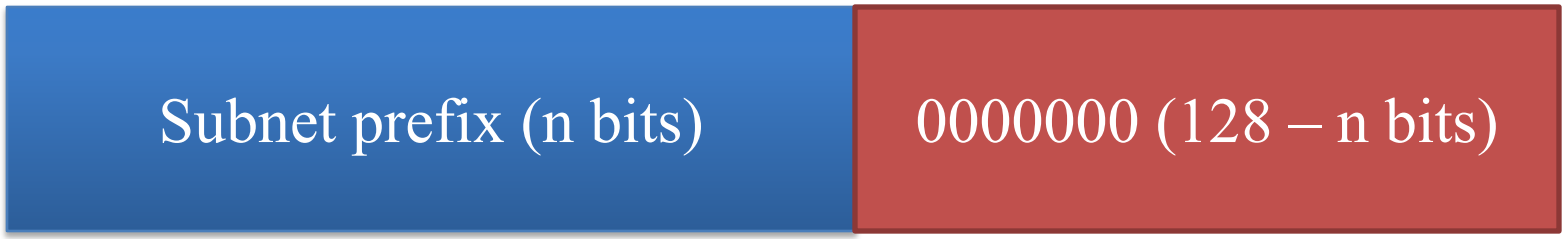

- Assigned to more than one interface
- All zero interface address
- Allocated from the unicast address space
- Ex: all root DNS servers

IP Tunnels

IP tunnels

- Tunnels
  - A technique used in many scenarios
    - VPN, IPv4-v6 transition, Mobile IP, Multicast, Non-IP forwarding, IPsec

What is a tunnel

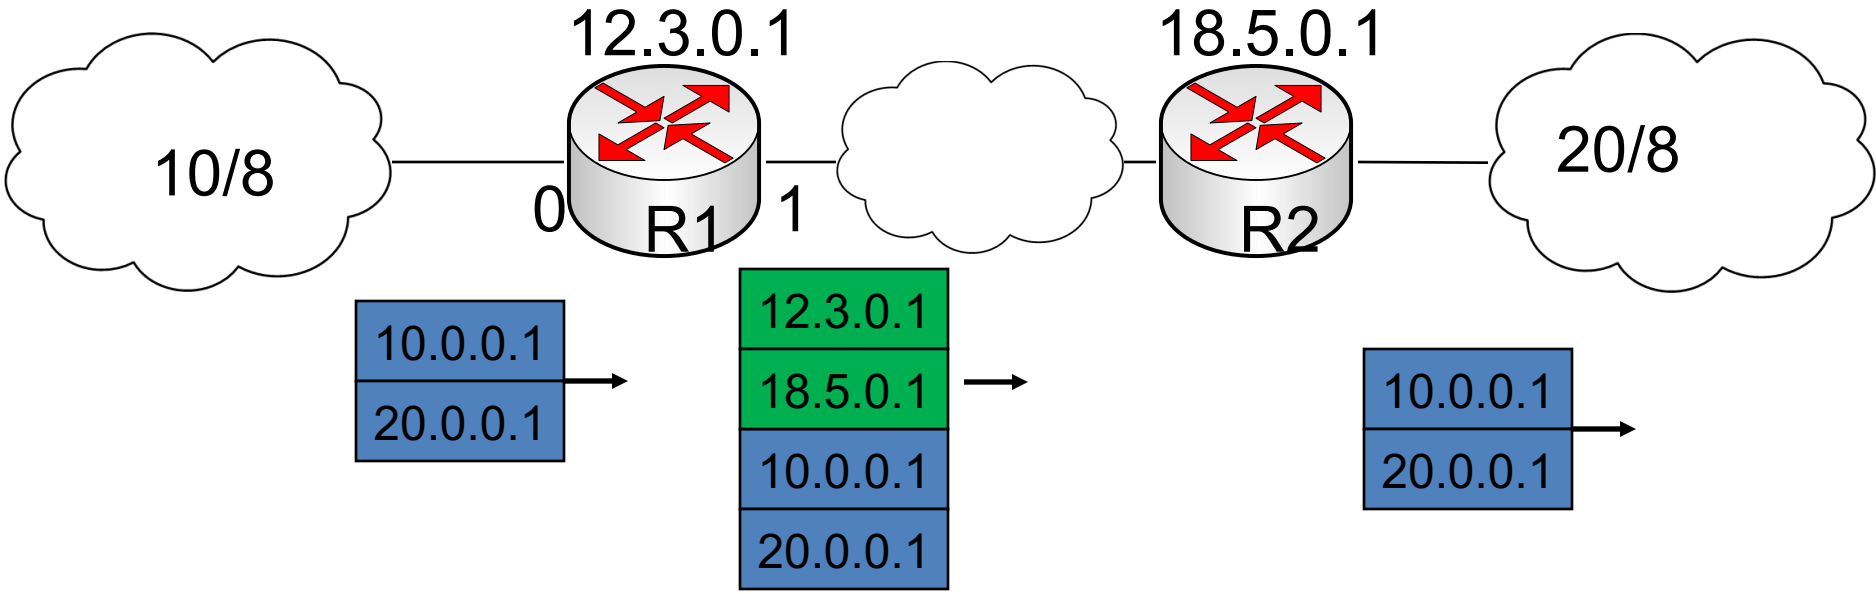

- A “pseudowire”, or a virtual point-to-point link
- The head router encapsulates a packet in an outer header destined to the tail router

Virtual interface

- A router adds a tunnel header for packets sent to a virtual interface

NetworkNum	nextHop
10/8	ether0
20/8	tun0
0/0	ether1

Tunnel applications

- Traversing a region of network with a different addressing format or with insufficient routing knowledge
- Building virtual private networks

FIGURE 3.26 An example of virtual private networks: (a) two separate private networks; (b) two virtual private networks sharing common switches.

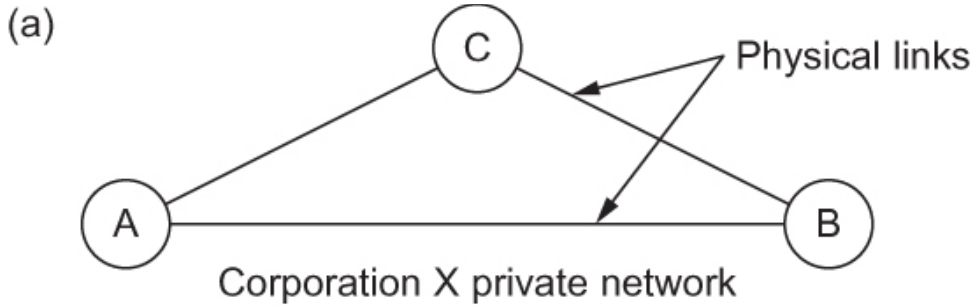

IPv4-v6 transition

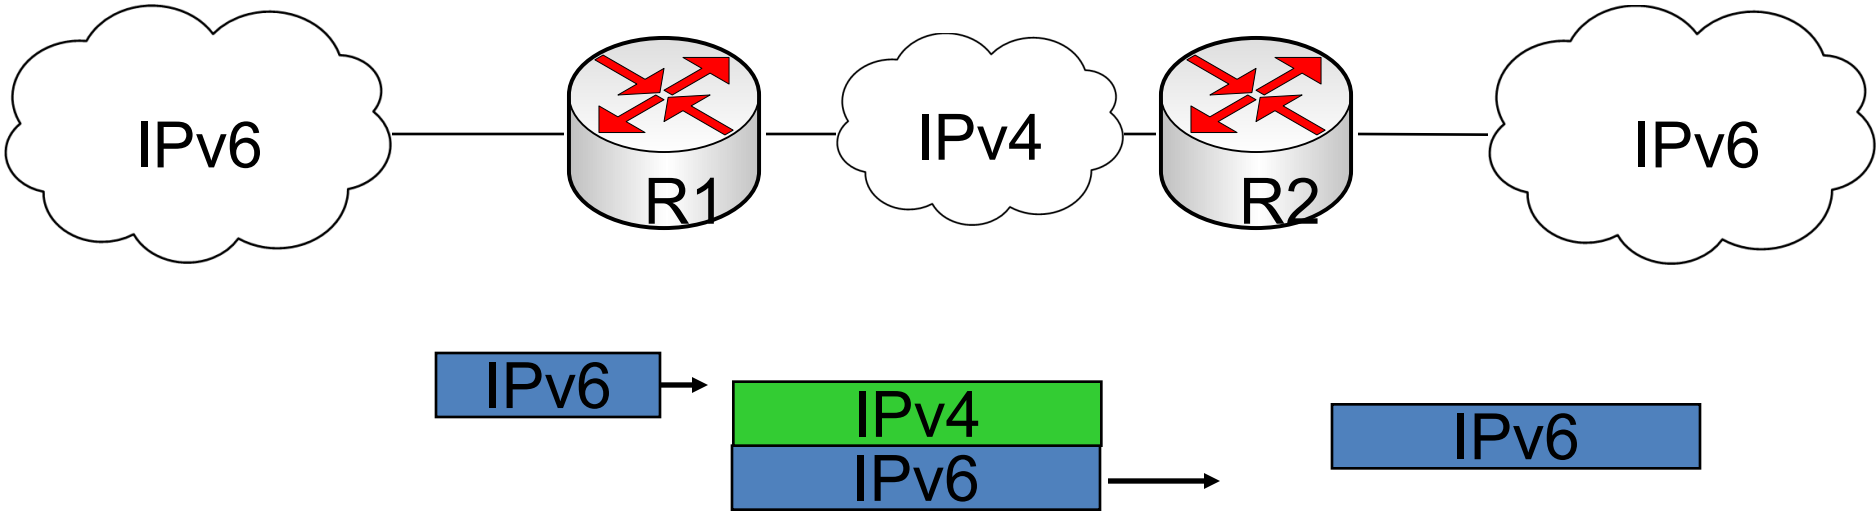

Generic Routing Encapsulation (GRE)

- Defined in RFC 2784 and updated by RFC 2890
- Can encapsulate any inner header

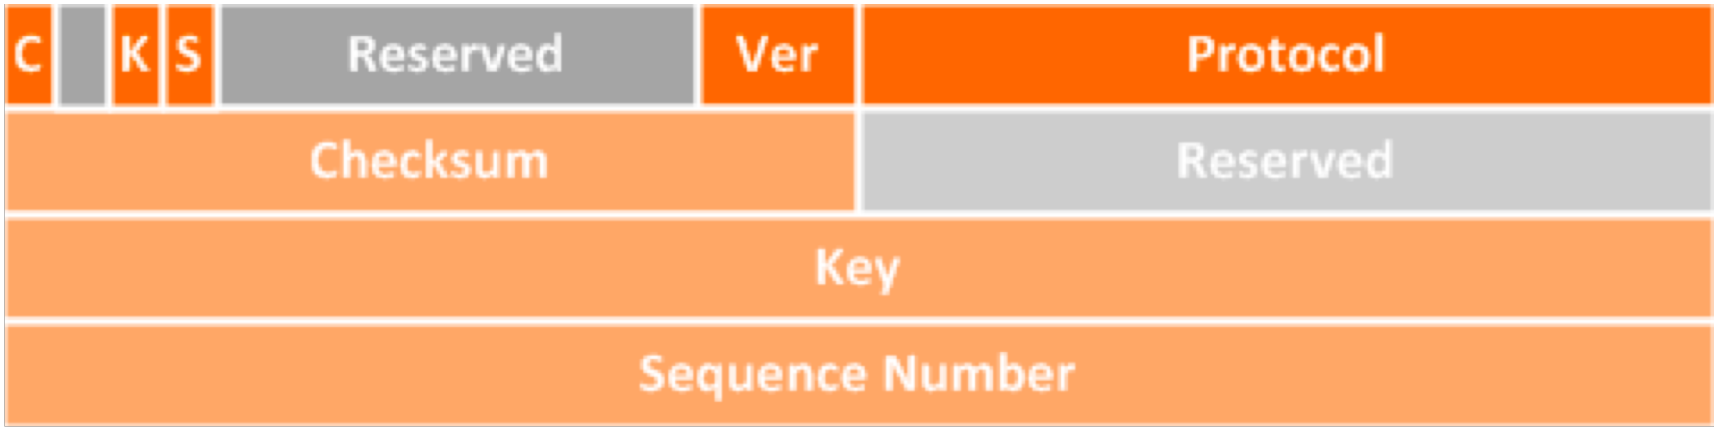

(brief) Midterm Review

Expectations

- Fundamental concepts
- Key algorithms / protocols

Midterm Policy

- Up to Feb 28's lecture
- Closed book/notes
- One page of your own note (letter-size)
  - Two sides notes are okay
- No Internet
- Calculator is allowed
- 75 mins

What we've learned

- Network architectures
  - Basic concepts, Internet architecture,
- Physical layer
  - Delay, bandwidth, and throughput
  - Delay bandwidth product
- Link layer
  - Coding/encoding, framing, error detection, reliable transmission
  - Multi-access links
  - Switching, bridges, ATM

What we've learned (cont.)

- Internetworking
  - Challenges, solutions
  - Classful vs classless IP addressing
  - IP forwarding, longest prefix lookup, ARP, ICMP
  - Dynamic routing protocols
    - Distance vector (RIP)
    - Link state (OSPF)
    - BGP
  - DHCP, and NAT

Common confusion

- KB, MB, etc.
- https://en.wikipedia.org/wiki/Kilobyte
- "In December 1998, the IEC addressed such multiple usages and definitions by creating prefixes such as kibi, mebi, gibi, etc., to unambiguously denote powers of 1024.[10] Thus the kibibyte, symbol KiB, represents $2^{10} = 1024$ bytes. These prefixes are now part of the International System of Quantities. The IEC further specified that the kilobyte should only be used to refer to 1000 bytes."

How to compute the size of sliding window

- Example from midterm

Construct the path to reach a destination

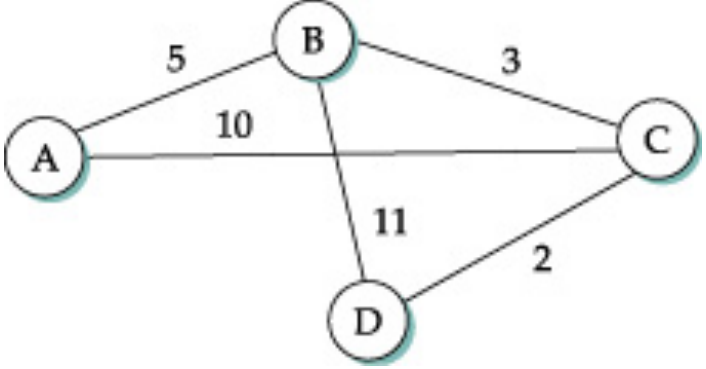

Step	Confirmed	Tentative
1	(D,0,-)	
2	(D,0,-)	(B,11,B), (C,2,C)
3	(D,0,-), (C,2,C)	(B,11,B)
4	(D,0,-), (C,2,C)	(B,5,C) (A,12,C)
5	(D,0,-), (C,2,C), (B,5,C)	(A,12,C)
6	(D,0,-),(C,2,C),(B,5,C)	(A,10,C)
7	(D,0,-),(C,2,C),(B,5,C), (A,10,C)	