Computer Networking

COMP 177 | Fall 2020 | University of the Pacific | Jeff Shafer

IPv4 - Part 1

Recap

Past Topics

- Overview of networking and layered architecture
- Wireshark packet sniffer
- Wired LAN and Ethernet
- Wireless LAN and WiFi

Today's Topics

- Internet Protocol v4

Why not just use Ethernet?

- Most computer systems use Ethernet networking
- Ethernet provides facilities to
  - Locate computers
  - Forward packets directly
  - Prevent loops
  - …
- What are the drawbacks of Ethernet for global communication?

Ethernet Drawbacks

- Locating computers
  - Do we really want to broadcast across the Internet?
- Preventing loops
  - Do we really want to rebuild an Internet-wide *spanning tree* whenever the topology changes?
  - Do we really want packets to live forever if loops remain?
- Unreachable computers
  - What happens if the destination is unreachable?
  - i.e., it doesn't exist, is turned off, is broken, …

Internet Protocol (IP)

- IP is the prominent *network layer* protocol
- IP provides for
  - *Universal addressing* of nodes using IP addresses
  - *Routing* packets using *routing protocols*
- IP headers include
  - *Source IP address*: the address of the ultimate sender
  - *Destination IP address*: the address of the ultimate receiver
- Routers use the destination IP address to determine next hop
- Using IP, a packet can be delivered from any arbitrary node to any other as long as the two ends are publicly available

Internet Protocol

- IP addresses are assigned to network interfaces (like link layer MAC addresses)
  - If a node has multiple network interfaces each with network layer service, that node needs multiple IP addresses
- IP addresses remain unchanged *end-to-end*
  - In contrast to link layer MAC addresses which change for every hop
  - Exceptions to the rule: NAT (Address translation)
- IP needs to be *scalable* to support the growing Internet
- IP is a *best-effort* service: no guarantee of delivery

IP for Interfaces

- IP addresses are assigned to *interfaces* (not machines)
- Each interface has both
  - MAC address (physical, *unchangeable*)
  - IP address (logical, *modifiable*)
- Hosts with multiple interfaces are called *multihomed*
  - Routers have at least two interfaces, each with an IP address
- Hosts usually have virtual interfaces as well
  - *Loopback* interface: used to talk between processes in the same machine
    - Can be used to test network applications locally on a host
  - *VPN* interfaces: Each end of VPN connection creates a virtual interface with an assigned IP address

An IP Datagram

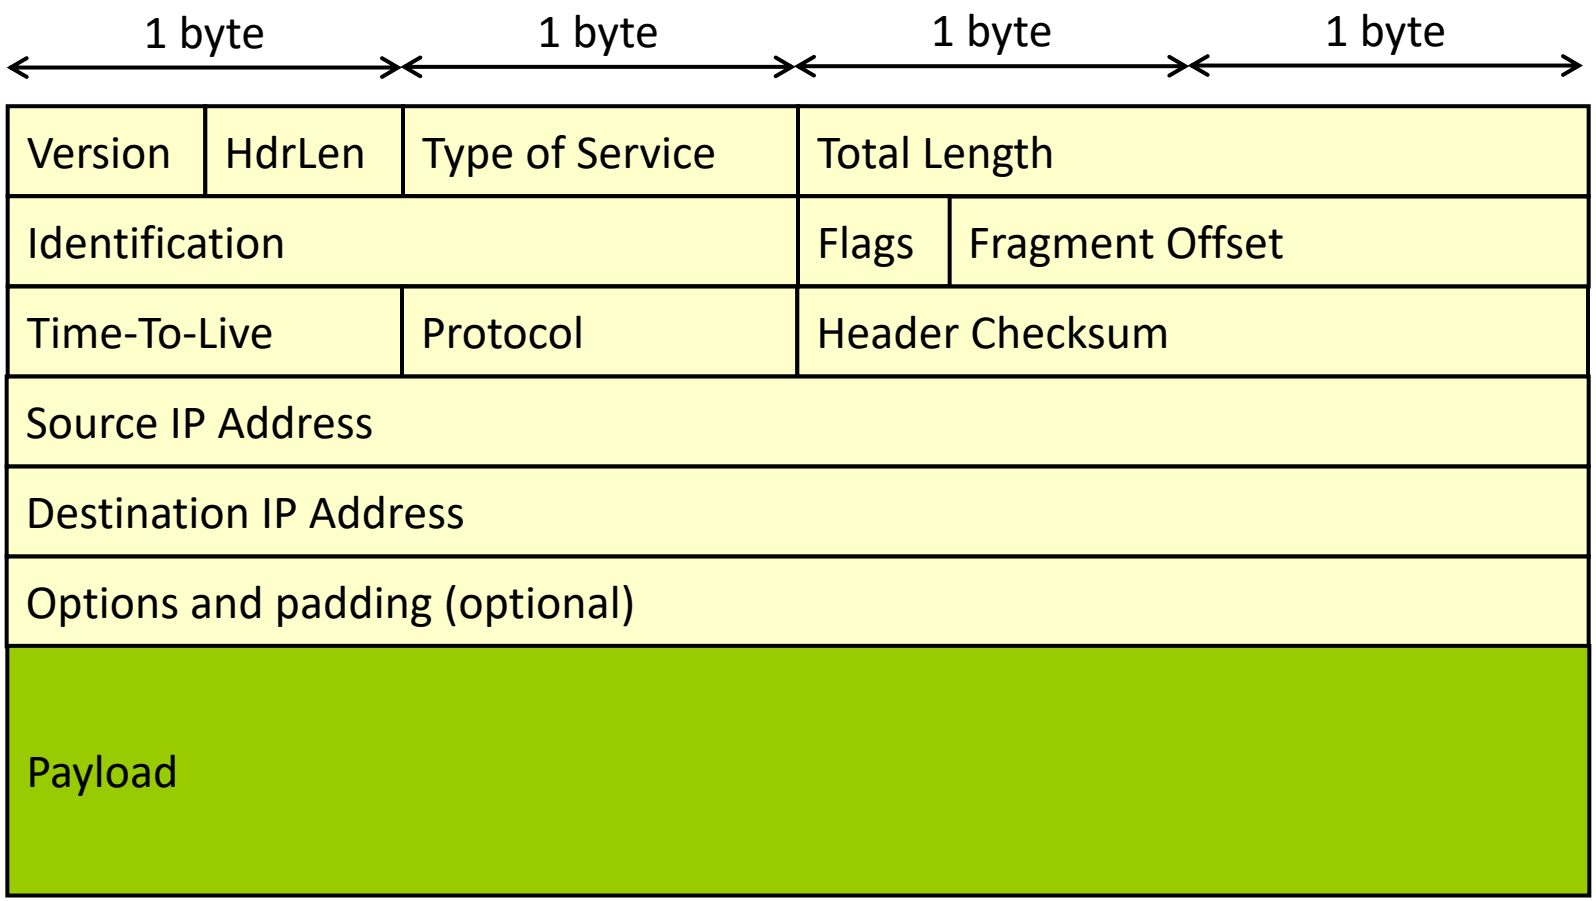

IPv4 Header Format

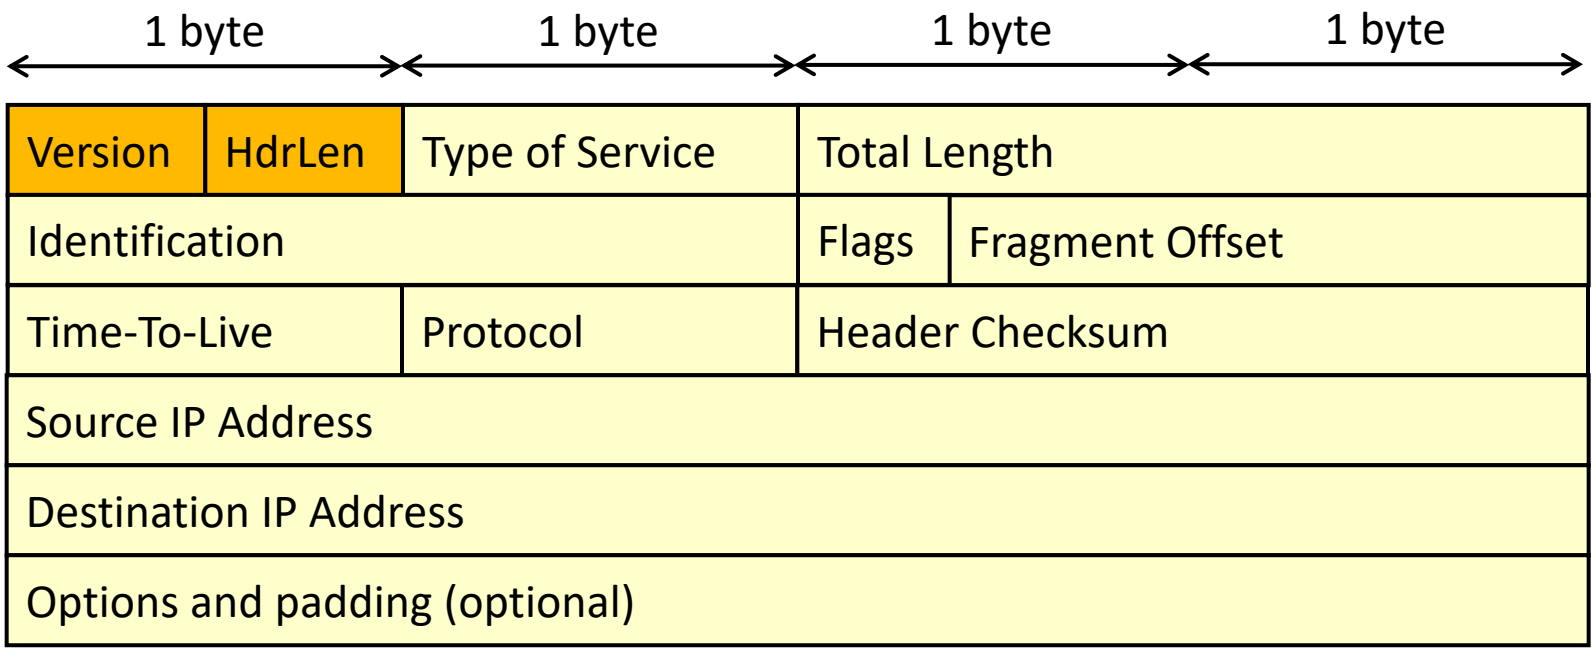

- Version: IPv4 or IPv6 (and other uncommon options)
- Header Length: Total length of IP packet header in 32-bit words

IPv4 Header Format

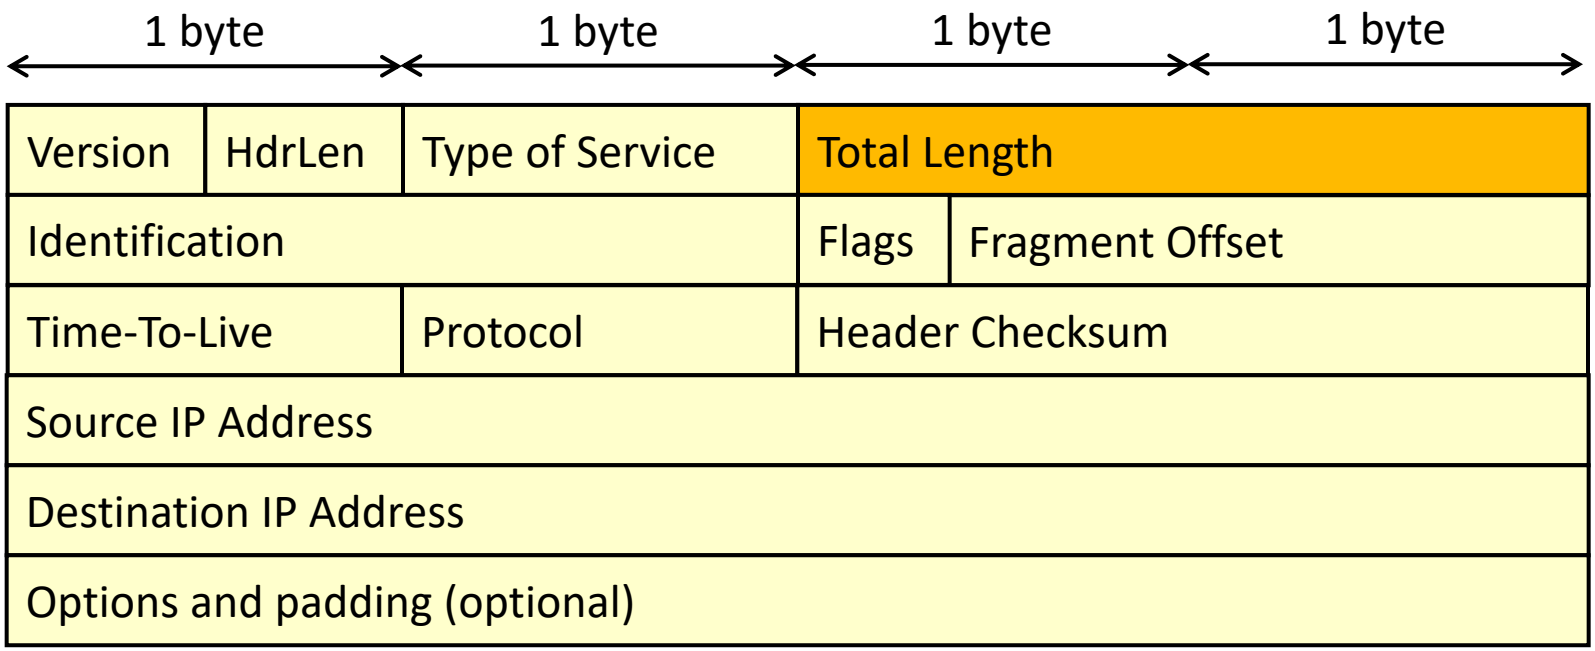

- Total length: 16 bits representing the whole IP packet (header + payload) in bytes
  - Allows a packet of up to 2^{16} bytes (64 KB)
  - This size is too big to communicate in Ethernet LANs. The solution is IP packet *fragmentation*

IPv4 Header Format

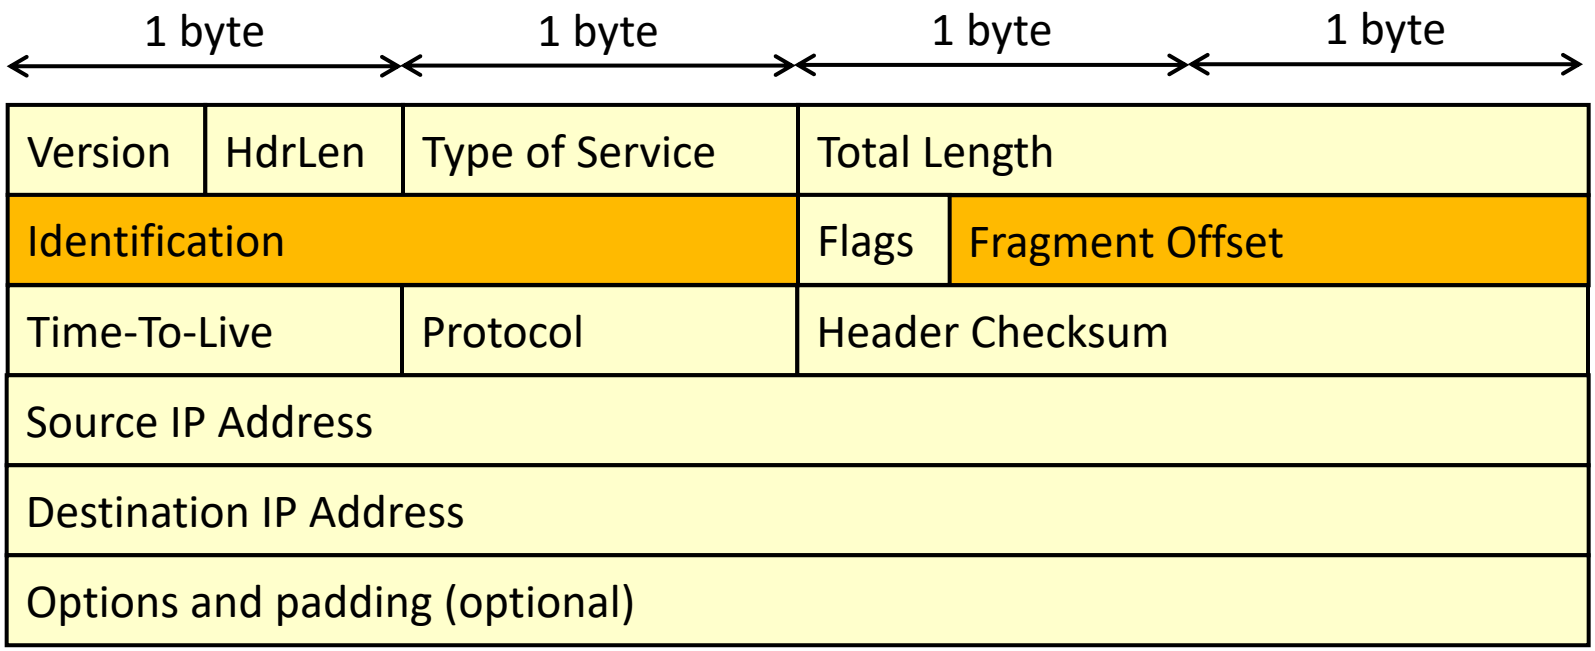

- Identification: 16 bits to identify an IP packet, useful for packet fragmentation
- Fragmentation offset: 13 bits representing the position of this fragment in the overall IP packet

IPv4 Header Format

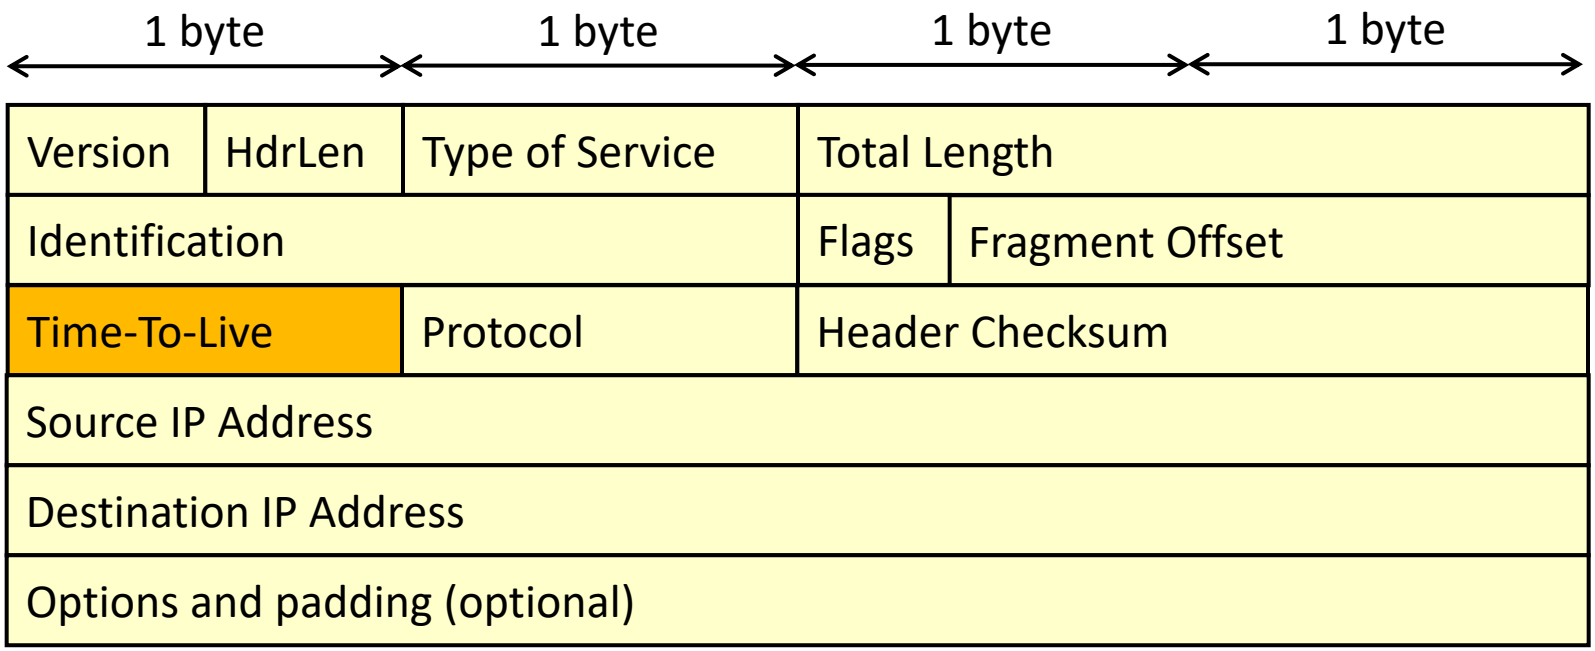

- Time to Live - “Hop count”
  - Decrement each hop
  - Discard datagrams with zero TTL

IP: Time-to-Live

- Sender sets a TTL value for each datagram
- Each router decrements the TTL
- When the TTL reaches 0
  - The router drops the datagram
  - The router sends an ICMP error (more later) to the sender
- Effectively a “maximum hop count”
- **Otherwise there could potentially be infinite loops!**

IPv4 Header Format

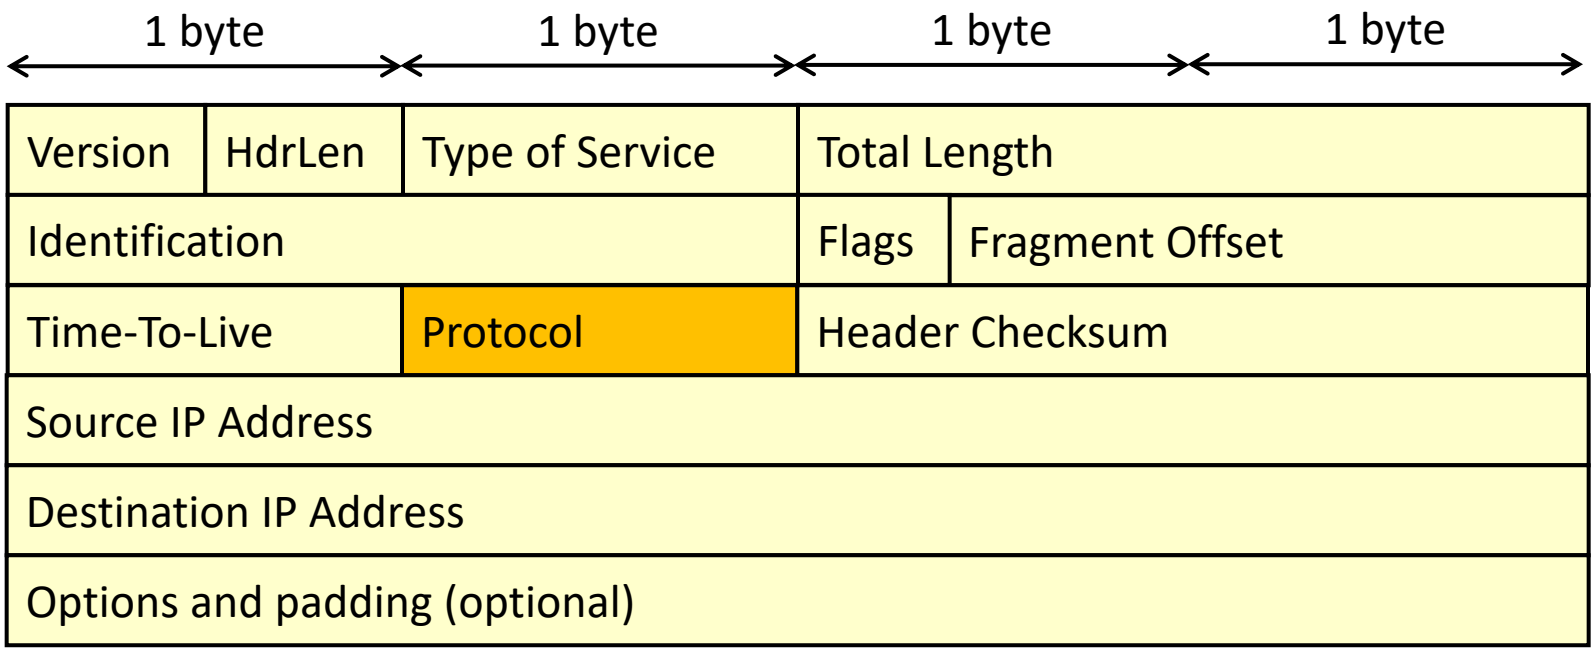

- Protocol: What is encapsulated in this IP datagram?
  - 1 = ICMP, 6 = TCP, 17 = UDP, etc…

IPv4 Header Format

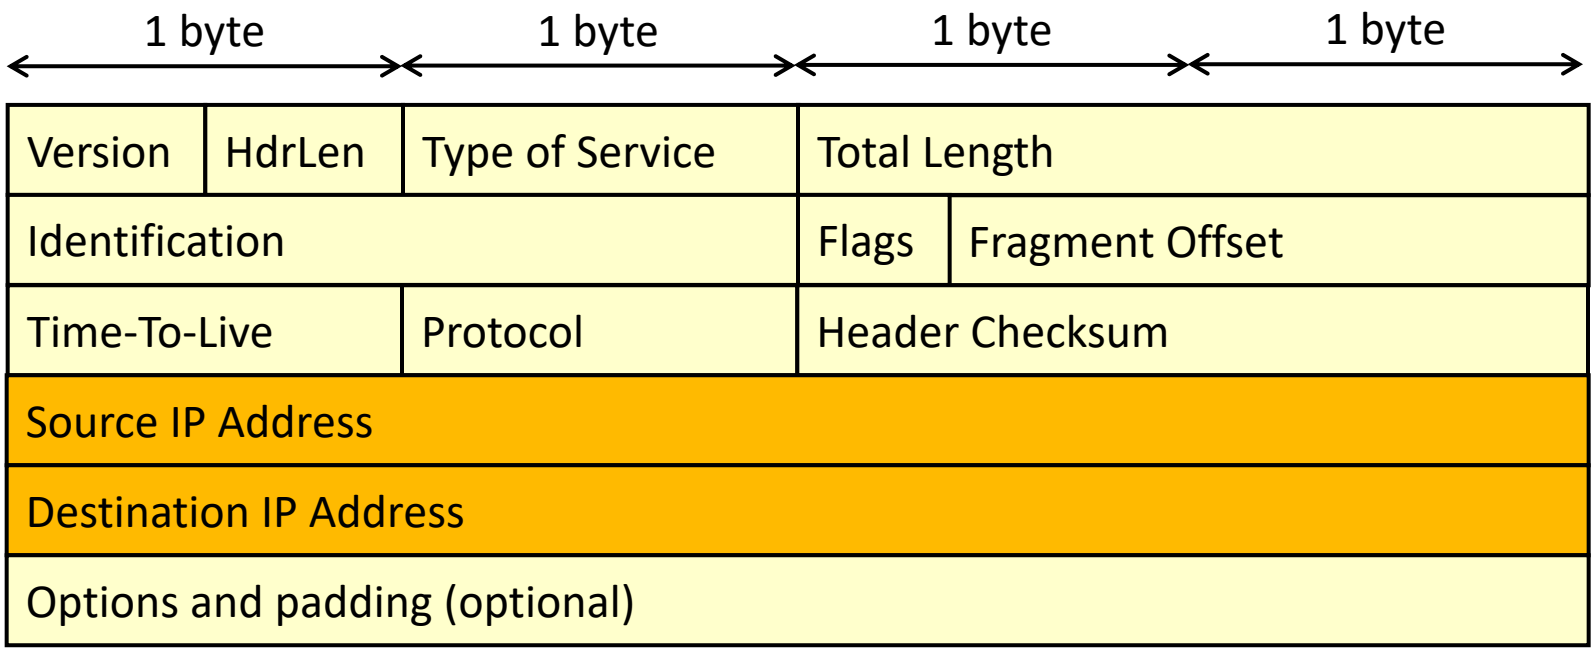

- Source IP address: 32 bits representing the logical address of the interface from which the packet originated
- Destination IP address: 32 bits representing the logical address of the interface to which the packet is destined

IP encapsulated in Ethernet

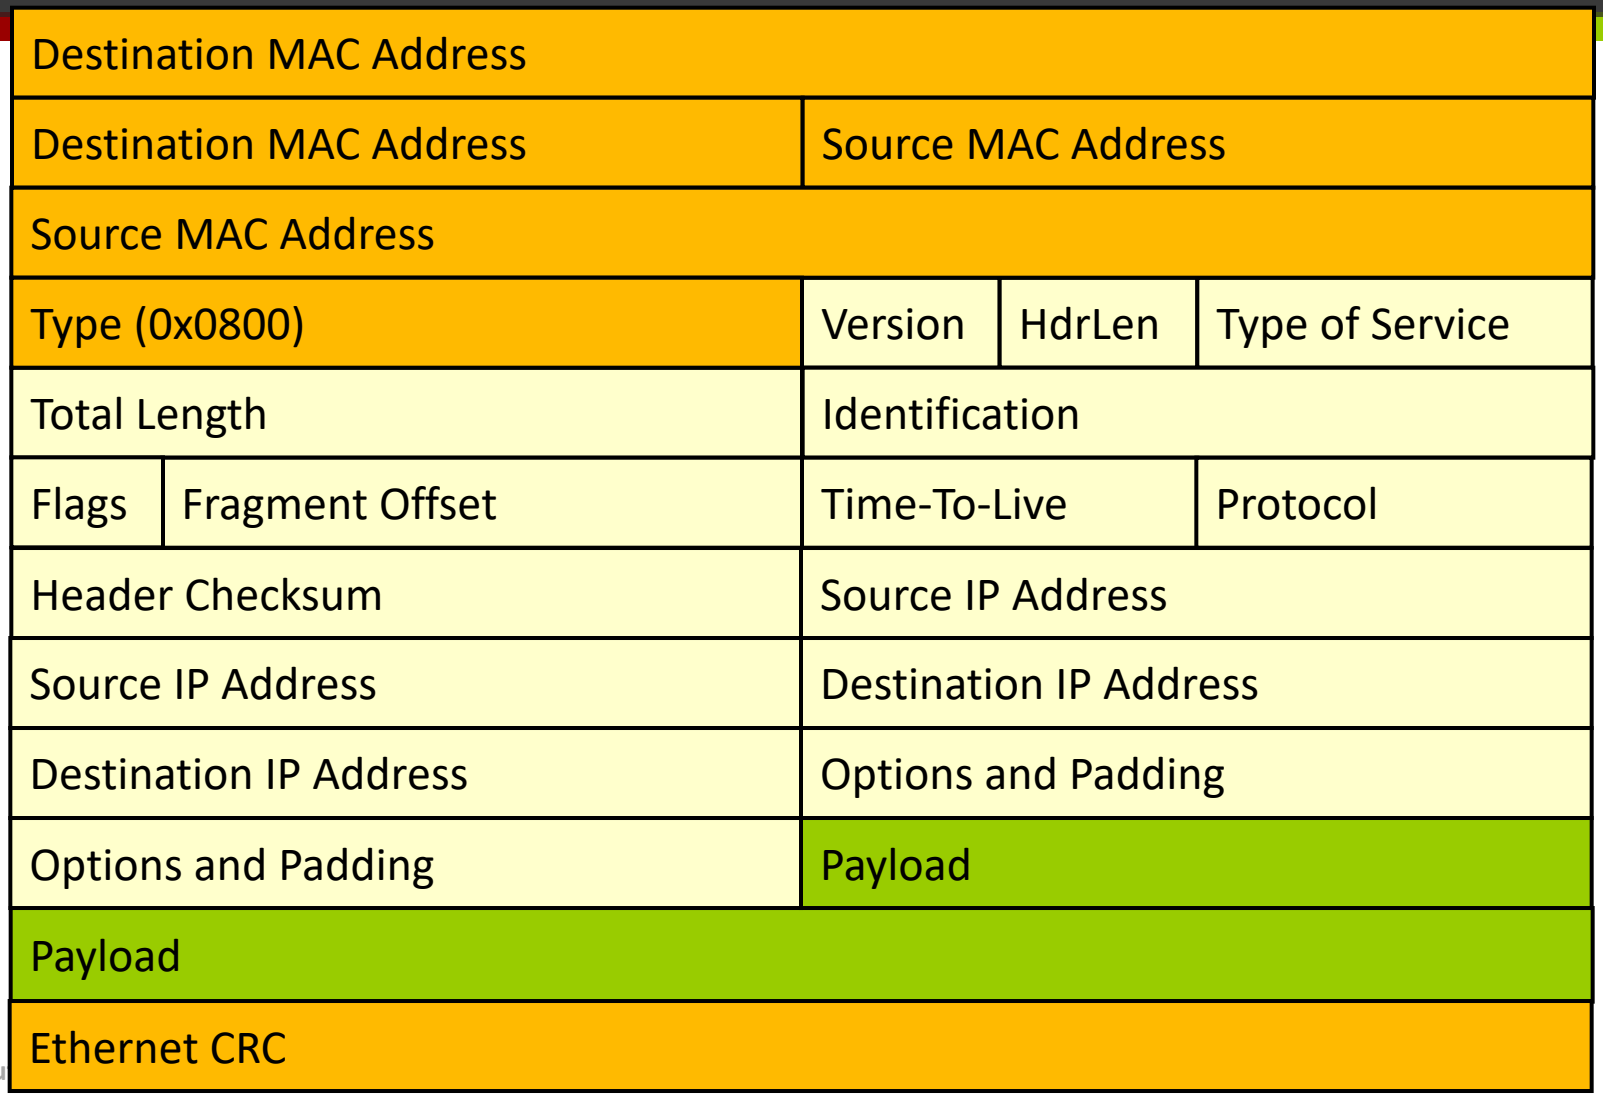

Routing Between LANs

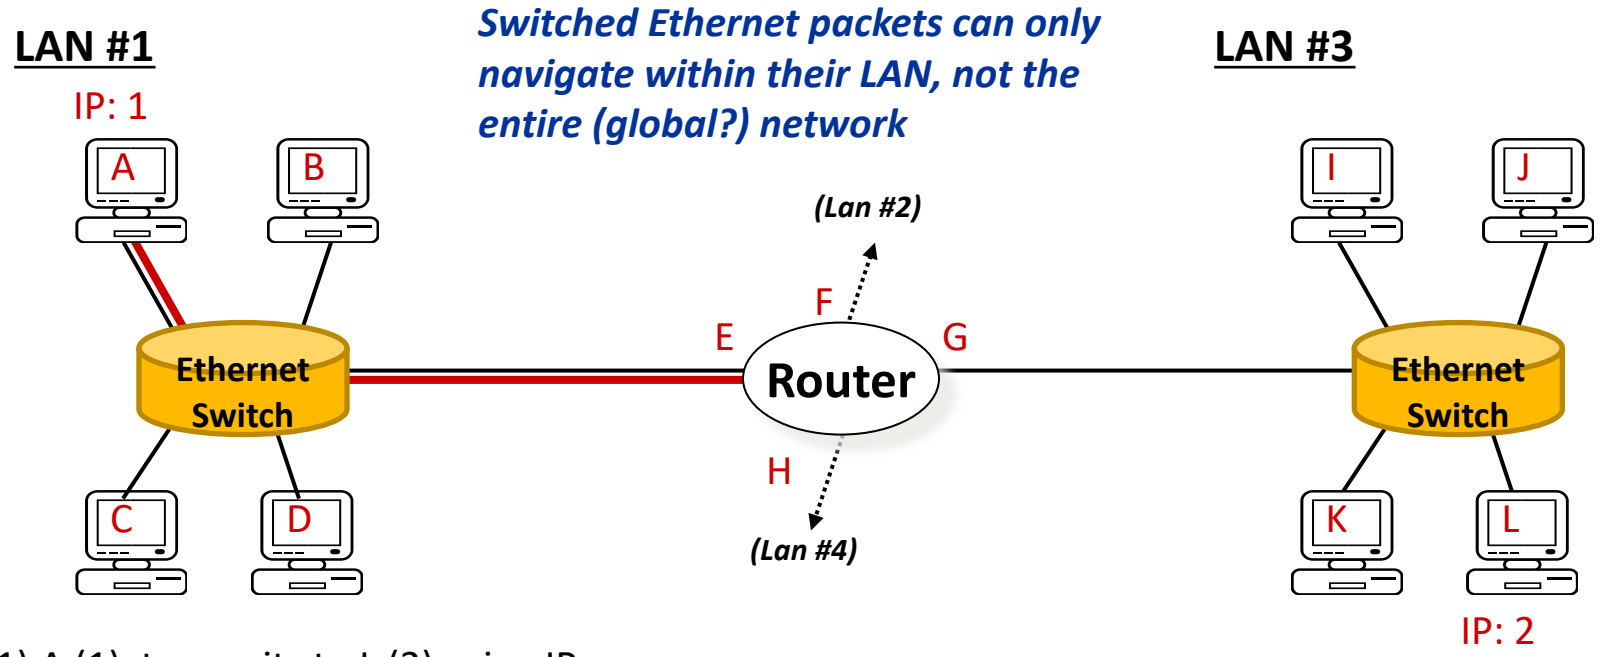

(1) A (1) transmits to L (2) using IP.
Ethernet frame destination is router

Frame:

EDA (E)	ESA (A)	0x0800	IPDA (2)	IPSA (1)

Ethernet Destination Address: MAC Addr of E
Ethernet Source Address: MAC Addr of A
IP Destination Address: IP of L
IP Source Address: IP of A

(2) Switch forwards frame to router

Routing Between LANs

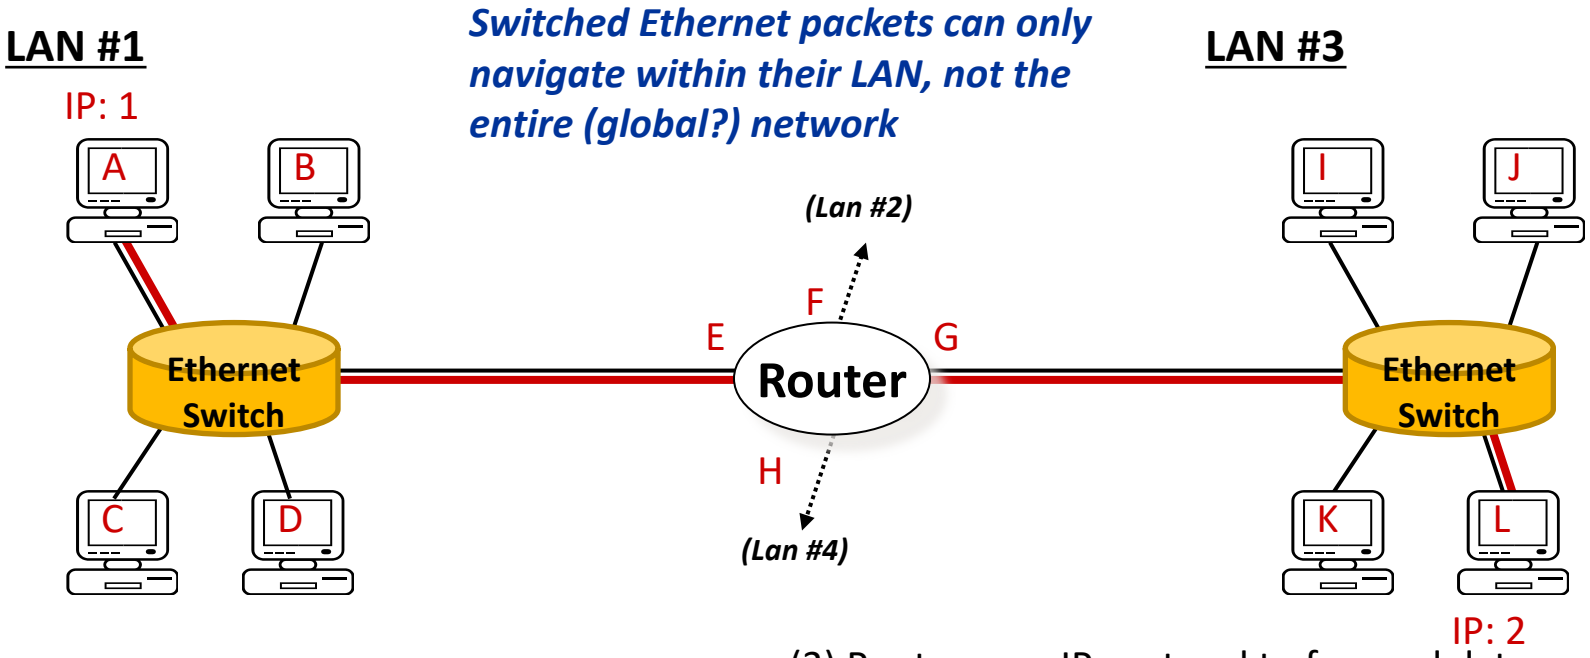

(3) Router uses IP protocol to forward data.
Ethernet: update src/dst/crc IP: update TTL/checksum

*Ethernet Destination Address: MAC Addr of **L***
*Ethernet Source Address: MAC Addr of **G***
IP Destination Address: IP of L
IP Source Address: IP of A

Frame:

EDA (L)	ESA (G)	0x0800	IPDA (2)	IPSA (1)

(4) Switch forwards frame to destination

IP Address Format

- An IP address consists of two parts:
  - *Network* part (prefix)
  - *Host* (interface) part (remainder)

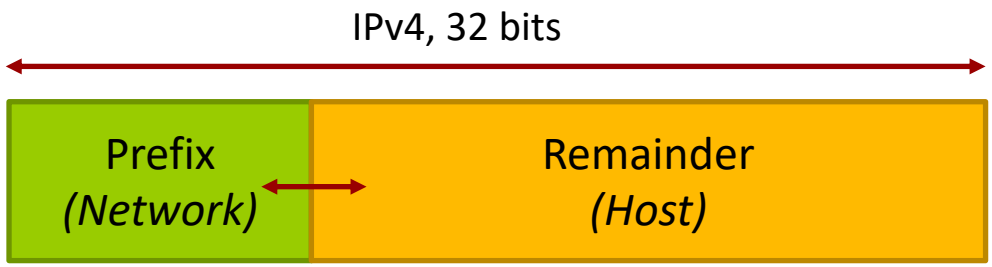

- The size of the prefix can **vary!**
  - /8 subnet = Prefix is 8 bits ("Class A")
  - /16 subnet = Prefix is 16 bits ("Class B")
  - /24 subnet = Prefix is 24 bits ("Class C")

IP Address Format

- The network part is assigned by the *ISP*
- The host part is usually configured by the *network administrator*
- As a result, IP addresses serve both as
  - End point identifiers
  - Logical locators in the network (not necessarily geographically)
- Note that MAC addresses are not locators, but only endpoint identifiers
  - *Knowing a MAC address does not help you locate that machine in the global Internet*

CIDR Addresses

- Traditional IP address classes are not very flexible
  - Class A networks can have ~16 million IP addresses
  - Class B networks can have ~64 thousand IP addresses
  - Class C networks can have up to 256 IP addresses (to be
  - exact 254)
- What if a customer needs ~1000 IP addresses?
- Solution: *Classless Inter-Domain Routing (CIDR)*
  - The division between the network prefix and remainder is **dynamic**
  - Allows more adaptable network sizes

Specifying Network Prefix

- Two ways to specify the network prefix and remainder
- **CIDRized** notation: Using `/n` at the end of IP address, where `n` is the length of the network prefix
  - Example: `147.126.5.125/16` refers to a class B address
  - Example: `9.126.5.125/8` refers to a class A address
  - *Easier for humans to understand!*
- **Netmask**: Using 1 for the location of bits that refer to network prefix
  - Example: `147.126.5.125` with netmask `255.255.0.0`
  - *This is how the computers (routers…) actually work*

Netmask vs. CIDRized notation

CIDR Notation	Netmask
200.1.130.98/27	255.255.225.224
200.1.130.98/28	255.255.255.240
200.1.130.98/14	255.252.0.0

https://www.tunnelsup.com/subnet-calculator/

Special IP Addresses

- Loop-back address: `127.0.0.1`
- Unrouted (private) IP addresses:
  - `10.0.0.0/8`
  - `172.16.0.0/12`
  - `192.168.0.0/16`

http://xkcd.com/742/

Closing Thoughts

Recap

- Today we discussed
  - Internet Protocol
  - IPv4 header format
  - Assigning IP addresses to interfaces
  - IP address classes, netmasks, and CIDR

Next Class

- More IPv4

Class Activity

CA.4 – IPv4 & Wireshark

Due tonight at 11:59pm