Introduction to Computer Networks

CS640

IP Introduction

https://pages.cs.wisc.edu/~mgliu/CS640/F22/

Ming Liu
mgliu@cs.wisc.edu

Today

Last lecture

- L2 reliable transmission

Today

- What functionalities do the IP layer provide?

Announcements

- Lab2 is due this Friday 10/14/2022, 11:59 PM
- Lab3 is released this Friday

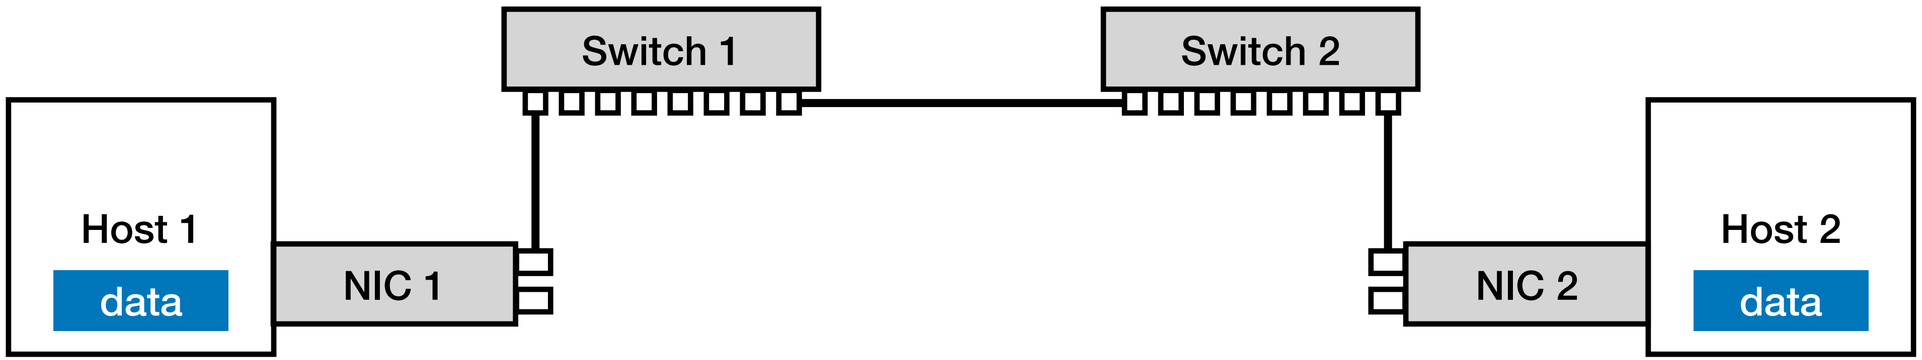
Switch 1
Switch 2
Host 1
data
NIC 1
NIC 2
Host 2
data

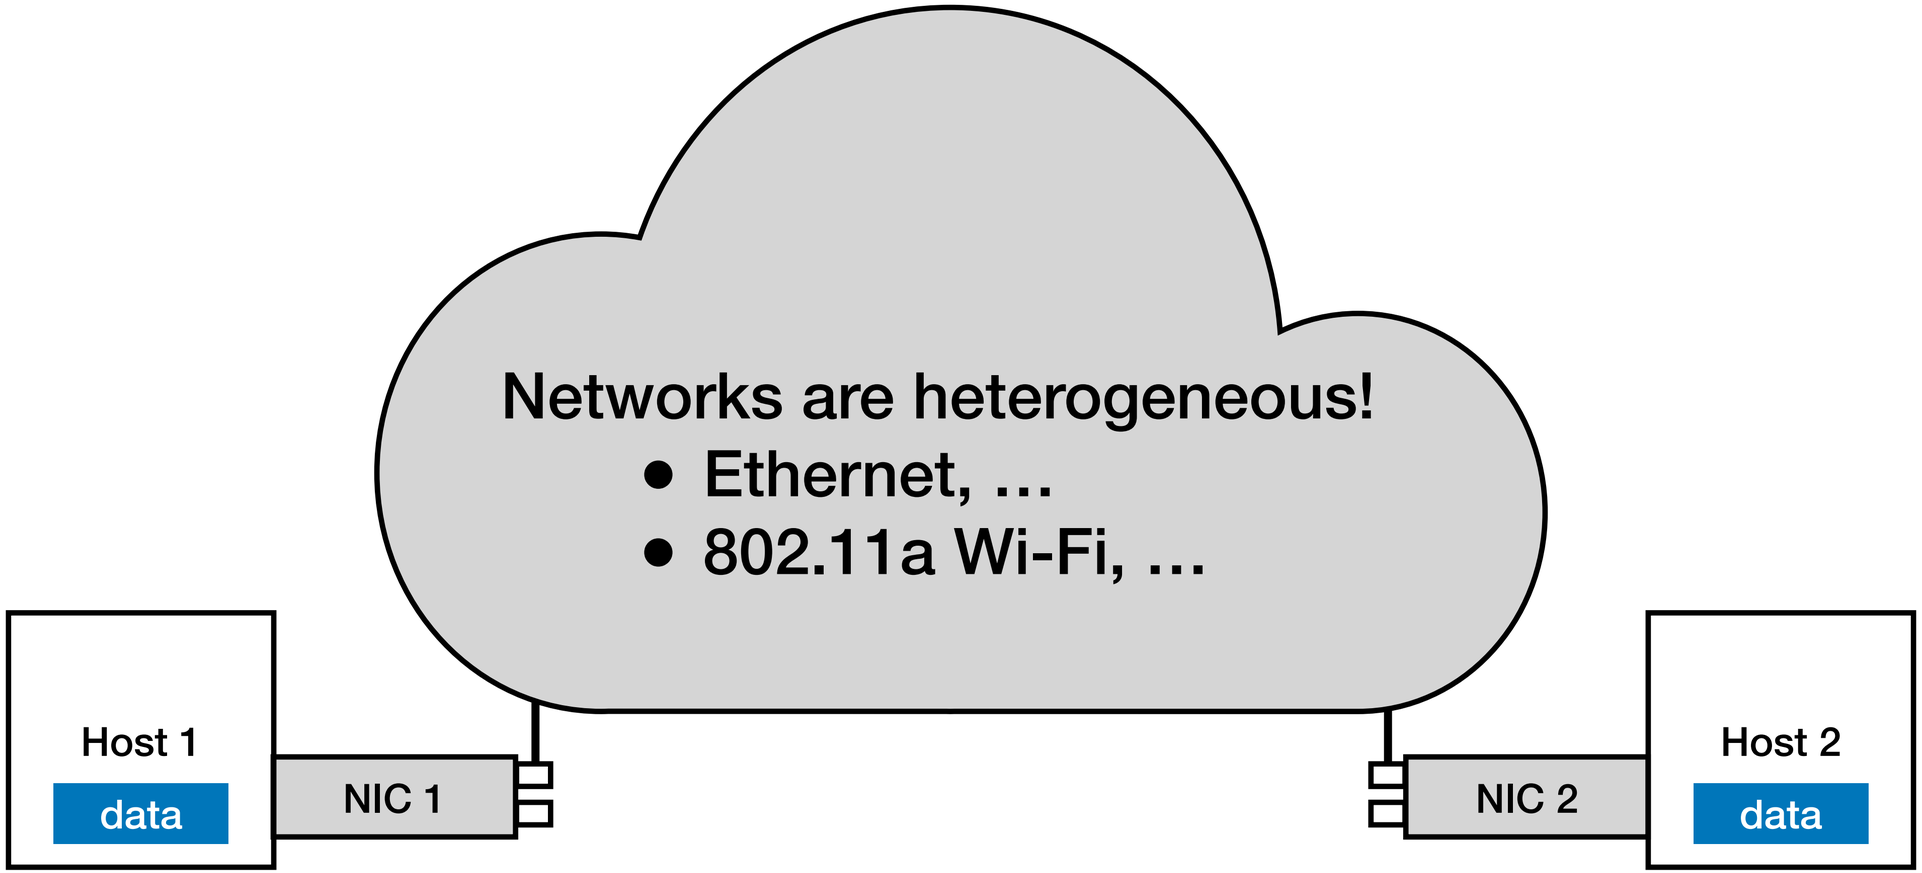
Networks are heterogeneous!
• Ethernet, …
• 802.11a Wi-Fi, …
Host 1
data
NIC 1
NIC 2
Host 2
data

Q: How to transmit a packet between two hosts in any network?

IP datagram: a unit of data in the IP/L3 layer

Q: How to transmit ~~a packet~~ between two hosts in any network?

Q: How to transmit an IP datagram between two hosts in any network?

Q1: How to address any host in any network?
Q2: How to deliver data for an arbitrary communication path?
Q3: How to improve the packet transmission efficiency given the unbounded scale?

Q: How to transmit an IP datagram between two hosts in any network?

A: Internet Protocol (IP)

- #1: run over all the entities in a collection of network
- #2: define the infrastructure that allows these nodes and networks to work as a single logical network

Hour Glass Model

Hide underlying L2 technologies from network apps

- Etherent, Wi-Fi, PPP, …

Support many different types of apps

- Email, browsing, streaming, …

Provide the minimal functionalities

- Two key functionalities: addressing and routing

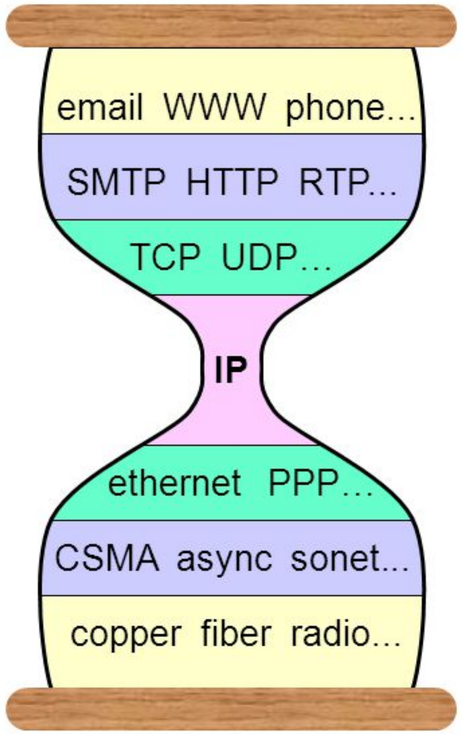

Q: Why IP is so powerful?

Q: Why IP is so powerful?

A: Best-effort host-to-host service model

Q: Why IP is so powerful?

A: Best-effort host-to-host service model

- #1: a unified header format

IP Packet Format

Version (4 bits)

- IP version number, default: 4

HLen (4 bits)

- # 32-bit words in header

TOS (8 bits)

- Type of service
- 6-bit DSCP (Differentiated service)
- 2-bit ECN (Explicit Congestion Notification)

IP Packet Format

Length (16 bits)

- # bytes in this datagram

Identification (16 bits)

- Sequence number
- Used by fragmentation

Flags + Offset (3+13 bits)

- Used by fragmentation

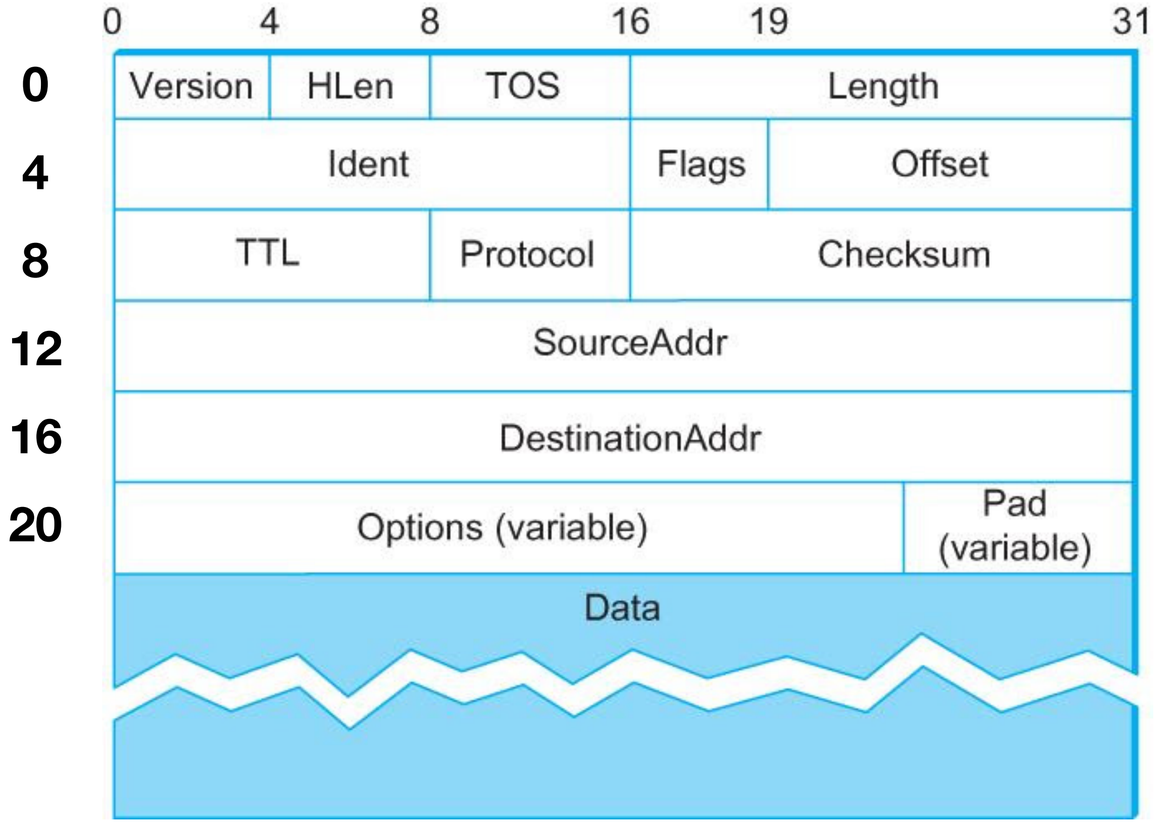

IP Packet Format

TTL (8 bits)

- # hops this datagram has traveled
- Decremented at every router

Protocol (8 bits)

- Demultiplex key (e.g., TCP=6, UDP=17)

Checksum (16 bits)

- The checksum of the IP header in terms of 16-bit words

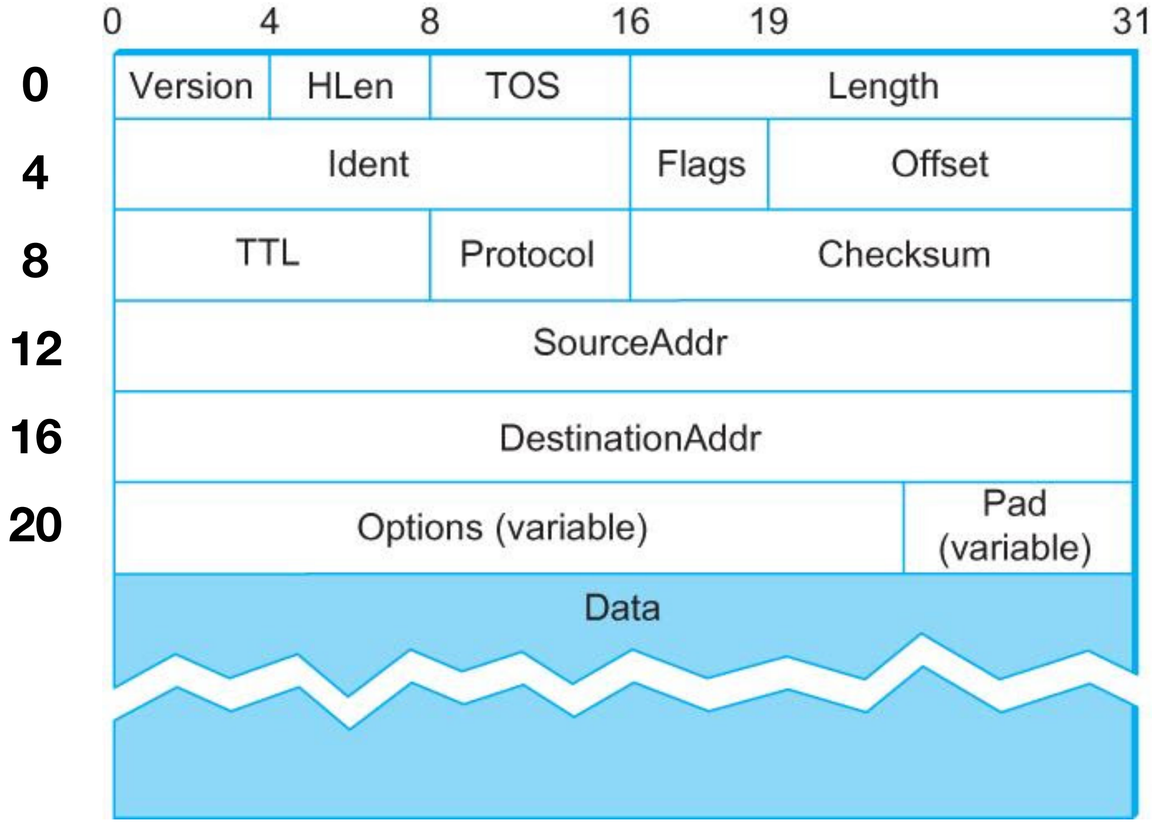

IP Packet Format

SourceAddr (32 bits)

- The address of the source host

DestinationAddr (32 bits)

- The address of the destination host

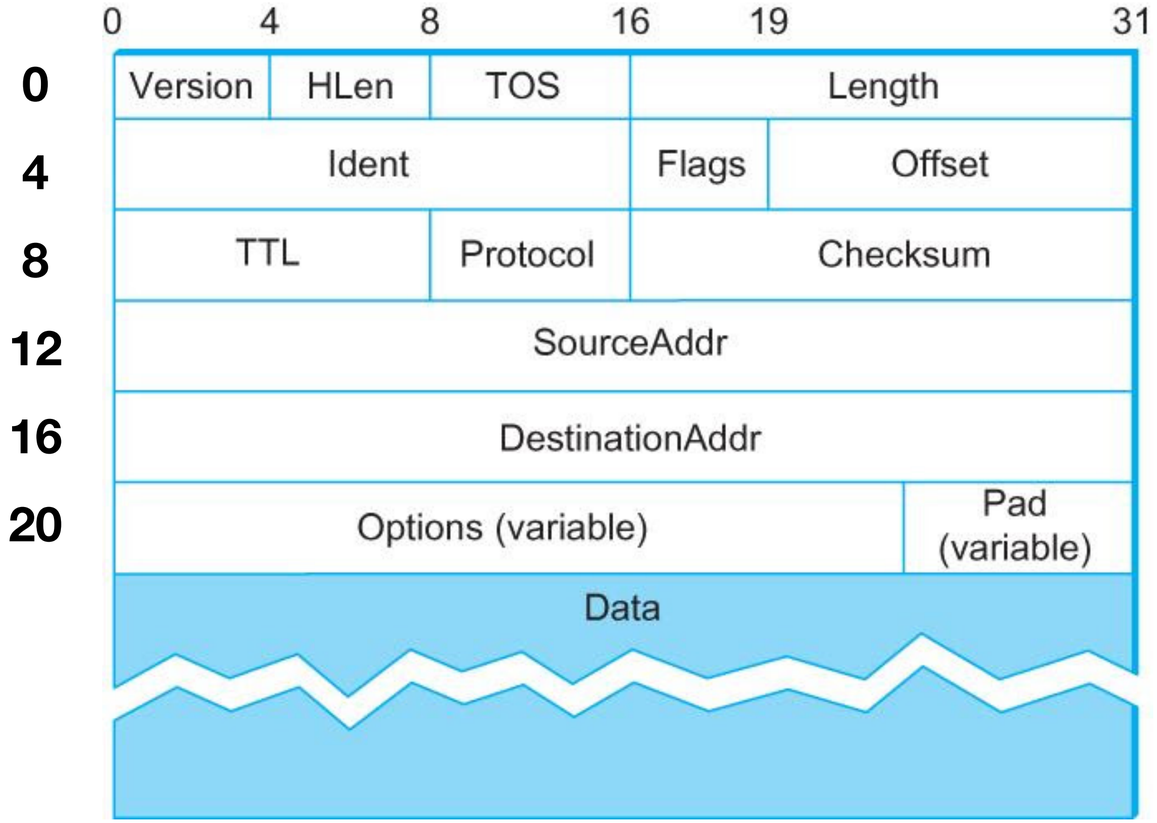

Data Transformation

Signals ⟷ Bits ⟷ Frames ⟷ IP datagrams

Data Transformation

Signals ⟷ Bits ⟷ Frames ⟷ IP datagrams

Ethernet frame

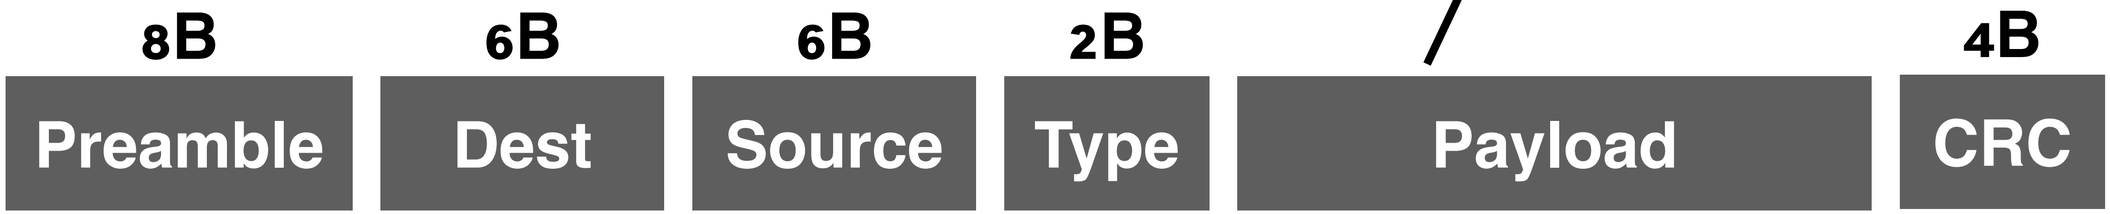

Data Transformation

Signals ⟷ Bits ⟷ Frames ⟷ IP datagrams

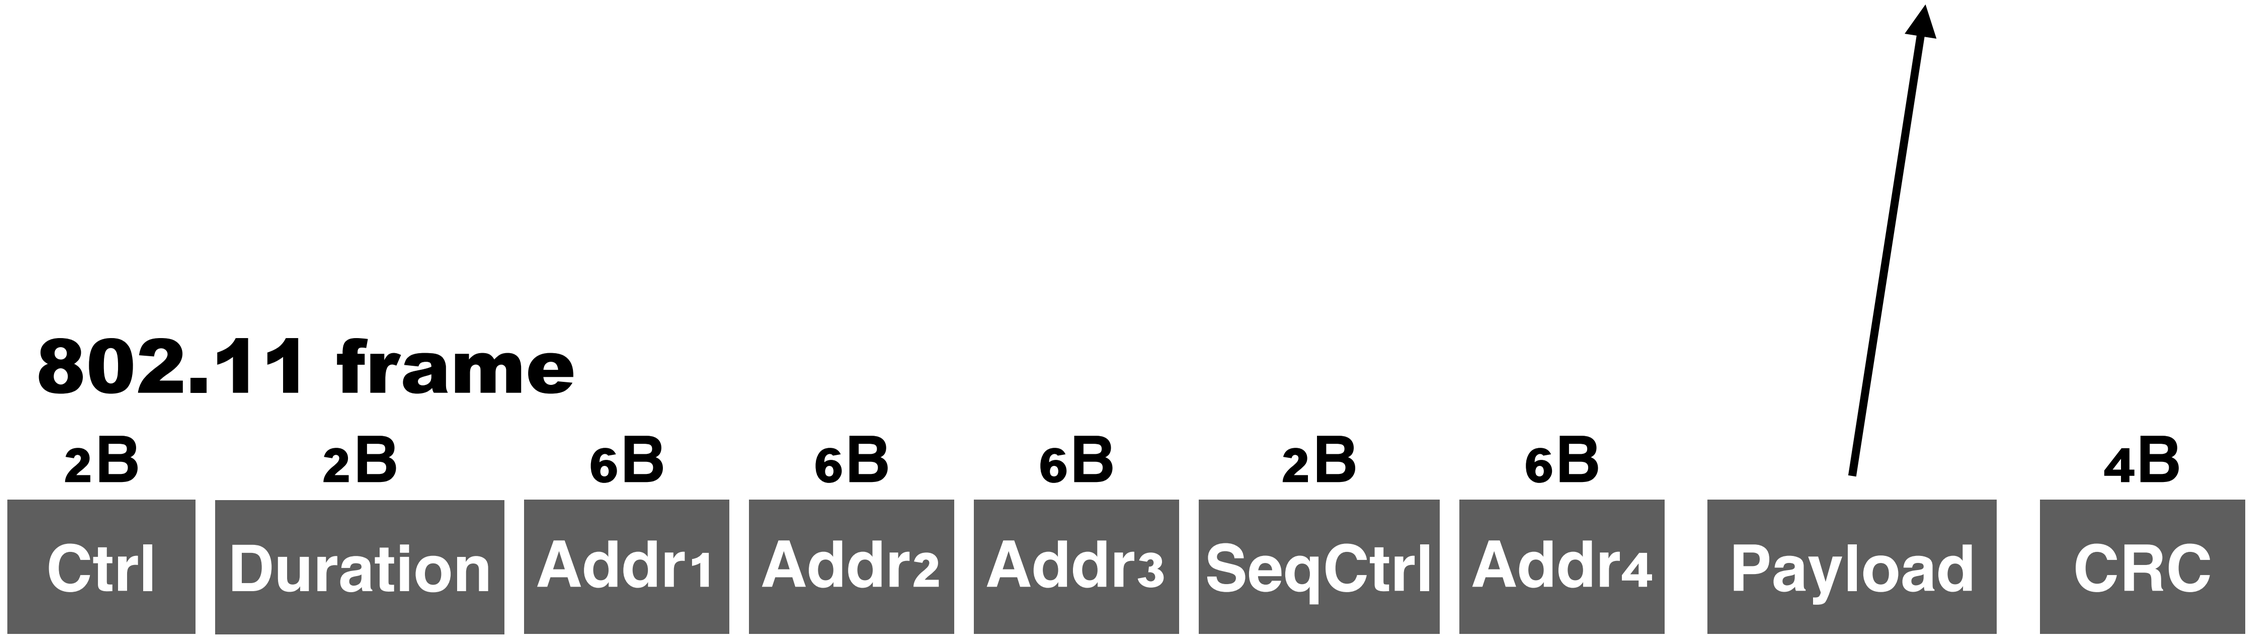

Q: Why IP is so powerful?

A: Best-effort host-to-host service model

- #1: a unified header format
- #2: support heterogeneous networks

Maximum Transmission Unit (MTU)

Different L2 networks define their own packet transmission limit

- Ethernet: 1500 bytes or 9000 bytes

Adapt the IP datagram to the underlying L2 frame

- #1: Fragmentation and reassembly
- #2: Synchronize the MTU

#1: Fragmentation and Reassembly

Breakdown the IP datagram when traversing a link using small-sized MTU

Strategy

- #1: Fragment when necessary (MTU < Datagram)
- #2: Avoid fragmentation at the source host
- #3: Re-fragmentation is possible
- #4: Delay reassembly until the destination host
    - Keep this functionality out of the network
- #5: Do not recover from lost fragments

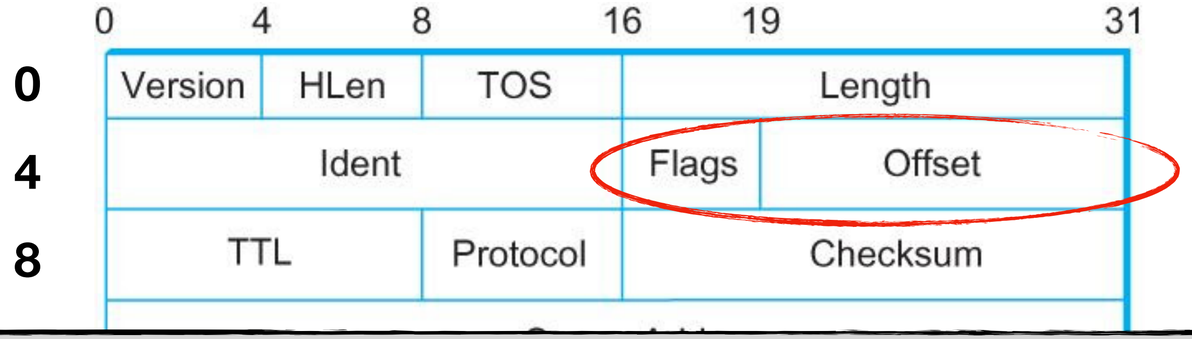

- Three-bit flags
  - bit 0: reserved; must be zero
  - bit 1: Don’t Fragment (DF)
  - bit 2: More Fragments (MF)
- Offset (13 bits)
  - Specify the offset of a particular fragment relative to the beginning of the original unfragmented IP datagram in units of eight-byte blocks

#2: Synchronize the MTU

Path MTU discovery

- Originally introduced by IPv4
- IPv6 delegate it to the endpoints

Key idea

- Set the Don't Fragment (DF) flag bit in the IP header
- Any device along the path whose MTU is smaller than the packet will drop it and send back an **Internet Control Message Protocol (ICMP)** message containing its MTU
- Source nodes then reduce their path MTU accordingly

Internet Control Message Protocol (ICMP)

A supporting protocol for the IP to handle errors

- Report the error status to the source host such that it can react accordingly
- Datagrams are not dropped blindly

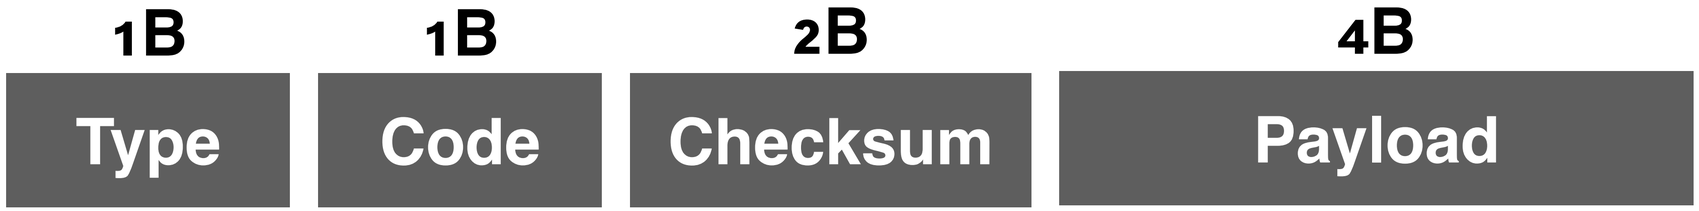

Issued by the router but processed by the host

- Lab3

Q: Why IP is so powerful?

A: Best-effort host-to-host service model

- #1: a unified header format
- #2: support heterogeneous networks
- #3: unreliable packet delivery without broadcasting

Q: Why IP is so powerful?

A: Best-effort host-to-host service model

- #1: a unified header format
- #2: support heterogeneous networks
- #3: unreliable packet delivery without broadcasting

  - Datagrams can be lost
  - Datagrams can be delivered out of order
  - Datagrams can be duplicated
  - Datagrams can be delayed for a long time

Router

A multiple-input multiple-output I/O device that performs forwarding based on the IP address

- Physical appearance is nearly the same as an L2 switch
- Key differences are functionalities

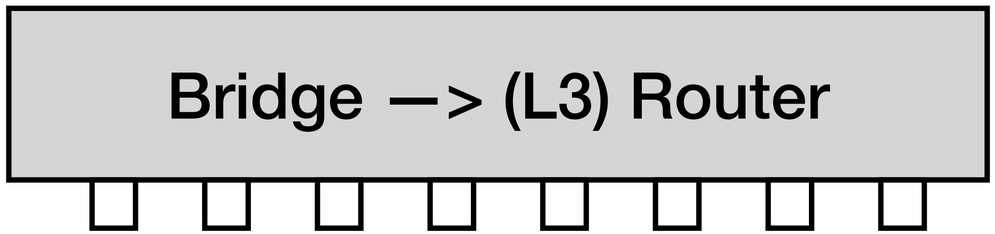

Forwarding in a Router

```
D = destination IP address
N = networkNum(D)

if N == Network Number of one of my interfaces
    Deliver packet over that interface
Else
    if N is in my lookup table
        deliver datagram to next Hop corresp. to N
    else
        deliver datagram to default router
```

IP Router v.s. Ethernet Switch (Incomplete!)

	IP Router	Ethernet Switch
Layering	Layer 3	Layer 2
Forwarding	Based on destination IP address	Based on destination Ethernet address
Packet drop	Speak the ICMP protocol	N/A

Terminology

1. Host
2. NIC
3. Multi-port I/O bridge
4. Protocol
5. RTT
6. Packet
7. Header
8. Payload
9. BDP
10. Baud rate
11. Frame/Framing
12. Parity bit
13. Checksum
14. Ethernet
15. MAC
16. (L2) Switch
17. Broadcast
18. Acknowledgement
19. Timeout
20. Datagram
21. TTL
22. MTU
23. Best effort
24. (L3) Router

Principle

1. Layering
2. Minimal States

Technique

1. NRZ Encoding
2. NRZI Encoding
3. Manchester Encoding
4. 4B/5B Encoding
5. Byte Stuffing
6. Byte Counting
7. Bit Stuffing
8. 2-D Parity
9. CRC
10. MAC Learning
11. Store-and-Forward
12. Cut-through
13. Spanning Tree
14. CSMA/CD
15. Stop-and-Wait
16. Sliding Window

16. Fragmentation and Reassembly
17. Path MTU discovery

Summary

Today's takeaways

#1: IP exposes the best-effort host-to-host service model

#2: The effectiveness of IP comes from a unified header format, the capacity to support heterogeneous networks, and unreliable packet delivery

Next lecture

- Efficient addressing