Internet Protocols by Example

Internet 101 revision in 90 minues

Philipp S. Tiesel – philipp@inet.tu-berlin.de
Slides credits to Anja Feldmann

Internet Protocols by Example

What happens when accessing a Web page:

- Application Layer
  - DNS, HTTP, TLS...
- Transport Layer
  - UDP, TCP
- Network Layer
  - IPv4 (+DHCP +ARP)
  - IPv6 (+auto configuration +neighbor discovery)

➢ See NPA material for references

DNS Example

- Host at cis.poly.edu wants IP address for gaia.cs.umass.edu

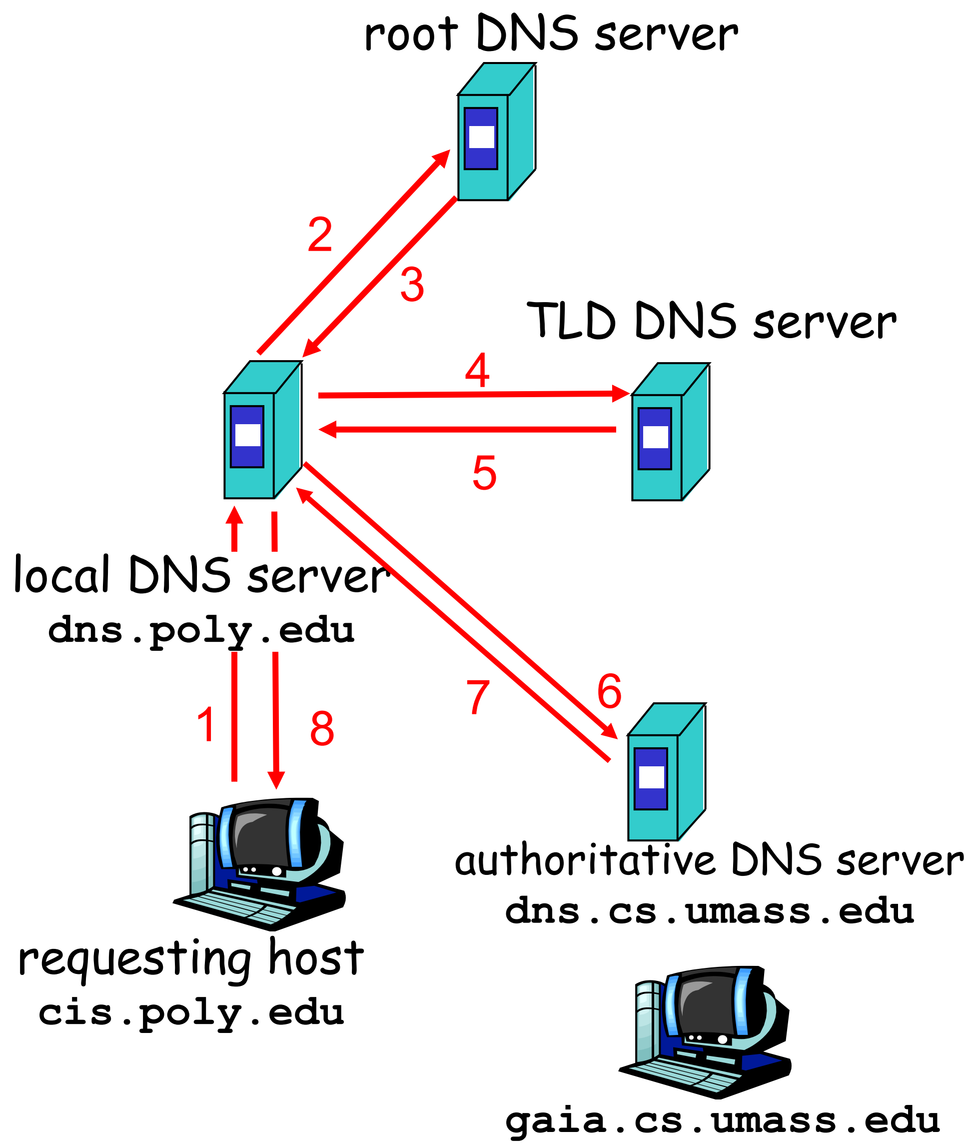

The HTTP Protocol

HTTP: Hypertext transfer protocol

- Application layer protocol for the Web
- Client/server model
  - *Client:* browser that requests, receives, "displays" web objects
  - *Server:* Web server sends objects in response to requests
- Also used as part of many other application layer protocol

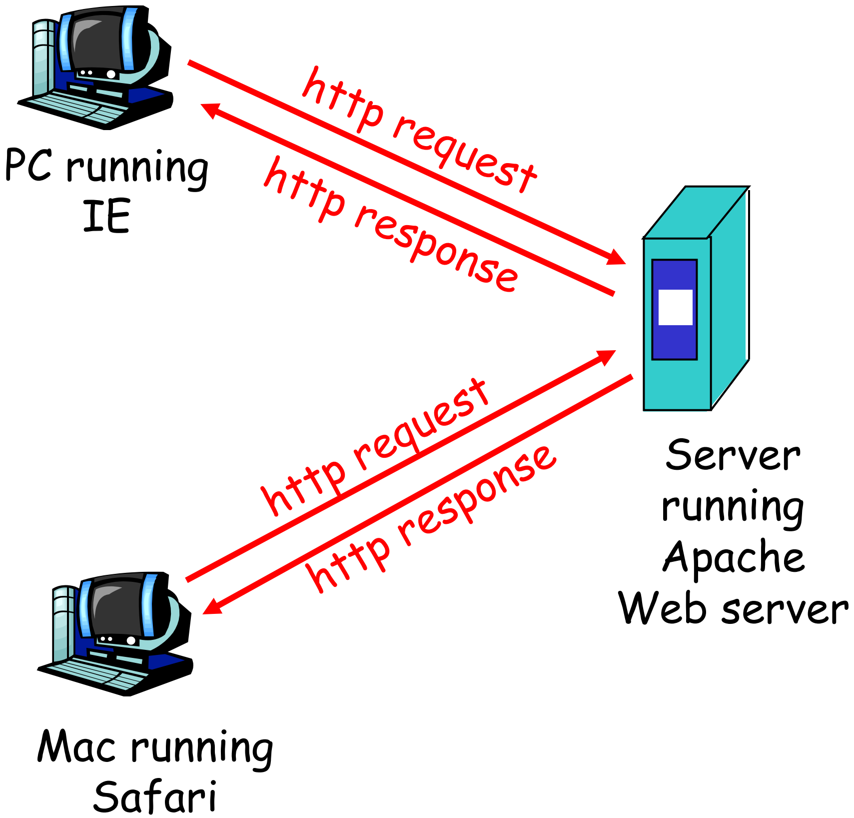

UDP: User Datagram Protocol RFC 768

- "Bare bones" Internet transport protocol
- "Best effort" service, UDP segments may be:
  - Lost
  - Delivered out of order to application
- *Connectionless:*
  - No handshaking between UDP sender, receiver
  - Each UDP segment handled independently of others

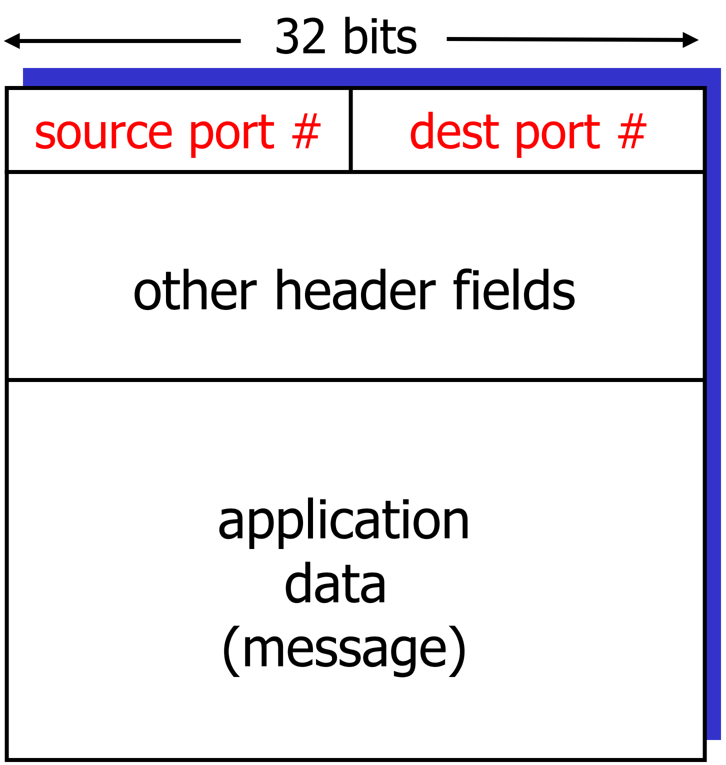

TCP/UDP segment format

TCP: Overview

RFCs: 793, 1122, 1323, 2018, 2581

- Point-to-point:
  - One sender, one receiver
- Reliable, in-order *byte stream:*
  - No "message boundaries"
- Pipelined:
  - TCP congestion and flow control set window size
- Full duplex data:
  - Bi-directional data flow in one connection
  - MSS: maximum segment size
- Connection-oriented:
  - Handshaking (exchange of control msgs) init's sender, receiver state before data exchange
- Flow controlled:
  - Sender will not overwhelm receiver
- Congestion controlled:
  - Sender will not overwhelm the network

TCP Segment Structure

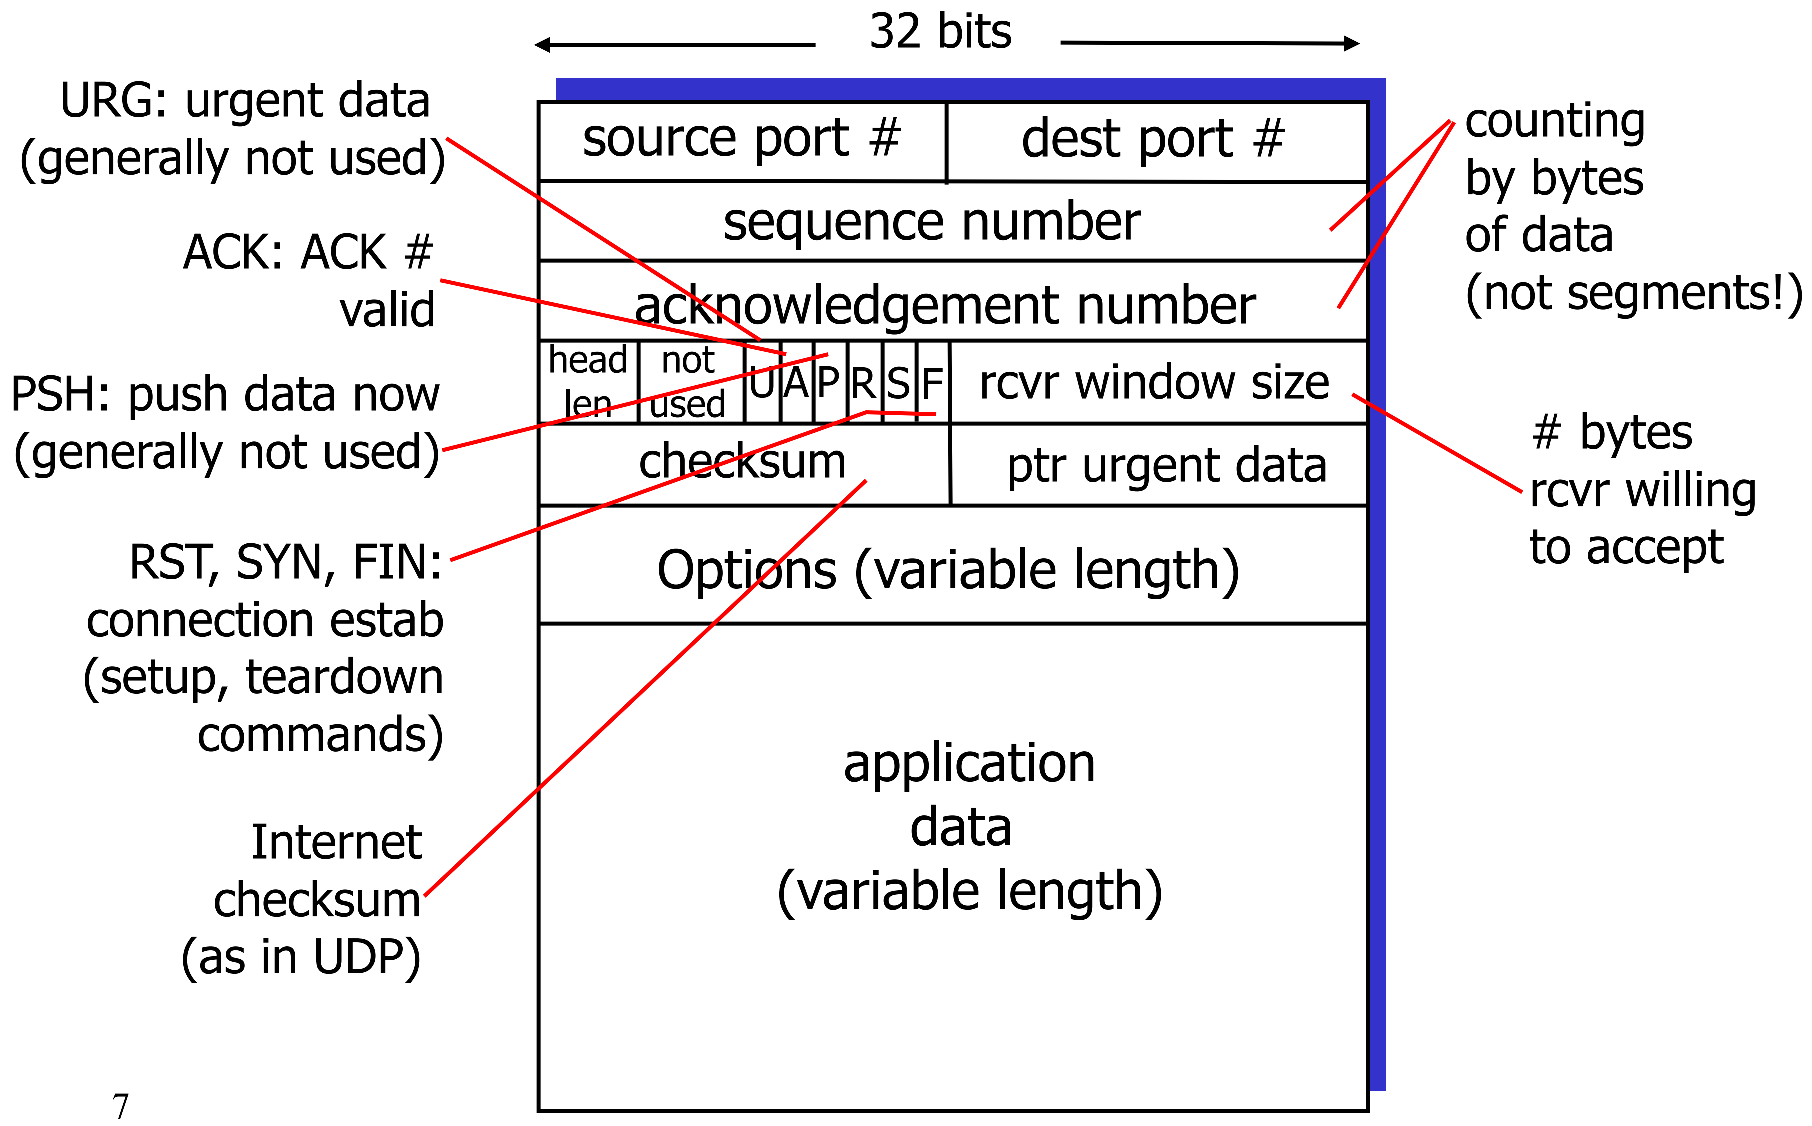

TCP state machine

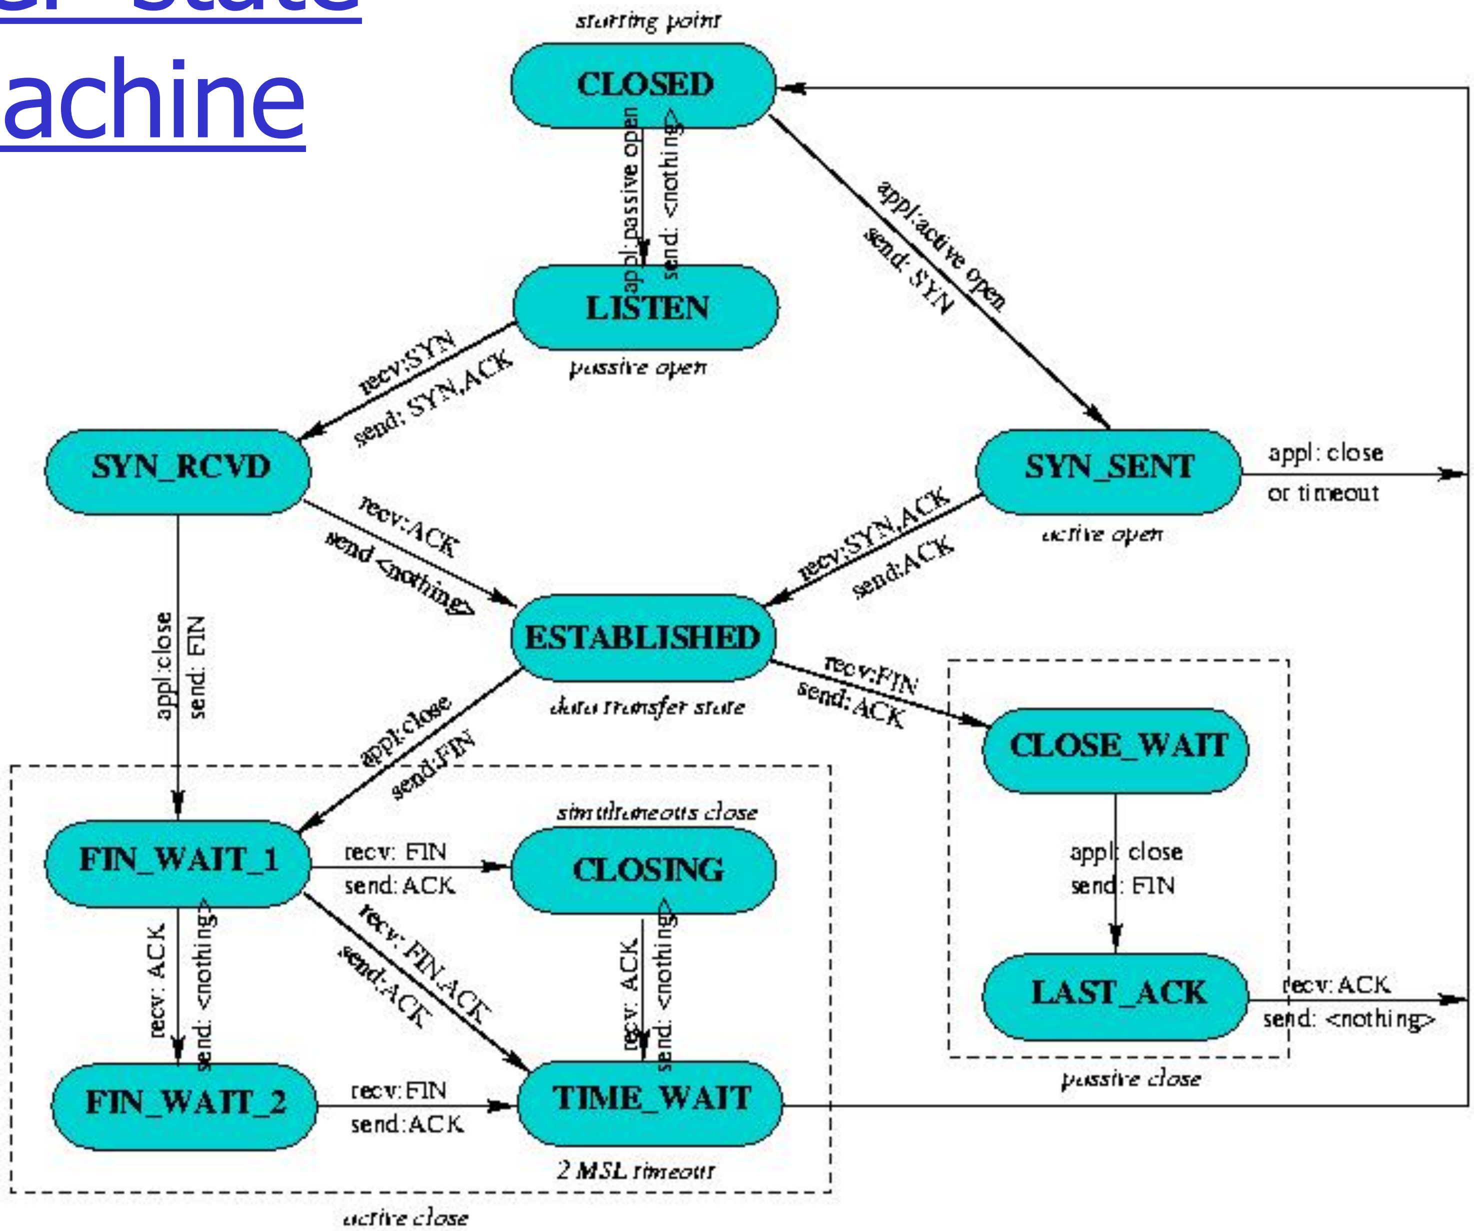

TCP Congestion Control

- End-end control (no network assistance)
- TCP throughput limited by rcvr window (flow control)
- Transmission rate limited by congestion window size, **`cwnd`**, over segments:

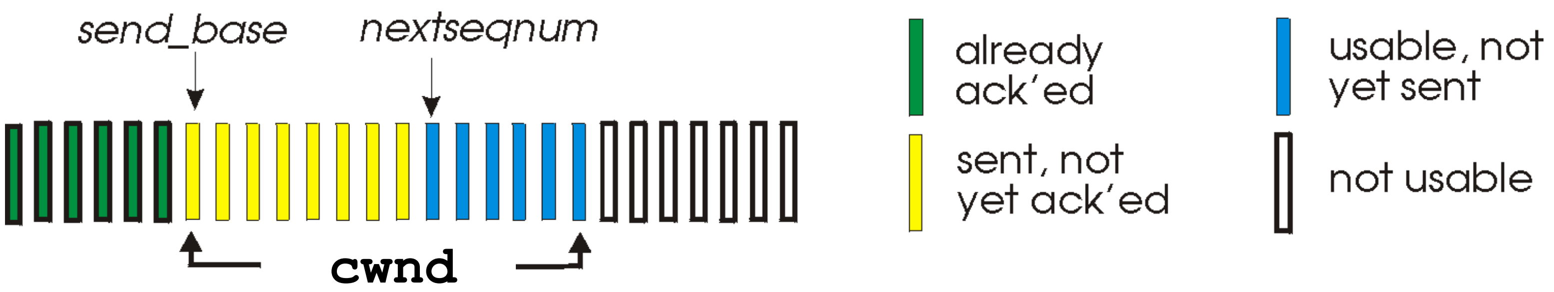

- w segments, each with MSS bytes sent in one RTT

IPv4 datagram format

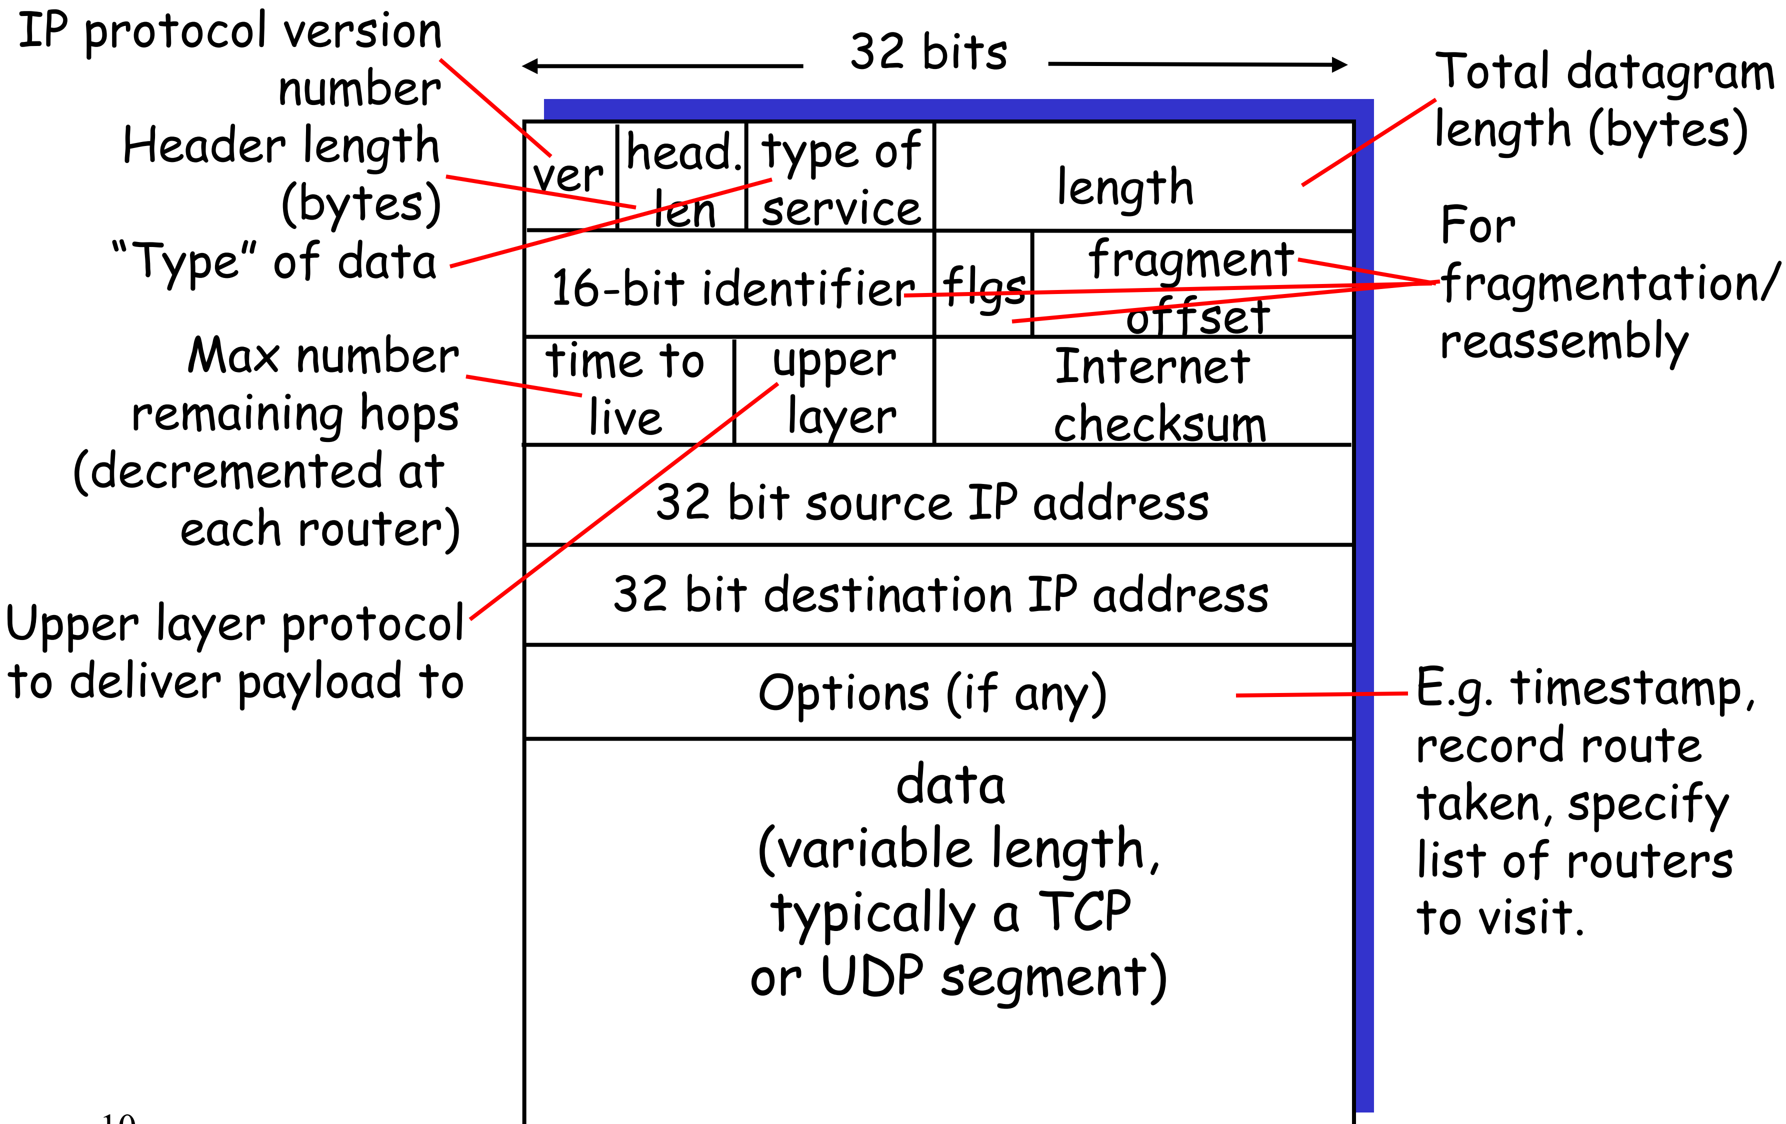

IPv6 datagram format

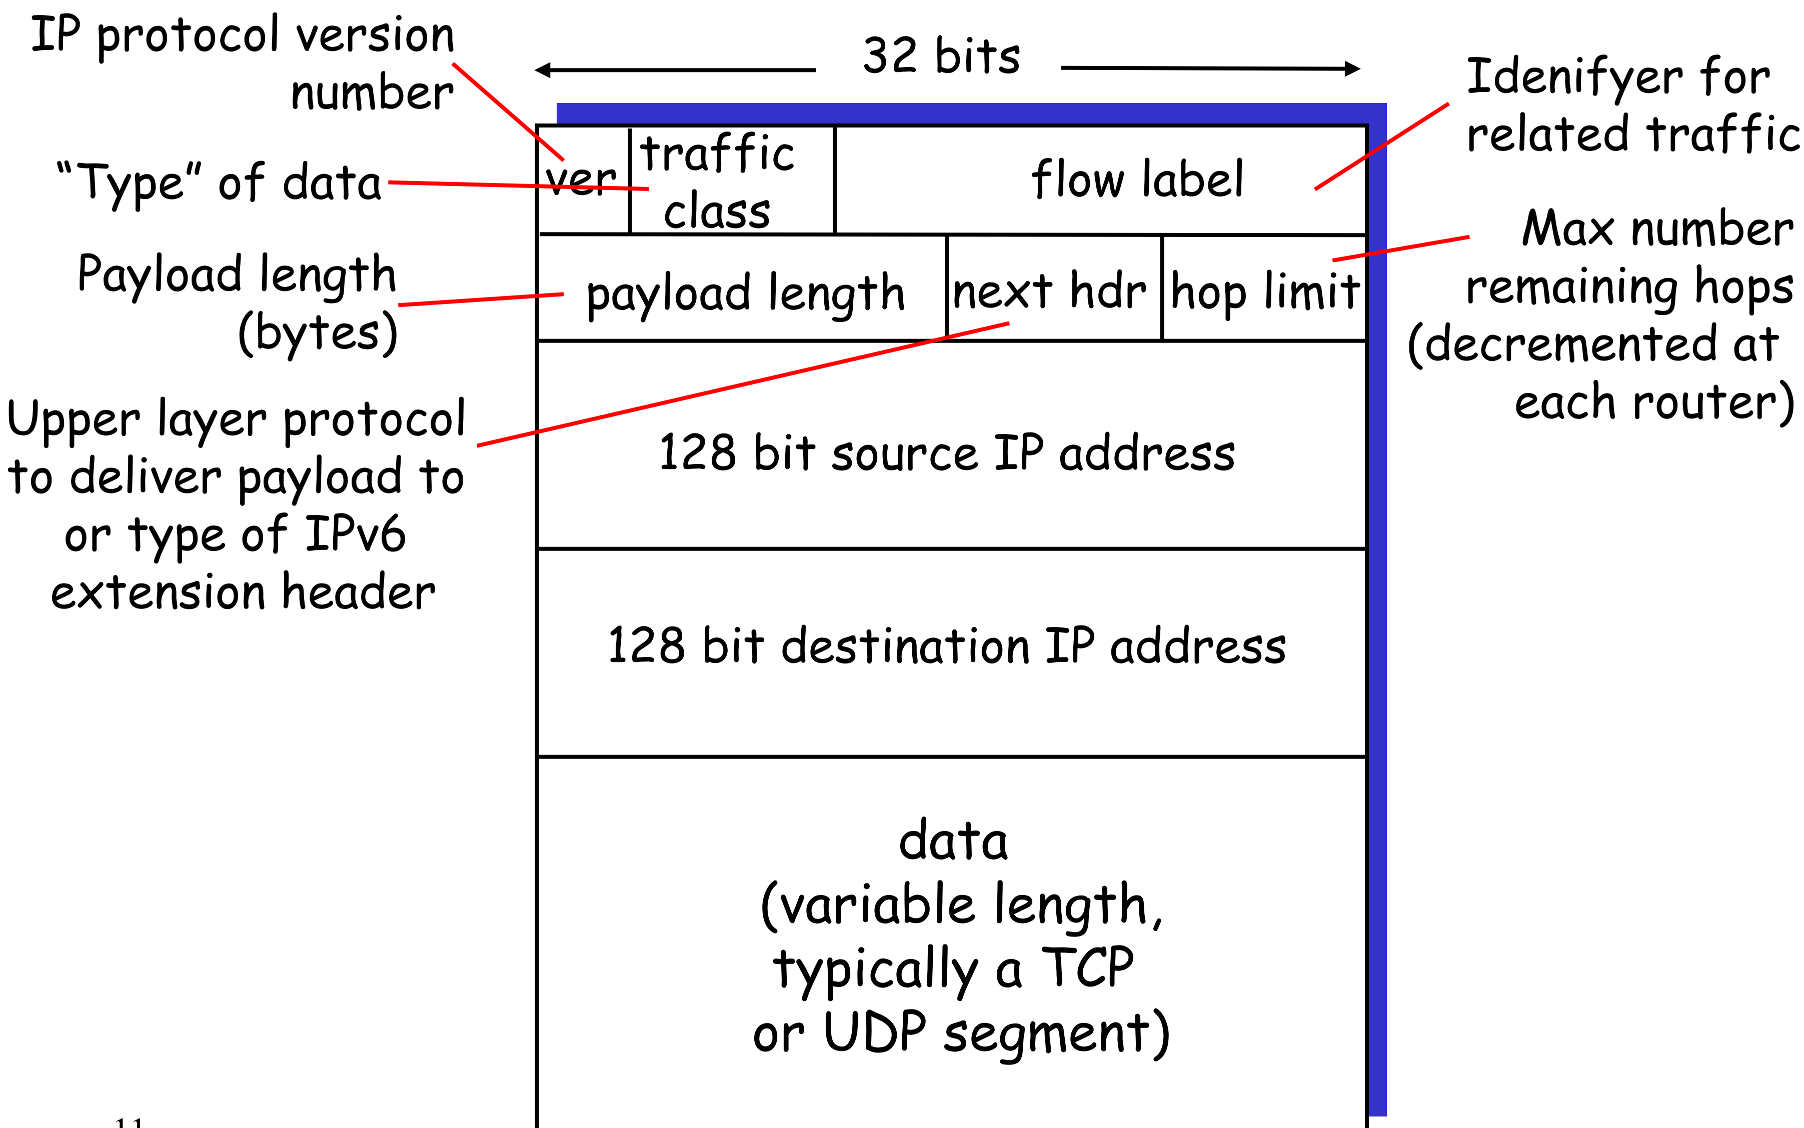

CIDR

CIDR: Classless InterDomain Routing

- Subnet portion of address of arbitrary length
- Address format: a.b.c.d/x, where x is # bits in subnet portion of address

subnet part | host part

11001000 00010111 00010000 00000000

200.23.16.0/23

How does a host get an IP addresses?

- Hard-coded by system admin

- DHCP / DHCPv6
  Dynamic Host Configuration Protocol
  - Request the address from a server
- IPv6 SLAAC
  Stateless Address Auto-Configuration
  - Router advertise the IPv6 prefix
  - Hosts add an Interface Identifier as Host-Part