IPv6 Details

CIS-3152, Spring 2013
Peter C. Chapin

General Goals

- Simplified header (relative to IPv4)
  - Routers should not have to compute a checksum.
    - In IPv4 there is a header checksum that needs recomputation at each step.
    - Underlying link protocol, or upper level transport application protocol and do error detection.
  - Many features relegated to “extension headers.”
    - Application only uses extensions that it needs.
    - Keeps header size reasonable.
- No (on network) fragmentation allowed.
  - Fragmentation in IPv4 has proved problematic.

Header

- IPv6 header is (nominally) 40 bytes.

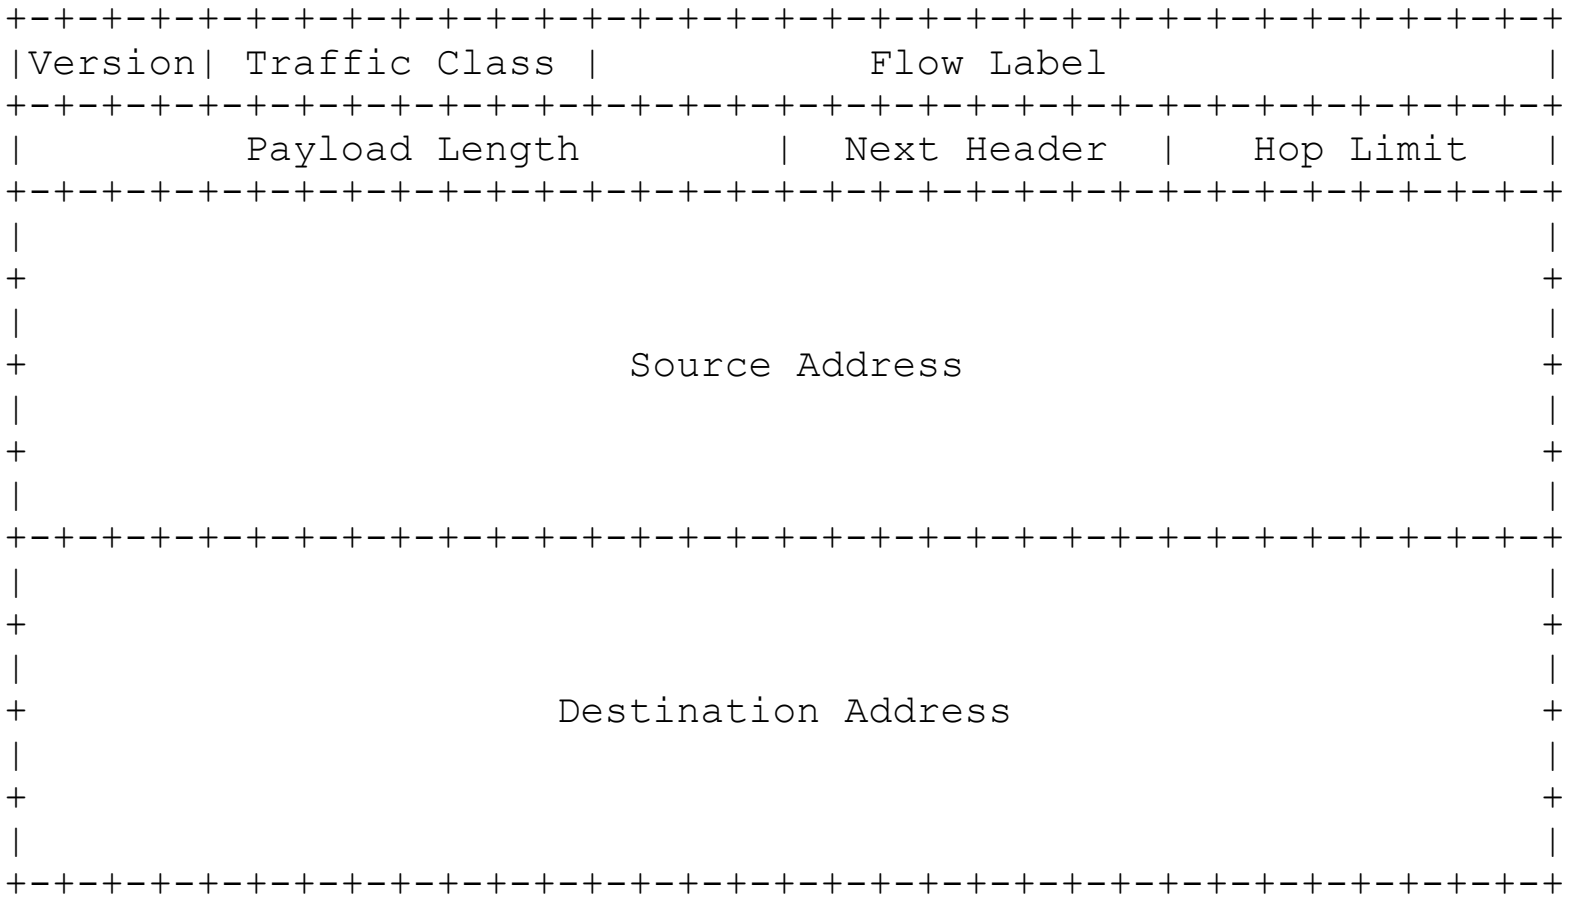

Hop Limit

- “Hop limit” field is similar to IPv4's “time to live.”
  - Number of routers the packet can pass through.
    - Limits packets from circulating indefinitely in an erroneous routing loop.
  - Field 8 bits so maximum number of hops just 255.
    - Is this enough?
    - IPv6 networks can contain a huge number of nodes.
- *Defn: The “diameter” of the network is the maximum number of hops between any pair of nodes.*

Diameter of the Internet?

- Nobody knows for sure.
- However... node count grows exponentially with the diameter.
  - Thus a large hop count field may be unnecessary
  - This was a debated topic in the design of IPv6.

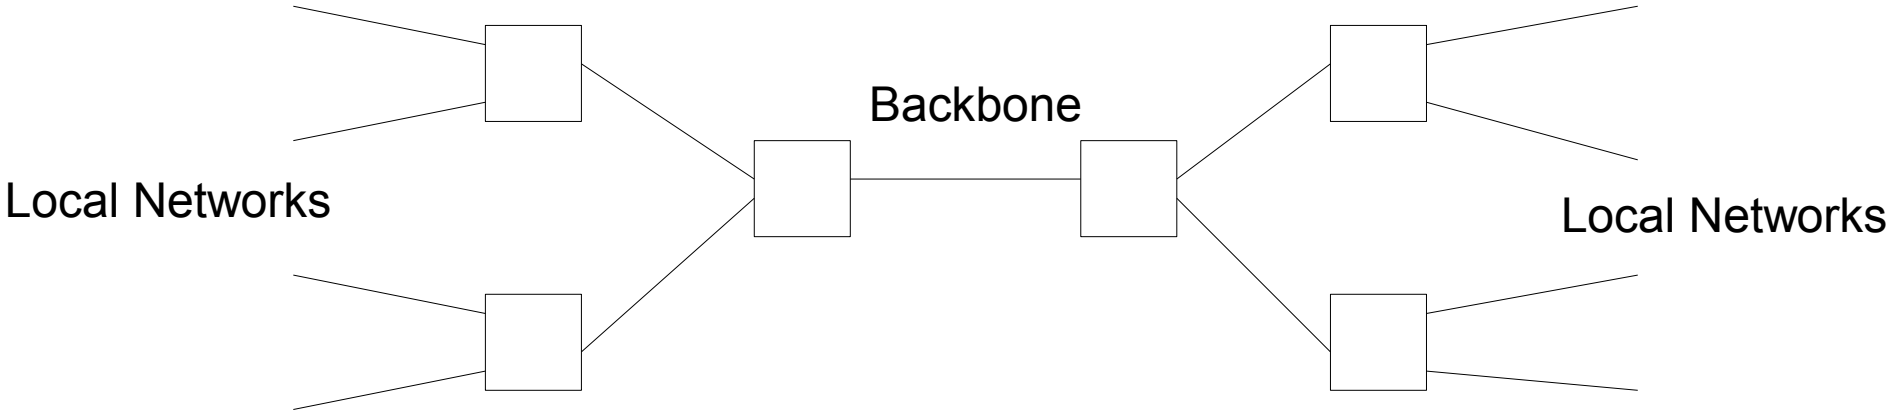

Flows

- Flows are an experimental feature (RFC-3697)
  - Not widely implemented (as far as I know)
  - Stream of packets designated as a “flow” by the source.
    - Often associated with a transport connection.
  - Defined by (source addr, dest addr, flow label)
  - Source must use flow label of zero by default.
  - Applications and transport protocols MUST have a means for setting the flow label.
- Intended to be used for QoS applications.

Extension Headers

- Various kinds are defined
  - Hop-by-hop (processed by routers).
  - Destination options (for all destinations listed in the routing header)
  - Routing (requested path through the network)
  - Fragment (source node *can* fragment packets)
  - Authentication (AH... part of IPsec)
  - Encapsulation Security Payload (ESP... part of IPsec)
  - Destination options (only for final destination)

Layout

This design allows for easy future expansion.
New extension headers can be defined at any time.

Extension Header Notes

- A few rules of interest... (see RFC-2460)
  - Headers only processed by destination
    - Except for “hop-by-hop options”... which must be first.
      - So routers don't have to dig around looking for it!
  - Headers must be processed in order given.
    - Some extensions may prohibit further processing.
  - Unrecognized header causes packet to be discarded.
    - An ICMPv6 message returned to sender.
    - “Don't process packets you don't fully understand.”

Option Headers Format

- Option Headers have a generic format.
  - Contain a variable number of “type-length-value” (TLV) encoded options.
  - New options can be defined later.

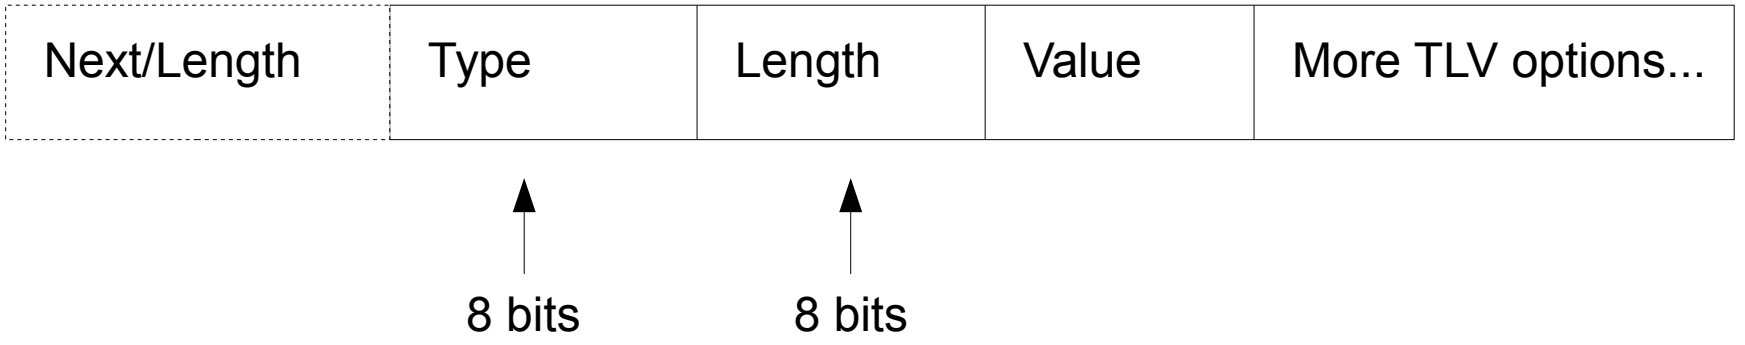

Options must be processed in order

Option Type Field

- Option type field has additional structure.

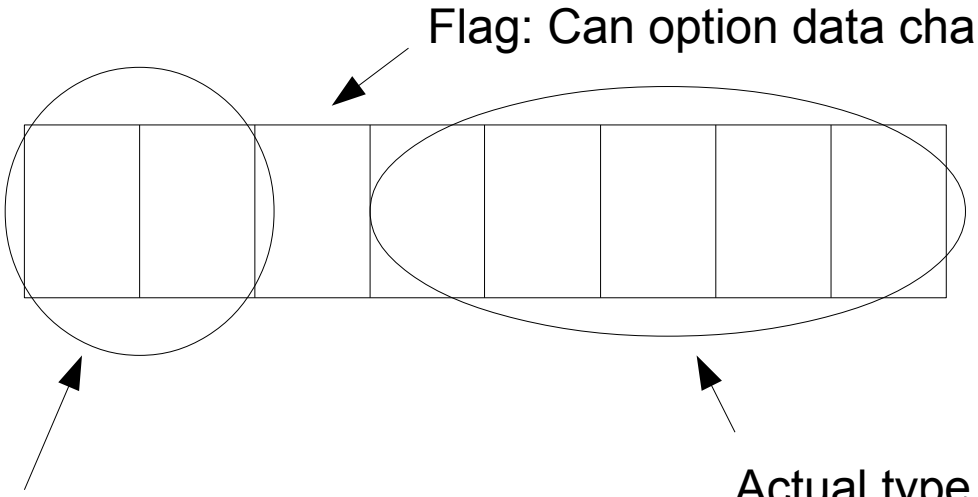

Neighbor Discovery Protocol

- Used for... (see RFC-2461)
  - Finding link layer addresses that correspond to an IPv6 address (like IPv4's ARP).
  - Finding routers on a given link.
  - Finding the link's prefix(es) (global addresses)
  - Link parameters
- ND thus combines the functionality of several separate IPv4 protocols.

Special ND Addresses

- Some ND functions are done before node has a normal address.
  - FF02::1 Link scope all nodes multi-cast (used to talk with all nodes on a particular link).
  - FF02::2 Link scope all routers multi-cast (used to talk with all routers on a particular link).
- Solicited node multi-cast (RFC-2373)
  - Suppose node addr = 4037::1:800:200E:8C6C
  - Solicited node addr = FF02::1:FF0E:8C6C
    - Prefix FF02::1:FF00:0000/104
    - Lower 24 bits from address above.

ND (ICMP) Message Types

- Router Solicitation
- Router Advertisement
  - Contains list of link prefixes
  - Options: hop limit, link MTU, etc, hosts should use.
- Neighbor Solicitation
  - Multi-cast to the solicited node address.
- Neighbor Advertisement
  - Unicast back to the requesting node.
- Redirect

Jumbograms

- Certain links support MTU values greater than 64 KiB.
  - Long fat pipes...
    - 10,000 mile connection at 80% speed of light, 10 Gbps.
      - ~67 ms transit time => ~670 million bits on the the wire.
      - ~56,000 traditional ethernet frames!
  - Using large packets reduces processing overhead
- Jumbograms implemented as a hop-by-hop option with a 32 bit payload length field.
  - No support needed for “normal” links.
- TCP and UDP need modification for this!