CSE561 – Internetworking

David Wetherall
djw@cs.washington.edu

Internetworking

- Focus:
  - Joining multiple, different networks into one larger network

- Service models
- Heterogeneity factors
- IPv4 and IPv6 formats
- Path MTU discovery
- Error reporting with ICMP
- Other glue: DHCP, ARP

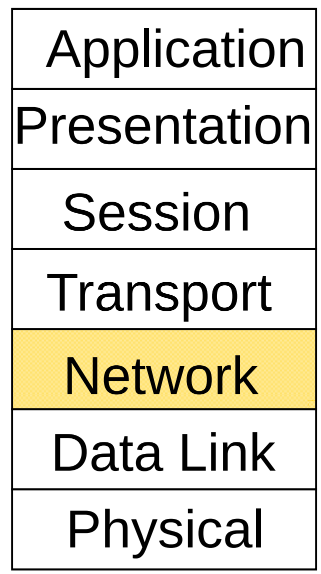

Network Service Models

- Datagram delivery: postal service
  - Also connectionless, best-effort or unreliable service
  - Network can't guarantee delivery of the packet
  - Each packet from a host is routed independently
  - Example: IP, switched Ethernet

- Virtual circuit models: telephone
  - Also connection-oriented service
  - Signaling: connection establishment, data transfer, teardown
  - All packets from a host are routed the same way (router state)
  - Example: MPLS, ATM, Frame Relay, X.25

- Q: How do we combine them?
- A: Not easily; with the lowest common denominator (IP)

Heterogeneity

- How else might networks differ?
  - Quality of service / priorities
  - Security
  - Maximum packet length

- How can we deal with these differences?
  - QOS: we’re screwed
  - Security: add it end-to-end
  - Packet lengths: path MTU discovery

IPv4 Packet Format

- Version is 4
- 32 bit addresses
- DiffServ field used to be TOS

- Header length is number of 32 bit words
- Limits size of options

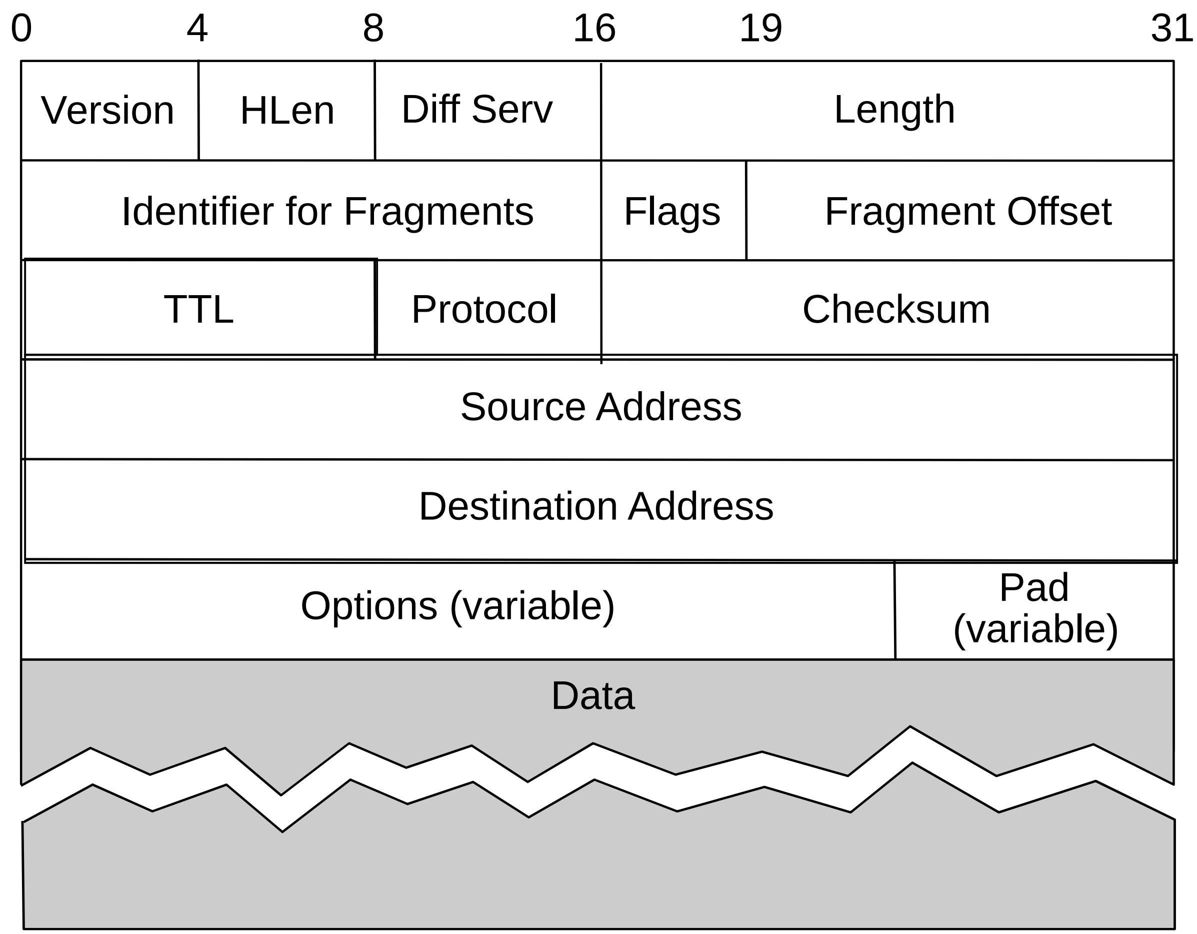

IPv6 packet format

- Version is 6
- 128 bit addresses
- Fields renamed and streamlined
- FlowLabel added
- Checksum gone

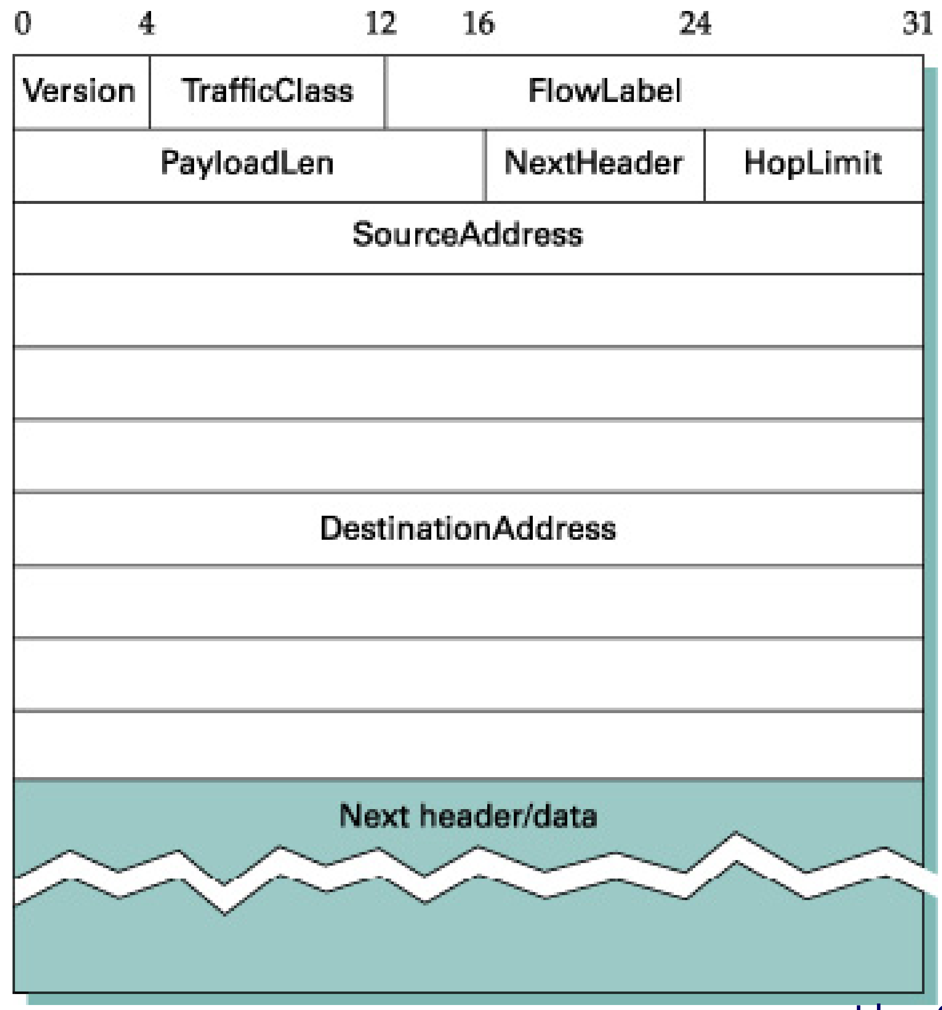

Fragmentation

- Sending small packets is wasteful, but don’t know a priori how large a packet will fit through the network
- One solution: network fragmentation
  - Network breaks large packets that are too large
  - Reassemble at destination (Why?)
  - Turns out to be bad (Why?)
- Better solution: discover largest packet for each a path (the “path MTU”) and tell the sender. (Downsides?)

Path MTU Discovery

- Path MTU is the smallest MTU along path
  - Packets less than this size don't get fragmented
- Hosts send packets, routers return error to host if packet too large
  - Use DF (Don't Fragment) header flag
  - Hosts discover limits, can fragment at source
  - Reassembly at destination as before
- Even better:
  - Host IP tells higher layer the right MTU to use; no fragmentation
  - At the cost of a layering violaiton

PMTU Games

- What is the best mechanism? Consider a path with M points of decreasing MTU and an initial MTU of L.

- How long would it take to find the PMTU if:
  1. Routers drop packet and return an error code to the sender
  2. Same, but to the receiver
  3. Routers now return the next hop MTU to the sender
  4. Same, but to the receiver
  5. Routers truncate packets

ICMP

- What happens when things go wrong?
  - Need a way to test/debug a large, widely distributed system
- ICMP = Internet Control Message Protocol (RFC792)
  - Companion to IP – required functionality
- Used for error and information reporting:
  - Errors that occur during IP forwarding
  - Queries about the status of the network

ICMP Generation

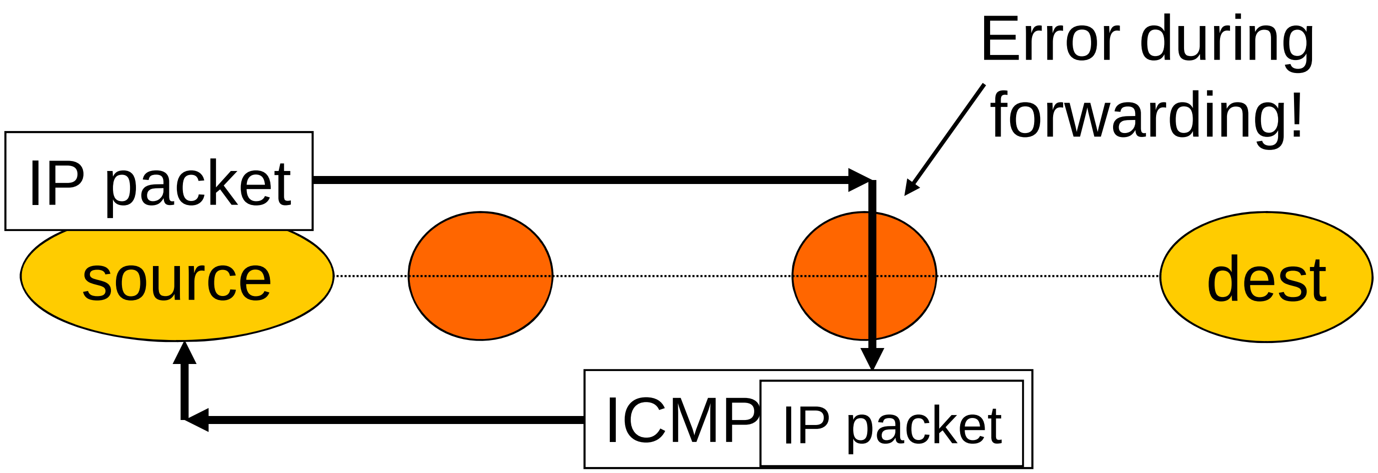

Common ICMP Messages

- Destination unreachable
  - "Destination" can be host, network, port or protocol
- Redirect
  - To shortcut circuitous routing
- TTL Expired
  - Used by the "traceroute" program
- Echo request/reply
  - Used by the "ping" program

- ICMP messages include portion of IP packet that triggered the error (if applicable) in their payload

Glue: Dynamic Host Configuration Protocol (DHCP)

- Q: How does a host get an IP address?
- A: DHCP, designed in 1993

- DHCP is widespread for the dynamic assignment of IP addresses, e.g., CSE, your cable company, …
- Host broadcasts a request; DHCP server responds with an IP
- Extensions:
  - Supports temporary allocation (“leases”) of IP addresses
  - DHCP client can acquire all IP configuration parameters

DHCP Interaction (simplified)

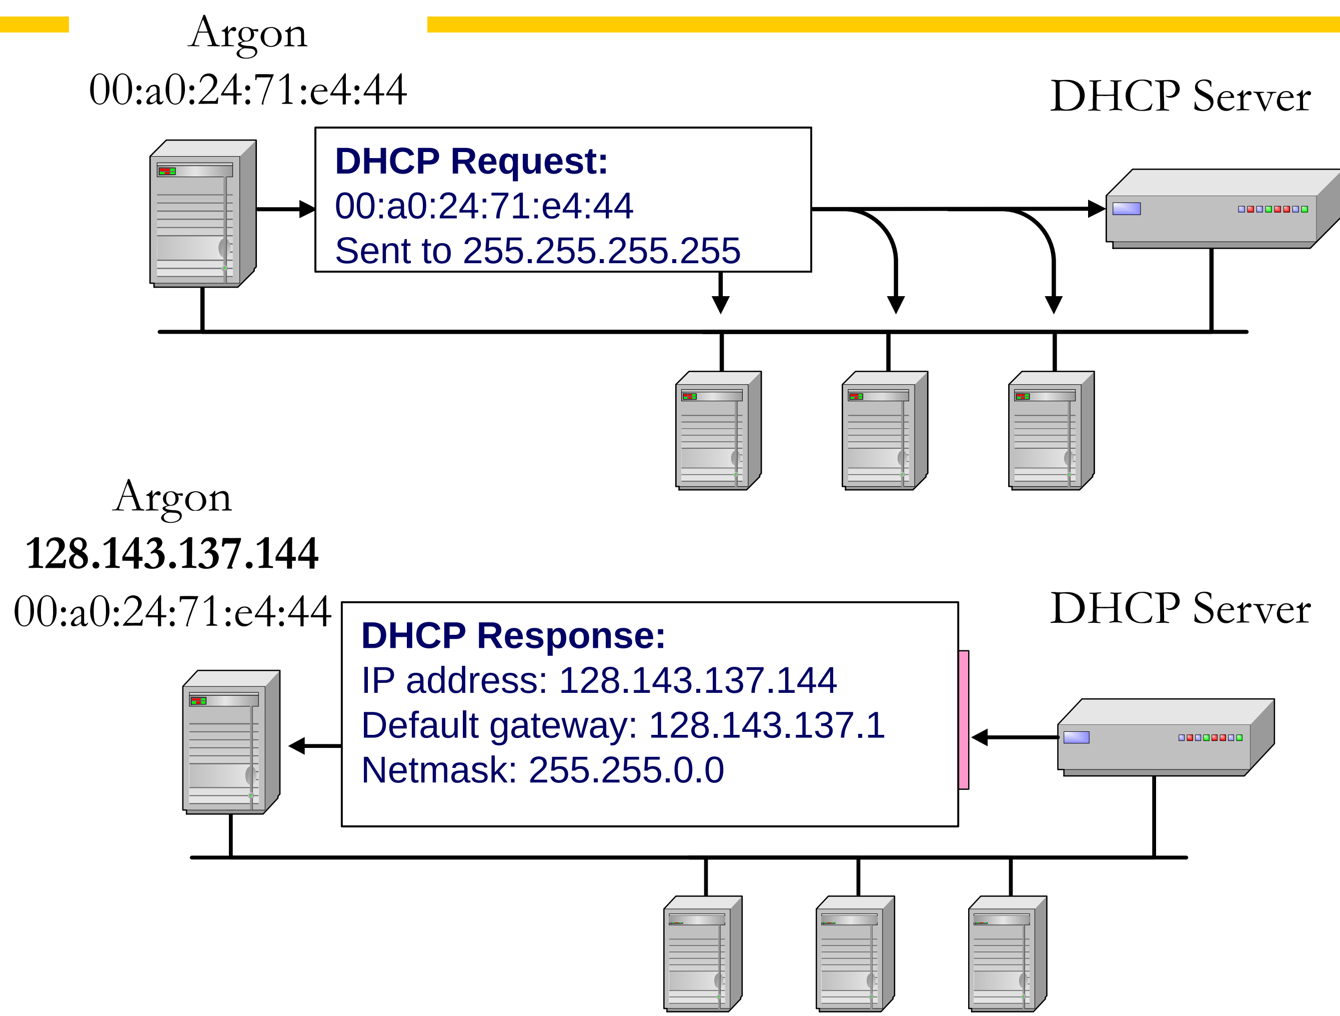

Address Resolution Protocol (ARP)

- Problem: We know a destination IP address, but how do we find the actual device on the LAN with that address?
- Solution: ARP maps local IP to local Ethernet addresses

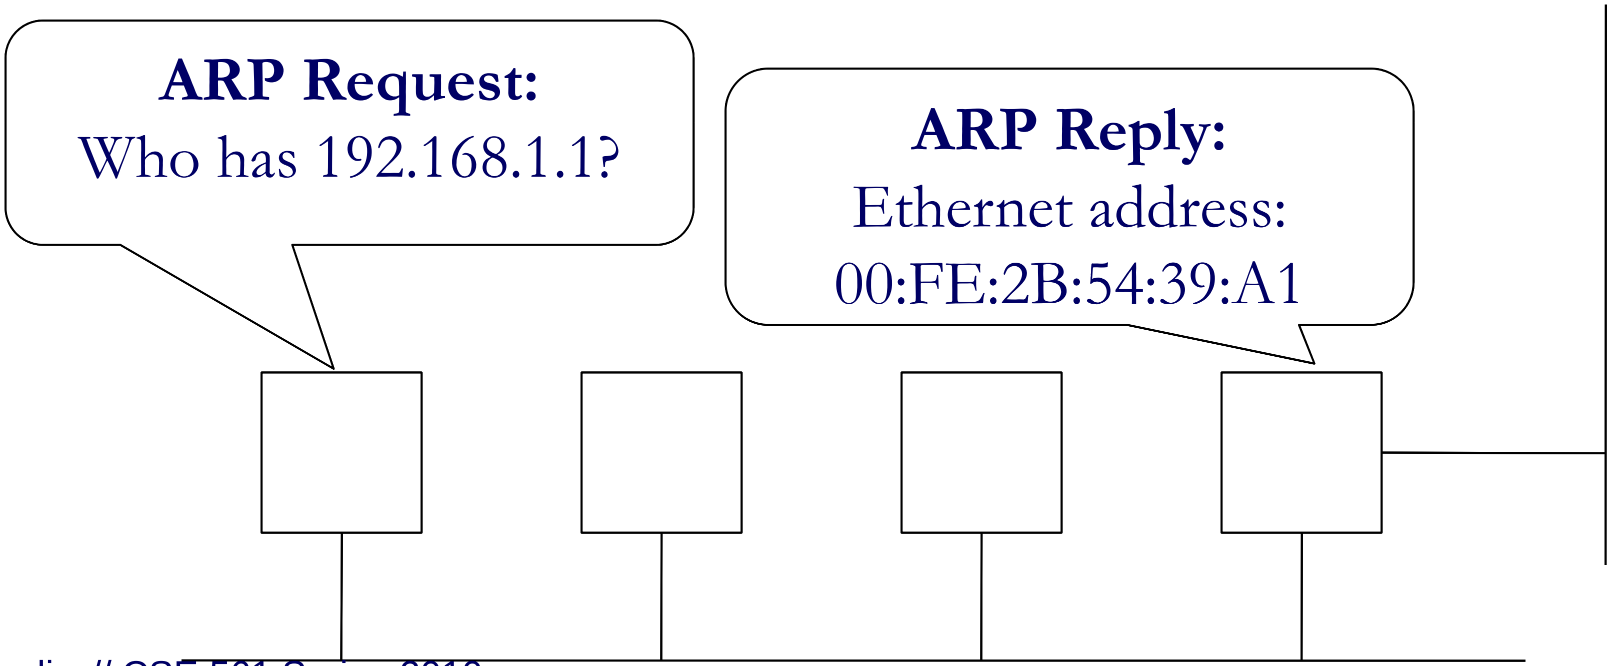

Clark paper questions

- So, why datagrams?
- Which of its goals does the Internet meet best?
- Which of its goals does the Internet meet least well?
- What is soft-state?
- What is fate-sharing?
- What are flows?