SC/CSE 3213 Winter 2013

L9: Control Protocols

Sebastian Magierowski
York University

Outline

- ICMP
- ARP
- DHCP
- NAT (not a control protocol)

Control Protocols

- IP is used to transfer data
- Network layer also contains protocols to help keep network operating
  – Or at least in some "control" of its behaviour

ICMP

- Internet Control Message Protocol
- Routers closely monitor the operation of their links
  – If something goes wrong they inform sender using ICMP
  – Encapsulated in IP
- About a dozen (non-deprecated) ICMP types
  – www.iana.org/assignments/icmp-parameters
- A few of the important ones…

Some ICMP Examples

- Type 3: Destination Unreachable
  - Sent (back to source) when router cannot locate the destination
  - Or if network allowing only small packets blocks delivery
- Type 5: Redirect
  - Packet seems to be routed incorrectly
    - Sender should update to a better route
- Type 12: Parameter Problem
  - An illegal value detected in IP header
    - A bug in router/host packet generation software
- Type 11: Time Exceeded
  - When IP packet's TTL field (time-to-live) drops to zero
    - Packets are looping
    - Counter values set too low

Traceroute

- Find route from local host to a remote host
- Time-to-Live (TTL)
  - IP packets have TTL field that specifies maximum # hops traversed before packet discarded
  - Each router decrements TTL by 1
  - When TTL reaches 0 packet is discarded
- Traceroute
  - Send UDP to remote host with IP TTL=1
  - First router will reply ICMP Time Exceeded Msg
  - Send UDP to remote host with IP TTL=2, …
  - Each step reveals next router in path to remote host

Traceroute Example: Home to CSE

```
BIT:Wes bit$ traceroute cse.yorku.ca
traceroute to cse.yorku.ca (130.63.92.157), 64 hops max, 52 byte packets
 1  192.168.1.254 (192.168.1.254)  1.996 ms  1.535 ms  1.536 ms
 2  d209-89-16-1.abhsia.telus.net (209.89.16.1)  25.814 ms  27.177 ms  29.906 ms
 3  173.182.202.129 (173.182.202.129)  26.653 ms  30.528 ms  29.346 ms
 4  chcgildtgr00.bb.telus.com (154.11.11.30)  63.154 ms  56.060 ms  56.194 ms
 5  * * *
 6  te0-7-0-18.ccr21.ord03.atlas.cogentco.com (154.54.87.189)  69.233 ms
    te0-7-0-0.ccr21.ord03.atlas.cogentco.com (154.54.83.229)  58.003 ms
    te0-3-0-1.ccr21.ord01.atlas.cogentco.com (154.54.29.17)  83.364 ms
 7  te0-2-0-2.ccr22.yyz02.atlas.cogentco.com (154.54.27.254)  81.562 ms
    te0-0-0-2.ccr21.yyz02.atlas.cogentco.com (154.54.6.153)  72.038 ms
    te0-1-0-3.ccr21.yyz02.atlas.cogentco.com (66.28.4.214)  82.603 ms
 8  te0-1-0-0.mpd22.yyz02.atlas.cogentco.com (154.54.43.166)  79.919 ms
    te0-2-0-0.mpd22.yyz02.atlas.cogentco.com (154.54.43.170)  81.081 ms
    te0-1-0-5.mpd22.yyz02.atlas.cogentco.com (66.28.4.58)  72.180 ms
 9  38.104.251.82 (38.104.251.82)  82.864 ms  81.696 ms  82.497 ms
10  york-hub-ut-hub-if-internet.gtanet.ca (205.211.94.42)  87.286 ms  82.603 ms  82.100 ms
11  yorku-york-hub-if-internet.gtanet.ca (205.211.95.134)  71.584 ms  71.696 ms  72.079 ms
12  core01-border.gw.yorku.ca (130.63.27.17)  81.488 ms  82.144 ms  88.743 ms
13  indigo.cs.yorku.ca (130.63.92.157)  83.161 ms  89.858 ms  85.132 ms
```

Home

Telus ISP

Cogent ISP

York

CSE

More ICMP Examples

- Type 0 & 8: Echo & Echo Reply
  - Used to see if a destination is alive, echo prompts echo reply
    - Employed by ping utility
- Type 13 & 14: Timestamp Request & Timestamp Reply
  - Like echo but arrival time of message and departure time of reply recorded
    - Useful for network performance measurement

Physical Addressing

- Internet uses IP addresses to direct messages
  - And encapsulates a packet in a frame
    - …a frame with a PHYSICAL ADDRESS
- But how do we know physical address of destination???
  - Physical addresses don't know anything about IP
- For example…
  - 192.32.65.7 wants to send to 192.32.65.5

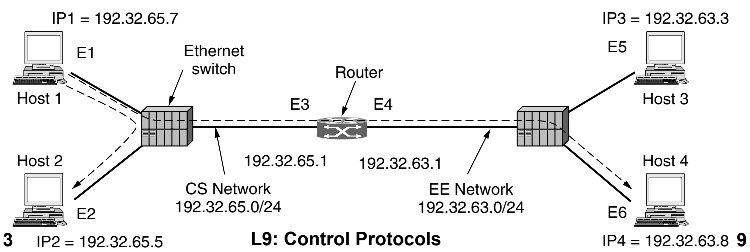

Map File?

- For example…
  - 192.32.65.7 wants to send to 192.32.65.5
- You can maintain a file that maps all IP addresses to all physical addresses in the network
  - Very complicated to maintain
- Come up with a protocol to automate the process
  - Address Resolution Protocol (ARP)

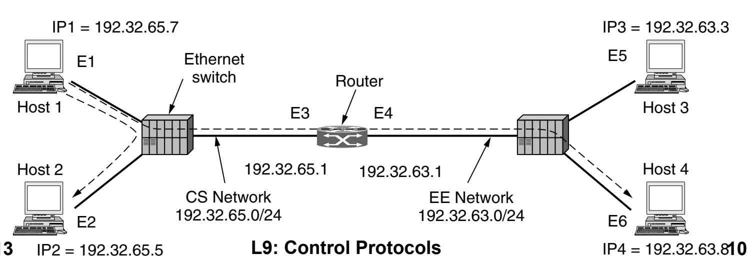

ARP

- Broadcast frame asking who owns destination IP address
  – Payload contains destination addr.: 192.32.65.5
- Each host on LAN checks if it has the requested address
  – Only Host 2 responds
  – Thus Host 1 learns 192.32.65.5 corresponds to E2 (put in ARP cache)
    - From broadcast E2 learns what IP address E1 corresponds to
- ARP handles this
  – Layer 2/3 protocol, implemented in device drivers

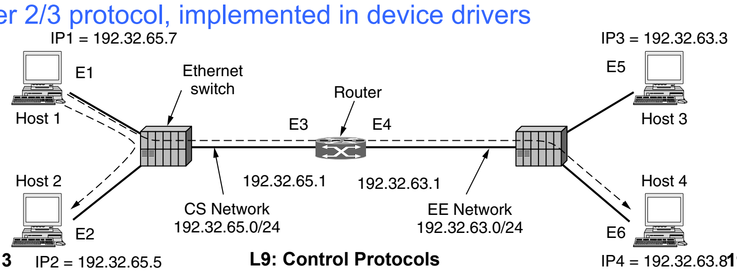

ARP Updates

- To keep fresh, ARP caches should time out every few minutes
- When a new machine is configured (obtains IP address) it should broadcast an ARP message looking for itself
  – This allows the remaining elements in the network to cache its IP/physical address relationship

ARP Internetworking

- What about sending to another LAN?
  - H1 to H4
  - H1 notices H4 IP address not in same network
    - H1 knows it must sent the packet to router
    - Router's IP address is automatically known lowest address in LAN
  - H1 uses ARP broadcast to discover E3
    - And hence can send the frame to the router
  - Router inspects IP and know which network to send the packet to
    - Will employ ARP to get H4's E6

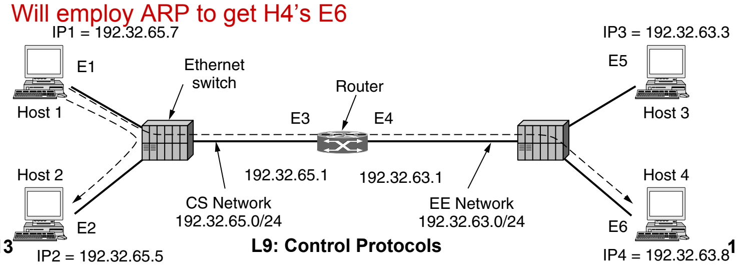

DHCP

- Dynamic Host Control Protocol
  - Application Layer, runs over UDP
- How do computers acquire an IP address?
  - New computer launches DHCP DISCOVER message
    - DHCP server (a must on every network) will receive this message (or router is configured to forward DHCP messages to known server)
  - DHCP server responds with DHCP OFFER message (containing assigned IP address)
    - DHCP DISCOVER packet contains senders Ethernet address
    - Allows DHCP server to know destination of OFFER message
- Hosts periodically ask for DHCP renewal to maintain address
  - A means of not running out of IP addresses
- DHCP also configures: mask, default gateway, DNS server

NAT

- Network Address Translation
- When you have a limited amount of IP addresses
- For example ISP might have 65,536 addresses
  - Easily handles 40,000 customers (with room to grow)
  - But what if each customer has on average 4 devices?
- Solution (until we get more addresses):
  - Let each customer in LAN have unique IP address
  - But anything leaving LAN appears as one constant IP address
- Three ranges of private IP addresses are available
  - 10.0.0.0 – 10.255.255.255/8 (16,777,216 hosts)
  - 172.16.0.0 – 172.31.255.255/12 (1,048,576 hosts)
  - 192.168.0.0 – 192.168.255.255/16 (65,536 hosts)

NAT @ Work

- Needs TCP or UDP
  - User sends with say 10.X.Y.Z from TCP port ABCD
  - NAT makes the association:
    - QRST: 10.X.Y.Z // ABCD
    - QRST just some TCP port
  - And sends out…
    - A TCP segment from port: QRST
    - Encapsulates in IP packet with the LAN’s one address, say: 198.60.42.12
  - Any response to 198.60.42.12 at port QRST…
  - …gets mapped back to:
    - 10.X.Y.Z at port ABCD