CS168 Fall 2014 Discussion 4

IP Addressing, IP Fragmentation, IPv4/IPv6… basically IP

Q0 - Warm Up

Find the binary representation, subnet mask, and address range of 192.168.0.0/13.

11000000 . 10101000 . 00000000 . 00000000
255.248.0.0
192.168.0.0 → 192.175.255.255

Which of the following addresses are part of this subnet?

123.100.0.5 192.128.69.5 **192.168.244.8** **192.175.100.0** 192.176.3.4

Q1 - IP Addressing

Berkeley.edu is the Provider AS for EECS, CS, Math, and Chemistry. Assume that the CIDR (Classless InterDomain Routing) addressing scheme is used.

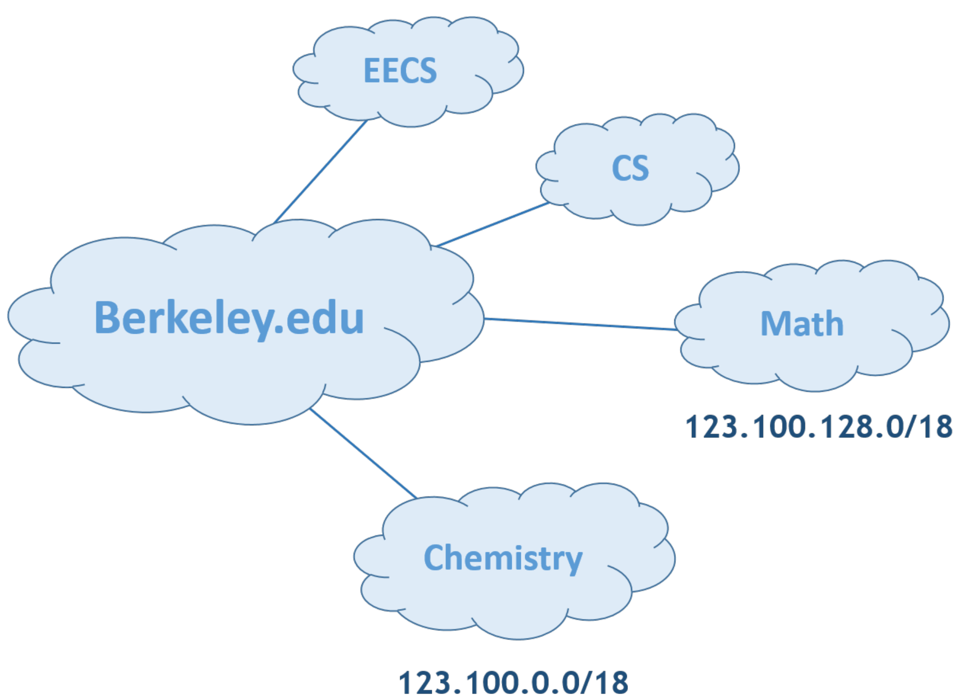

a) What range of addresses does Math hold? How many addresses are in this range?

123.100.128.0 → 123.100.191.255, 2^{14} addresses

b) 123.100.192.0/18 is reserved for EECS and CS. Assign equal halves of this address space to the two departments.

EECS:
01111011 . 01100100 . 11000000 . 00000000
123.100.192.0/19

CS:
01111011 . 01100100 . 11100000 . 00000000
123.100.224.0/19 (or vice versa)

c) What is the longest prefix for Berkeley.edu that encompasses all of Chemistry, Math, EECS and CS?

Take the first N bits common to all three prefixes, and take this as the network portion for Berkeley.edu

EECS & CS: 01111011 . 01100100 . 11000000 . 00000000
Math: 01111011 . 01100100 . 10000000 . 00000000
Chemistry: 01111011 . 01100100 . 00000000 . 00000000

The first diverging bit is the 17th bit. Thus, the answer is **123.100.0.0/16**.

(Alternatively, observe that the total size of the Berkeley AS must be at least $(3 * 2^{14}) > 2^{15}$, so you need at least 16 bits for the host portion. Taking the first 16 bits of any of the prefixes, you have 123.100.0.0/16)

d) You want to start a new department Floriology, but you foresee that no more than 50 people will enroll. Assuming one address per person, what prefix would you assign to it?

Floriology: Need 64 addresses
Must allocate from unclaimed chunk: 123.100.64.0/18
123.100.64.0/26

e) Your friend came up with the brilliant idea of starting yet another (slightly redundant) department, Mathematical Floriology (123.100.64.0/29), which is ***multi-homed*** from the existing Math and Floriology departments.

Why might it be a good idea for Mathematical Floriology to be multi-homed, instead of directly attached to only Math or Floriology?

In case one of Math or Floriology goes down (but not both).

How does this affect Berkeley.edu?

Berkeley.edu must now remember that 123.100.64.0/29 is attached to both Math and Floriology (i.e. one extra route to keep track of). This makes the forwarding tables of the routers in Berkeley.edu bigger. Breaks aggregation!

Q2 - IPv4

The following is the structure of an IP header, taken directly from your textbook.

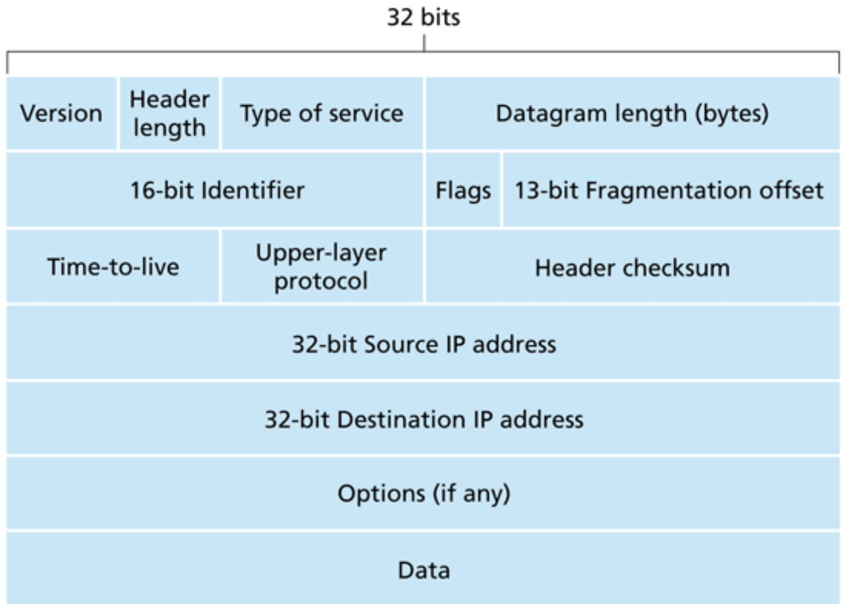

a) Which header fields must be updated before the router sends out a packet?

TTL, checksum

b) Suppose there is a bug in the IP router such that it no longer updates the time-to-live field. What problems might this cause?

If there is a loop in the network's routing configuration, then this packet will loop forever (until the incorrect routing configuration changes).

c) Suppose vendor A designs its routers such that it no longer updates the checksum. Its rationale is that end points commonly compute their own checksum anyway. What problems might this cause?

Per the end-to-end argument, lifting checksums from the network when end points already compute their own is a reasonable choice. However, not all router vendors ignore checksums, so if vendor A routers don't update checksums on an IP level, then other routers may think that the packet is corrupt even when it is not, resulting in a lot of dropped packets.

d) To accelerate packet forwarding, vendor A decides to always parse the last 4 bytes of the header for the destination address. What problems might this cause?

What happens when there are options?

Q3 – IP Fragmentation

Maximum Transmission Unit (MTU) is the size of the largest packet that a link can carry. Host A sends an **600 byte** IP packet (including header) to Host B, which is fragmented along the way. Assume the typical IP header length of 20 bytes.

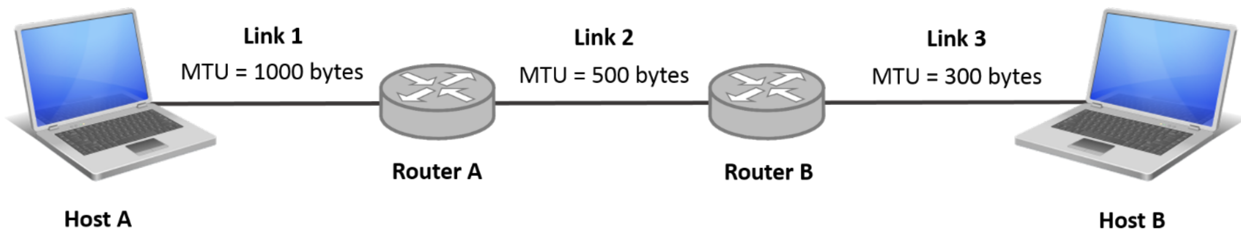

Recall that fragmentation flags occupy 3-bits in the format of R|DF|MF, which means <*reserved*, *don't fragment*, and *more fragments coming*>. Fragmentation offsets are in terms of 8-byte units.

a) The packet fits within the MTU of Link 1 and arrives at Router A. What are the resulting fragments that traverse Link 2? For each fragment, identify the total length (including header), flags, and offset.

Original payload = 600 byte total length – 20 byte header = 580 bytes

F1: total length = (20 byte header + 480 byte payload) = 500 bytes, flags = 001, offset = 0
F2: total length = (20 byte header + 100 byte payload) = 120 bytes, flags = 000, offset = 480/8 = 60

b) The fragments arrive at Router B. What are the resulting fragments that traverse Link 3?

F1a: total length = (20 byte header + 280 byte payload) = 300 bytes, flags = 001, offset = 0
F1b: total length = (20 byte header + 200 byte payload) = 220 bytes, flags = 001, offset = 280/8 = 35
F2: total length = (20 byte header + 100 byte payload) = 120 bytes, flags = 000, offset = 480/8 = 60

Observe that fragmentation offsets in F2 stays the same.

c) Why is the MF flag needed?

Packets can arrive out of order; MF flag tells the end host which fragment is the last one.

d) Why can't we just number our fragments instead of keeping track of fragmentation offsets?

How else would you order fragments of fragments? (Consider the further fragmentation of F1)

f) IP fragmentation is removed altogether in IPv6. Why might this be?

Fragmentation and re-assembly is time-consuming. The alternative of a router dropping the oversized packet, then sending a "Packet Too Big" error message back to the host is simpler and enables faster forwarding.