CS118 Final

TOTAL POINTS

85.5 / 140

QUESTION 1

1 20 pts

1.1 1.1 (HTTP) 2 / 2

- **- 0 Not graded (free credit)**

1.2 1.2 (Ethernet) 0 / 2

- + 2 no correct statements
- + 2 only select d or e
- **+ 0 otherwise**

1.3 1.3 (HTTP) 2 / 2

- + 2 select e
- **+ 1.5 select e & one of other choices**
- **+ 0.5 Adjustment for choice (d). It is not correct (it is used by DNS), but we accept it for the grading purposes**
- + 1 select e & and two more choices
- + 0.5 select e & and more than two more choices
- + 0 Incorrect

1.4 1.4 (TCP) 0.5 / 2

- - 0.5 a
- **- 0.5 b**
- **- 0.5 c**
- **+ 2 d**
- **- 0.5 e**
- + 0 Nothing marked

1.5 1.5 (Multicast) 1 / 2

- + 2 select a
- **+ 1 select a &(b|c)**
- + 0 otherwise

1.6 1.6 (Mobile-IP) 3 / 3

- **+ 1 PA: everybody else knows. permanent id**
- **+ 1 CA: temporary id**
- **+ 1 foreign network assigns CA**
- + 0 wrong answer

1.7 1.7 (Mobile-IP-2) 3 / 3

- **+ 3 ip-ip tunneling**
- + 2.5 Using VPN
- + 0 wrong answer
- + 2 no clear explanation.

1.8 1.8 (NAT) 0.5 / 4

- - 0 Any of the following: use relay, port redirection, D-NAT, reverse NAT, uPnP
- - 1.5 Correct direction, but incorrect/incomplete explanation
- **- 3.5 Good attempt, but incorrect**
- - 4 Missing
- - 2 Correct in general, but no specifics given

QUESTION 2

2 TCP 20 pts

2.1 2.1 TCP graph 5 / 8

- - 0 Correct
- - 1.5 Wrong CWND in slow start phase
- - 1.5 Wrong CWND (pkt 4, 5)
- **- 1.5 Wrong SSthresh after packet lost**
- - 1.5 3 RTT (RTO) is not considered or wrong RTO
- - 1.5 Wrong CWND (pkt 22nd)
- **- 1.5 Pkt 48 lost should trigger time out**
- - 1 incorrect seq number

2.2 2.2 Number of RTTs 3 / 4

- - 0 Correct (18~24)
- **- 1 16~17 or 25~26**
- - 0 13~15 or 27~29
- - 0 ~12 or 30~

2.3 2.3 TCP Facts 2 / 8

- - 0 Correct:

(1) Maximum payload size is 65,495 bytes or 2^16 - 41 (or 2^16)

(2) Sequence number range is 0 to 2^32 -1

(3) 2 * (2^16 -1) or 2 * (65,536-1) or 131,072 - 2

(4) 2^16 - 1

- 2 **(1) The maximum TCP payload size is wrong**
- 2 (2) The sequence number range is wrong
- 2 **(3) Number of distinct applications is wrong**
- 2 **(4) Number of applications is wrong**

QUESTION 3

3 OSPF 20 pts

3.1 3.1 Routing table 2 / 8

- 4 400 ms
- 4 **12 packets**
- 0 Correct
- 2 **Partial credit if answer is 200 ms**
- 2 Partial credit if answer is close to 12

3.2 3.2 When detect 4 / 4

- 3 HELLO message
- 1 30 sec
- 0 **Correct**

3.3 3.3 When converge 0 / 4

- 2 If mention longest path or 4 hops
- 4 **Incorrect**
- 0 Correct

3.4 3.4 Why Intra-AS 4 / 4

- 2 I point is correct: 1, global topology; 2, policy
- 4 Incorrect/blank
- 0 **Correct**

QUESTION 4

4 BGP 20 pts

4.1 4.1 Path1 5 / 5

- 0 **Correct**
- 2 One incorrect path (violates no-valley policy)
- 4 More than one incorrect path, violation no-valley policy
- 0.5 Incorrect/not specified preference
- 5 Incorrect/missing

4.2 4.2 Path2 4 / 5

- 0 Correct
- 1 **One valid path missing**
- 2 Two valid paths missing
- 3 More than two paths missing
- 1 Incorrect/missing preferred path
- 2 1 or more invalid paths listed
- 5 Incorrect/missing

4.3 4.3 Adjust local preference 1 / 5

\+ 3 Changing or adding providers, connecting to new peers or similar

\+ 3 Setting local preference and/or adjusting next hop

\+ 3 Updating AS-PATH (including more of self or similar) / updating cost

\+ 1 **Good attempt, but incorrect**

\+ 0 Incorrect/missing

4.4 4.4 Adjust remote preference 1 / 5

\+ 3 Changing or adding providers, connecting to new peers or similar

\+ 3 Setting local preference and/or adjusting next hop

\+ 3 Updating AS-PATH (including more of self or similar) / updating cost

\+ 1 **Good attempt, but incorrect**

\+ 0 Missing or incorrect

QUESTION 5

5 Network 20 pts

5.1 5.1 Datagram sequence 6 / 8

\+ 8 9 items, deduct 1 point for each wrong item

\+ 0 if all items are correct, but the order is wrong, then 4

- 2 if there are frames for IP reply, deduct 2

\+ 0 Totally wrong

\+ 6 **Point adjustment**

5.2 5.2 Number of broadcasts 4 / 6

\+ 2 **for s1, 2 times**

\+ 2 for s2, 3 times

\+ 2 **for s3, 1 time**

\+ 0 all wrong

5.3 5.3 Routing table example 4.5 / 6

\+ 2 1 subnet correct

\+ 4 2 subnets correct

\+ 6 **3 subnets correct**

\+ 0 Incorrect / not a valid routing table

- 1 Wrong/missing gateway

- 1 Missing masks
+ 0 0.75 point for each host that can be reached
- **0.5 missing one host 192.168.2.4 or 192.168.3.1**
- **1 wrong/missing interface**

QUESTION 6

6 IP 20 pts

6.1 6.1 Garbled IP 1 / 5

- **1 Missing that IPv4 IHL is wrong (should be 5)**
- **1 Missing that IPv4 Total Length is wrong (should be 1020)**
- **1 Missing that IPv4 flags are wrong (first bit should be 0)**
- 1 Missing that IPv4 TTL is wrong (max value 255)
- **1 Missing that IPv6 Payload Length is wrong (should be 1000)**
- 0 Wrong explanation about correct value
- **0 Correct field identified as wrong field**
- 0 Correct

6.2 6.2 IPv4 fragmentation 3 / 5

- **0.5 Wrong payload length (First packet payload length not divisible by 8. First packet payload length should be 776, or total length = 796; second packet payload length should be 224, or total length = 244)**
- **0.5 Wrong offset (round up error; not count in 8 bytes)**
- 1 Second packet MF is not 0 (or student answered 001) / First packet MF is not 1 (or student answered 000)
- **1 Did not mention ID/wrong ID**
- 1 Did not mention fragment offset
- 0 Correct

6.3 6.3 IPv6 fragmentation 5 / 5

+ **2 Positive answer**
+ **3 Sender split payload and use extension header**
+ 1 fragmentation using 6to4
+ 0 Incorrect

6.4 6.4 Not valid IPs 4 / 5

- 2.5 Not choosing a
- 2.5 Not choosing h
- **1 Choose up to three other options**
- 2 Choose three more other options
- 0 Correct

QUESTION 7

Misc 20 pts

7.1 CSMA/CD 4 / 4

- **0 Correct Answer (1/32)**
- 2 Miss by 1 error (eg. ans is 1/31 or 1/16) or Partially correct (eg. correct window)
- 4 Incorrect or Missing Answer

7.2 DHCP 0 / 4

- 0 2 correct reasons
- 1 1 correct, 1 partially correct reason
- 2 1 correct reason or 2 partially correct reasons
- 3 1 partially correct reason
- **4 Both reasons incorrect or missing**

7.3 ARP 4 / 4

- **0 2 Correct reasons**
- 1 1 correct, 1 partially correct
- 2 1 correct reason or both partially correct
- 3 1 partially correct reason
- 4 Both incorrect reasons

7.4 RTS/CTS 4 / 4

- **0 Answer is no with correct explanation**
- 2 Answer is no but with incorrect or missing explanation
- 2 Correct explanation but wrong deduction
- 4 Incorrect or Missing Answer

7.5 Traceroute 3 / 4

- 0 Correct
- **1 Correct but did not explain the use of TTL increment/expiration**
- 1 Correct but did not explain the use of ICMP/Port Unreachable message
- 4 Incorrect or Missing answer

CS118
Spring 2017 Final Exam

2 hour 50 minutes
Close book and closed notes,
except a SINGLE piece of paper as a cheat sheet.

NO use of any device except calculators.

- This exam has pages, including this cover page. Do all your work on these exam sheets. NO EXTRA PIECES OF PAPER WILL BE ALLOWED.
- Cross out all the scratch work that you do not want to be counted as part of your answer before you submit the exam.
- Be *specific, clear,* ***concise*** in your answers, and *explain* your answers.
- When the answer to a problem is not immediately clear, do not simply dump everything, relevant or irrelevant, on the paper. Irrelevant answers may lead to point-deduction as they show the lack of understanding of the problem.

Problem 1 (20 points)

1.1. Circle one or more statements about HTTP 1.1 that are CORRECT:

(a) The HTTP response messages never have an empty message body.
(b) Two distinct HTTP requests can be carried in a single TCP segment of a non-persistent HTTP connection
(c) Pages **http://site1/page1** and **https://site1/page2** can be retrieved over the same persistent HTTP connection.
(d) Pages **http://site1/page1** and **https://site2/page3** (site1 and site2 has same IP) can be retrieved over the same persistent HTTP connection.
(e) The "Date:" field in the HTTP response message header indicates when the object in the response was last modified.

1.2. Circle one or more statements about Ethernet interconnection that are CORRECT:

(a) Hosts are aware of presence of switches.
(b) An Ethernet network cannot detect collisions until it has computed a checksum over the frame.
(c) Ethernet switch actively sends ARP request to learn MAC addresses.
(d) Ethernet switches learn addresses by looking at the destination address of packets as they pass by.
(e) Ethernet hubs learn addresses by looking at the addresses of packets as they pass by.

1.3. Which protocol is NOT used when Bob uses his new laptop to request the CNN Web page (using HTTP)?

(a) DNS	(c) DHCP	(e) SMTP
(b) IPv4 or IPv6	(d) UDP	(f) TCP

1.4. Circle one or more statements about TCP that are TRUE?

(a) The size of the TCP receiver window (rwnd) never changes throughout the duration of the connection.
(b) The number of unacknowledged bytes can never exceed the size of the receiver window (rwnd).
(c) TCP provides in-sequence and best-effort delivery.
(d) A packet with SYN flag set cannot carry data segment.
(e) The maximum size of the congestion window (cwnd) equals to the maximum sequence number.

1.5. Which of the following statements about multicast are true?

(a) 225.0.0.5 is a valid multicast group.
(b) Using IGMP membership queries, router discovers knows how many members are in each multicast group.
(c) IGMP snooping is used by routers to manage multicast groups.

1.6. What is the difference between a permanent address and a care-of address? Who assigns a care-of address?

- permanent address: mobile's address in home network used to always reach mobile device
- care-of address: mobile's address obtained from a visited network as device moves

Care-of-address is given by the visited network

1.7. How does a phone with US number roaming in China can access google services bypassing the Great China Firewall which drops all IP packets destined to any of Google's servers?

China's firewall acts as a hidden terminal issue so all of Google's servers are hidden from your phone. So if google & your phone can access a router B, then one can simply redirect through B to get to google. Or one can use a VPN. It encrypts any packets prior to entering public internet, so firewall will not detect any google packet [illegible] & let it go.

1.8. How can a user call in to Skype host behind a NAT router? →UDP NAT on page 1 col 1

NAT router is ~~was to conserve IP addresses~~ to hide networks behind them. You would use port numbers in router to identify hosts in Network, so send UDP datagrams to initiate the call in to skype host.

like HW,

Problem 2 (20 points)

Consider a TCP Reno (i.e., one that implements fast retransmit and recovery) flow that has exactly 50 segments to send. Assume that during the transmission, exactly four packets are lost: the 4th, 5th, 22nd, and 48th; no other losses occur. Initial CWND is 1 MSS and initial slow start threshold (SSthresh) is 8 MSS. Consider fixed value of RTT and RTO = 3*RTT.

2.1 Using the graph below, plot the approximate evolution of the congestion window as each segment is sent. Each tic of x-axis represents a single round trip time (RTT) and each tic of y-axis indicates 1 MSS.

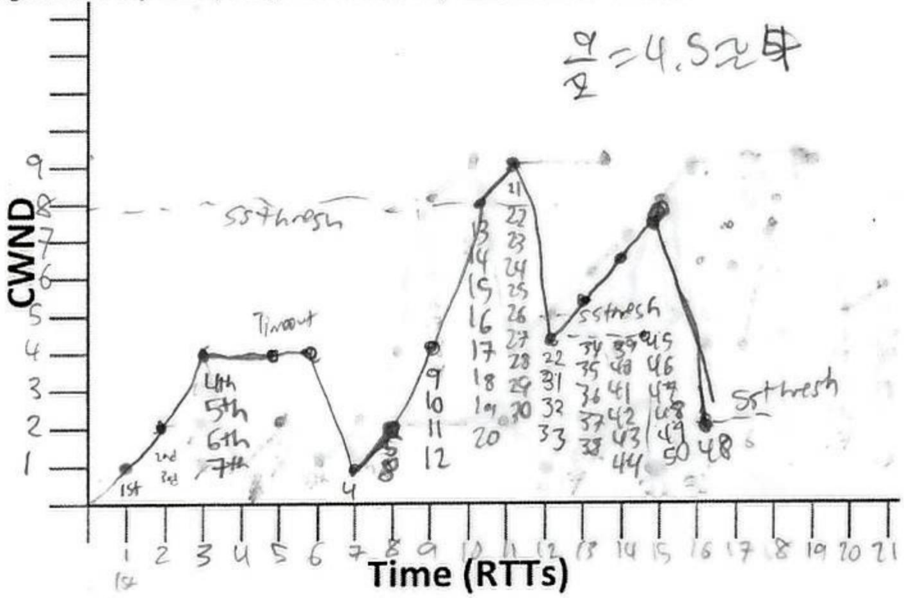

2.2 How long (in number of RTTs) will it take to send all segments and receive all corresponding ACKs?

16 RTTs

2.3 Finish the following statements about TCP protocol:

The maximum size of a TCP payload is 512 bytes

The range for TCP sequence numbers is $0 - 2^{32}$

For a computer with two IP addresses (e.g., one for wireless and one for wired), there could be 2 maximum number of distinct TCP server applications.

To prevent anybody else to start a TCP server application, one need to start at least 2 number of applications, each creating one socket, binding, and listening on a single port.

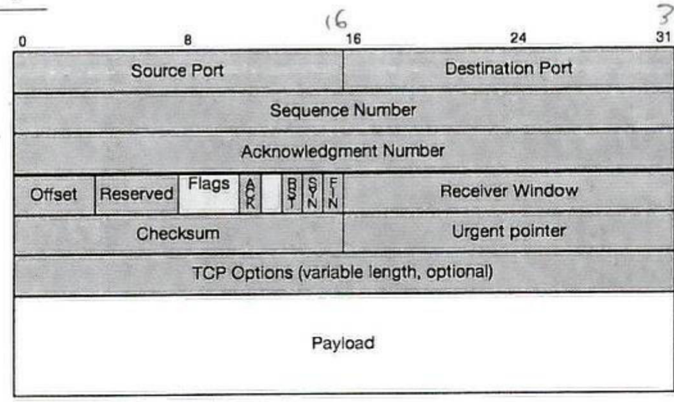

Link state

Problem 3 (20 points)

Consider a network with 9 routers connected as a grid in the figure. The routers use OSPF routing protocol. The numbers above each link indicate link costs. When a router has to choose between two or more equal cost paths to the same destination, it breaks the tie by picking the one with the lower node ID (in alphabetic order).

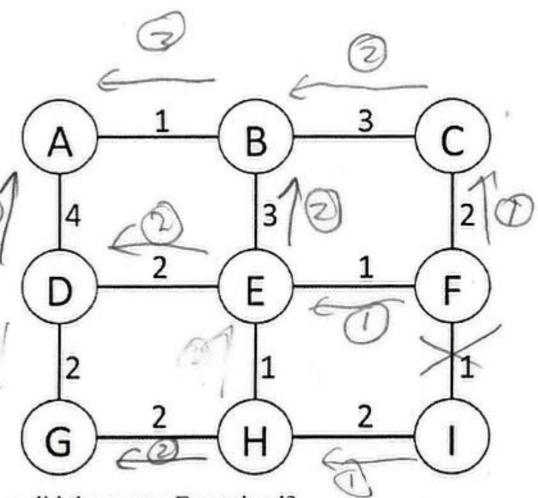

Assume that,

- Initially (T=0) the routing tables of all routers are **empty**.
- Propagation delay between any two connected routers is **200 ms.**
- First routing update messages will be sent out at T=0.
- Ignore processing, queuing, and transmission delays.

3.1 Assume that there is no link failure and no packets are lost. How long did it take to make the converged routing table of router E? How many packers did the router E received?

#	N'	A	B	C	D	E	F	G	H	I
0	e	∞	3,e	∞	2,e	[illegible]	1,e	∞	1,e	∞
1	ef	∞	3,e	3,f	2,e	[illegible]	1,e	∞	1,e	2,f
2	efh	∞	3,e	3,f	2,e			3,h	1,e	2,f
3	efhd	6,D	3,e	3,f	2,e			3,h		2,f
4	efhdi	6,D	3,e	3,f				3,h		2,f
5	efhdib	4,b	3,e	3,f				3,h		
6	efhdibc	4,b		3,f				3,h		
7	efhdibcg	4,b						3,h		
8	efhdibcgD	4,b								

8 packets
1600 ms

200ms • 8 packets = 1600ms

3.2 Assume that a link between router F and router I is down after every router has converged routing table. How and approximately when (relative time) F and I will detect this failure?

- will detect it with controlled flooding. F&I send Hello msg to each other & will timeout = no response so link failed.
- will wait 30 seconds

3.3 How long would it take for the routing to converge again (relative to failure detection, assuming F and I detected simultaneously)?

3 hops • 200 ms = 600 ms → If I & F flood simultaneously.
If F & I go one at a time, then (3+2 hops) • 200ms = 1000ms
(see diagram for marks on hops)

3.4 OSPF is considered an intra-AS routing protocol. Please least at least two reasons why people don't use OSPF instead of BGP for global routing?

- Global routing usually crosses AS, requiring inter-AS routing protocol (for outside subnet/AS)
- OSPF uses Link state but BGP uses distance vector, latter of which is more efficient for global routing & updating all routing tables

Problem 4 (20 points)

Assume we have an AS topology annotated with AS relationship as shown in the figure:

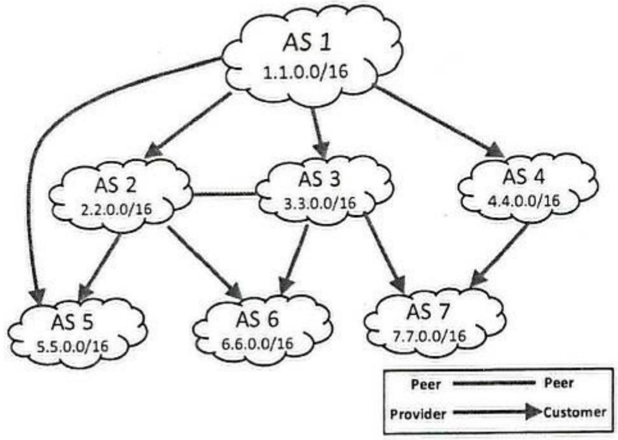

4.1 In order to reach destination prefix 4.4.0.0/16 in AS4, list all valid path(s) AS2 can take. Among these path(s), which valid path does AS2 prefer the most?

Most prefered : → AS2 → AS 1 → AS 4

~~AS2 → AS 3 →~~ AS1 → AS4

4.2 In order to reach destination prefix 5.5.0.0/16 in AS5, list all valid path(s) AS7 can take. Among these path(s), which valid path does AS7 prefer the most?

- AS7 → AS3 → AS1 → AS2 → AS5
- AS7 → AS4 → AS1 → AS2 → AS5
- AS7 → AS3 → AS1 → AS → 5
- ~~AS7 → AS3 → AS2 → AS5~~
- AS7 → AS4 → AS1 → AS5 → prefered most

4.3 List at least two mechanisms how AS7 can adjust its preferred path to reach 5.5.0.0/16 in AS5

1. eBGP
2. OSPF

4.4 List at least two mechanisms how AS5 can adjust AS7's preferred path to reach 5.5.0.0/16?

1. iBGP
2. OSPF

Problem 5 (20 points)
Consider three LANs interconnected by two routers, as shown in the figure.

For each question assume that:
- all ARP tables (ARP caches) on routers and hosts are initially empty
- all switch tables on switches are empty too.

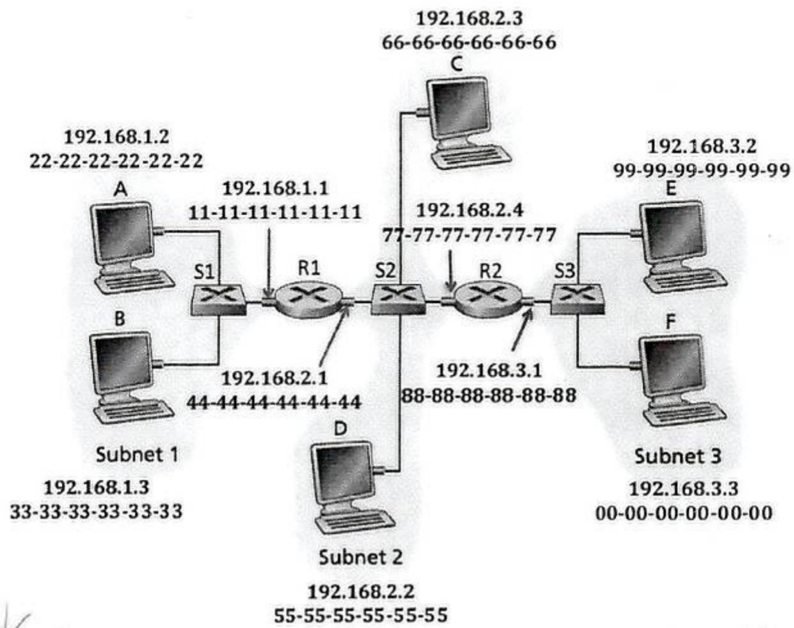

5.1 When sending an IP datagram from Host E to Host B. Write a sequence of sent Ethernet frames that hosts and routers will generate (ignore switches for this question). (dest, src)

write src & dest MAC

* Note: since MAC addresses are unique, will shorten XX-...-XX to simply XX for all problems

1. E sends (FF, 99) Broadcast
2. R2 sends (99, 88)
3. E sends (88, 99) data
4. R2 sends (FF, 77) broadcast
5. R1 sends (77, 44)
6. R2 sends (44, 77) data
7. R1 sends (FF, 11) Broadcast
8. B sends (11, 44)
9. R1 sends (44, 11) data
10. B sends (99, 44) reply

5.2 Host A, B, C, and F each sends an IP datagram to Host D in that order with long enough time gaps in between. Assume routers and hosts create entries in ARP cache ONLY when ARP response is received. How many times have switch S1, S2, and S3 broadcast a data frame?

= response back! ✓ = count for broadcast

S1: 2 broadcasts
S2: 2 broadcasts — 5 total
S3: 1 broadcast

3 broadcasts

Broadcast = Flood w/ FF-...-FF to get MACs

5.3 Show an example of the routing table on router R1 that allows IPv4 communication between all hosts and routers.

Source R1:

Dest	IPv4
A	22
B	33
C	66
D	55
E	99
F	00
R2	77

Dest	Gateway	Mask
192.168.1.2	192.168.1.1	/31
192.168.1.3	192.168.1.1	/32
192.168.2.4	192.168.2.1	/30
192.168.2.3	192.168.2.1	/32
192.168.2.2	192.168.2.1	/32
192.168.3.2	192.168.2.1	/31
192.168.3.3	192.168.2.1	/32

$2^{16} = (2^8)^2 = ((2^4)^2)^2 = ((16)^2)^2$ 16 × 16 = 96 + 160 = 256

Problem 6 (20 points) The figure below shows a potentially garbled IPv4 and IPv6 datagrams.

(1)

4	6	0	1200
4859		100	0
576	TCP	checksum	
192.168.1.5			
131.179.128.5			
1000 bytes of data			

192 ≠ 00000011

(2)

6	0	~~12345~~
1200	TCP	60
2606:6000:628c:1000:4884:3afa:e8e4:3d73		
2607:f010:2e8:228::ff:fe00:152		
1000 bytes of data		

0xA

6.1 Based on your knowledge of the valid ranges of values for each of these fields, identify which, *if any*, contain erroneous values. The IPv4 and IPv6 header formats are shown at the end of this problem for your reference. Please provide a brief explanation for each of all the erroneous values you identified.

- TTL ≠ 576 since max TTL for 8 bits is $2^8 - 1 = 255$ sec
- Largest amount of data that can be sent is 500 bytes, not 1000, at least without fragment flag (which is 0)
- IPv6 datagram should not say TCP protocol. That field is actually the next header location, protocol not in header.

6.2 Suppose that those datagrams need to be sent through a link with MTU 800 bytes, which would require IP fragmentation. How is fragmentation achieved in IPv4 case (highlight the mechanisms, and calculate values of the fields)?

- Fragmentation achieved by setting fragmentation flag to 1 & setting the offset to 8 Byte units. Router fragments IP packets if next link has smaller MTU.

1. First will send 780 bytes + 20 bytes of header (780 data bytes)
2. Then send 220 bytes with 20 more bytes of header (240 total)

6.3 Is fragmentation possible in IPv6 case? If so, describe the mechanism (you do not need to calculate any values for this question)

It is possible to fragment in IPv6, but it is done with extension headers. Payload is divided into pieces with a new base header for each fragment

6.4 Circle one or more IPv4 and IPv6 addresses that are NOT valid

(a) 131.179.260.70
(b) 1
(c) 0x83.0xb3.12.0260
(d) 0x7F000001
(e) 2607:f010:2e9:e:9424:82d4:b103:5e9e
(f) ::ffff:131.179.196.70
(g) 2607:f010:3f9::1003
(h) 2607::3f9::1003

0	8		16	24–31
Version	IHL	DS fiels	ECN	Total Length
Identification			U D M	Fragment Offset
Time to Live		Protocol		Header Checksum
Source Address				
Destination Address				
Options + Padding (optional)				
Payload				

IHL: IP Header Length, U: Unused, D: Don't fragment, M: More fragments

0	8	16	24–31
Version	Priority	Flow Label	
Payload Length		Next Header	Hop Limit
~~Time to Live~~	~~Protocol~~	~~Header Checksum~~	
Source Address			
Destination Address			
Payload			

Problem 7 (20 points)

7.1 In CSMA/CD, after the fifth collision, what is the probability that a node chooses K = 4?

m = 5 so $2^5 - 1 = 32 - 1 = 31$ so range of probability is: {0, 1, ... 31}
32 possible #s → $P(K=4) = \frac{1}{32}$

7.2 DHCP discover, DHCP offer, DHCP request, and DHCP ACK all sent in broadcast IP datagrams. Give at least two reasons why offer, request, and ACK don't use unicast addressing.

- DHCP is how your computer gets an IP address, which must broadcast to quickly receive answers & responses as to where your IP dest might be. unilateral only connects to 1 node at a time whereas broadcast floods neighboring nodes requiring less hops.
- Much faster & more efficient broadcasting over unicast
- Can cache entities much earlier

7.3 ARP requests are sent within a broadcast frame and ARP responses in unicast? Give at least two reasons why ARP responses are unicasted.

- ARP caches the [IP : Mac] pairs so it only needs to send to destination
- Easier on bandwidth & less congestion

7.4 Suppose the IEEE 802.11 RTS and CTS frames were as long as the standard DATA and ACK frames, respectively. Would there be any advantage to using the CTS and RTS frames? Why or why not?

There would be no advantage in using these frames. These are used in Collision Avoidance, allowing sender to reserve a channel to avoid collisions of long frames. Sender sends RTS to AP & gets CTS in response to clear to send. The benefit was that RTS & CTS were small so that latency is low & not having large frames collide. Using these frames would at best reduce the number of packets sent since ACKs & Data are with CTS & RTS but at the risk of higher chance of collisions, and wasting bandwidth.

7.5 Explain how 'traceroute' command works.

Source sends series of UDP packets to destination, when nth packet arrives at nth router, the router discards the msg & sends to source an ICMP msg. Source calculates RTT if it receives it. If not, it waits 5 seconds for response & gives up after. Traceroute does this 3x per hop. Source stops ping if UDP pack arrives at destination & destination returns "Port unreachable"