Carnegie Mellon
Computer Science Department.
15-744 Spring 2008 Midterm

Name: ______________________________

Andrew ID: ______________________________

INSTRUCTIONS:

There are 10 pages (numbered at the bottom). Make sure you have all of them.

Please write your name on this cover and put your initials at the top of each page in this booklet **except the last.**

If you find a question ambiguous, be sure to write down any assumptions you make.

It is better to partially answer a question than to not attempt it at all.

Be clear and concise. Limit your answers to the space provided.

Question	A	B	C	D	E	F	G
Points	/ 27	/ 15	/ 12	/ 16	/ 16	/ 12	/ 2

A Short Answer and Multiple Choice

1. In a high-performance router, copies of the forwarding table are kept in:
   A. The switching fabric.
   B. The output ports.
   C. The input ports.
   D. High-performance routers dont use forwarding tables.

Solution: 4pts - (c) The input ports

2. We discussed two types of QoS architectures, Integrated Services and Differentiated Services.
   (a) What kind of isolation (per-flow vs. per-aggregation) does IntServ offer on its performance guarantees?

   Solution: 1pt - IntServ offers per-flow guarantees and isolation.

   (b) What kind of isolation (per-flow vs. per-aggregation) does DiffServ offer on its performance guarantees?

   Solution: 1pt - DiffServ offers per-aggregation guarantees and isolation (all packets of type X, regardless of flow, receive the same QoS).

   (c) Which is a more scalable architecture and why?

   Solution: 2pts - DiffServ. IntServ requires per-flow state in all routers. In DiffServ, core routers only have to store state per class of traffic that is being differentiated.

3. Consider the figure below, where a link has a capacity of 10Mbps. Three flows, f_1, f_2, and f_3, all send traffic over this link at rates r_1, r_2 and r_3 respectively.

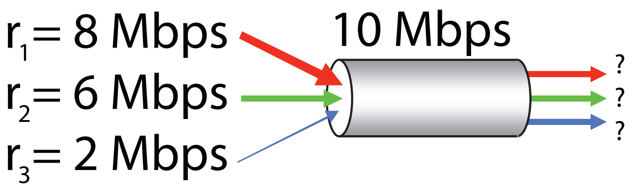

Let r_1 = 8Mbps, r_2 = 6Mbps, and r_3 = 2Mbps. Under max-min fairness allocation, what outgoing rate does each flow receive through this link?

Solution: 4pts, -1pt for minor error

First, we start with what is fair: 10Mbps / 3 = 3.33Mbps. This can satisfy f_3 only, so we give the full 2Mbps to f_3, deduct 2Mbps from the capacity, and continue.

Now the capacity of the link is 8Mbps, and the fair allocation would be 4Mbps (8Mbps/2). Since both r_1 and r_2 are greater than 4Mbps, we give the remaining capacity as a fair share of 4Mbps each to f_1 and f_2.

Thus f_1 gets 4Mbps, f_2 gets 4Mbps, and f_3 gets 2Mbps under max-min fairness.

4. Harry Bovik's first job assignment at Acme is to allow Acme users to wander the U.S. and still be able to access Acme computing resources.

   (a) Harry initially installs a DHCP client on all the users' laptops. The users complain that when they are roaming the U.S., their TCP connections are reset whenever they change location. Why do they have this problem?

   Solution: 2 points - TCP binds to the src addr, src port, dest addr and dest port. As a result when the dest addr changes - the connection is no longer valid.

   (b) Harry addresses the above problem by using Mobile IP instead of DHCP. Many routers in the network check the source address on a packet and drop packets that do not arrive on the interface that the router would use to forward packets to that source address. This is also known as a Reverse Path Forwarding (RPF) check. Assume that this technique was widely deployed in the Internet. Why would RPF create a problem for users roaming away from the Acme network?

   Solution: 3 points - Packets sent from the mobile host have the home address of the host in the src address field. However, it is not appropriate for the location that the mobile host is transmitting from. Therefore, routers will drop packets sent by the mobile host.

5. Answer the following about input-queued routers.

   (a) What is head-of-line blocking?

   Solution: 2 points
   Assume two flows enter the same input interface, where flow 1 is destined to a congested output interface O_1 and flow 2 is destined to an uncongested interface O_2. Since all packets in the input interface have to wait until O_1 frees up, regardless of flow, flow 2 traffic is blocked by flow 1 when flow 1's packets are at the head of the input queue and waiting for congestion at O_1 to go away. This is called head-of-line blocking.

   (b) How can you eliminate head-of-line blocking without requiring a signficantly faster router backplane?

   Solution: 3 points
   One solution for eliminating this is to have Virtual Output queuing Each output receives has a separate input queue, and packets are forwarded from each input queue according to some discipline (e.g. Round Robin). Even if flow 1's input queue does not empty to due congestion, flow 2's packets will still be processed.

6. Chiu and Jain showed in their paper that AIMD was the only choice of increase/decrease algorithms that converged to a fair and efficient utilization. *Briefly* explain why XCP is able to get away with using a multiplicative increase algorithm?

   Solution: 5 points

   Because XCP uses MIMD for efficiency and a separate AIMD fairness controller to allocate the increase *between* flows.

B BGP Routing

7. In the network depicted below, circle all of the paths that packets may take under valley free routing between a pair of clients.

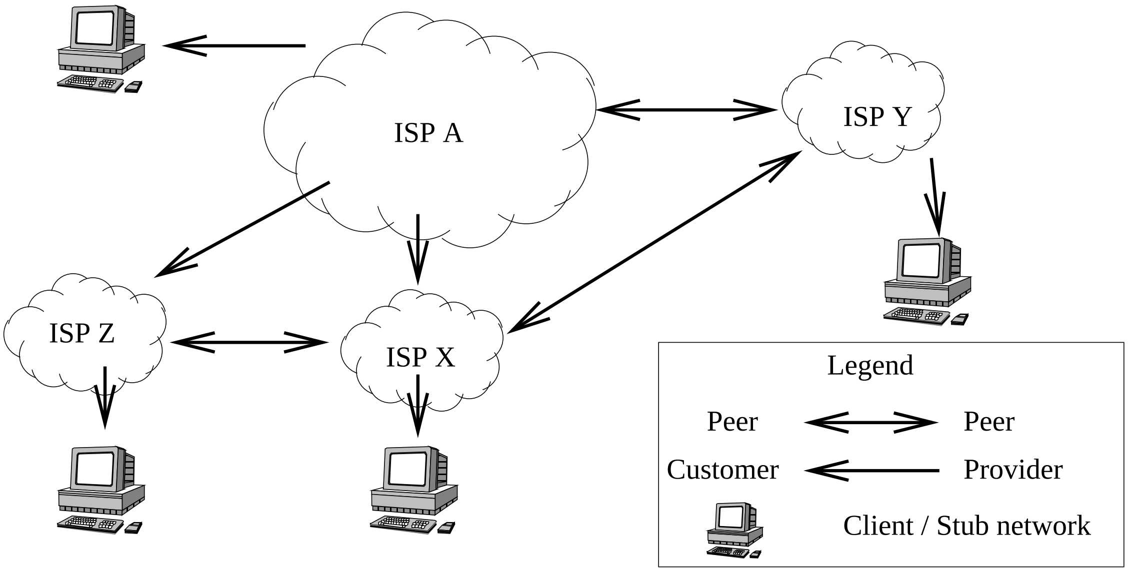

A. Client ⟶ ISP Z ⟶ ISP X ⟶ ISP A ⟶ Client
B. Client ⟶ ISP X ⟶ ISP Y ⟶ ISP A ⟶ ISP Z ⟶ Client
C. Client ⟶ ISP Y ⟶ ISP X ⟶ Client
D. Client ⟶ ISP Z ⟶ ISP A ⟶ ISP Y ⟶ Client
E. Client ⟶ ISP Z ⟶ ISP X ⟶ ISP Y ⟶ Client

Solution: Total: **max 5 pts**

(C)(**3 pts**) and (D)(**3 pts**) are correct paths.

3 pts deducted for each incorrect answer

8. Why doesn't BGP suffer from the count-to-infinity problem as does RIP?

Solution: 4pts - BGP is a path vector protocol, so when a speaker hears a BGP announcement containing itself in the path, it will know that this path should not be used. This solves the count-to-infinity problem that distance-vector protocols like RIP have, since nodes in RIP are not aware that the advertised cost is associated with the (now invalid) path that it advertised previously.

9. Suppose there are two feuding ASs, X and Y, who are opposed to carrying traffic for each other.

   (a) Using only BGP, is it possible for AS X to implement a policy stating that traffic outbound from my AS should not cross Y? Why or why not?

   > **Solution:** 3pts - Yes, assuming it has an alternative route. The policy is: prefer paths that don't contain Y. Of course, if the only path to a destination contains Y, it can not reach that destination without going through Y.

   (b) Using only BGP, is it possible for AS Y to implement a policy stating that, I dont want to carry traffic from X to my customers? Why or why not?

   > **Solution:** 3pts - Not in general. Traffic from a neighbor might be both from good ASes and X. BGP can only accept all traffic by advertising a route or deny all of it by not advertising.

C ETX-OR

Consider the wireless network pictured below. Assume that links experience Bernoulli losses. The labeled edges indicate the *combined* delivery ratio (i.e., the probability that a packet is successfully received in the forward direction and that the acknowledgement is received in the reverse direction). If there is no edge, assume that no packets make it through.

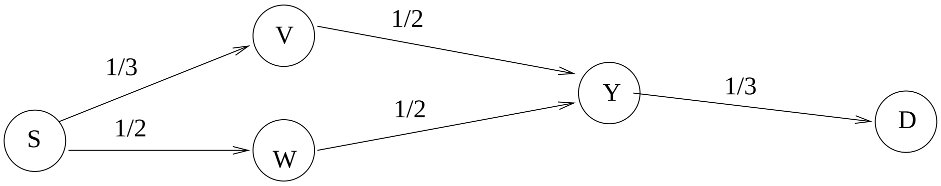

10. Assume that the link layer performs retransmissions. What is the expected number of transmissions to send one packet from S to D using the ETX metric along the path S-W-Y-D?

Solution: 6 points.

$$\begin{aligned} \text{path ETX} &= \sum_{hops} \text{link ETX} && (1)\\ &= \sum_{hops} \frac{1}{\text{Link loss rate}} && (2)\\ &= 2+2+3 && (3)\\ &= 7 && (4) \end{aligned}$$

11. What is the expected number of transmissions needed to send a packet from S to D using ExOR? Assume that there are enough packets in the batch so that the overhead of ExOR headers and its batch maps is insignificant, and that batch maps are received 100% reliably (i.e. they do not suffer from the shown loss rates).

Solution: 6 points. To analyze ExOR, split it into two parts: 1) the number of transmissions needed to successfully reach Y; (2) the number of transmissions from Y to D.

Assuming losses on the paths are independent, when S broadcasts, there is a 1/3 chance that the packet reaches V and a 1/2 chance that the packet reaches W. Regardless of which of these nodes gets the packet, ExOR would then require expected 2 hops to reach Y. So the ETX is the expected number of hops to reach *either* V or W.

$$\begin{aligned} P[VorW] &= 1 - P[!Vand!W] && (5)\\ &= 1 - P[!V]P[!W] && (6)\\ &= 1 - \frac{1}{2}\frac{2}{3} && (7)\\ &= \frac{2}{3} && (8) \end{aligned}$$

The path ETX with ExOR is thus $1.5 + 2 + 3 = 6.5$.

D Wireless MAC

12. Consider the following topology of wireless laptops A, B, C and D. The dotted lines indicate the range of wireless transmissions from each node. For example, B is within range of A, A & C are within range of B, B & D are within range of C and only C is within range of D.

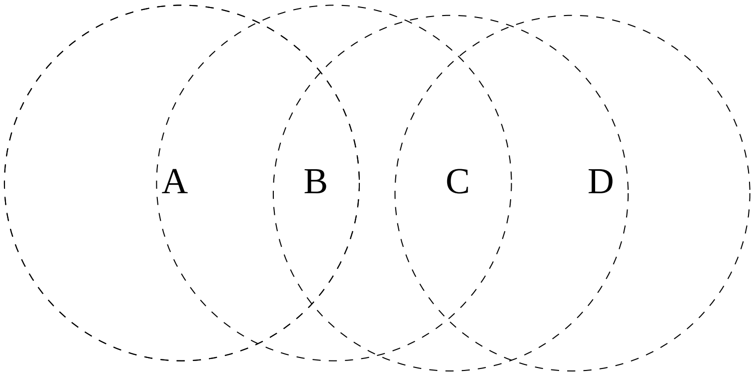

Assume that each node uses an RTS/CTS based MAC protocol (i.e. like MACAW)

(a) If C is sending B an RTS, why does A know not to transmit?

Solution: 4pts - A hears the CTS

(b) If B is sending data to C, why does D know not to transmit?

Solution: 4pts - D heard the CTS from C

(c) Using the nodes above, give an example of the hidden terminal problem.

Solution: 4pts - If A wants to transmit to B and C wants to transmit to D, the transmissions will clobber each other, even though A and C cannot hear each others transmissions.

(d) Irene Packet is considering implementing a walkie-talkie service for her wireless PDAs. Her program largely uses small packets to avoid delaying any voice. Should Irene use RTS/CTS for her deployment? Why?

Solution: 4pts - No. RTS/CTS is primarily to permit collision resolution to finish quickly. The overhead of RTS/CTS isn't worth it for really small packets.

E TCP, RED, and Fair Queueing

13. Recall that TCP's throughput depends on the RTT, the packet size (MSS), and the loss rate (p). Two connections experience the same loss rate and use the same packet size. Connection A has RTT 5ms, and connection B has RTT 10ms. Express the throughput of connection B ($tput_B$) in terms of the throughput achieved by connection A ($tput_A$), or indicate if there is no relationship between the two.

    Solution: 4 points $tput_B = \frac{tput_A}{2}$

14. Suppose that two TCP connections A and B traverse a network consisting only of RED gateways. They both traverse the same congested gateway, and do not traverse any other congested gateways. Assume that the flows have the same RTT and packet size.

    Express the throughput of connection B ($tput_B$) in terms of the throughput achieved by connection A ($tput_A$), or indicate if there is no relationship between the two. Give a one-sentence explanation of your answer.

    Solution: 4 points $tput_B = tput_A$. The two flows experience the same loss rate, and have the same RTT.

15. Now consider a network that uses routers with fair queueing. This is *not* weighted — all weights are the same.

    (a) Two connections share the same congested gateway and have no other congested gateways. Connection A has RTT 5ms, connection B has RTT 10ms. Express the throughput of connection B ($tput_B$) in terms of the throughput achieved by connection A ($tput_A$), or indicate if there is no relationship between the two.

    Solution: 4 points $tput_B = tput_A$

    (b) Two connections traverse the same congested gateway, but *also* traverse some other unshared congested gateways. Express the throughput of connection B ($tput_B$) in terms of the throughput achieved by connection A ($tput_A$), or indicate if there is no relationship between the two.

    Solution: 4 points Without knowing the other flows in the network, you cannot find a relationship between the two.

F TCP Congestion Control

In the graph shown below, we have plotted the congestion window for a particular TCP connection over its lifetime. Avg Cwnd is the average of the cwnd variable over the congestion avoidance (linear mode) part of the graph. The value of Avg Cwnd is 1100KB. Answer the questions below based on this graph.

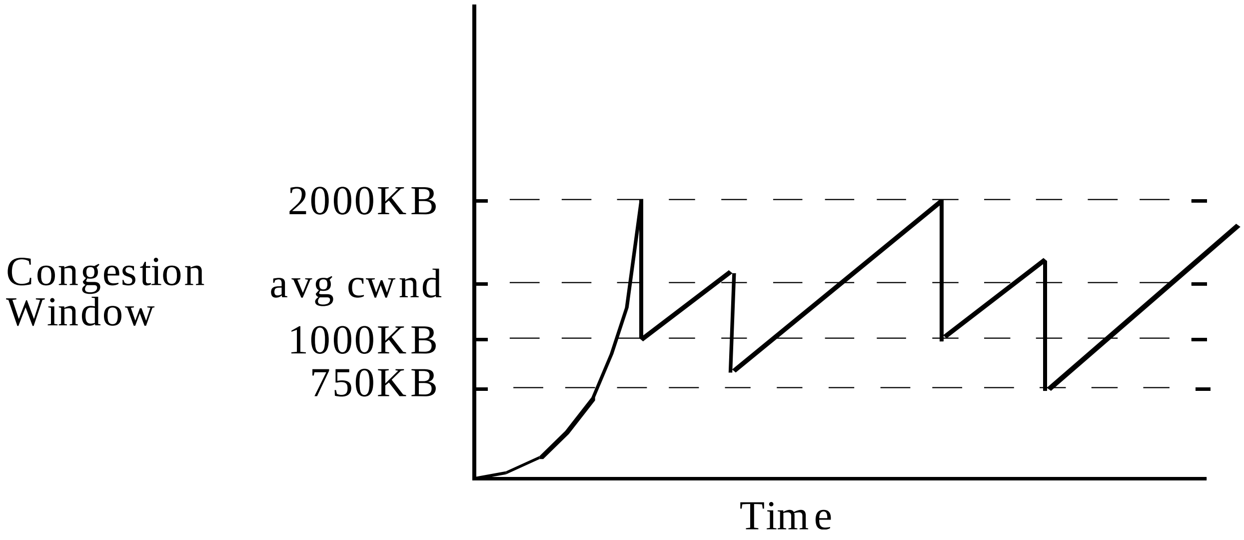

16. Assuming that the connection is using something like TCP NewReno or SACK, mark an X on the Time axis of the congestion window plot whenever a loss event occurred.

> **Solution:** 4pts - Where there is a dip, there should be an X

17. Assume the following: 1) The propagation RTT of the link connecting the source and destination was 100ms and 2) This was the only TCP connection on this link. After the TCP connection leaves slow start, what is the maximum speed link that this TCP connection could have kept continuously fully utilized?

> **Solution:** 8pts - 750KB/100ms = 7500KBps = 60000Kbps = 60 Mbps
>
> A common mistake was to put 1100KB/100ms as the answer (4pts given). However, when the cwnd drops below 1100KB for some time, we are not keeping a 88Mbps pipe full. The lowest cwnd after slow start is 750KB, and so since the cwnd is always above this value, a pipe that is 750KB/100ms = 60Mbps will be kept full. Note that we do not assume any particular buffer sizing along the pipe, so that we may have enough buffering to rise to a CWND of 2000KB before experiencing loss.

Survey!

G 2 Free Points for Tearing Off Page: Anonymous Feedback

List one thing you liked about the *class* and would like to see more of or see continued (any topic - lectures, homework, projects, discussion site, topics covered or not covered, etc., etc.):

List one thing you would like to have changed or have improved about the class: