Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.829 Computer Networks: Fall 2005

Quiz I

There are 18 questions and 8 pages in this quiz booklet. Answer each question according to the instructions given. You have **85 minutes** to answer the questions.

If you find a question ambiguous, be sure to write down any assumptions you make. **Be neat and legible.** If we can't understand your answer, we can't give you credit!

Use the empty sides of this booklet if you need scratch space. You may also use them for answers, although you shouldn't need to. *If you do use the blank sides for answers, make sure to clearly say so!*

Note well: Write your name in the space below AND your initials at the bottom of each page of this booklet.

THIS IS AN "OPEN NOTES, OPEN PAPERS" QUIZ.
NO OTHER MATERIALS, NO PHONES, NO COMPUTERS, NO LAPTOPS, NO PDAS.
MAKE SURE YOU'VE READ ALL THE INSTRUCTIONS ABOVE!

Do not write in the boxes below

1-4 (xx/13)	5-7 (xx/12)	8-11 (xx/11)	12-14 (x/10)	15-16 (x/7)	17-18 (x/7)	Total (xx/60)
13	10	4	8	7	5	47

Name: DAN PORTS

I Network architecture

1. [2 points]: In the 1970s, the designers of the original Internet protocols rejected translation gateways in favor of a universal network layer, IP. Yet, the current Internet has a large number of translation gateways. Why?

(Answer legibly in the space below.)

(2) The shortage of IPv4 addresses prompts the use of translation gateways so that not every device needs its own IP address that can be addressed by the rest of the world.

2. [4 points]: Which of the following statements is true about application-level framing (ALF)?

(Circle True or False for each choice.)

(4)

A. $\boxed{\textbf{True}}$ / **False** ALF advocates that application data units and protocol data units be the same.

B. **True** / $\boxed{\textbf{False}}$ TCP is an example of the ALF design principle.

C. $\boxed{\textbf{True}}$ / **False** ALF advocates that messages be delivered to the application in whatever order they appear, even if that order is different from the order in which they were sent.

3. [5 points]: A TCP Tahoe sender may sometimes retrannsmits packet that have already reached the receiver, even if no ACKs are ever lost. Using an example, explain why that could happen?

(Answer legibly in the space below.)

(5) suppose A initially sends a window of four packets:

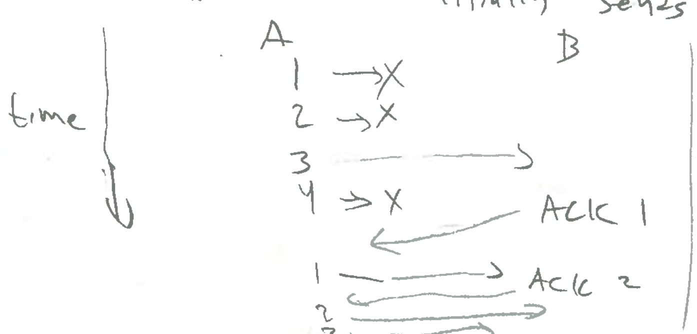

Only 3 is received. B sends an ACK causing A to lower its window and resend 1. When B ACKs this, A might resend both 2 and 3 if its window is wide enough, since it doesn't know 3 already arrived.

4. [2 points]: In the current Internet, autonomous systems (ASes) run an IGP and may also run iBGP. Give one reason why ASes don't just introduce routes to external destinations into their IGP, eliminating iBGP.

(Answer legibly in the space below.)

(2) IGPs typically don't scale as well to the large number of routes that BGP handles, esp. when they change frequently

Initials:

Dan Ports

II Interdomain Internet Routing

5. [3 points]: Studies have found that Internet routing can be highly asymmetric: packets from A to B may go through different cities compared to packets from B to A. Give one reason why this behavior occurs in today's Internet.

(Answer legibly in the space below.)

1

ASes will express their preferences (via LOCAL PREF) about which routes to take, for outgoing traffic, e.g. due to customer/peer/provider status or other reasons. Too vague But why does behavior really occur?

6. [4 points]: Interdomain Internet routing is often non-transitive: i.e., there may be no usable end-to-end path between addresses A and C, even though there is some B for which the paths from A to B and from B to C are working. Assume that the only policies implemented are for customer-provider transit and for mutual peering. Give an example showing how transitivity may be violated in this routing system.

(Answer legibly in the space below.)

4

Suppose B buys transit from A to C. Then there is a route A ⇒ B, and B ⇒ C, but B will not broadcast its B → C route to A.

7. [5 points]: Nick Feamster's lecture and the Gao/Rexford paper showed that BGP is not a "safe" routing protocol in general. Suppose we restrict each autonomous system (AS) to:

A. Advertise routes from its customers to all neighboring ASes.

B. Advertise routes from its peers and providers only to its customers.

Suppose that the provider graph has no cycles (i.e., no AS's ancestor is also its descendant, in a graph where each node is an AS and a directed edge signifies a customer-provider relationship). Assume that each AS can pick an arbitrary path from the set of paths available to it, unconstrained by the "customer > peer > provider" rule.

Is BGP guaranteed to be safe under these conditions? Explain your answer.

5

(Answer legibly in the space below.)

No. Consider Fig 4 in Gao/Rexford, where AS 3 prefers the provider route (2 0) over the customer route (0). This satisfies the restrictions, but

(on back)

Initials:

Dan Ports

III End-to-end Congestion Control

Alyssa P. Hacker feels that the multiplicative decrease rule in TCP's AIMD congestion control causes too much of a window size reduction to be useful for applications like streaming video. She is interested in developing an algorithm that doesn't have the same drastic reduction. She generalizes AIMD as follows:

Upon no congestion, every round-trip time (R), increase the window size, w_t, as:

$$w_{t+R} = w_t + \alpha/w_t^k.$$

Upon congestion, as signaled by a duplicate ACK, decrease the window size as:

$$w_{t+\delta t} = w_t - \beta w_t^l.$$

As in TCP, the window size reduces at most once during each RTT. α and β are constants > 0. Observe that $k = 0, l = 1$ is what TCP AIMD implements, and that when $l < 1$, the reduction is less drastic than multiplicative decrease.

In these questions, assume that $w_t >> 1$, that the packet loss rate is small, and that the RTT (R) does not change appreciably during a connection. Connections are long-running bulk data transfers.

8. [2 points]: What are k and l for TCP during its *slow start* phase, assuming multiplicative decrease upon congestion?

(Answer legibly in the space below.)

2

TCP doubles its window size each RTT during slow start and will halve it upon congestion. So

(no congestion) $w_{t+R} = 2w_t = w_t + \alpha/w_t^k \Rightarrow k = -1$ ✓

(congestion) $w_{t+R} = 1/2\, w_t = w_t - \beta w_t^l \Rightarrow l = 1$ ✓

9. [1 points]: In Alyssa's congestion control scheme, by how much should the sender increase the window size *on each ACK* to approximate the specified increase rule?

(Answer legibly in the space below.)

1

There will be w_t ACKs per RTT, so the sender should increase the window size by

$$\frac{\frac{\alpha}{w_t^k}}{w_t} = \frac{\alpha}{w_t^{k+1}}$$

Initials:

DAN PORTS

(1)

10. [3 points]: Sketch w_t as a function of time schematically, showing how the window increases and reduces upon congestion. For this sketch, assume that $k > 0$ and $0 < l < 1$.

(Answer legibly in the space below.)

On back

11. [5 points]: Call a congestion control scheme "TCP-friendly" if its throughput is the same as TCP AIMD at a given loss rate. What are the conditions on k and l such that Alyssa's congestion control scheme is TCP-friendly?

Hint: There are at least two ways to solve this problem. The easier method is to think about the "drift" from the average window size.

(Answer legibly in the space below.)

Assume steady state, with max throughput C packets/RTT. Whenever the window size exceeds C, congestion occurs and the window size is reduced by $\beta w_t^l \approx \beta C^l$; thereafter, it increases by α / w_t^k per RTT.

X

D

Initials:

DAN PORTS

IV Router-assisted congestion control

3

12. [3 points]: Instead of the standard drop-tail scheme, some people have suggested that packets at the *head* of a queue be dropped when a queue is full and a new packet arrives. Why might this scheme be better than drop-tail?

(Answer in the space below)

This causes packets to be dropped sooner, meaning that the sender will notice faster and react accordingly, e.g. by reducing its window size.

13. [4 points]: How can an XCP router estimate the number of *distinct* connections observed over time T while maintaining only a constant amount of state (independent of the number of flows)? Explain clearly.

You may assume that each XCP sender correctly inserts its congestion window size ($cwnd_i$) and round trip time (rtt_i) in each packet. Assume that the size of a packet from connection i is s_i.

(Answer legibly in the space below.)

2

The throughput for a connection i is $\frac{\cancel{s} \; s_i \cdot cwnd_i}{rtt_i}$ since $cwnd_i$ packets are sent each RTT. So from each packet, the XCP router can get the throughput for that flow. By computing an average of this, and dividing the total throughput by the average, it can estimate the number of flows.

in XCP, cwnd in bytes, but ok

but how ??

Not a true statement

1, 2, 3, 4, 100

Avg $\approx$ 100 but #flows = 5 !

14. [3 points]: Both CSFQ and XCP use the idea of "dynamic packet state". Both protocols claim to achieve some form of fairness. But CSFQ provides isolation between flows even if the sources are greedy, whereas XCP does not. What is the reason for this difference in their properties?

(Answer legibly in the space below.)

3

XCP relies on the source to lower its cwnd in response to the router's feedback in order to achieve congestion control while CSFQ enforces this without source cooperation by dropping packets at the routers.

Initials:

Dan Ports

V Packet Pair Congestion Control

Some researchers have proposed a congestion control algorithm called *packet pair*. Here, the sender maintains a transmission rate, r, in packets per second. The sender executes the following loop:

```
do {
    send two packets back-to-back;
    // each pair of packets has ''first'' and ''second'' pkt
} every (2/r) seconds
```

On each ACK, the sender executes the following:

```
if (ACK for first packet of a sent pair) then
    firsttime <-- current_time;
else {
    // all pkts have same size
    rate <-- packet_size/(current_time - firsttime);
    r <-- alpha * rate + (1 - alpha) * r;
}
```

On detecting a lost packet, the sender cuts r by one-half. Assume that the ACKs don't experience any loss or queueing.

15. [3 points]: Suppose that a connection using the packet pair algorithm runs in a network with other traffic where the bottleneck router implements drop-tail queue management and FIFO scheduling. Does the packet pair scheme perform correct congestion control? Explain.

(Answer legibly in the space below.)

No. With FIFO scheduling, the back-to-back packets will be processed in sequence at the router, and so the measured rate is the ~~max throughput~~ link bandwidth for all flows in the router. The source's rate will increase exponentially to this, but packets will be lost as soon as the rate exceeds the router's per-flow ~~throu~~ capacity.

16. [4 points]: What router scheduling scheme and/or queue management method guarantees that packet pair will properly implement congestion control? Explain.

(Answer legibly in the space below.)

With fair queuing the packets will be transmitted by the router such that each flow sees throughput corresponding to its fair share. This means that the algorithm will estimate this as 'rate' and hence its send rate will converge to this (with multiplicative decrease if it overshoots)

7

Initials: DAN PORTS

VI Internet QoS

17. [3 points]: Consider a network shared by n instances of an application whose utility function, U, is shown below. U depends on the rate available to the application. Suppose the goal of the network designer is to maximize the aggregate utility of the applications running in the network (ignoring any "unhappiness" caused to applications turned away). Assume that each instance of the application gets a rate of C/n, where C is the capacity of the network.

How many instances of the application should the network admit to maximize the aggregate utility? Answer this question by showing how to get this number on the graph below.

U is convex, so admission control is necessary. The max will be in the convex region

Maximize $nU(C/n)$

$\Rightarrow n \frac{dU}{dC_n}(C/n) \cdot \frac{-n^2}{C} \cdot U\left(\frac{C}{n}\right) = 0$

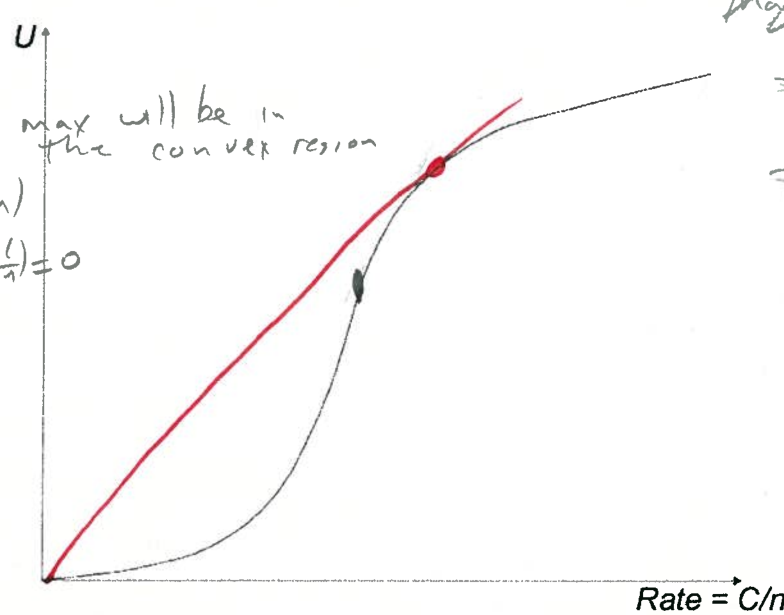

1

18. [4 points]: Consider the behavior of QoS based on explicit reservations and QoS based on implicit information carried in packets, in a network in which routes may change at any time.

4

A. Which type of scheme is oblivious to a change in route?

(Answer legibly in the space below.)

✓ Implicit information since the information is in the packets,

B. Which type of scheme is more likely to provide end-to-end bandwidth guarantees in practice, assuming routes don't change?

(Answer legibly in the space below.)

✓ Explicit reservations.

Initials:

Dan Ports