College of Computing

Georgia Institute of Technology

CS 7260: Internetworking Protocols and Architectures: Spring 2007

Quiz II

There are 14 questions and 10 pages in this quiz booklet (including this page). Answer each question according to the instructions given. You have **85 minutes** to answer the questions.

If you find a question ambiguous, write down any assumptions you make. **Be neat and legible.** If I can't understand your answer, I can't give you credit! There are three pretty challenging questions (clearly marked); you may want to look through the whole quiz and save those for last.

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 ENCRYPTED WIRELESS TRAFFIC.
MAKE SURE YOU'VE READ ALL THE INSTRUCTIONS ABOVE!

Initial here to indicate that (1) you've read the instructions and (2) you agree to abide by the Georgia Tech Honor Code:

The last page has easy bonus questions, which you can answer outside of the allotted time. Rip the last page off of your quiz for five bonus points. Turn it in anonymously if you like. You won't get the five points if you don't tear off the page (this is to make certain you've read this far ;).

Do not write in the boxes below

1-5 (xx/20)	6-8 (xx/15)	9-12 (xx/20)	13-14 (xx/10)	Bonus (xx/5)	Total (xx/70)

Name:

I Warmup

1. [4 points]: Which of the following statements about the Chord lookup service are true?
(Circle ALL that apply)

- **A.** If nodes are evenly spaced around the Chord ring, each node is guaranteed to service a roughly equal number of lookups.
- **B.** Without finger tables, the lookup protocol will still be able to find the node that stores a certain key.
- **C.** Chord's lookup algorithm requires only $O(\log n)$ iterative lookups, which necessarily bounds the latency required to resolve any particular query.
- **D.** According to the Chord protocol described in the paper, each node maintains a pointer to both its successor and its predecessor node.

Answer 1 The answer is (B) and (D). (A) is false because the query load may be unbalanced. (C) is false because nodes may be spread out topologically/geographically, even if they are close in keyspace. ■

2. [4 points]: Which of the following statements about BitTorrent are true?
(Circle ALL that apply)

- **A.** BitTorrent clients' attempts to download the rarest piece first attempts to ensure that clients have chunks of the file to trade with each other.
- **B.** Per BitTorrent's "tit for tat" strategy, a client that has never uploaded a block to any of its peers is unable to download any blocks itself.
- **C.** BitTorrent provides no incentives for clients who have completely downloaded a file to remain in the swarm as seeders.
- **D.** Chopping the BitTorrent file into smaller chunks will reduce the likelihood of free-riding (or, in game theory terms, that a client will "defect").

Answer 2 The answer is (A) and (C). (D) is false because the tit-for-tat punishment strategy only hits a boundary condition at the last block; the size of the blocks don't actually matter. ■

3. [4 points]: You decide you need to design a very fast-spreading worm to build your botnet army that spreads by exploiting a known software vulnerability. Which of the following statements are true about the speed at which your worm spreads?
(Circle ALL that apply)

- **A.** Preferentially scanning for vulnerable hosts on the local subnet over hosts on remote networks may speed up the spread of the worm.

Initials:

B. Pre-emptively scanning address space for vulnerable hosts before releasing the worm may speed up the spread of the worm.

C. The rate at which owners of machines patch the vulnerability may affect the speed at which the worm spreads.

D. None of the above.

Answer 3 (A), (B) and (C) are all true. This question comes directly out of the paper by Weaver *et al.* ■

Initials:

4. [4 points]: Since the study of the Jung *et al.* DNS paper was performed in 2001, the use of DNSBLs has increased substantially. If you were to re-perform this study, what differences would you necessarily see as a result of this increased use of DNSBLs?

(Circle ALL that apply)

A. An order of magnitude increase in DNS traffic to the DNS root servers.

B. A smaller ratio of DNS lookups for some hostname to TCP connections to the hostnames being looked up.

C. A longer "tail" in the distribution of names being looked up (i.e., lookups for more IP addresses).

D. None of the above.

Answer 4 (B) and (C) are true. (A) is false; caching of DNSBL nameservers would likely prevent the load on the DNS root servers from going very high. ■

5. [4 points]: Which of the following are true about layer 3 BGP/MPLS VPNs? (For the rest of the question, we simply call these "VPNs").

(Circle ALL that apply)

A. Multiple distinct VPNs can number their networks from the same address space.

B. Traffic traversing one VPN cannot exhaust the resources of another VPN that is using the same physical infrastructure.

C. The size of the routing table at the provider edge will continue to increase as the number of VPNs increases.

D. The size of the routing tables in the provider's "core" will continue to increase as the number of VPNs increases.

Answer 5 (A) and (C) are true. (B) is false because BGP/MPLS VPNs do not inherently provide any resource isolation. ■

Initials:

II Potpourri

6. [9 points]: Suppose that you have access to a /16 network that you are using as a "network telescope". You observe SYN/ACK traffic coming from 1,000 different source IP addresses to 2,000 destination addresses in the telescope.

a. Estimate the total number of attack victims.

b. Give your best estimate for the total number of attackers.

c. For each of these estimates, give *one* possible source of inaccuracy.

(Answer legibly in the space below.)

Answer 6
(a) $2^{16} * 1,000$
(b) $2^{16} * 2,000$
(c) The answer to (a) could be preferential scanning or hosts that explicitly know to avoid the telescope. (b) could be inaccurate if attacking hosts don't pick source IP addresses at random. ■

7. [3 points]: Explain how fake Web clicks on Web advertisements from "click bots" can be used to generate revenue (1) for the person who hires the click bot and (2) for the advertising network. Explain why advertising networks have little incentive to stop click bots.

(Answer legibly in the space below.)

Answer 7 A person who hires a bot could set up "fake" web sites that get paid by the ad network for hosting ads that get clicked on. The advertising network also makes money (per click) from distributing the ads on behalf of the advertiser. ■

8. [3 points]: Give two reasons why a RON can allow end hosts to react more quickly to end-to-end path failures than Internet routing protocols would. Give an example failure scenario that might not be detected by an Internet routing protocol at all but could be detected by RON (*hint:* think about IP prefixes and aggregation).

(Answer legibly in the space below.)

Answer 8 A RON can probe/react more quickly than BGP routing protocols can converge (15 minutes). It might also detect "soft failures" (e.g., lots of congestion) that a routing protocol may not notice, and routing protocols are known to not detect/reflect all failures anyway. One possible scenario would be a subnet within a larger prefix becoming unreachable. ■

Initials:

III Selfish, Wireless Routing

Consider the wireless network shown below, with a source, destination, and two intermediate nodes which are within wireless range of each other. Assume that s is *not* within range of d.

The annotations on each edge indicate the *success probability* of transmitting a packet on that link. (Assume that this success probability shown already incorporates the probability of the data being received by the sender *and* the ACK being received by the receiver. In other words, it is the probability that no retransmission is needed.)

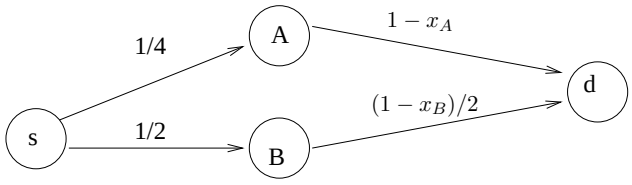

9. [4 points]: Using the ETX metric, what is the expected number of transmissions for sending traffic from s to d via A and the expected number of transmissions for sending traffic from s to d via B in terms of x_A and x_B.

(Answer legibly in the space below.)

Answer 9
Via A: $4 + \frac{1}{1-x_A}$
Via B: $2 + \frac{2}{1-x_B}$

■

10. [6 points]: Suppose that there is one (1) unit of traffic flow that is infinitely divisible, and the source has the option of dividing this traffic to send any fraction to each of node A or node B (x_A and x_B, respectively).

Give values for x_A and x_B such that all traffic is routed (i.e., $x_A + x_B = 1$) and such that the allocation of traffic constitutes a "Nash flow" (i.e., neither of the two flows are envious of the other's success probability).

(Answer legibly in the space below.)

Answer 10
$4 + \frac{1}{1-x_A} = 2 + \frac{2}{1-x_B}$

One possible solution is to set $x_A = x_B = 1/2$. In this case, the expected number of transmissions along each path is 6. ■

Initials:

11. [4 points]: After reading the EXOR paper, Ray Diot recognizes that perhaps s might have a higher probability of successfully transmitting packets by exploiting broadcast. Suppose now that s sends a packet and relies on either A or B to receive it.

What is the probability that either A or B receives this packet within a single transmission?

(Answer legibly in the space below.)

Answer 11
$1 - (3/4)(1/2) = 5/8$ ■

12. [6 points]: (Hard!) Suppose that transmission schedule is such that A and B each get one chance to send along a packet to d before s will retransmit the packet.

What is the probability that a single packet can be transmitted from s to d *without requiring* s *to retransmit the packet*?

(Answer legibly in the space below.)

Answer 12 There are three cases where the scenario could succeed:

- **A.** only A gets the packet from s (with probability $(1/4)(1/2) = 1/8$);
- **B.** only B gets the packet from s (with probability $(3/4)(1/2) = 3/8$;
- **C.** both get the packet from s (with probability $(1/4)(1/2) = 1/8$)

(note that this matches the answer from the previous question). Now, compute the conditionals for each case, and the probability of success of the first try will be: $(1/8)(1/2) + (3/8)(1/4) + (1/8)(1 - (1/2)(3/4)) = 15/64$ ■

Initials:

IV Incentives in BitTorrent

Ben Bitdiddle has decided that BitTorrent's "tit for tat" strategy is too permissive, and that clients that freeride should be choked for a longer period of time.

He estimates that clients in BitTorrent that free-ride for a single "round" can download data faster than they could if they were sharing with other clients, which makes them extremely happy. He also estimates that there is some "negative" cost to not getting a block of data in a certain time block. He expresses the game theoretic matrix as follows:

	Upload	FreeRide
Upload	(1,1)	(0,6)
FreeRide	(6,0)	(-1,-1)

Suppose that each player has a discount value, δ, for receiving blocks in a later round. So, for example, if a player were to cooperate for all time, the payoff would be $1 + \delta + \delta^2 + ...$ In a sense, smaller values of δ mean that a client is more "impatient" and wants to cheat to get blocks of the file more quickly.

13. [5 points]: Ben first decides that his new BitTorrent client will punish a client that free rides for one round by choking it for the rest of time. For what values of δ would clients be motivated to cheat regardless?

Answer 13

$$\begin{aligned} \frac{1}{1-\delta} &< 6 - \frac{\delta}{1-\delta} \\ \delta &< 5/7 \end{aligned}$$

Many people forgot to take into account that in tit-for-tat, both parties will end up choking each other in the "punishment phase". ■

14. [5 points]: Alyssa P. Hacker tells Ben that his new BitTorrent client is simply too harsh and that it should only choke a client for 5 rounds. Write down the inequality that must hold in order for clients to cooperate under this shorter punishment phase. *You do not need to solve for δ in a closed form.*

Answer 14

$$\frac{1}{1-\delta} > 6 - \left(\frac{1-\delta^k}{1-\delta}\right)$$

■

Initials:

Some people answered this slightly differently ($1 + \delta + \delta^2 + \ldots > 6 - \delta - \ldots$), which is equivalent. Again, many folks forgot the punishment phase (the negative terms on the right side).

Initials:

V Bonus: Anonymous Course Feedback

This page is anonymous. Rip this off from your exam, and turn it in separately if you like. You'll get five points for simply ripping off the last page of the exam, but I'd prefer if you fill it out and hand it in in a separate stack.

What was your favorite paper from the reading list? Your least favorite paper?

What paper do you know of that you think should be added to the syllabus next year?

What would be the best way to improve this course in the future?

Initials: