Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.829 Computer Networks: Fall 2005

Quiz II

There are 19 questions and 9 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

1-6 (x/20)	7-9 (x/8)	10-13 (x/12)	14-16 (x/9)	17-19 (x/11)	Total (xx/60)
18	8	11	8	8	53

Name: DAN PORTS

I Wireless networks

1. [4 points]: Wired Ethernet and 802.11 (WiFi) LANs both use carrier sense (CSMA) for channel access. Unlike in a wired Ethernet, however, 802.11's link layer uses link-layer acknowledgments (ACKs). Give *two reasons* why 802.11 uses link-layer ACKs, but not wired Ethernet?

(Answer legibly in the space below.)

802.11 cannot reliably detect collisions at the sender like Ethernet can, since the transmitting radio signal usually is powerful enough to drown out reception of other transmitting radios. Also, the hidden terminal problem means that node A's transmission and node C's transmissions could cause ~~the~~ a collision at node B even though A and C cannot hear each other.

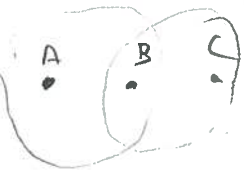

Also, bit errors causing packet corruption are more likely in 802.11 than wired Ethernet.

4

2. [4 points]: Alyssa P. Hacker reads the ExOR and MRD papers and hits upon a brilliant idea to combine the techniques in these papers to improve the performance of a multi-hop wireless network. Explain how you might be able to take advantage of erroneous transmissions in an ad hoc network to deliver error-free packets to a destination. For simplicity, assume that there is only one destination for all packets in the network.

(Answer legibly in the space below.)

One could take advantage of the fact that multiple receivers will hear the same packet though some may have errors. We could have nodes retransmit their packets ~~based on the~~ in order based on their proximity to the receiver as in ExOR, but have do this even if they received a packet with bit errors (detected by CRC). Then, each node in sequence will retransmit the packet only if they have not heard a node closer to the destination retransmit it or that transmission contained errors. If they hear a closer node transmit first, with errors, they can use this to correct their own bit errors as in MRD.

4

Initials: DRKP

3. [2 points]: The Multi-Radio Diversity (MRD) scheme described in the Miu *et al.* paper tends to perform better when bit errors occur in bursts, rather than when they are spread uniformly over a frame. Why?

(Answer legibly in the space below.)

(2)

MRD assumes that bits are partitioned into blocks, and chooses blocks from the two packets to reassemble and check the CRC on. The cost of this is 2^n checks, where n is the number of blocks. If bit errors are in bursts, we can use a larger block size (smaller n); otherwise we need a block size of 1.

4. [5 points]: Consider a model similar to Shepard's analysis of the aggregate impact of far-away transmissions on the signal-to-noise ratio (SNR) of a spread-spectrum system. Suppose nodes are distributed at uniform density, ρ, in three dimensions (i.e., there are ρ nodes per unit volume). Suppose also that the signal strength attenuates as $\frac{1}{d^4}$, where d is the distance between sender and receiver. Assume no limit on the total number of nodes in the network. Calculate the SNR for communication between a pair of nodes separated by distance R_0 (the expected distance between a node and its nearest neighbor), as a function of the density ρ.

(Answer legibly in the space below.)

Signal strength at the receiver is proportional to R_0^{-4}

To estimate noise, we consider the contribution from nodes at distance r, using a spherical shell of radius r and thickness dr. The number of nodes in this shell is $4\pi r^2 \rho\, dr$, so the noise contribution is $\frac{4\pi r^2 \rho\, dr}{r^4}$. With no limit on total nodes we must integrate this from R_0 to infinity:

(4)

$$\int_{R_0}^{\infty} 4\pi r^{-2} \rho\, dr = 4\pi\rho \int_{R_0}^{\infty} r^{-2}\, dr = 4\pi\rho \left[-r^{-1}\right]_{R_0}^{\infty}$$

$$= 4\pi\rho \left[-\frac{1}{\infty} + \frac{1}{R_0}\right] = 4\pi\rho R_0^{-1}$$

Then SNR = signal / noise = $R_0^{-4} / 4\pi\rho R_0^{-1} = \boxed{\frac{R_0^{-3}}{4\pi\rho}}$ ✓

How does R_0 relate to ρ?

Initials: DRKP

Opt E. Mizer is interested in understanding the performance of her 802.11 network as a function of the sending node's bit rate. She picks a sender-receiver pair in the network, selects a bit rate, R, and plots the probability of successful packet delivery as a function of distance between sender and receiver in her network. She obtains the graph shown below.

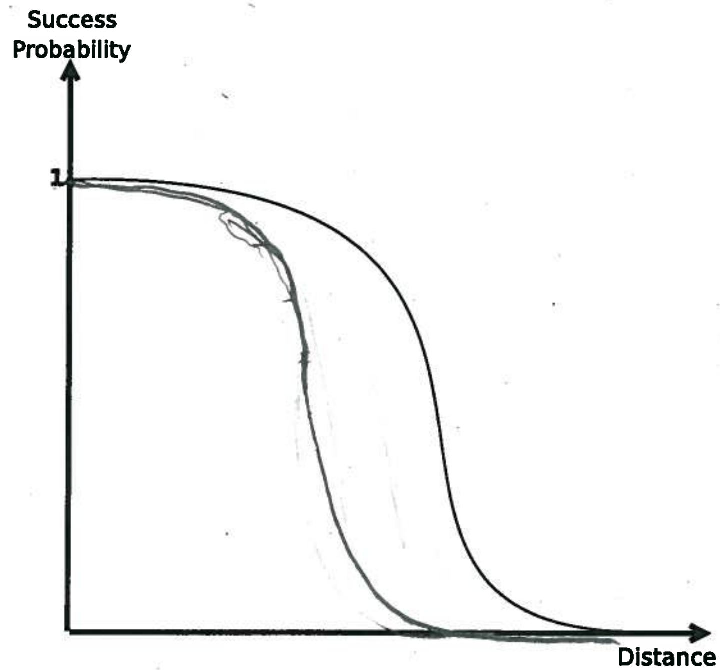

(2) ✓

5. [2 points]: She now increases the 802.11 bit rate, and keeps everything else in the network and the environment unchanged. *On the graph above*, show what curve she is likely to observe.

6. [3 points]: At some distance d between sender and receiver, suppose that the packet loss rate to the receiver is p_i when the sender sends packets at a bit rate R_i ($i = 1, 2$). Link-layer acknowledgments and retransmissions are turned on, and all link-layer ACKs are lost with probability q. Describe how you can use this information to determine which bit rate will maximize data delivery throughput at this distance.

(Answer legibly in the space below.)

Throughput $\propto$ bit rate / # transmissions. ✓ The probability of a retransmission being necessary is $p_i + (1-p_i)q$: a packet must be resent if it is lost, or it's received and the ack is lost. So this is the expected number of retransmissions, so we must choose i to maximize

$$\frac{R_i}{p_i + (1-p_i)q}$$

(2) ✗ Not quite correct.

Initials: DAKP

III DHTs and DHT-based Systems

10. [2 points]: Suppose that a new node joins a correctly working Chord DHT with N nodes and K keys. What is the expected number of keys that have to be moved for Chord's consistent hashing assignment to be maintained? Assume that each key is stored at exactly one node.

(Answer legibly in the space below.)

1

$\frac{K}{N+1}$

$O(\frac{K}{N})$ in expectation (and w.h.p. if $O(\log n)$ virtual nodes are used). The new node's key can be treated as random, so the keys of all nodes will be approximately evenly distributed and the node gets 1/nth of the keys. These are the only ones that need to be moved.

11. [4 points]: Chord's stabilization procedure runs periodically. This distributed procedure attempts to ensure that each node n's successor is set to the smallest node ID larger than n in the system. Node n contacts its successor and asks it for its predecessor, and if the value returned is not n, then various actions are taken. The pseudocode is shown below (it is the same as in the Chord paper, except we don't bother with fixing the fingers here because it isn't relevant to the question).

```
procedure STABILIZE()
  // runs periodically at each node
  x ← successor.predecessor();
  if x ∈ (n, successor) then{
    successor ← x;
  }
  // Tell successor about n; successor updates its predecessor if necessary
  successor.notify(n);
```

Describe (with a picture that shows nodes, their IDs, and successor pointers between nodes) an initial configuration for which the stabilization procedure given above *will not* cause the nodes to have correct successor pointers, even when the procedure runs an infinite number of times.

(Answer legibly in the space below.)

4

A partitioned ring: if for some reason ~~the successor pointers from~~ the nodes in two separate subsets of the ring each have no pointers to nodes in the other subset, then two consistent rings may be formed, but ~~the~~ stabilization can't fix the partition and merge them into one ring.

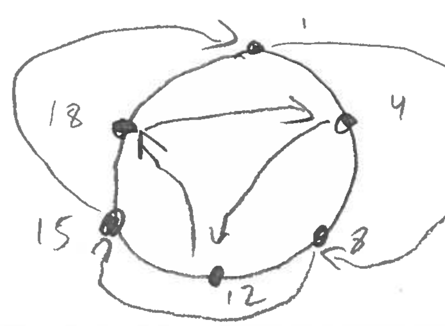

Initials: DKP

12. [4 points]: Suppose that Calvin Butterball runs a Web server from his DOA-enabled laptop. The laptop has EID e. When he sits down at Greenbacks (a coffee shop with a WiFi hot-spot), he connects via wireless to a DOA-enabled NAT ("NEB" in the paper) with IP address X and EID f, and obtains a private IP address, Y, on the Greenbacks network. How should various bindings be updated in the components of the DOA system (the DHT, the NEB, the laptop, etc.) to make the Web site on Calvin's laptop globally reachable (assume that Calvin invokes no other intermediaries)?

(Answer legibly in the space below.)

4

The laptop uses the same EID it did before. It updates the mapping in the DHT so that EID now resolves to the NEB's EID, and tells the NEB its private IP address and EID so the NEB can add this to its EID → private IP table. Now clients know to contact the NEB to reach the laptop's EID, and the NEB can forward this packet to the laptop's IP.

13. [2 points]: The Internet Indirection Infrastructure (*i*3) and the Delegation-Oriented Architecture (DOA) both use a DHT to separate the location and identity of end hosts. But they use the DHT in different ways. What is the main difference in how they each use a DHT?

(Answer legibly in the space below.)

2

DOA uses the DHT to store mappings from EID to IP (perhaps via delegation); a client uses the DHT to find the appropriate IP address to use to send the packet to the destination.
i3 uses the DHT nodes as a forwarding intermediary; a sender routes the packets for a destination IP to the node found by a Chord lookup on the ID, and that node forwards it on to the destination.

Initials: DRKP

IV DDoS

Alyssa P. Hacker obtains an IP address space allocation of 2^{22} bits to start her own ISP. Even before she sets up her network, she finds that her ISP is receiving TCP SYN-ACK segments (i.e., responses to TCP SYNs) from the IP address of web.mit.edu. These SYN-ACK packets don't have spoofed source addresses. Alyssa concludes that she is seeing the effects of some denial-of-service attack.

14. [2 points]: What specific attack is probably happening against web.mit.edu?

(Answer legibly in the space below.)

1

TCP SYN – flooding

+ spoofed IP.

15. [3 points]: Suppose the attackers are picking IP addresses completely at random from the IPv4 address space in attacking web.mit.edu. Alyssa observes a TCP SYN-ACK rate from the site of 100 segments per second. Estimate the rate at which web.mit.edu is experiencing attacks.

(Answer legibly in the space below.)

3

Alyssa has $\frac{2^{22}}{2^{32}} = \frac{1}{2^{10}}$ of the address space. So we estimate web.mit.edu is experiencing attacks at a rate of $100 \cdot 2^{10}$ segments/sec $= 102,400$ segments/sec

16. [4 points]: Alyssa tries to get to the site from her 1 Mbit/s access link and finds that the response to Web requests from web.mit.edu is slower than it should be. She correctly attributes the slow response to the on-going attack. Suppose that web.mit.edu is connected to the Internet using a 1 Gigabit/s link and that TCP SYNs and SYN-ACKs are 40 bytes in size.

Which of the following is a possible reason why Alyssa is unable to get prompt responses from the site.

(Circle True or False for each choice.)

4

A. True / False Network congestion at web.mit.edu's access link. 40 bytes · 1024 · 100/s = 4 MB/s

B. True / False Network congestion at Alyssa's ISP. 40 bytes · 100/s = 4 KB/s

C. True / False Kernel resource exhaustion at web.mit.edu.

Yes, due to SYNs.

Not enough to saturate lines

Initials: DRKB

V Pot Pourri

17. [3 points]: Consider an $N \times N$ switch that implements virtual output queueing. Assume that all input and output links have the same speed. Which of the following is true of crossbar scheduling algorithms in such a switch?

(Circle True or False for each choice.)

\

A. True / False If the switch has speedup of N, then it does not need much packet buffering at the inputs.

B. True / False To achieve 100% throughput with one iteration of the PIM algorithm may sometimes requires a greater switch speedup than achieving 100% throughput with iSLIP.

18. [5 points]: Which of the following statements is true of the IP multicast?

(Circle True or False for each choice.)

A. True / False It allows any host to send packets to the multicast group.

B. True / False If the multicast routing protocol always works correctly, then no packet sent to the group will ever be sent more than once on the same link.

C. True / False In some IP multicast routing protocols, routers must maintain state proportional to the number of receivers subscribed to each group.

D. True / False When a host sends a "join" message for a multicast group to its multicast-capable router, that router may create some multicast forwarding state. This state has to be refreshed periodically.

19. [3 points]: How does the TVA capability scheme thwart a "denial of capability" attack caused by hosts making a large number of capability requests?

(Answer legibly in the space below.)

Requests are rate-limited so they cannot use more than 5% of link bandwidth; this prevents the link from being overwhelmed with request traffic. To prevent DoS due to request processing time, the routers use fair queuing for requests, using network (AS) path identifier tags to select the queue.

3

End of Quiz II!

Initials: DRKP