Carnegie Mellon

Computer Science Dept.
Carnegie Mellon University

15-744 Fall 2004

Final

Name: _____

INSTRUCTIONS:

There are 13 pages (numbered at the bottom) make sure you have all of them.

Please write your name on this cover and at the top of each page in this booklet.

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

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

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

For multiple-choice CIRCLE ALL ANSWERS THAT APPLY.

SCORING:

A	B	C	D	E	F	G	H	I	Total
/17	/14	/13	/15	/16	/14	/14	/18	/4	/125

ALL MULTIPLE CHOICE WERE MAX 5PTS AND MIN 0PTS

A. Potpourri

1. Briefly define the full queues problem in the context of queue management and explain why it occurs. (4pts)

   TCP increases in speed until queues fill and losses occur – thereby, keeping queues in a full state

2. For each of the following strategies indicate wheter it solve the problems of full queues or lockout. (4pts, -0.5 per mistake)

Protocol	Full Queues	Lockout
drop tail	**N**	**N**
drop random	**N**	**Y**
drop head	**N**	**Y**
random early drop	**Y**	**Y**
fair queueing	**Y**	**Y**

3. .Why would/wouldnt you run BGP over TCP ? Give one reason for and one reason against. (4pts)

   for: session oriented, reliable delivery
   against: don't want to do congestion control over a routing protocol, tcp level attacks

4. Network measurements have shown which of the following to be true.

   a) On TCP connections, the probability of the packet after a lost packet also being lost is identical to the probability of any packet being lost. (-3)

   b) ISP topologies are accurately modelled by a random graph with the correct degree distribution. (-3)

   c) Packet pair techniques typically overestimate the speed of multi-channel links. (+2)

   d) That many of an ISPs BGP routing policies can be inferred from observations at other parts of the network. (+3)

B. CDN, HTTP, Web caching

5. The key idea behind Summary Cache is to have a cache keep track of the content of its sibling caches.

   e) Why does Summary Cache use bloom filters instead of just keeping the list of files on these sibling caches? (3pts)

   reduce the size of the index

   f) What drawback does this approach have? (2pts)

   false positives, queries may return yes even when the object does not exist

6. Bovik is trying to figure out a scheme his clients should use, so that given a URL U, they can find the appropriate CDN node to fetch the content from. Bovik starts with a hash function h that takes a string and maps it to a number $\alpha \in [0, 1)$. He also comes up with two different schemes for mapping URLs to different nodes.

   SCHEME 1: Each node is assigned a value, $h(\text{node}_i)$. When a client needs to fetch a URL with $h(\text{URL}) = \beta$ and has to decide which replica to query, it picks node_i, such that the absolute value of the difference between $h(\text{node}_i)$ and $h(\text{URL})$ is minimum.

   SCHEME 2: Let there be m CDN nodes in all; sort them using the $h(\text{node}_i)$ values. If the rank of a node is i, $(0 \leq i \leq m\text{-}1)$, it is responsible for storing all URLs that map to the interval $[i/m, (i+1)/m)$.

   Despite the fact that SCHEME 2 balances load slighly better, Bovik claims that a CDN with a large number of nodes (that occasionally crash and are later repaired) should choose SCHEME 1 over SCHEME 2. WHY? (4pts)

   less traffic overhead .. in scheme 2 addition/removal of a node can cause multiple pairwise exchanges because the intervals change with inc/dec in m, whereas in scheme 1 there are at most 2 nodes that have to re-sync their state

7. Suppose you retrieve index.html from cnn.com and find that it has 3 embedded images that have been akamaized. Assume that your browser does not use persistent connections and that your DNS cache is empty. Assume that the only TTLs used for DNS are 1 day and 1 minute. In retrieving the 3 images (not the html file), how many connections will your browser make to: (5pts, 1 per part)

   The original content provider 0

   the DNS root/gTLD server 1 to resolve akamai;s ns

   the Akamai high-level DNS server 1 to direct to a region

   the Akamai low-level 1 – the 1 minute TTL is sufficient for the entire transfer

   the closest Akamai server 3 one per image

C. DNS

8. Based on the Jaeyeon Jung (MIT) DNS measurement paper, what do you expect will be the DNS hit-rate when all web transfers use persistent connections? (4pts)

   0 % , 80% can arise from 4-5 embedded objects per page itself

9. Which of following is true about DNS? (5pts)

   a) The NS-record for a name usually has a larger TTL than the A-record for the name. (+2pts)

   b) While resolving the A-record for the name www.cs.cmu.edu, a name server that does not have the A-record cached may not need to go to the root or gTLD server. (+2pts)

   c) While resolving the A-record for the name www.cs.cmu.edu, suppose a name server contacts the name server for cmu.edu. Now, if the query packet sent here is lost, a new lookup originating from the root is initiated. (-4pts)

   d) One of the implications of the Jaeyeon Jung (MIT) DNS measurement study is that setting very low ttl's is unlikely to have an adverse impact on DNS performance (+2pts)

10. Why is reversing the IP address necessary in DNS reverse name queries? (4pts)

    to make the name most specific to least specific

D. P2P & DHTs

11. There are three forms of P2P lookup algorithms: centralized (napster), flooding-based (gnutella) and routing-based (DHTs). Which of the following statements is true about these algorithms? (5pts)

    a) Flooding-based and centralized systems can support much richer queries (regular expressions, wildcards) than routing-based systems. (+3pts)

    b) Routing-based systems are more scalable than flooding-based systems since they produce less traffic per search. (+3pts)

    c) Routing-based systems ensure that a client finds the copy of a file that is closest to it in the network. (-3pts)

    d) Ring-based DHTs are not as flexible as other DHTs since they only allow routing in one dimension. (-3pts)

12. Imagine a Chord system using 4-bit ids. Let there be 4 nodes participating with IDs 0, 3, 9, 11 and 12, . Use the table below to fill in the finger table for node 0. Please assume clockwise data assignment. (6pts, -1 per error)

ID pointed to	Node storing ID
1	**3**
2	**3**
4	**9**
8	**9**

13. Using the above Chord ring, what path would a request starting at node 0 take to find data item 12? (4pts)

    0 → 9 → 11 → 12
    9's finger table has IDs 10,11,13, 1

E. Mobile routing/transport/MAC

14. Consider the following topology of wireless laptops A, B, C and D.

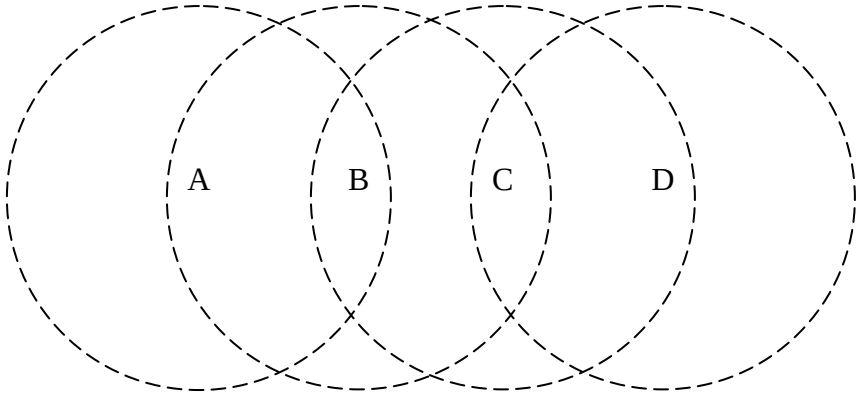

The dotted lines indicate the range of wireless transmissions from each node. For example, B is within range of A, A & C are within range of B, B & D are within range of C and only C is within range of D. Assume that each node uses an RTS/CTS based MAC protocol (i.e. like MACAW) (6pts, -2pts per error)

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

A hears the CTS

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

D heard the RTS from C

c) If B finishes sending and D was waiting to send to C, how does D know it's time to try again?

RTS contains the duration of the data transmission

d) Harry Bovik is considering implementing a walkie-talkie service for his wireless PDAs. His program largely uses small packets to avoid delaying any voice. Should Harry use RTS/CTS for his deployment? Why?

RTS/CTS is primarily to prevent collision resolution from finishing quickly

15. Which of the following are true about pure link-layer, tcp-aware link-layer (Snoop) and split-connection (itcp) approaches to wireless TCP? (5pts)

a) Snoop creates soft state at the base station (+3pts)

b) Snoop essentially just splits the TCP connection into two: one from wired to base station and one from base station to wireless host (-3pts)

c) If a wireless host moves to a new base station, the pure link layer approach requires that the TCP handshake must be performed again. (-3pts)

d) Split connection protocols impose hard state at the base station, which becomes a problem during handoffs. (+3pts)

16. Which of the following are true about IP routing to mobile hosts? (5pts)

a) Many Mobile IP solutions suffer from triangle routing. (+3pts)

b) When using Mobile IP, TCP connections must be reset when a mobile hosts moves between different IP networks. (-3pts)

c) Since DSR uses source routes, any node that overhears a packet transmission learns a valid path through the network.. (+3pts)

d) Geographic ad hoc routing requires O(logN) state at each node. (-3pts)

F. Multicast

17. IP multicast enables group communication, but requires enhancements to hosts and routers. Which of the following are true?

    a) Without truncation in a DVMRP network, packets transmitted with a large TTL are flooded throughout the entire network until routes are pruned back. (+3pts)

    b) A given multicast group can have only one sender (-3pts)

    c) Per-source multicast trees provide better latency than shared trees (+3pts)

    d) The Internet address registry assigns IP multicast addresses to groups (-3pts)

    e) Senders must join the multicast groups that they are transmitting to. (-3pts)

18. In SRM, the timers for requests and replies are randomized. What purpose does this randomization serve? (4pts)

    Helps in suppressing duplicate NACKs

19. Harry Bovik claims that even though overlay multicast is more popular than IP multicast, it uses more network *bandwidth* resources. Is he correct? Why? (5pts)

    Yes, overlay multicast must only use endpoints – as a result it cannot choose as optimal split points in the delivery tree.

G. Intserv/Diffservh

Harry Bovik is implementing QoS support in his network. He is a little confused by some of the concepts - it's your job to help him out.

20. All of the traffic in Harry's network has a utility (U) function like the one shown below. Would you recommend that Harry deploy IntServ in his network? Why or why not? (3pts)

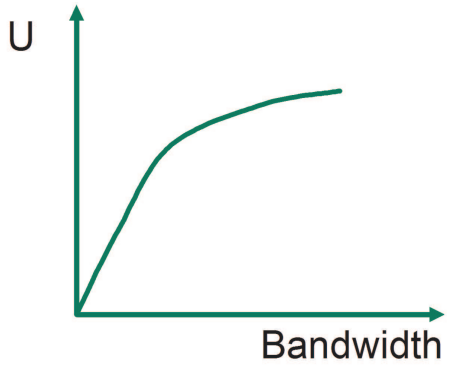

no .. concave function => admit more flows

21. Recall that the Integrated service (IntServ) architecture considers three types of service. For each of the 3 types of service, Harry needs to know whether policing at the edges to drop or mark packets is necessary or not. Explain each of your answers in one sentence. (5pts, -2 per error)

    guaranteed

    flow isolation is already guaranteed so no need for policing

    predicted

    need policing to prevent misbehaving flow from being affecting others

    best effort

    no need for policing

22. Harry realizes that IntServ is overloading his routers and he decides to use a DiffServ-based solution for his network. He decides that he only wants to support intolerant & rigid applications in this setup. These applications mark their packets as high priority (in-profile). He would like to setup RIO such that none of in-profile packets are dropped until all out-of-profile (best-effort) traffic is dropped. Draw the appropriate RIO drop rate graphs below. Also, clearly identify the relationship between max_in & max_out, min_in & min_out, Pmax_in & Pmax_out. (6pts)

If you want strict priority, min_in > max_out.

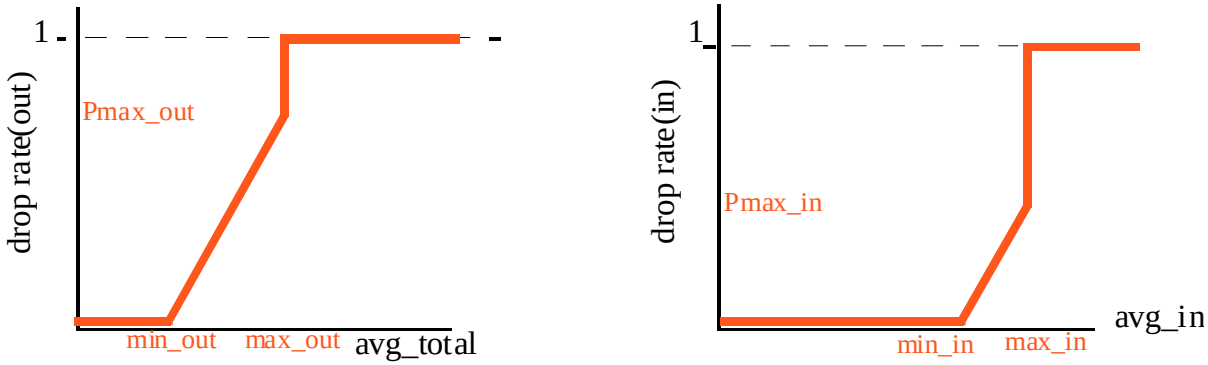

H. Security

23. Harry Bovik obtains a class A sized address space (2^24 addresses) to start his own ISP. Before he even installs any machines he notices that his ISP is receiving 600 SYN/ACK packets from www.google.com every minute. Assume that all packets are delivered, all solicit a response, and any attackers using spoofed addresses generate them at random..

    a) Describe the attack that is going on. (3pts)

    syn flooding of google using random addresses some which fall within boviks space

    b) What is the current rate of attack (3pts)

    (600/min)/(60min/sec) * 2^32/2^24 = 256 * 10 = 2560 probes/sec

A common technique in denial of service attacks is to insert a fake/bogus address in the IP source address field. One technique to identify the real source of such packets is to use node-sampling. Each router inserts its IP address into a single location in the packet header (possibly overwriting a previous entry) with some fixed probability P.

24. In node-sampling, how does a victim identify how far away a particular router is along the path to the attacker? (4pts)

    p(1-p)^(n-1)

25. One attack on the node sampling technique is to always insert a false router name into the header.

    a) How can you adjust P to still traceback to a malicious node? (2pts)

    p > 0.5

    b) Why might this solution make it take much longer to trace back to this node? (2pts)

    1-p becomes small .. further away routers need to get more packets

26. In the edge-sampling algorithm proposed by Savage et al., the routers in the network can use different marking probabilities. Why does it not require all routers to use the same marking probability like the node sampling algorithm? (hint: look at question 24) (4pts)

    distance is carried explicitly in each packet

I. Freebies

27. Name one paper you would eliminate from the second half of the semester.

28. Name one topic you wish the class had covered in greater depth.

29. Did Srini have a baby boy or baby girl at the beginning of the semester?

boy

THE END!!