MIDTERM TAKE-HOME EXAM

Csci5221: Foundations of Advanced Networking

Spring 2007

Prof. Zhi-Li Zhang

Last Name: **First Name:**

Student Id.

Instructions:

1. This is a **take-home** exam. **It is due by 5pm Wednesday Feb 28, 2007.** Please either email or hand in your exam to Prof. Zhang by the due time. (Please slide your exam under the door if Prof. Zhang is not in his office.)
2. Please remember to put down your name and student id. before you handle in. *Please work on your exam independently (i.e., by yourself, without discussing with any other students)! If you get your answer from another source, please cite and credit the source!*
3. There are *five* questions plus one *bonus* question. Each question may have several sub-questions. The number of points for each question is given in parentheses. There are 100 points total (including the bonus points).
4. Partial credit is possible for an answer. Please try to be as concise and make your exam as neat as possible. We *must* be able to read your handwriting in order to be able to grade your exam.
5. Good luck. Enjoy!

---

1. End-to-End Argument and Multicast (15 points total.

a. (5 points) What is the *end-to-end* argument?

b. (10 points) Based on your understanding of the end-to-end argument, do you think that "multicast" should be a functionality that is best placed at end systems, or within the network? Argue under what circumstances (e.g., in terms of group sizes, number of groups, bandwidth requirements, etc.) it is better to implement multicast at the end systems vs. within the network.

2. Implementation of Leaky Bucket Traffic Shaper/Regulator (15 points total.)

Please provide a detailed algorithmic description (e.g., in psuedocode) for efficiently implementing leaky bucket traffic shaper/regulator (ρ, σ) where ρ is the long-term average rate, and σ is the burst size. It is assumed that the system clock (say, of the line card processor) is as fast as the line speed. In other words, if the link speed is C bits per second, then the clock cycle is $1/C$ seconds. You can invoke a system call `clock()` with an input parameter t (i.e., `clock(t)`) which will generate an *interrupt* after t clock cycles. In your description, you need to pay attention to what states that need to be maintained, what need to be done when a packet arrives, when should a packet be queued, and when a queued packet is scheduled to depart the traffic shaper/regulator.

3. QoS Theory and FIFO Queueing (10 points total.)

Suppose we have n flows, where each flow i conforms to an arrival curve $\alpha_i(t) = \min\{\rho_i t + \sigma_i, Pt + M\}$, $P > \rho_i$, $\sigma_i > M$. (In other words, $x_i(\tau + t) - x_i(\tau) \leq \alpha_i(t)$ for $t \geq 0$, where $x_i(s)$ is the cumulative amount of bits generated by flow i up to time s. Note that P and M are the same for all n flows.) These n flows share a FIFO queue served by an FIFO scheduler with a line capacity of R, where $\sum_{i=1}^{n} \rho_i < R < nP$. Namely, the FIFO scheduler services the *aggregate* traffic of the n flows with a service curve $\beta(t) = Rt$.

a. In order to guarantee that there is no packet loss for any flow, what is the minimum buffer size we need to allocate for the FIFO queue?

b. What is the worst case delay bound that we can provide for each flow?

4. TCP Modeling and Performance (15 points.)

Suppose that we have two end hosts, S, and R, each having two separate network interfaces. Using these interfaces, S and R are connected to each other via two *disjoint* paths, P_1 and P_2. Let RTT_i and p_i be the (steady state average) round trip time and packet loss rate of path P_i, $i = 1, 2$. Suppose we need to transfer a large file (say, a video) between S and R. Consider the following two scenarios. (In the following we assume that all packets are of the same size.)

a. In scenario 1, we transfer the file over a single TCP connection between S and R. However, the network layer of S would automatically perform load-balancing between the two paths by sending one packet to P_1 and the next packet to P_2. Hence in this scenario, TCP is unaware that the packets are sent using two different paths with differing RTT's and packet loss rates. In other words, a single congestion window is maintained, and is adjusted based on the aggregate RTT and packet loss rate over the two paths. What will be the (steady state) throughput of this TCP connection (thus the application) using the approximate TCP throughput formula given in the lecture notes. Briefly justify your answer.

b. In scenario 2, we transfer the file by establishing two separate TCP connections, one using path P_1 and another using path P_2, and perform "round-robin" load-balancing at the application layer by sending one packet over one TCP connection, and the next packet over the other TCP connection. (In this case, the network layer does not perform load balancing!) Note that in this scenario, each TCP connection will maintain its own congestion window, and adapts it based on the corresponding path characteristics. What will be the (steady state) throughput of the application using the approximate TCP throughput formula given in the lecture notes. Briefly justify your answer.

c. Suppose $p_1 = p_2$. Which scenario yields higher throughput?

5. Packet Classification (25 points total.)

Consider the classifier (i.e., a prioritized list of classification rules) given the following table. Answer the following questions.

Rule	F1	F2	Action
R_1	10*	00*	deny
R_2	0*	10*	send to eth0
R_3	1*	1*	send to eth1
R_4	01*	0*	send to eth0
R_5	0*	1*	send to eth1
R_6	*	1*	deny
R_7	*	*	send to eth2

a. Please draw the *hierarchical tries* data structure for the above classifier.

b. Please draw the *grid-of-tries* data structure for the above classifier.

c. Given a packet with $F1 = 100000$ and $F2 = 010000$, what action will be taken on this packet?

d. Prove that any W-bit range G (i.e., $G = [l, u] \subset [0, 2^W - 1]$) can be represented by at most $2(W - 1)$ prefixes.

6. Multicast Revisited: Hybrid Implementation (*bonus question, 20 points total.*)

Suppose after your graduation, you joins an Internet service provider company (say, Google). The service provider wants to provide IPTV (e.g., streaming live TV programs) to users on the Internet. Within its own network (which, say, geographically spans the entire US), its routers support native source-specific IP multicast. In order to save bandwidth within its own network, it is desirable to take advantage of IP multicast to carry the multimedia streams *within its own network*. However, since users reside outside its own network, video streams have to be delivered to individual users as *unicast* streams. You are therefore asked to come up with a design that combines IP multicast within the provider's network with IP unicast outside the provider's network.

To make things more concrete, let S be a video server which streams live TV programs to users, and G_i, $i = 1, 2, \ldots, m$ be special "gateway" servers that reside between the provider's network and the rest of the Internet and "map" or convert the multicast streams within its own network to appropriate unicast streams to individual users. Suppose that a user accesses a IPTV program by accessing a website of the service provider, and given each user's IP address, it is mapped to one of the (closest) gateway servers. Internally (i.e., within the provider's network) each IPTV program is mapped to one multicast group (i.e., an IP multicast address). Given these assumptions, please *outline* a design focusing on the following two issues. (Note that you do not need to provide a lot of specific details, just the key ideas!)

a. *"Signaling" and Membership Management Issue:* note that in the standard IP multicast, users join a multicast group via IGMP. However, here since users do *not* reside within the provider's network, they (or rather, their end hosts) cannot use IGMP to join a multicast group. Instead, the "gateway" servers need to join the appropriate multicast groups on behalf of the users. We therefore need a "signaling" mechanism for the web server (upon receiving a user's HTTP request for an IPTV program) to convey to appropriate "gateway" server that a user wants to join a particular multicast group.

b. *Multicast/Unicast Stream Mapping and Delivery Issues:* the video server S will deliver the TV programs in multicast packets (i.e., the destination addresses of the packets are IP multicast address). In other words, the video server S is oblivious of individual users. Briefly describe the operations at the "gateway" servers, paying in particular the states that must be maintained the "gateway" servers for mapping and duplicating each multicast stream into multiple (say, UDP) unicast streams to be delivered to individual users.