EECS489 Computer Networks, Final Exam (Winter 2007)

Instructions: You are allowed to use two sheets of notes (front and back). This exam is closed book, no computers are allowed. You can use a calculator. Read the entire exam through before you begin working. Work on those problems you find easiest first. Read each question carefully, and note all that is required of you. Please keep your answers clear and concise, and state all of your assumptions carefully.
Please write out the details of how you reach your final answer in order to get partial credit.

You are to abide by the University of Michigan/Engineering honor code. Please sign below to signify that you have kept the honor code pledge.

Honor code pledge: I have neither given nor received aid on this exam.

Name: _____

Signature: _____

Uniqname: _____

Problem 1 [20 pts]: Security

In network communications, there are several desirable security properties. For example, *confidentiality* is the property that the original plaintext message cannot be determined by an attacker who intercepts the cyphertext-encryption of the original plaintext message.
Another important property is *message integrity*. This means that the receiver can detect whether the message sent (regardless if it was encrypted) was altered in transit.

[3 pts] (a) For the two properties of confidentiality and message integrity, can you have one without the other? Justify your answer. (If your answer is yes, show an example of where one can exist without the other. If you answer is no, explain why one would imply the other.)

[3pts] (b) *Digital signatures* is an electronic signature used to authenticate the identity of the sender of a message. Assume that you want to send a message M to your lawyer and give him/her the assurance that it was unchanged from what you sent (message integrity) and it is really from you (message authenticity). Describe how you can achieve this using public key encryption.

[3pts] (c) You have achieved message integrity and authenticity in (b). Describe how you can achieve confidentiality by slightly modifying your solution in (b).

[3pts] (d) To prevent replay attacks (you do not want an attacker to re-send the same message to your lawyer), describe a simple modification to your solution in (c).

[4pts] (e) Suppose we have a very short secret s (e.g., a single bit or even a Social Security number), and we wish to send someone else a message m now that will not reveal s but that can be used later to verify that we did know s. Explain why $m = MD5(s)$ or even $m = E(s)$ with RSA encryption would not be secure choices and suggest a better choice.

[4pts] (f) suppose two people want to play poker over the network. To "deal" the cards they need a mechanism for fairly choosing a random number x between them; each party stands to lose if the other party can unfairly influence the choice of x. Describe such a mechanism. (Hint: You may assume that if either of two bit strings x_1 and x_2 are random, then the exclusive-OR $x = x_1 \oplus x_2$ is also random. Another hint: try to use your solution from (e))

Problem 2: Network Mobility Support [10 pts]

[8pts] (a) Mobile IP support allows mobile hosts to be reached via their original IP addresses associated with their home networks. This is accomplished mainly with the help of the home agent which tunnels the traffic to the foreign agent before traffic finally reaches the mobile host. **Now consider the correspondent also to be mobile.**
Sketch below the additional network-layer infrastructure that would be needed to route the datagram from the original mobile use to the (now mobile) correspondent. Show the structure of the datagrams (in terms of src and dst IP address and any packet encapsulation) between the original mobile user and the (now mobile) correspondent using the figure below. (You need to show at least the packet from the mobile user to the correspondent and the reply packet from the mobile to the correspondent.)

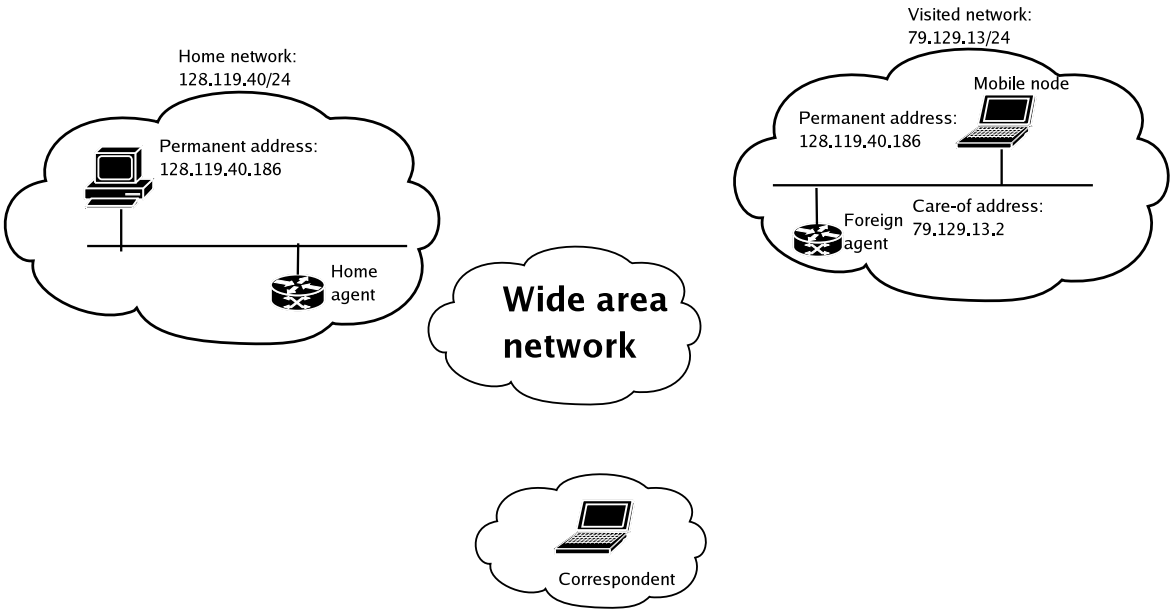

[2pts] (b) How would mobility support for both hosts (the correspondent and the mobile) affect end-to-end delays of datagrams between the source and the destination?

Problem 3: TCP and Routing [20 pts]

[3 pts] (a) TCP waits until it has received three duplicate ACKs before performing a fast retransmit. Why do you think the TCP designers chose not to perform a fast retransmit after the first duplicate ACK for a segment for a segment is received?

[3 pts] (b) Host A is sending an enormous file to Host B over a TCP connection. Over this connection there is never any packet loss and the timers never expire. Denote the transmission rate of the link connecting Host A and the Internet by Rbps. Suppose that the process in Host A is capable of sending data into its TCP socket at a rate Sbps, where $S = 10 \times R$. Further suppose that the TCP receive buffer is large enough to hold the entire file, and the send buffer can hold only 1% of the file. What prevents the process in Host A from continuously passing data to its TCP socket at rate Sbps?

[4pts] (c) Consider a datagram network using 8-bit host addresses. Suppose a router uses longest prefix matching with the following forwarding table:

Prefix match	Interface
00000000	0
01000000	1
10000000	2
11000000	3

For each of the four interfaces, give the associated range of destination host addresses and the number of addresses in the range.

[3pts] (d) Suppose two TCP connections share a path through a router R. The router's queue size is 6 segments; each connection has a stable congestion window of 3 segments. No congestion control is used by these connections. A third TCP connection now is attempted, also through R. The third connection does not use congestion control either. Describe a scenario in which, for at least a while, the third connection gets none of the available bandwidth, and the first two connections proceed with 50% each. Does it matter if the third connection uses slow start? How does congestion avoidance by the first two connections help solve this?

[7pts] (e) You are an Internet Service Provider; your client hosts connect directly to your routers. You know some hosts are using experimental TCPs and suspect some may be using "greedy" TCP with no congestion control. What measurements might you make at your router to establish that a client was not using slow start at all? If a client used slow start on startup but not after a timeout, could you detect that?

Problem 4: Multimedia Networking [10 pts]

[4 pts] (1) What is the difference between end-to-end delay and packet jitter? What are the causes of packet jitter?

[2 pts] (2) Why is a packet that is received after its scheduled playout time considered lost?

[2 pts] (3) Suppose we send into the Internet two IP datagrams, each carrying a different UDP segment. The first datagram has source IP address A1, destination IP address B, source port P1, and destination port T. The second datagram has source IP address A2, destination IP address B, source port P2, and destination port T. Suppose that A1 is different from A2 and P1 is different from P2. Assuming that both datagrams reach their final destination, will the two UDP datagrams be received by the same socket? Why or why not?

[2 pts] (4) Given the answer to (3), how would a host communicate with two separate clients at the same UDP port? (Hint: in CS, "a level of indirection" tends to be very powerful).

Problem 5: [20 pts] Wireless Networks

[2 pts] (1) Ethernet protocol uses CSMA/CD for multiple access control. When a collision is detected, how does Ethernet back off to avoid future collisions? What is the advantage of this scheme compared to random backoff?

[2 pts] (2) Why doesn't collision detection work well in wireless networks? (describe a scenario where this doesn't work.)

[9 pts] (3) Suppose nodes A and B are on the same 10Mbps Ethernet segment and the propagation delay between the two nodes is 225 bit times. Suppose A and B send frames at the same time, the frames collide, and then A and B choose different values of K in the CSMA/CD algorithm. Assuming no other nodes are active, can the retransmissions from A and B collide?

Assuming A chooses K=0, B chooses K=1, basically we need to find out whether A's retransmission will reach B before B's scheduled retransmission time.

[2 pts] (4) Why are acknowledgments used in 802.11 but not in wired Ethernet?

[2 pts] (5) In the CSMA/CA protocol, a station that successfully transmits a frame begins the protocol for a second frame at step 2 where its adapter has to sense the channel idle for 96 bit times before it can start transmitting again. What is the rationale behind this? Why not allow the station transmit the second frame immediately, if the channel is sensed idle?

[3 pts] (6) The idea of "snooping TCP" (shown below) is to buffer all unacknowledged packets to the destination mobile host and perform local retransmission in case of packet loss on the wireless link. Why is this beneficial? Would this work for IPsec traffic, what about SSL traffic?

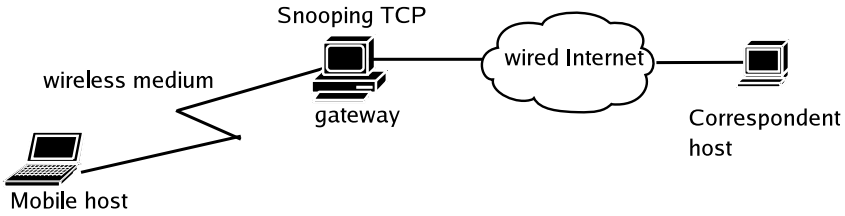

Problem 6: [20 pts] P2P and NAT

[4 pts] (1) Give two reasons why native IP multicast is not widely deployed today.

[6 pts] (2) Give an advantage and a disadvantage of application-layer multicast (ALM). Describe two metrics for evaluating the overhead of application-layer multicast.

[10 pts] (3) Given the following topology, with delay marked at each link. Assuming node A wishes to send data to all other nodes. *Mark* on each of the other three topologies how native unicast, native IP multicast, and ALM would deliver the content from A to B,C,D. For ALM, there are multiple options, pick the one that has the lowest overhead (based on the metrics in (2).

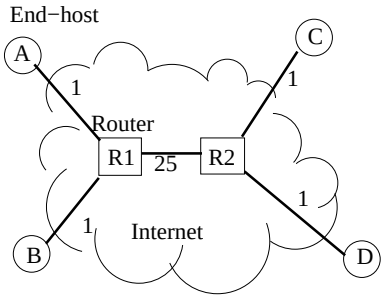

Topology

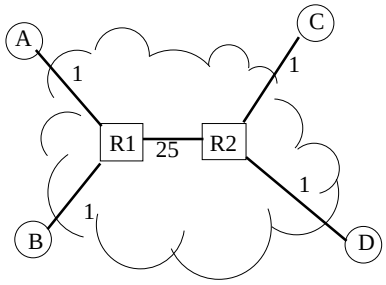

Unicast

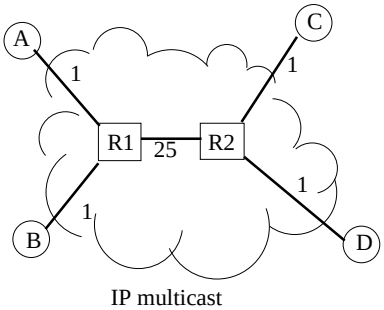

IP multicast

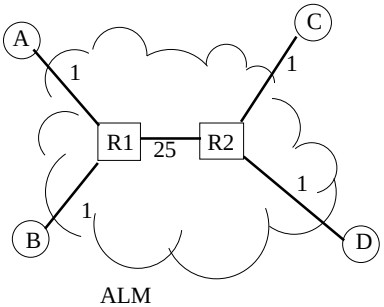

ALM

Extra Credit [10 pts]

[5 pts] (a) *Available bandwidth probing*: Explain whether available bandwidth estimation can be accurately validated using bulk TCP throughput. If so, explain why such estimation would be accurate. Otherwise, explain why such estimation would be inadequate.

[5 pts] (b) *Packet-pair based probing*: Given an H-hop path defined by the sequence of capacities $P = \{C_0, C_1, ..., C_H\}$. Two packets of size L are sent back-to-back from the source to the sink. These packets are the *packet pair*. The *dispersion* of the packet pair is the interval from the instant the last bit of the first packet is received at a certain path point to the instant of the last bit of the second packet is received at that point. The dispersion is $\Delta_0 = \tau_0 = L/C_0$ after the source, and let it be $Delta_i$ after link i. When the packet pair reaches the sink, the dispersion is $Delta_H$ and the receiver computes a bandwidth estimate $b = L/Delta_H$.

Assuming no cross traffic is present, what does b measure, i.e., what kind of bandwidth does b represent?

Under what condition can we claim that dispersion Δ_i cannot be lower than the dispersion at the previous hop Δ_{i-1}. What kind of delay would cause the dispersion to decrease at a later hop?