B1. Networks (25 points total)

Four packet-switched nodes P, Q, R, and S are connected by three full-duplex communication links as shown below:

P--------Q -------R --------S

The links P-Q and R-S have data rates of 10 Mbps each, while the link Q-R has a data rate of 1 Mbps. The distance between each pair of neighboring nodes is 800 km. The signal propagation speed is 2×10^8 meters/sec. You may assume there is no processing delay at any of the nodes.

(a) (3 points) What is the total end-to-end delay for sending exactly one packet of size 5000 bytes from P to S?

(b) (6 points) Suppose 100 packets of size 5000 bytes each are transmitted continuously from P. What is the total end-to-end delay for sending this sequence of packets from P to S?

(c) (8 points) Suppose 500 Kbytes of data is to be transmitted in a long continuous stream of packets from P to S. You are given two options:

Option A: Use a packet size of 5000 bytes, and double the data rate of link R-S.

Option B: Reduce the packet size to 1000 bytes, and keep the data rate of link R-S unchanged.

Which of these two options would you prefer if your objective is to achieve (i) a higher throughput, (ii) a lower total end-to-end delay? Support your answer with appropriate facts and/or reasoning.

(d) (8 points) A large file is to be transmitted from node P to node S. The MTU of the network that contains links P-Q and R-S is 5000 bytes, while the MTU of the network that contains link Q-R is 2000 bytes. The data is transmitted via TCP with each segment having a TCP header of size 20 bytes. IPv4 with fragmentation is being used with an IP header of size 20 bytes. Each link also uses a link layer header of size 20 bytes. What is the percentage header overhead on each of the three links? Assuming that P transmits all the data continuously with no constraints of window size or congestion mechanisms, what is the average application layer throughput seen by the end-users?

B2. Networks (25 points total)

A user on host "stimpy.eecis.udel.edu" receives a file from a server on host "www.cnn.com". The connectivity between these hosts is shown in the figure below. The conventional Internet protocol suite is being used.

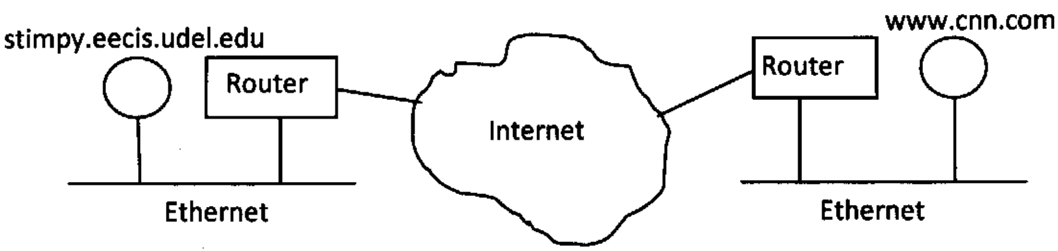

(a) (3 points) Identify the three primary means by which each host and/or the host interface is identified.

(b) (4 points) Name the protocols that are used to map between these three different methods of host identification. At what protocol layers do these mappings occur?

(c) (6 points) What is the rationale for using three different identification methods to identify hosts on the Internet? List the benefits of using these three different identification methods and of mapping between them. What are the disadvantages of using this scheme?

(d) (6 points) Explain the concepts of "route aggregation" and "longest prefix matching". What are the advantages of each of these concepts in the context of Internet routing?

(e) (6 points) How does longest prefix matching support route aggregation? Answer this question by considering a hypothetical network that uses route aggregation but does not use longest prefix matching. What problems would occur in such a network?

B3. Networks (25 points total)

(a) (10 points) Explain the difference between a unicast, anycast, multicast, and broadcast destination address.

(b) (5 points) Flooding is one means for local information to be distributed to all other nodes in a network. In a naive implementation of flooding (herein called "pure"), when a packet arrives on an incoming link of node X, the node X transmits a copy of the packet out every one of its outgoing links. Discuss the main positive and negative feature(s) of pure flooding.

(c) (10 points) In practice, the negative features of pure flooding outweigh the positives. However several improvements can modify pure flooding to make the resulting "controlled" flooding a worthwhile means of distributing local information. For example, controlled flooding is used in peer to peer protocols (e.g., Gnutella) and link state routing protocols (e.g., OSPF, IS-IS). Describe how controlled flooding works.

B4. Networks (25 points total)

(a) (5 points) In the context of reliable data transfer protocols, what is a checksum?

(b) (6 points) Explain how a checksum is used.

(c) (4 points) The Ethernet checksum is physically located at the end of the Ethernet protocol data unit (PDU). Is there any advantage of putting it there rather than say at the beginning of the PDU?

(d) (4 points) Ethernet uses CRC-32 "100000100110000010001110110110111" to compute its 32 bit checksums. In general what does this binary number represent?

(e) (6 points) Consider the simpler CRC-3 "1011". Show the math involved to use this CRC-3 to compute and append an appropriate checksum to the data "1111".