

MidTerm Exam

Oct 12, 2010

NAME: _____

1. **Modeling Ethernet Networks** (20 points) The papers by Alefiya Huessein and Mohamed Hafez discussed the tradeoffs between modeling Ethernet networks in varying levels of detail. While they discuss somewhat different approaches to this, they both talk about three different levels of detail.
- (a) Fully Abstract. Speed of light is assumed to be infinite, and each node will transmit packets in order with no collisions and no backoffs.
 - (b) Partially Abstract. Speed of light is assumed to be infinite, but each station will randomly back off if any other station is transmitting.
 - (c) Fully Detailed. Accurate speed-of-light models, completely accurate model of the Ethernet behavior.

Discuss briefly the pros and cons of the three approaches, and mention when each of the approaches might be the right choice.

2. Modeling Protocol Stack Processing Delays (20 points)

The *GTNetS* simulator, as well as most other popular network simulators, ignore delays in the protocol stack processing, and ignore any routing lookup delays in routers. In other words, a packet will be presented to the layer 4 processor at exactly the same time as it arrived at the layer 2 processor. Further, a packet will be transmitted on an output link at the same instant as it arrived on some other input link (assuming the output link is not busy).

(a) Comment on the loss of accuracy in the simulation results due to this simplifying assumption.

(b) What would be the affect on simulator performance if we removed this simplifying assumption and in fact modeled these delays?

3. **Tuning RED for Web Traffic** (20 points)


(a) Give a critique of Jeffay's experimental method in the *Tuning RED* paper. Are these results representative of RED performance in the Internet? How could this experiment have been improved?

(b) The measurement study by Bruce Mah is the basis for both Jeffay's study, and the Web Browser model in *GTNetS*. Give at least two criticisms of this model.

4. **Lab Project 1** (20 points)

In lab project1, we measured the overall TCP throughput as a function of several independent variables. The graph below shows throughput as a function of the TCP segment size, for a fixed bottleneck queue limit of 8000 bytes. Notice that performance goes down as the segment size increases.

Suppose we specified the queue limit in *packets* instead of in bytes. For a fixed queue size in packets (say for example 64 packets), what would the below graph look like in this case? Explain your rationale for the predicted behavior.



5. Measuring Web Browser Performance (20 points)

The graph below shows a cumulative distribution function of web response time for several different bottleneck queue limit values. These results were obtained using the GTNetS simulator with an experimental setup similar to Jeffay's in the Tuning Red paper.

Give a plausible explanation why the line for the 30 packet queue limit appears to give the best results for smaller response times and gradually gets to be the worst for larger response times. In other words the lines for the 30 packet limit and the larger packet limits cross over at various points in the graph. Why would this happen?

