

NETWORKING AREA EXAM – Spring 2002

This exam consists of 9 questions. Please answer exactly 6 of them. Typed answers are preferred but don't spend a lot of time type-setting equations, hand written equations are fine.

If you have any questions during the exam please contact Mostafa Ammar, ammar@cc.gatech.edu, 404 894 3292 (Office), (770) 432 2070 (Home).

1. In a high speed switch that uses a crossbar design, a controller must determine which inputs can send to which outputs in each data transfer cycle. The paper by Turner and Yamanaka states that the controller must solve a maximum matching problem.
 - (a) Explain how one constructs the appropriate instance of maximum matching, using an example and then stating the general principles.
 - (b) How does the problem change if the switch supports multicasting?
2. The host anycasting service described in RFC 1546 provides stateless best effort delivery of an anycast datagram to at least one host. Describe the design of a reliable anycast packet delivery service.
3. Using a transit-stub topology and making reasonable assumptions about inter- and intra-domain routing protocols, demonstrate that the routing protocols can construct the following:
 - (a) A forwarding path that is *not* a shortest hop count path.
 - (b) A forwarding path that remains within a single domain and is *not* a shortest hop count path.
 - (c) A forwarding path that goes a long physical distance even though the endpoints are in close physical proximity (i.e., physically close together).
4. The following defines a network load and resource sharing condition called *continuously backlogged fair sharing* (CBFS). Let n be the number of sessions and r be the total bandwidth of the link. In CBFS,
 - (a) (Fair Sharing) Each session has equal weight, that is, for any $1 \leq i < j \leq n$, $\phi_i = \phi_j$.
 - (b) (Continuously Backlogged) Each session has a packet arrival at time 0. Also, for any $t > 0$ and $1 \leq i \leq n$, $A_i(t) \geq \frac{r}{n}t$. Here $A_i(t)$ is the amount of session i traffic that has arrived during the interval $[0, t]$.

The following defines another such condition *CBFS+*. In *CBFS+*, the j th packet ($j \geq 2$) in session i arrives just at the time the $(j - 1)$ th packet finishes service under the GPS scheduler. In other words, each packet in session i arrives just in time to satisfy the CBFS condition.

Answer the following two questions:

- (a) Given a packet arrival instance that conforms to CBFS condition, does Virtual Clock always generates the same service schedule as WFQ? Explain why or why not? Here tie-breaking (e.g., on GPS virtual finish times) does not need to be considered.
- (b) Given a packet arrival instance that conforms to *CBFS+* condition, does *WF²Q* always generates the same service schedule as WFQ? Explain why or why not? Again, tie-breaking does not need to be considered.

5. Our context here is strictly the current Internet built on top of TCP/IP protocols. For each of the following security functions, explain whether or not the end-to-end arguments are applicable to it.
 - (a) Authentication
 - (b) Integrity
 - (c) Secrecy
 - (d) Nonrepudiation
 - (e) Traceback of Distributed Denial of Service Attacks

Make sure you are working with the correct definitions of these functions. A handout with these definitions is available in the NTG Lab.

6. This question deals with the use of queuing theory to understand the behavior of packet-switched networks.
 - (a) By way of a concrete example, show how a queueing theoretic model can be used to evaluate the end-to-end delays packets may encounter in a packet-switched network.
 - (b) What assumptions would be required for one to obtain “closed form” equations describing packet delays in your model and how realistic are these assumptions.
 - (c) Comment on the suitability of using such queueing models in the design and evaluation of the Internet today.
7. One of the main requirements for schemes for multicasting of video over the Internet is that they deal with receiver heterogeneity.
 - (a) Explain what is meant by *receiver heterogeneity*. How may one model such heterogeneity quantitatively.
 - (b) List and explain a comprehensive set of evaluation criteria that are important to consider in evaluating a video multicast scheme’s functions designed to deal with receiver heterogeneity.
 - (c) Briefly describe the layering and replication schemes and evaluate them using the criteria you discuss above.
8. This question deals with the issue of state maintenance in routers within the Internet:
 - (a) Describe two networking contexts where it is desirable for a router to maintain state.
 - (b) There in general a distinction made between “Hard” and “Soft” state. Describe the distinguishing features of each of the states.
 - (c) Which type of state is in general preferred within the Internet and why? Do you agree with this preference, why or why not?
 - (d) Assume you want to systematically and quantitatively want to evaluate the difference between Hard and Soft State, describe the process by which you would set out to perform such an investigation.
9. This question deals with TCP Congestion Control
 - (a) What is an “ambiguous acks” and why does it require special treatment. Explain how and why TCP reacts to the reception of an ambiguous ACK.

- (b) Modern TCP congestion control includes “fast retransmit” and “fast recovery” procedures, describe how these work and explain the reasoning behind their incorporation in TCP.
- (c) As originally defined TCP (RFC793) did not include any congestion control and avoidance mechanisms. One can argue that incorporating congestion control functions within TCP has made it difficult to re-use the same functions in other transport protocols. Assume that the Internet continued with the simple TCP in RFC 793 and the choice was made to implement *end-to-end* congestion control functions in a separate protocol. Devise such a protocol that would be deployed only at end systems and that would work with the plain RFC793 TCP (and any other transport protocol) to perform congestion control functions. The protocol should be architecturally independent from TCP (i.e. operating at the layer below or above TCP or operating “beside” TCP, akin to the relationship of ICMP and IP). At a minimum your description should include the modified end-system protocol stack and a description of how the protocol would work and the properties it would achieve.