Remarks: Please keep the answers compact, yet precise and to-the-point. Long-winded answers that do not address the key points are of limited value. Binary answers that give little indication of understanding are not good either. Time is not meant to be plentiful. Make sure not to get bogged down on a single problem.

PROBLEM 1 (40 pts)

(a) In what sense is TCP's Slow Start "slow," and why can "fast" be considered a more accurate characterization? What is the rationale behind Slow Start? What is likely to happen when TCP congestion control mediates transfer of small files in non-congested networks? What role does TCP's Congestion Avoidance play that complements the role of Slow Start? What is the reason behind Congestion Avoidance employing linear increase/exponential decrease instead of, say, linear increase/linear decrease? Explain using the technical meaning of "congestion" as decrease in throughput when load becomes excessive.

(b) Suppose a new company which owns and manages its own computing systems, including IP addresses leased from its ISP, decides to use a domain name, newkiddglobal.com, along with a web server www.newkidotbglobal.com, and email addresses for people working at the compancy. The company will run primary and secondary DNS servers, ns1.newkiddglobal.com and ns2.newkiddglobal.com, respectively. What information must the company convey to its DNS registrar (e.g., GoDaddy.com) so that the latter can initialize DNS with the .com TLD so that newkiddglobal.com and its services can be accessed through the Internet? What responsibility falls unto the company—not its registrar—to set up so that DNS can resolve names related to newkidotbglobal.com? Suppose after several months the company decides to use a more powerful machine with hostname s1.newkiddglobal.com and IPv4 address A as its web server. Describe two methods for achieving this transition. Explain what determines how long it takes for the Internet to resolve www.newkiddglobal.com from its old IPv4 address to A.

PROBLEM 2 (30 pts)

(a) What are the main distinguishing features of intra-domain routing protocols OSPF, RIP, and IS-IS? Why has Purdue utilized RIP whereas large ISPs prefer OSPF and IS-IS? Suppose a company Z is an autonomous system (AS) and has two transit providers, AS X and AS Y, through which it connects to the Internet. Suppose the service level agreement (SLA) with provider X is such that it costs less during daytime to send/receive IP packets whereas it costs less during nighttime to send/receive IP packets through Y. How would Z's border router running BGP and connected to border routers belonging to X and Y need to behave so that Z's total Internet access cost can be minimized? Explain what types of BGP messages Z's border router would need to communicate to its counterparts—and in what sequence—to achieve the goal.

(b) Suppose three coffee shops, J, K, and L are located adjacent to each other along the same side of a street where K is between J and L. Each coffee shop operates its own 802.11 WLAN hot spot—same make of WiFi router connected wireline to an ISP—that customers use to access the Internet. The three coffee shops are close to each so that a wireless station associated with J can sense a station associated with L, and vice versa. By happenstance, the owners of the three coffee shops configured their basic service sets (BSSs) to use the same channel. What is likely to happen to the throughput experienced by customers of the three coffee shops, and why? Suppose coffee shops J and L moved away from each other along the street so that wireless stations belonging to the two hot spots cannot sense each other. The coffee shop in the middle, K, remains within carrier sense distance of J and L. Compared to the first scenario, what is likely to happen to the throughput experienced by customers of the three coffee shops? What is a simple solution to the problem?

PROBLEM 3 (30 pts)

(a) Routers on the Internet that are equipped with IETF DiffServ capability, by default, run with the capability disabled. What makes it difficult for an ISP from enabling DiffServ on its routers and exporting platinum, gold, silver, bronze services to its customers who access services on the global Internet? In what way can DiffServ scheduling that utilizes WFQ (weighted fair queuing) be viewed as an improvement over priority scheduling? Does marking the TOS field of IPv4 packets with a better service class (e.g., platinum) imply that they receive improved service over packets whose TOS field is marked with a lesser service class (e.g., silver) when using WFQ? What about when differentiated services is implemented using priority queuing? Explain your reasoning.

(b) In lab5 we implemented a pseudo-real-time audio streaming client/server application using UDP. Unlike Problem 1 of lab4 where we built a reliable file transfer app using UDP in which lost packets were retransmitted, in lab5 we did not implement reliable transport. Audio packets that were lost remained lost. Suppose in a different version of lab5, the client reads from the audio buffer one audio packet (e.g., 4 KB) every 100 msec (not 313 msec) which is written to the client's audio device. Hence γ is fixed at 10 pps (packets per second). Suppose that the audio streaming service operates in a network environment where packet losses occur infrequently, and if a loss occurs, the likelihood of a retransmitted packet also being lost is extremely low.

In the modified version of lab5, the sender (i.e., server) sets a timer to 2×RTT before transmitting an audio packet. If a positive ACK is not received within the timeout, the audio packet is retransmitted. Assume RTT is known or accurately estimated, and one-way delay (e.g., from server to client) is RTT / 2. The target buffer level at the receiver (i.e., client), Q^*, which specifies the amount of prefetched audio data, is measured in unit of packets. Assuming congestion control method D is able to maintain buffer level close to Q^*, how large must Q^* be in relation to RTT so that retransmitted packets are likely to arrive at the client before its deadline for playback on the audio device? Explain your logic and derivation. In lab5, we implemented the receiver's FIFO buffer as a circular queue. What simple data structure is suited to implement the modified protocol where lost audio packets are retransmitted and put in their proper sequence at the client before playback?

BONUS PROBLEM (10 pts)

What is the main performance improvement feature implemented by HTTP/1.1 compared to version 1.0? Why does the feature improve HTTP performance? What is the HOL (head of line) problem of HTTP/1.1? What method has been used by HTTP/1.1 to help address the problem? Why is the method considered insufficient? Describe the approach used by HTTP/2 to mitigate the HOL problem.