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 utility. 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) What is the technical definition of *congestion*? Does packet loss imply congestion? What is the aim of congestion control? How is room temperature control using a thermostat related to congestion control?

(b) Internet routing follows a user-optimal, i.e., selfish, approach to selecting paths as opposed to system optimal routing that seeks to minimize all users' total delay. What are two detrimental consequences of the selfish design choice? Why, in spite of the negatives, is a user-optimal approach employed? Does user-optimal routing imply that an end user actually gets to select a path of his/her choosing?

(c) Why is UDP better suited for real-time multimedia transport than TCP? Would TCP be adequate for pseudo-real-time streaming? For real-time multimedia streaming, why would one not dispense with UDP and use raw IP instead?

(d) What are the top two solutions—the second one, in turn, has two subsolutions—that have prevented 32-bit IPv4 addresses from becoming depleted without migrating to a 128-bit IPv6 address infrastructure? How have these solutions affected IP packet routing/forwarding? Which one is relevant for global Internet routing and which one is pertinent to home networking?

PROBLEM 2 (30 pts)

(a) How do TCP SYN flooding and ICMP flooding achieve denial-of-service at a target server? From an attacker's perspective, which requires less effort? Which is more taxing on the victim assuming an equal number of packets arrive at the victim? Which might be easier to defend against? Sketch a solution that could be carried out at a border router.

(b) Why are ad hoc WLANs a bad idea when the number of wireless stations becomes large? Use a $\sqrt{n} \times \sqrt{n}$ grid topology where n is the number of nodes to argue the case. Is the primary limitation inherent to wireless networks or is it shared by (ad hoc) wireline networks? What feature, if any, is specific to wireless networks? Should a large residential community in Lafayette connect as an ad hoc WLAN if one of the homes has a 1 Gbps Internet connection and charges other homes a flat fee of $5 for "unlimited" Internet access?

(c) Describe the steps that TCP undertakes to tear down a full duplex connection. How is this related to the 2-person consensus problem and why is TCP's solution a hack? What may go wrong as a result? Is this a serious practical problem? How is TCP's solution for connection set-up different from connection termination, in fact, being less of a hack?

PROBLEM 3 (30 pts)

(a) Suppose aliens from a distant galaxy have abducted you (a danger that comes with taking CS536) to employ you as their new telecommunication czar in hopes of reigning in network resource wastage that is crippling their economy. Upon arrival, you find that `AoL` (`Aliens-on-Line`), the state run network access provider, is operating a huge dial-up modem pool, one per customer. You also find, through measurement, that dial-up requests from users are approximately independent, following an exponential distribution be^{-bt} between successive request arrivals. Each dial-up session lasts at most 10 seconds (aliens surf & talk rather quickly). You conduct a customer survey and find that customers are willing to tolerate a small blocking probability δ in return for a significant price reduction that may be obtained by reducing the number of high-speed modems `AoL` has to maintain. Describe the steps that you would undertake to estimate the number of modems you have to operate assuming there are n customers.

(b) We know that congestion control method D, $d\lambda/dt = \varepsilon(Q^* - Q) - \beta(\lambda - \gamma)$, is superior to method B (linear increase/exponential decrease), implemented and benchmarked in one of the homeworks in the context of pseudo-real-time audio streaming. What are the technical challenges associated with implementing method D for the "black box" Internet? What are some non-technical barriers that would need to be overcome to implement method D? Come up with a measurement-oriented variant of D, call it method E, that may be feasibly implemented for the global Internet.

BONUS PROBLEM 4 (10 pts)

Measurement studies of the Internet have revealed that network connectivity (inclusive information networks such as the World Wide Web) and file sizes obey a power-law distribution. What are the performance consequences—good and bad—of these power-law phenomena when compared to a network environment where network connectivity and file sizes are exponential? How does power-law connectivity impact information search on the Internet (aka why `google` has been beating the competition)?