CS 536 (Park) **Final** **Dec. 12 (Mon.), 2016**

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 addition to routing at the LAN layer using MAC addresses, two major complementary systems determine packet routing on the global IP Internet. What are they and what roles do they play? What is their main difference? Provide an example as a packet travels from one IP device to another on the Internet. Why is there a need for LAN layer routing? What is a key concern and how is it addressed?

(b) Suppose we are using HTTP (version 1.1) on port 80 using GET to fetch a file /index.html at www.purdue.edu. Describe the format of a generic request message discussed in class. What happens if we use version 1.0? Suppose the response specifies status code 200 (OK) for success with header line containing "set-cookie: name=xyz". Assuming the HTTP client accepts cookies, what happens to the GET request when a second request for the same URI is issued later in the day? How can cookies benefit a service provider? How can a browser that uses caching employ if-modified-since to reduce response time and network traffic? How can an HTTP server use ETag to aid a caching HTTP client improve performance?

(c) What feature of wireless networks in indoor environments distinguishes it from outdoor environments? Suppose mobile device A is closer in proximity to a WLAN access point than another mobile device B. By the (inverse) quadratic distance rule of transmission power, can we conclude that A will achieve better reception/transmission? What actions may device B undertake (if feasible) to improve its throughput?

(d) Internet data traffic is known to be significantly more bursty than telephony traffic. What causes burstiness that persists even at larger time scales? Why is streaming of compressed video traffic bursty? Why is providing quality of service to video streaming apps more challenging than voice traffic? What are the two methods that today's routers support for providing quality of service (QoS) to IP traffic? These features, by default, are disabled. What are the challenges associated with enabling them to provide QoS-sensitive services on the Internet?

PROBLEM 2 (30 pts)

(a) Describe the role of TCP's two main congestion control components: congestion avoidance and slow start. In what sense can "slow start" be viewed as a misnomer? The bulk of Internet traffic is comprised of TCP sessions that transport files. Given the heavy-tailed ("many mice, a few elephants") property of real-world file sizes, what contributions do the two TCP components make when transporting files? What is TCP's fast retransmit and what is its rationale?

(b) IEEE 802.11n/g WLANs use OFDM to transmit bits at the link/physical layer even though CSMA (i.e., DCF) is used at the link layer as the MAC protocol to achieve sharing of bandwidth. What is the rationale for not using OFDMA—a contention-free MAC protocol—that assigns individual users their own orthogonal carrier frequencies in place of CSMA? What is the role of OFDM in WLANs considering that contention-based CSMA is being used to mediate bandwidth sharing?

(c) What is QoS routing and why is it considered algorithmically a difficult problem? In practice, is this a significant concern? Both OSPF and IS-IS implement shortest-path routing. What is their main difference? How is RIP different from OSPF? Which protocol was used at Purdue for the longest time? Why is this protocol not suitable for intranets of tier-1 transit providers? What type of autonomous system is Purdue? What does that mean? With respect to global Internet routing, how is Purdue known to the outside world? Provide one specific address block (we noted three in class).

PROBLEM 3 (30 pts)

(a) What is the 2-party consensus problem? Describe the intuition (not a formal argument or proof) behind why there is no solution to the problem. How does the problem manifest itself in TCP connection termination? How does TCP go about addressing the problem? Give an example of a typical sequence of events that lead to tearing down of a TCP connection. In the example, what may go wrong that makes TCP's "solution" a heuristic? How are application programmers exposed to this issue? How is consensus for connection set-up handled differently?

(b) We discussed four congestion control methods in the context of pseudo real-time multimedia streaming. What was Method A's problem and how did Method B address the problem? What issue did Method B face that it could not resolve? What was the key idea behind Method C? What was the innovation behind Method D that made it outperform Methods B and C? TCP's congestion avoidance mechanism is based on Method B. Why is adopting Method D for TCP file transport (not multimedia streaming) not a simple porting matter? Discuss your reasoning.

BONUS PROBLEM (10 pts)

What is the Internet's end-to-end paradigm? How does it manifest itself in the context of reliable data transport? What about secure data transport (e.g., confidentiality)? How about in the context of IPv4 fragmentation?