Homework 1 for CSE 561 (Autumn 2004; Wetherall)

Due: Wed, Nov 17, 2004, at the beginning of class. (Out: Mon, Nov 8, 2004.)

1. Peterson 1.12 (in 2nd edition: 1.10)

2. Bandwidth, latency, and Internet measurement. In the Internet, a fixed size error indication packet is returned from a router to the sender if the TTL reaches zero during forwarding. Consider an Internet path composed of N links with differing bandwidth and delay that are connected by routers. Explain, giving a procedure and expressions, how a source can use the round-trip time of packets to infer the bandwidth and delay of each link along the Internet path. You can assume that the rate and delay are the same in both directions along a link and that routing is symmetric; make other simplifying assumptions that you need.

3. Checksums. Consider the 16-bit Internet checksum as given in Peterson 2.4.2. (Also see RFC 1071 for interest.) Characterize the error detection capabilities of this checksum against both burst and random bit errors. Explain what errors will always detected and why. Explain what errors will sometimes be detected and the probability with which they will be detected and why. Make simplifying assumptions as you need them.

4. Peterson 2.13 (in 2nd edition: 2.11)

5. Link-State Routing Convergence. This question explores how quickly link-state routing converges to stable routes after a failure. Assume a simplified network in which all links have equal propagation delay and all routers process messages equally fast. Consider the case of a single link failure that does not partition the network.

a) Describe the state of network connectivity immediately after the failure, before the nodes have dealt with it.
b) How long might it take from the instant of actual failure until all nodes have connected, loop-free routes to all destinations if we are using arbitrary link costs? What network factors does your answer depend on?
c) How long might it take from the instant of actual failure until all nodes have connected, loop-free routes to all destinations if we are using shortest-hops as our cost metric? What network factors does your answer depend on?
d) How would your answers above differ if we were using distance vector rather than link-state?

6. BGP Path Vector Convergence. This question explores how quickly a path vector protocol converges after a failure. Assume a simplified network in which N ISPs are meshed to each other only at one Internet Exchange point, all links connecting the ISPs have equal propagation delay, and all routers process messages equally fast. Assume BGP is selecting the shortest route, and when there are equal length routes it is breaking ties by choosing the lowest next hop address. Consider the case of a single link failure elsewhere in the network that causes ISP 1 to lose reachability to an address P.

a) Before the failure, what route is an ISP X using to reach P and what other routes is it hearing from its neighboring ISPs?

b) After the failure and ISP 1 sends a withdraw message for the path 1 to reach P, what route will an ISP X be using to reach P?
c) How long will it take until all ISPs realize that they have no valid route to P?
d) How would you change the protocol to fix this problem to provide more rapid convergence?

7. Sliding windows. Consider a standard sliding window protocol. The size of the sliding window affects the transfer rate that transport protocols such as TCP can achieve.

a) Give an expression for the throughput as a function of the size of the sliding window. What would this throughput be in practice if stop-and-wait were used on a cross-country connection (100ms round trip time)?
b) TCP's maximum receive buffer size for flow control is 64K without the use of extensions. (Defaults are typically 16 or 32K too.) At what throughput will flow control become the limiting factor for cross-country connections? What sliding window size would be needed to support a 1Gbps cross-country connection?
c) How does TCP slow-start change your answer for the expression in part a)?

8. TCP Sawtooth. TCP uses a shadow congestion window, cwnd, to control congestion within the network. The size of this window ideally follows the TCP "sawtooth" of additive increase with multiplicative decrease after a single packet loss.

a) What is the steady-state relationship between transmission rate and packet loss for this sawtooth?
b) What is the maximum packet loss rate needed to sustain transfers of 1Gbps?
c) Why does this make high-speed TCP transfers difficult to achieve?

9. Peterson 6.14a (in 2nd edition: 6.13a)

10. Guaranteed services. Consider the network topology below. A is a flow with an average rate of 1Mbps and bursts of up to 500KB. B is a flow with an average rate of 5 Mbps and bursts of 100KB. Describe how you would set the network up to provide rate and delay guarantees for flows A and B by using a general procedure that holds for any topology. What conditions must hold?

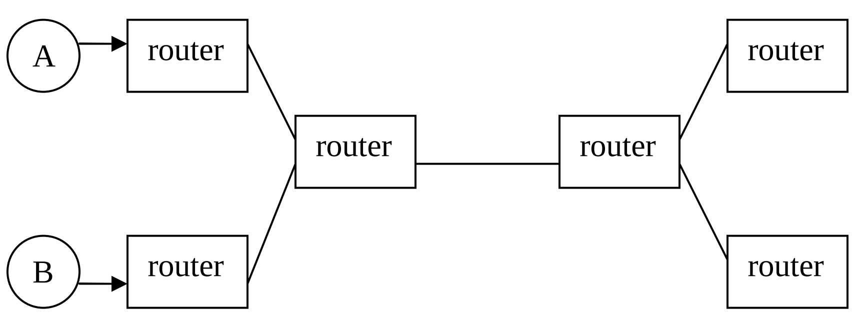

——END——