Example sheet 5

TCP

Network Performance—DJW—2011/2012

Question 1. This question is about using relay nodes in a TCP connection. Sometimes we have a choice: either use a single TCP connection from the source to the destination, or use one TCP connection from the source to the relay plus another TCP connection from the relay to the destination. In the latter case, the end-to-end throughput is the lesser of the rates of the two TCP connections, as determined by the TCP throughput equation.

(i) State the TCP throughput equation, and explain the terms in it.

(ii) Consider the following scenario. Between me and my destination there is a noisy wireless link with packet drop probability 15% and round trip time 1ms, and a backbone link with packet drop probability 1% and round trip time 200ms. What throughput do I get with a single TCP connection? What throughput do I get if I use a relay between the two links?

Question 2. The drift model for the window size $w(t)$ of a TCP flow, with constant packet drop probability p and constant round trip time RTT, is

$$\frac{dw(t)}{dt} = \frac{1}{\mathsf{RTT}} - p\frac{w(t)^2}{2\mathsf{RTT}}.$$

Let $x(t) = w(t)/\mathsf{RTT}$ be the throughput. Find a drift equation for $x(t)$.

Question 3. ScalableTCP has been proposed as an alternative to TCP. It uses a congestion window w like TCP, but it has different rules for increasing and decreasing the window. When it receives an ACK it increases w by α; when it detects a drop it cuts the window by βw.

(i) Find the drift model for Scalable TCP. Hence find the throughput equation for ScalableTCP.

(ii) Consider a link with n TCP flows and n ScalableTCP flows, where all flows have common round trip time RTT. Sketch a graph which shows total TCP throughput and total ScalableTCP throughput as a function of packet drop probability p. At what value of p do they share the link fairly?

Question 4. This question is about buffer size for Internet routers. Consider a collection of n TCP flows sharing a single bottleneck link. Theoretical study has proposed two extremes: very small buffers, and very large buffers, for the queue at the bottleneck link. Both extremes should in theory achieve near-full utilization, i.e.

$$y \approx C$$

where y is the sum of the throughputs of all of the n flows, and C is the link capacity.

The two extremes of buffer size differ in their contribution to delay. The total round trip time RTT for a TCP flow is made up of two parts: $\mathsf{RTT} = \mathsf{PT} + \mathsf{QT}$ where PT is the propagation delay (determined by the distance and the speed of light) and QT is the queueing delay (determined by buffer size).

(i) If the buffer is small, then $\mathsf{QT} \approx 0$. Suppose that all n flows have identical PT. Write down the TCP throughput formula, and calculate p.

(ii) If the buffer is large, then the rule of thumb says to set the buffer size so that $\mathsf{QT} \approx \sum \mathsf{PT}_i/n$, where PT_i is the propagation delay of the ith flow. Suppose again that all n flows have identical PT, and QT is set according to the rule of thumb. Calculate p. How does it compare to the drop probability with small buffers?

The throughput of a TCP flow with round trip time RTT is proportional to $1/\mathsf{RTT}$. Suppose there are $n/2$ flows with $\mathsf{PT} = 2a$ and $n/2$ flows with $\mathsf{PT} = 4a$.

(iii) With a small buffer, what fraction of capacity ends up being allocated to the flows with short RTT?

(iv) With a large buffer, what is the recommended QT? Then, what fraction of capacity ends up being allocated to the flows with short RTT?

(v) Which would you recommend, small buffers or large buffers? Why?

Question 5. The diagram below shows the paths from sender to receiver for each of three sets of TCP flows, sharing two links. At a link with capacity C and incoming traffic y, the packet loss rate is $p = (y - C)/y$ if $y > C$, or $p = 0$ if $y \le C$. The return paths have negligible delay and loss.

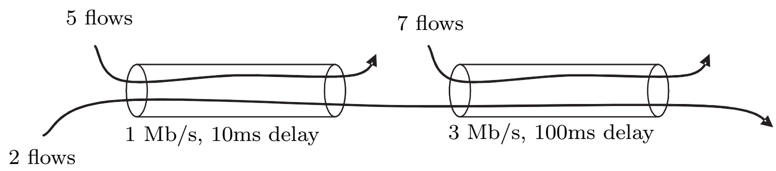

(i) Write down the fixed point equations for this system. Describe an iterative procedure to calculate the flow rate of each of the TCP flows. Execute this procedure.

(ii) Write down drift equations for the flow rates of each of the three flows. (The drift equations depend on the current packet drop probabilities, which in turn depend on the current flow rates.) Simulate this drift model.

(iii) Explain why the solution of part (i) agrees with the fixed point in part (ii).

Question 6. The drift model for the transmit rate of a TCP flow is

$$\frac{dx(t)}{dt} = \frac{1}{\mathsf{RTT}^2} - p(t)x(t)\frac{x(t)}{2}$$

where $x(t)$ is the throughput and $p(t)$ is the packet drop probability at time t, and RTT is the round trip time. It has been suggested that a more accurate drift model is

$$\frac{dx(t)}{dt} = \frac{1}{\mathsf{RTT}^2} - p(t - \mathsf{RTT})x(t - \mathsf{RTT})\frac{x(t)}{2}, \tag{1}$$

and that a suitable function for the packet drop probability is

$$p(t) = \frac{(1-\rho(t))\rho(t)^B}{1-\rho(t)^{B+1}} \quad \text{where} \quad \rho(t) = \frac{x(t)}{C}. \tag{2}$$

(i) Justify equations (1) and (2). Explain what B and C represent.

(ii) Simulate the drift model for (1) & (2) for $\mathsf{RTT} = 0.3$sec, $C = 10$pkt/sec and $B = 10$pkt, running your model for at least 30 seconds of simulated time. Repeat with $B = 50$pkt. Plot $x(t)$ for your two simulations. What do you observe?