

Example sheet 7

TCP

Network Performance—DJW—2010/11

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. 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) State the window increase and decrease rules for TCP.
- (ii) State the drift model for TCP. Find the drift model for Scalable TCP and explain your working.
- (iii) State the throughput equation for TCP. Find the throughput equation for ScalableTCP and explain your working.
- (iv) 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 3. 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: $RTT = PT + QT$ 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 $QT \approx 0$. 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 $QT \approx \sum PT_i/n$, where PT_i is the propagation delay of the i th 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 $1/RTT$. Suppose there are $n/2$ flows with $PT = 2a$ and $n/2$ flows with $PT = 4a$.

- (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 4. 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 $y \leq C$. The return paths have negligible delay and loss.



Describe an iterative procedure to calculate the flow rate of each of the TCP flows. Execute this procedure.

Question 5. The drift model for the transmit rate of a TCP flow is

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

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

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

and that a suitable function for the packet drop probability is

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

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

Question 6. What is meant by the statement “TCP has a utility function”? [Do not waffle. Simply state equations and theorems.]

Suppose two users with different RTT are sharing a link. Sketch the utility curves for the two users. Use your sketch to explain why one user gets higher throughput than the other.