Answer THREE of FIVE questions
Marks for each part of each question are indicated in square brackets
Calculators are permitted

1. Write an essay on fixed points in network modelling. You should

   (a) explain what is meant by the fixed point of a drift model, and give an example of its application to network modelling;

   (b) describe a simple simulator, e.g. written in Excel, for finding a fixed point of a drift model;

   (c) describe the fixed point method for solving simultaneous equations, and give an example of its application to network modelling;

   (d) explain the difference between stable and unstable fixed points.

   Be precise and concise.

[Total 33 marks]

2. (a) Explain what is meant by a processor-sharing link.

[4 marks]

   (b) Describe a Markov process model for a processor-sharing link. Draw the state space diagram, and explain carefully how you derived the transition rates.

[10 marks]

   (c) Calculate the equilibrium distribution of the above Markov process model, and find the mean number of active jobs.

[10 marks]

   (d) State Little's law.

[4 marks]

   (e) Use Little's law to calculate the mean job completion time in a processor-sharing link.

[5 marks]

[Total 33 marks]

3. In a fully-connected network running the Dynamic Alternative Routing algorithm, the probability $B$ that a given link has all its circuits busy can be found by solving the fixed-point equation

$$B = E\left[\frac{1}{\mu}\big(\lambda + 2\lambda B(1-B)\big), C\right].$$

(a) Describe the routing algorithm that gives rise to this equation.

[3 marks]

(b) Define the terms $E$, $\mu$, $\lambda$ and $C$ in the equation. Explain carefully how the equation is derived.

[10 marks]

The Dynamic Alternative Routing algorithm is meant to make the network robust to link failures. We will now investigate what happens when a link fails.

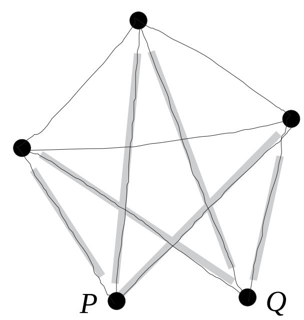


Suppose that $P$ and $Q$ are two of the $n$ nodes in the network, and the link between $P$ and $Q$ fails. We would expect that the links which terminate at $P$ or $Q$ should become slightly busier, as traffic between $P$ and $Q$ shifts to alternative paths. Label a link *red* if it terminates at $P$ or $Q$; there are $2(n-2)$ of these. Label the other links *green*. Let $B_{\text{red}}$ be the probability that a red link has all its circuits busy, and let $B_{\text{green}}$ be the probability that a green link has all its circuits busy.

(c) Derive fixed-point equations for $B_{\text{red}}$ and $B_{\text{green}}$. Explain your reasoning. *[Hint. First consider a red link, say between P and some other node M, and list all the different routes that might use this link. Then consider a green link, and do the same.]*

[20 marks]

[Total 33 marks]

4. Consider a TCP flow. Let $w(t)$ be its congestion window at time $t$, let the round trip time be RTT, and let the packet drop probability be $p$. The drift equation for $w(t)$ is

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

and the throughput equation is

$$x = \frac{\sqrt{2}}{\mathrm{RTT}\sqrt{p}}.$$

(a) Describe the algorithm by which TCP adjusts its congestion window. Write down the drift equation, find its fixed point, and derive the throughput equation. [8 marks]

UCL is part of a research effort to develop a multipath version of TCP. The idea is that each flow might have several subflows, each with its own congestion window. We will investigate congestion-control algorithms for a TCP with two subflows. Let the congestion windows for these two subflows be $w_A(t)$ and $w_B(t)$. Let the two subflows experience packet drop probabilities $p_A$ and $p_B$. Assume that the round trip time is equal to RTT on each subflow.

In order that multipath TCP should increase its total window just as much as regular TCP, it has been proposed that when subflow $A$ receives an ACK it should increase its congestion window by $1/(w_A(t)+w_B(t))$, and when it detects a drop it should cut its congestion window by $w_A(t)/2$; and subflow $B$ should behave similarly.

(b) Write down drift equations for $w_A(t)$ and $w_B(t)$ under this proposal. Explain how to find the fixed point (though you need not work through the algebra). [6 marks]

A second proposal, which aims to match the increase and the decrease of regular TCP is as follows. When subflow $A$ receives an ACK it should increase its congestion window by $1/(w_A(t)+w_B(t))$, and when it detects a drop it should cut its congestion window by $(w_A(t)+w_B(t))/2$; and subflow $B$ should behave similarly.

(c) Write down drift equations for $w_A(t)$ and $w_B(t)$ under this second proposal. [6 marks]

(d) Sketch the drift diagram for the second proposal, for the case $p_A = p_B$ and for the case $p_A < p_B$. Your diagrams for both cases should have two axes, $w_A$ and $w_B$, and arrows to indicate the drift in $w_A$ and $w_B$. Identify the fixed point or points. [13 marks]

[Total 33 marks]

5. Consider a FIFO queue operating in discrete time. In every timeslot, the probability that a packet arrives is $p$, and the probability that a packet is served is $q$. No more than one packet can arrive per timeslot, and no more than one packet can be served. Arrivals and services are independent, and successive timeslots are independent.

   (a) The total queue size can be modelled by a Markov chain. Draw its state space diagram, and write down the probability of each transition. *[Hint. Remember that this is a Markov chain, not a Markov process, so the sum of transition probabilities out of any given state must equal 1.]*

   [6 marks]

   (b) Find the equilibrium distribution. Show that the probability that the queue is empty is $1-\rho$, where $\rho = [p(1-q)]/[q(1-p)]$. Find the mean queueing delay.

   [12 marks]

   Now consider a wireless network with three nodes, operating in slotted time. In each timeslot, the probability that a packet arrives at node $i$ is $p_i$. If nodes 1 and 2 attempt to transmit at the same time then there is interference and neither node manages to send a packet; likewise for nodes 2 and 3. In order to limit collisions, the network designer proposes that node $i$ should transmit with probability $r_i$, if it has any packets queued.

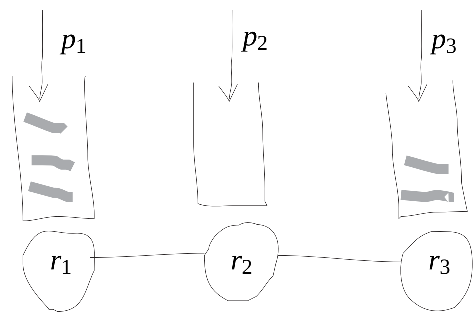


   (c) Assuming that node 1 has packets queued, it will successfully transmit a packet if (i) it attempts to transmit, and (ii) either node 2 has no packets queued, or it has packets queued but does not attempt to transmit. Write down a formula for the probability that node 1 successfully transmits a packet. Write down similar formulae for the other two nodes. Derive a set of fixed-point equations for computing the probability of successful transmission. *[Hint. Make sure you have as many equations as you have unknowns. Take $p_i$ and $r_i$ as known.]*

   [15 marks]

   [Total 33 marks]

*A system designer could use your equations to calculate the mean queueing delay at each of the three queues. She could then choose the probabilities $r_i$ so as to minimize the overall average queueing delay.*

END OF PAPER

1. I have taken measurements of file sizes served by my web server. Here is a sample of file sizes, in bytes:

$$1427, 4713, 1409, 186, 4315, 1108, 431, 4215.$$

(a) Explain how to plot an empirical distribution function, and draw a rough sketch of the empirical distribution function for this sample.

[6 marks]

When I plot the empirical distribution function of the full dataset using a logarithmic y-axis, i.e. when I plot $\log \text{prob}(X \geq x)$ as a function of $x$, it seems to follow two straight line segments, one from $x = 0$ to 1024, one from $x = 1024$ up.

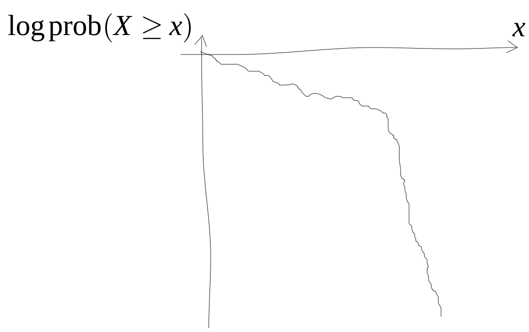


I propose fitting the distribution

$$\log \text{prob}(X \geq x) = \begin{cases} -\lambda x & \text{if } x \leq 1024 \\ -\lambda x - \mu(x - 1024) & \text{if } x > 1024 \end{cases}$$

This distribution has density function

$$f(x) = \begin{cases} \lambda e^{-\lambda x} & \text{if } x \leq 1024 \\ (\lambda + \mu) e^{-(\lambda+\mu)x + 1024\mu} & \text{if } x > 1024 \end{cases}$$

(b) Find formulae for the maximum likelihood estimators of $\lambda$ and $\mu$.

[15 marks]

(c) Give pseudocode for a random number generator that generates random samples from this distribution.

[12 marks]

[Total 33 marks]

2. (a) Describe briefly the processor-sharing model for TCP.

[3 marks]

(b) Describe a Markov process model for a processor-sharing link. Draw the state space diagram, and add labels showing the transition rates.

[5 marks]

(c) Calculate the equilibrium distribution and the mean number of active jobs.

[10 marks]

The processor-sharing model is an idealization of how TCP allocates bandwidth. One way to refine this model is to use the TCP throughput equation in conjunction with a formula for packet drop probability, to calculate the throughput that each TCP flow should get.

(d) Write down the TCP throughput equation, and the formula for packet drop probability at a FIFO queue with buffer size $B$, and explain briefly how to use them to calculate TCP throughput.

[10 marks]

(e) Modify the transition rates in your state space diagram to take account of this refinement.

[5 marks]

[Total 33 marks]