Coursework 4

Advanced random processes
Network Performance—DJW—2009/10

This coursework is worth 4.5% of your final grade. You should hand it in to Computer Science reception by noon on 22 January 2010. You may use any computer language you wish. The mark scheme awards more marks for correct and well-explained reasoning than it does for computer code and final answers. In the questions that ask you to write computer code, you should explain what your code does and how it works, as if to a colleague who can program but who doesn't understand any of the maths; you should also include the source code as an appendix. If you produce computer output that does not make sense, but you do not comment on the problem or explore the reason for the nonsensical output, you will lose marks.

In this coursework, you will use advanced analytical techniques to understand the behaviour of the multipath TCP that you simulated in Coursework 2.

Little's Law specifies a relationship between the average number of active jobs in a system, the arrival rate of jobs to the system, and the average length of time that a job spends in the system. Little's Law is a theorem about long-run averages; it does not apply exactly to simulations of finite duration. It is interesting to explore how accurate Little's Law is, for typical simulation runs.

Question 1. Write down a relationship derived from Little's Law between the average number of active flows from source A and the average completion time of these flows, in the setting of Coursework 2. Use the simulator to test this relationship. Repeat for sources B and C.

Hint. It is up to you to choose parameter values that illustrate the point you wish to make. It is also up to you to choose how to present your results.

Here is a standard fact about Poisson processes, known as *thinning*. If an arrival process is a Poisson process with rate λ, and if each incoming job is labelled "green" with probability p and "red" with probability $1-p$ independently of other jobs, then the arrival process of green jobs is a Poisson process of rate λp, and the arrival process of red jobs is an independent Poisson process of rate $\lambda(1-p)$.

Question 2. Using the concept of thinning, explain why the `unipath` system behaves like four independent processor-sharing links. Using theoretical results that you have learnt, calculate the values of x for which link 1 is stable. Repeat for the other links. Also, find the average flow completion time for each of the three sources. Validate your findings with the simulator.

Hint. It is up to you to decide what simulations to run, and how to present your results. For every theoretical claim you make, you should present a simulation to back it up, and you should explain why it is that the simulation backs up your theoretical claim.

Question 3. Let (n_A, n_B, n_C) be the number of active flows from each of the sources A, B and C at some point in time. What are the possible values that (n_A, n_B, n_C) can take? Write out the state space diagram, and find the transition rates.

In order to make calculations simpler, we shall (for Question 4 only) work with a simpler system that has only two sources. Specifically, let there be two sources A and B, and three links 1, 2 and 3 with capacities $C_1 = 2.5$ Mb/s, $C_2 = 4$ Mb/s and $C_3 = 2$ Mb/s. Suppose that jobs arrive at the two sources according to two independent Poisson processes with arrival rates $(\lambda_A, \lambda_B) = (0.4x, 0.6x)$ jobs per second, and that job sizes are all Exp(1) Mb. Suppose that jobs from source A can use links 1 and 2, and that jobs from source B can use links 2 and 3. The `uncoupled` algorithm works as described in Coursework 2. The `coupled` algorithm can be described more simply: when there are n_A flows active from source A and n_B from source B, then each flow from source A gets throughput

$$\frac{1}{n_A}\min\Big[\max\Big(\frac{n_A}{n_A+n_B}(C_1+C_2+C_3),C_1\Big),C_1+C_2\Big]$$

and each source B flow gets throughput

$$\frac{1}{n_B}\min\Big[\max\Big(\frac{n_B}{n_A+n_B}(C_1+C_2+C_3),C_3\Big),C_3+C_2\Big].$$

Question 4. *Answer parts (i)–(iii) below for both the* `coupled` *and* `uncoupled` *algorithms, and for a range of values of x in the range 5–15. You do not need to repeat your answer for every value 5, 6, 7, ..., 14, 15: instead, pick out two or three values of x that illustrate the typical behaviour.*

(i) Find formulas for the drift in n_A and the drift in n_B.

(ii) Draw the drift diagram. *Hint. Your diagram should have two axes, one for n_A and one for n_B.*

(iii) Pick an arbitrary starting pair of values for (n_A, n_B), and solve the drift model numerically. Draw the solution on your drift diagram. Repeat for several other starting values.

(iv) From your drift diagrams, for what values of x is `uncoupled` stable? For what values of x is `coupled` stable?

(v) Use the simulator to validate your findings about stability.

It has been suggested that the fixed point method might be used to estimate the average number of active flows from each of the three sources, for the `coupled` and `uncoupled` algorithms. In applying the fixed point method, we shall compute the average number of active flows at source A, $\bar{n}_A$, assuming that the numbers of active flows at sources B and C are constant and are given by $\bar{n}_B$ and $\bar{n}_C$ respectively. Then we will compute $\bar{n}_B$, assuming that the numbers of active flows at sources A and C are constant and are given by $\bar{n}_A$ and $\bar{n}_C$ respectively. Similarly for $\bar{n}_C$.

In order to calculate $\bar{n}_A$ given $\bar{n}_B$ and $\bar{n}_C$, we need to write out a state space diagram for n_A, compute the equilibrium distribution, and then find the mean of the distribution. Under the assumption that the numbers of active flows at B and C are constant, the state space diagram for n_A is very simple, and it is easy to find the equilibrium distribution using the detailed balance method. Although the state space diagram has infinitely many states, we can find a reasonably good approximation to the equilibrium distribution by counting only say the first 500 states. The mean of the distribution is just $\bar{n}_A = \sum_i i\mathbb{P}(n_A = i)$.

Question 5. *Answer parts (i)–(iv) below for both the* `coupled` *and* `uncoupled` *algorithms, and for a range of values of x.*

(i) Draw out a state space diagram for n_A, the number of active flows at source A, assuming that the numbers of active flows at sources B and C are $\bar{n}_B$ and $\bar{n}_C$ respectively. What are the transition rates? Repeat for the other two sources.

(ii) Program a function that computes $\bar{n}_A$ given $\bar{n}_B$ and $\bar{n}_C$.

(iii) Use the iterative fixed point method to compute $\bar{n}_A$, $\bar{n}_B$ and $\bar{n}_C$. Explain carefully what you are doing.

(iv) Use the simulator to validate your findings.