Coursework 2

Due 18 Jan 2008

Network Performance—DJW—2007

This coursework will study fair bandwidth sharing in a simple network. Consider the following network, consisting of two links and three traffic routes. Both links have the same capacity C Mb/s.

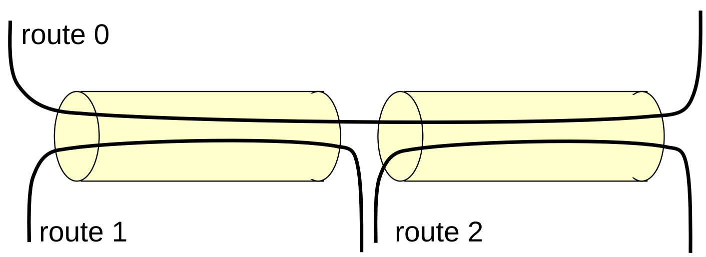

If there are n_0 flows active on route 0, n_1 on route 1, and n_2 on route 2, then TCP will allocate to each of the flows on route 0 the throughput

$$\frac{C}{n_0 + (n_1^2 + n_2^2)^{1/2}}$$

and the residual capacity on each link $i = 1, 2$ is shared equally between the flows on route i. This concept of fairness has been generalized by Mo and Walrand [1995], who give the name α-fairness to the allocation in which each of the flows on route 0 gets

$$\frac{C}{n_0 + (n_1^\alpha + n_2^\alpha)^{1/\alpha}}. \tag{1}$$

Interesting values of α are

$\alpha = 0$	Maximum throughput, often seen as maximally efficient
$\alpha = 1$	Proportional fairness, introduced by Kelly [1997]
$\alpha = 2$	TCP fairness
$\alpha = \infty$	Max-min fairness, often seen as perfect fairness

When $\alpha = 0$ or $\alpha = \infty$, the allocation formula (1) is replaced by its formal limit:

$$\text{if } \alpha = 0, \quad \frac{C}{n_0 + n_1 + n_2} 1_{n_1 = 0 \text{ or } n_2 = 0}$$

$$\text{if } \alpha = \infty, \quad \frac{C}{n_0 + \max(n_1, n_2)}$$

You will be required to simulate this system. Assume that jobs arrive as Poisson processes, and the rate of arrivals of jobs on route i is ν_i, where $\nu_0 = 0.5$, $\nu_1 = 0.3$, $\nu_2 = 0.4$ jobs/sec. Assume that all job sizes are exponential, with mean size 1 Mbit.

1. Let $\alpha = 2$. Vary C. At each value of C, measure the mean number of active jobs on each of the three routes. What do you find? For what values of C is the system stable?

2. Pick C such that the total mean number of active jobs you measured in question 1 is around 50–100. Now vary α, including the cases $\alpha = 0, 1, 2, \infty$.

How does α affect the mean number of active jobs, and how does it affect stability?

3. Repeat question 2, but using different distributions for job sizes. Include the case where all jobs are size 1 Mbit, and the case where job sizes are heavy-tailed with mean size 1 Mbit. What do you observe?

Hints. *You should plot appropriate graphs of your results. Refer to Section 0.3 of lecture notes for a discussion of what makes good graphs. Remember to plot error bars.*

The questions direct you to simulate the system for certain parameter values. These are minimum requirements; if you find that extra simulations will help you write a more complete account of your observations then run these extra simulations.

In your answers, you should not merely report on your simulation results. You should also discuss any theory which can put your results in context, in particular the theory for the $M/G/1$ processor-sharing queue. In addition, you should comment on the implications of your findings for network design.

You should provide the source code for your simulator, and also a one-paragraph description of its general structure. Explain how long you leave it to run-in, and how you detect when the system is unstable. Describe any sanity checks you have performed to check that the simulator is working correctly, e.g. setting the capacity of one of the links to be very large and comparing the output to the $M/M/1$ processor-sharing queue.

I was able to run each individual simulation in under 10 minutes. Do not spend too much time running hundreds of very long simulations: you will be marked for the quality of your explanations, not the quantity of CPU cycles.

References

F. P. Kelly. Charging and rate control for elastic traffic. *European transactions on telecommunications*, 1997. URL `http://www.statslab.cam.ac.uk/~frank/pf/`.

Jeonghoon Mo and Jean Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 1995. URL `http://walrandpc.eecs.berkeley.edu/Papers/moIEEE.pdf`.