COMPSCI 742 S2 C Assignment 2

Department of Computer Science

The University of Auckland

Due Sunday 18 September 05, 11:59 pm

This assignment will contribute 1/3 of your coursework mark, and 10% to your overall course mark.

Submit your assignment via the DropBox, either in PDF (preferred), or in MS Word format. Assignments in other format will *not* be accepted or marked!

TCP flavours, Email Queue statistics, M/M/k/n queue simulation

1. **TCP Flavours** [8 marks]

   Over the years, a number of different *flavours* of TCP have appeared. They include *Tahoe* (1988), *Reno* (1990), *New-Reno* (1995), *SACK* (1996-98) and *Vegas*. *Vegas* has not been widely implemented; the others were discussed in FS96, i.e.

   > *Simulation-based comparisons of Tahoe, Reno and SACK TCP,*
   > Kevin Fall, Sally Floyd, CACM, vol 26, pp5-21, 1996.

   That paper is available on the 742 *resources* page.

   (a) Read the FS96 paper, then write short summaries (about two paragraphs) summarising *Tahoe, Reno* and *New-Reno.* Your summaries should make it clear which of the algorithms explained in RFC 2581 are implemented in each flavour. [6 marks]

   (b) Find some information (by searching the web) about the *Vegas* flavour. Write a short summary (again, about two paragraphs) of it, and comment on why it has not been widely implemented. [2 marks]

2. **Email queue statistics** [6 marks]

   The University's email system runs several servers, with load-sharing arrangements to spread the incoming email messages across four servers. The servers accept incoming email messages, and send them to a suitable email forwarder.

   On the 742 *Assignments* page you will find a file called `a2.log`. It's a log file with one record for each arriving email message to one of the servers. Each record begins with three fields:

   - Inter-arrival time (i.e. time since previous message), seconds
   - Arrival time, Unix seconds
   - Sender address

   The sender address is UA for messages originating inside the University, and '*' for messages originating outside.

CONTINUED

(a) Produce three plots (on the same graph) showing the probability distribution for all the records, the UA records, and the '*' records. *You'll need to count the number of records for each interarrival time so as to produce a histogram. Then you can divide each count by the total number of records to convert them to probability estimates.* [3 marks]

(b) Assuming that the interarrival time distribution is exponential ($f(x) = \lambda e^{-\lambda x}$), make an estimate of λ for the '*' records interarrival times. Plot your exponential function on the same scales as your '*' probability distribution so as to demonstrate how well it fits the data. [3 marks]

3. **M/M/k/n queue simulation** [1 mark]

On the 742 *Assignments* page you will find a file called `mmkn.tcl`. It's an ns script, providing a simple simulation of an M/M/k/n queueing system.

(a) Consider an M/M/1/20 system with $\lambda = 2.3$ and $\mu = 3.0$
Using the formulae from our lecture notes, calculate the average number of requests in the system, and the average length of its non-empty queue. [2 marks]

(b) Run the mmkn script using the parameter values above. Compare its 'average number of requests in the system' with the one you calculated above. Comment on any difference you observe. [1 mark]

(c) Modify the mmkn.tcl script to compute the average length of its non-empty queue. Compare it with the one you calculated above. [3 marks]