Assignment #1

CS5262 Multimedia Networking and Systems

Department of Computing Science, National Tsing Hua University Hsin Chu, Taiwan

Due at 10:00 a.m. September 30, 2011. Please turn in hardcopies before the lecture starts. See course website for grading policies, especially about late submissions.

The following questions are from [KR08] Kurose and Rose, Computer Networking: A top-down Approach Featuring the Internet, Addison Wesley, 2008.

1) Consider the procedure describe in Sec. 7.3 for estimating average delay d_i. Suppose that $u = 0.1$ Let $r_1 - t_1$ be the most recent sample delay, let $r_2 - t_2$ be the next most recent sample delay, and so on.
   a) For a given audio application suppose four packets have arrived at the receiver with sample delay $r_4 - t_4$, $r_3 - t_3$, $r_2 - t_2$, and $r_1 - t_1$. Express the estimate of delay d in terms of the four samples.
   b) Generalize your formula for n sample delays.
   c) For the formulate in part b, let n approach infinity and give the resulting formula. Comment on why this averaging procedure is called an exponential moving average.
2) Given that a CDN does not increase the amount of link capacity in a network (assuming the CDN uses existing links to distribute its content among CDN nodes), how does a CDN improve the performance seen by hosts? Give an example.
3) See Fig. 1, which is similar to Figures 7.22–7.25. Answer the following questions.
   a) Assuming FIFO service, indicate the time at which packets 2 through 12 each leave the queue. For each packet, what is the delay between its arrival and the beginning of the slot in which it is transmitted? What is the average of this delay over all 12 packets.
   b) Now assume a priority service, and assume that odd-numbered packets are high priority, and even-numbered packets are low priority. Indicate the time at which packets 2 through 12 each leave the queue. For each packet, what is the delay between its arrival and the beginning of the slot in which it is transmitted? What is the average of this delay over all 12 packets?

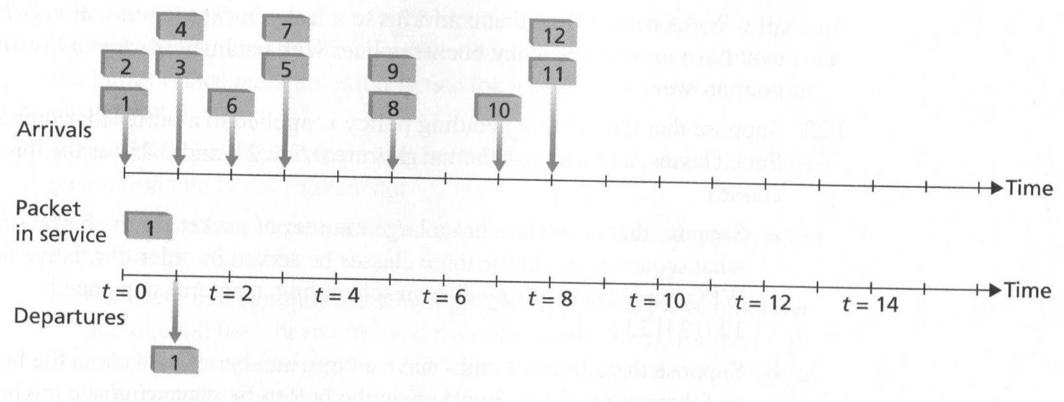

Fig. 1. Figure for Problem 3

c) Now assume round robin service. Assume that packets 1, 2, 3, 6, 11, and 12 are from class 1, and packets 4, 5, 7, 8, 9, and 10 are from class 2. Indicate the time at which packets 2 through 12 each leave the queue. For each packet, what is the delay between its arrival and its departure? What is the average delay over all 12 packets?

d) Now assume weighted fair queueing (WFQ) service. Assume that odd-numbered packets are from class 1, and even-numbered packets are from class 2. Class 1 has a WFQ weight of 2, while class 2 has a WFQ weight of 1. Note that it may not be possible to achieve an idealized WFQ schedule as described in the text, so indicate why you have chosen the particular packet to go into service at each time slot. For each packet what is the delay between its arrival and its departure? What is the average delay over all 12 packets?

e) What do you notice about the average delay in all four cases: FCFS, RR, priority, and WFQ?

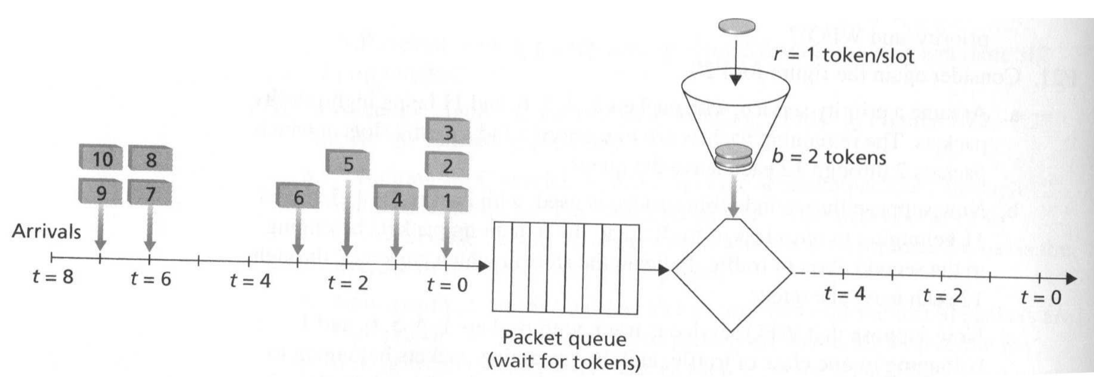

Fig. 2. Figure for Problem 4

4) See Fig. 2, which shows a leaky bucket policer being fed by a stream of packets. The token buffer can hold at most two tokens, and is initially full at $t = 0$. New tokens arrive at a rate of one token per slot. The output link speed is such that if two packets obtain tokens at the beginning of a time slot, they can both go to the output link in the same slot. The timing details of the system are as follows.

   a) Packets (if any) arrive at the beginning of the slot. Thus in the figure, packets 1, 2, and 3 arrive in slot 0. If there are already packets in the queue, then the arriving packets join the end of the queue. Packets proceed towards the front of the queue in a FIFO manner.
   b) After the arrivals have been added to the queue, if there are any queued packets, one or two of those packets (depending on the number of available tokens) will each remove a token from the token buffer and go to the output link during that slot. Thus packets 1 and 2 each remove a token from the buffer (since there are initially two token) and go to the output link during slot 0.
   c) A new token is added to the token buffer if it is not full, since the token generation rate is $r = 1$ token/slot.
   d) Time then advance to the next time slot and these steps repeat.

   Answer the following questions:

   a) For each time slot indicate which packets appear on the output after the token(s) have been remove from the queue. Thus, for the $t = 0$ time slot in the example above, packets 1 and 2 appear on the output link from the leaky buffer during slot 0.
   b) For each time slot, identify the packets that are in the queue and the number of tokens in the bucket, immediately after the arrivals have been processed (step a above) but before any of the packets have passed through the queue and removed a token. Thus for the $t = 0$ time slot in the example above, packets 1, 2, and 3 are in the queue, and there are two token in the buffer.

5) Consider the leaky bucket policer that policies the average rate and burst size of a packet flow. We now want to police the peak rate p as well. Show how the output of this leaky bucket policer can be fed into a second leaky bucket policer so that the two leaky buckets in series police the average rate, peak rate, and burst size. Be sure to give the bucket size and token generation rate for both leaky buckets.