The University of Queensland
School of Information Technology and Electrical Engineering

Semester Two, 2011

COMS3200 – Tutorial 11 - Solutions

Question 1

It is possible to define flows on either a host-to-host basis or on process-to-process basis.

a. Discuss the implications of each approach to application programs.

b.IPv6 includes a FlowLabel field for supplying hints to routers about individual flows. The originating host is to put here a pseudorandom hash of all the other fields serving to identify the flow; the router can thus use any subset of these bits as a hash value for fast lookup of the flow. What exactly should the FlowLabel be based on for each of these two approaches?

a) From the application perspective, it is better to define flows as process-to-process. If a flow is host-to-host, then an application running on a multi-user machine may be penalised (by having its packets dropped) if another application is heavily using the same flow. However, it is much easier to keep track of host-to-host flows; routers need only look at the IP address to identify the flow. If flows are process-to-process (i.e. end-to-end), routers must also extract the TCP or UDP ports that identify the endpoints. In effect, routers have to do the same demultiplexing that is done on the receiver to match messages with their flows.

b) If flows are defined on a host-to-host basis then *FlowLabel* would be a hash of the host-specific information; that is the IP addresses. If flows are process-to-process then the port numbers should be included in the hash input.

Question 2

The transmission schedule (Table 1) for a given flow lists for each second the number of packets sent between that time and the following second. The flow must stay within the bounds of a token bucket filter. What bucket depth does the flow need for the following token rates? Assume the bucket is initially full.

a. 2 packets per second

b. 4 packets per second

a) If we start with an empty bucket but allow the bucket volume to become negative (while still providing packets), we get the following table of "indebtedness": at T=0, for example, we withdrow 5 packets and deposit 2.

Table 1:

Time, secs	0	1	2	3	4	5
Bucket volume	-3	-6	-5	-3	-7	-6

We therefore need an initial bucket depth of 7, in order not to run out of tokens at T=4. The bucket with depth 7 will not overflow.

b) if we do the same thing as above we get

Table 2:

Time, secs	0	1	2	3	4	5

Table 2:

Bucket volume	-1	-2	1	5	3	6

A bucket depth of 2 will therefore accommodate T=1. If we start with an initially full bucket of depth 2, we get

Table 3:

Time, secs	0	1	2	3	4	5
Bucket volume	1	0	*2*	*2*	0	*2*

Entries in italic represent filing the bucket and therefore losing the overflowing tokens. It turns out that due to this truncation we get to the bottom at T=4 as well. However the depth of 2 does suffice.

Question 3

Suppose a router has accepted flows with the TSpecs shown in Table 2, described in terms of token bucket filters with token rate r packets per second and bucket depth B packets. All flows are in the same direction and the router can forward one packet every 0.1 second.

a. What is the maximum delay a packet might face?

b. What is the minimum number of packets from the third flow that the router would send over 2.0 seconds, assuming that the flow sent packets at its maximum rate uniformly?

a) If the router queue is empty and all three flows dump their buckets at the same time, the burst amounts to 15 packets for a maximum delay of 1.5 sec. Since the router can keep up with packets due to steady-state traffic alone, and can drain any earlier bucket dumps faster than the buckets get refilled, such a burst is in fact the maximum queue.

b) In 2.0 seconds the router can forward 20 packets. If flow 1 sends an initial burst of 10 at T=0 and another single packet at T=1, and flow 2 sends 4 at T=0 and 2 at T=1, that amounts to 17 packets in all. This leaves a minimum capacity of 3 packets for flow3. Over the long term, of course, flow 3 is guaranteed an average of 8 packets per 2.0 seconds.

Question 4

Suppose an RSVP router suddenly loses its reservation state but otherwise remains running.

a. What will happen to the existing reserved flows if the router handles reserved and nonreserved flows via a single FIFO queue?

b. Eventually the receivers on these flows will request that their reservations be renewed. Give a scenario in which these requests are denied.

a) If the router was initially combining both reserved and nonreserved traffic into a single FIFO queue, then reserved flows were not getting genuine service guarantees before the loss. After the loss the router is still handling all traffic via a single FIFO queue; the only difference is that all traffic is now considered nonreserved. The state loss should therefore make no difference.

b) Suppose new reservations from some third parties reach the router before the periodic refresh requests are received to renew the original reservations; if these new reservations use up all the reservable capacity the router may be forced to turn down the renewals.

Question 5

Which ATM traffic classes can cause congestion in switches? Which traffic classes are capable of creating short term congestion and which could create longer lasting congestion?

All traffic classes except for Constant Bit Rate (CBR) may create congestion in ATM switches as data rate for these classes is not static. Congestions created by VBR traffic will, however, be short leaved (short term congestion) because the congestion will be created if several flows send traffic bursts at the same time. When the bursts are finished (and their length is described in the flow specfication) the congestion will ease. Short time congestion can be solved by buffering cells in switches. The ABR traffic class which has less predictable behaviour may create a long term congestion and therefore special methods for congestion avoidance are necessary.

Question 6

Give an argument why a leaky bucket algorithm should allow just one packet per tick, independent of how large the packet is.

The router has to do approximately the same amount of work queueing a packet, no matter how big it is. There is little doubt that processing 10 packets of 100 bytes each is much more work than processing one packet of 1000 bytes.

Question 7

Consider the user of differentiated services with expedited forwarding. Is there a guarantee that expedited packets experience a shorter delay than regular packets? Why or why not?

There is no guarantee. If two many packets are expedited, their channel may have even worse performance than the regular channel. To provide some guarantees, the network would need to provide admission control for expedited traffic at the edges of the network.