CSE 473S, Fall 2008

Solution for Assignment 1

Task 1. The logical topologies of the following paths were established by identifying the longest prefix match for the destination IP address of datagrams in each port starting from the port attached to the source.

Path from *n2p1* to *n1p1*:

n2p1, link (1 Gbps), *NPR2* (ports 1-2), link 2-4 (1 Mbps; actual: 0.683 Mbps), *NPR1* (ports 4-1), link (1.5 Mbps; actual: 1.366 Mbps; buffer size: 100 KB), *n1p1*.

Path from *n2p4 to n1p1*:

n2p4, link (1 Gbps), NPR2 (ports 4-3), link (100 Mbps; actual: 99.718 Mbps), NPR1 (ports 3-1), link (1.5 Mbps; actual 1.366 Mbps; buffer size: 100 KB), *n1p1*.

Path from *n1p1* to *n2p1*:

n1p1, link (1 Gbps), *NPR1* (ports 1-4), link 4-2 (1Gbps), NPR2 (ports 2-1), link (1 Gbps), *n2p1*.

Path from *n1p1* to *n2p4*:

n1p1, link (1 Gbps), NPR1 (ports 1-4), link 4-2 (1 Gbps), NPR2 (ports 2-4), link (1 Gbps), *n2p4*.

Task 2. When the paths are empty, *ping* reports small RTTs around 0.09 ms. According to *iperf* reports, *n1p1* receives X_{exp} = 2.4 MB from *n2p1* and Z_{exp} = 2.39 MB from *n2p4*.

The screen shot for the bitrate of traffic coming out of port 1 of NPR1 shows that (a) the bitrate rises from 0 to about 0.7 Mbps when *n2p1* starts transmitting, (b) the bitrate rises from about 0.7 Mbps to about 1.4 Mbps when *n2p4* starts transmitting, (c) the bitrate decreases from 1.4 Mbps to about 0.8 Mbps when n2p1 stops transmitting, (d) the bitrate decreases from about 0.8 Mbps to 0 when *n2p4* stops transmitting.

The screen shot for the queue size at link coming out of port 1 of NPR1 shows that (a) the queue size grows from 0 to about 100,000 bytes when *n2p4* starts transmitting; the growth lasts for about G_{exp} = 6 seconds at rate g_{exp}; (b) the queue drains from about 100,000 bytes to 0 when *n2p1 s*tops transmitting; the drainage continues for around D_{exp} =1.5 seconds at rate d_{exp}.

Task 3. When *n2p1* starts transmitting at the rate 2 Mbps at time 0, the transmission is bottlenecked at link 2-4 that has capacity C_{24} = 0.683 Mbps. Traffic beyond this rate is queued and then discarded at port 2. Monitoring bitrate of link 2-4 confirms this. The bitrate on link 3-3 stabilizes at R = 0.8 Mbps. In fact, when we change the forwarding table to route all traffic with destination address that matches 192.168.1.32/28 to port 2, we no longer see any traffic on link 3-3, and link 2-4 has become the new bottleneck. However, for our analysis below, we will use the original topology without this modification in forwarding table.

When *n2p4* starts transmitting at rate R = 0.8 Mbps after delay t = 20 seconds, the combined load of *n2p1* and *n2p4* is $C_{24} + R$ and exceeds the capacity of link going out of port 1 of NPR1, C. Hence, while the bitrate on link going out of port 1 of NPR1 stabilizes at the link capacity C = 1.366 Mbps, the queue at link going out of port 1 of NPR1 starts growing at rate $g_{an} = C_{24} + R - C = 117$ Kbps and fills up the buffer of size B = 100,000 bytes completely after delay $G_{an} = B / (C_{24} + R - C) = 6.84$ seconds.

When *n2p1* stops transmitting at time T = 40 seconds, the transmission rate of *n2p4* alone is too small to saturate the link going out of port 1 of NPR1. Hence, the queue at link going out of port 1 of NPR1 starts draining at rate $d_{an} = C – R = 566$ Kbps and empties the buffer of size B = 100,000 bytes completely after delay $D_{an} = B / (C – R) = 1.41$ seconds. Then, the queue size at link going out of port 1 of NPR1 stays near 0, and the bitrate on link going out of port 1 of NPR1 stabilizes at R = 0.8 Mbps.

When *n2p4* stops transmitting at time T + t, the bitrate at link going out of port 1 of NPR1 goes to 0.

To predict X_{an} and Z_{an}, the overall amount of data received by *n1p1* from *n2p1* and *n2p4* respectively, one needs to make an assumption how losses are distributed between the senders when their combined transmission overloads link going out of port 1 of NPR1. We assume that each datagram has the same probability to be discarded. Then, the loss rate on the overloaded link going out of port 1 of NPR1 for both senders is $L = (C_{24} + R - C) / (C_{24} + R) = 8\%$.

Under the above loss distribution assumption, we compute:

$X_{an} = C_{24} * t + (1 - L) * C_{24} * (T – t) = 3.26$ MB

$Z_{an} = (1 - L) * R * (T – t) + R * t = 3.84$ MB

Several factors contribute to the differences between the experimental and analytical results.

First, the *iperf* rates are for application data. With headers added by lower layers, the sources inject traffic into their network paths at higher bitrates. In particular, R is higher than 0.8 Mbps. This explains why $g_{exp} > g_{an} = C_{24} + R - C$ and $d_{exp} < d_{an} = C – R$ (or equivalently why $G_{exp} < G_{an}$ and $D_{exp} > D_{an}$). Similarly, the header overhead increases the loss rates and thereby results in lower amount of delivered data: $X_{exp} < X_{an}$ and $Z_{exp} < Z_{an}$.

Second, delay t before the transmission from *n2p4* starts is not exactly 20 seconds in the experiment. The longer the transmissions from *n2p1* and *n2p4* overlap, the smaller become the amounts of data received from the senders by *n1p1*.

Third, the distribution of losses between the sources could be different from the assumed in the analysis. This would change the distribution of data delivered from the sources.

Fourth, the analysis does not account for data buffered at the overloaded links. The buffered data is delivered after excessive transmission ceases. Accounting for the buffered data would increase the overall amounts of delivered data.