

CSE 473S – Introduction to Computer Networks Fall 2009

Lab Assignment 1

1 Preliminaries

Before proceeding with the assignment, make sure that you have all required background knowledge and tools:

1. Attend the ONL class of the course.
2. Begin reading the *NPR Tutorial* (see the sidebar at www.onl.wustl.edu). Here is a short guide to the sections in the tutorial: *Scan Packet Processing*; *Read The Remote Laboratory Interface*; and *Begin to scan Filters, Queues and Bandwidth*.
3. Request your personal ONL account.
4. After receiving the account, log in and follow the steps on the web page accessible through the link *Getting Started*.

2 Pre-run Analysis

Acquire an ONL configuration file from www.cse.wustl.edu/~jain/cse473-09/ftp/assignment_1.exp. Create local directory `~/onl_dir/` and put this file into this directory. The configuration file sets up a topology enabling two senders at hosts `n2p1` and `n2p4` to communicate with their receivers at host `n1p1` as shown in Figure 1.




Figure 1: Configuration for Lab 1

Start the RLI and open the configuration file. Note that, the RLI requires Java 1.6 (actually JRE 1.6) or above to be installed in your computer. Three windows are supposed to appear and contain:

- Configured network topology
- Display monitoring the bitrate of traffic going out of port 1 of *NPR1*
- Display monitoring the queue size at output port 1 of *NPR1*

2.1) Find out the size of this buffer by clicking on port 1, opening the Queue Table menu, and looking up the value in the *threshold(bytes)* entry for *queue id* 300 in the lower *Queues* table. Report the size of the buffer.

2.2) Use the RLI to examine the forwarding tables at the ports of *NPR1* and *NPR2*, and determine the logical topology for the configured packet paths (i.e., sequences of ingress and egress ports) from *n2p1* and *n2p4* to *n1p1* and back from *n1p1* to the senders. To find the forwarding table of a port, select the port in the RLI and proceed into the *Route Table* menu. Report the sequences of ingress and egress ports for the 4 paths.

2.3) Determine the capacity for each of the links on these four paths by selecting the corresponding port, opening the *Queue Table* menu, and looking up the value in the *Port Rate* entry. Assume that the links from *n2p1* and *n2p4* to *NPR1* are 1 Gbps. Report the capacities.

3 Run-time ONL Experiment

Before you can make an ONL reservation or run an experiment, you must build an SSH tunnel to the testbed. The *Getting Started* page gives instructions on how to do this. You can use a Unix command line or SSH tool (e.g., PuTTY). If you are unsuccessful after a few attempts, see your course TA for help. Note: You do NOT need to build the tunnel if you only plan to read or change the configuration file.

You must make a reservation for the hardware resources for a time period when you plan to run an ONL experiment. You can make the reservation days (or even weeks) in advance. You do NOT need to make a reservation if all you are doing is constructing or examining an experiment (configuration) file.

To make a reservation for this assignment, start the RLI and open the configuration file *assignment_1.exp*. Then, select *File => Make Reservation* from the *File* menu and respond to the prompt by requesting a reservation for 60 (sixty) minutes. The RLI will respond with a message about the reservation. You can also check the reservation through the ONL web page by logging into the ONL website and selecting *My Reservations* in the sidebar.

When your reservation time comes, reestablish the SSH tunnel and run the RLI. Open the configuration file *assignment_1.exp*. Request the configured topology by selecting *Commit* in the *File* menu. It might take several minutes for ONL to commit the requested resources. The RLI provides a Progress Bar that shows the progress of the *Committing* process and ends with the message *Commit Completed*. Once the topology is realized (e.g., links turn from dashed lines to solid ones, and hosts turn from light to dark blue), check how the resource commitment affected the link capacities; i.e., see what actual link capacities that the links will have in your experiment.

ONL hosts have two network interfaces: control network and data network. Both interfaces have an IP address and a name. A control network interface has a name of the form *onlXXX* (the full name is *onlXXX.arl.wustl.edu*) where *XXX* is some integer, and an IP address of the form 10.0.A.B. A data network interface has a name of the form *nXpY* where *X* is the NPR number and *Y* is the port number, and an IP address of the form 192.168.X.Y where *X* is the NPR number and *Y* is 16 times one more than the port number.

Click on the hosts *n2p1*, *n2p4* and *n1p1*, and determine the control network name and IP address and the data network IP address of each host. Log into *n1p1* through its *onlXXX* address (e.g., `onlusr> ssh onl070`) and execute the command `iperf -s -u` to start the UDP receiver. Similarly log into *n2p1* and *n2p4* through their *onlXXX* addresses and start the UDP senders: first, execute the command `iperf -c n1p1 -u -b 2m -t 40` at *n2p1*; about 20 seconds later, execute the command `iperf -c n1p1 -u -b 0.8m -t 40` at *n2p4*.

The first sender transmits at rate 2 Mbps, and the second sender transmits at rate 0.8 Mbps. Each of the transmissions lasts for 40 seconds. After the senders stop transmitting, execute the command `ping -c 10 n1p1` at both *n2p1* and *n2p4* to measure round-trip times (RTTs) for small packets when the paths

serve no other traffic. The RTTs include all propagation and processing delays of the paths.

Monitor the bitrate of traffic going to input port 4 of *NPR1* by opening the Monitoring menu, selecting Add Monitoring Display, clicking port 4, proceeding to the RXBYTE submenu, and choosing rate.

3.1) How much traffic do you see on link 2-4 during your experiment? What causes this traffic? Is your answer consistent with the logical network topology that you inferred before?

3.2) Change the forwarding table of port 4 of *NPR2* by adding an entry that suggests forwarding to port 2 all traffic with addresses that match prefix 192.168.1.32/28. Does the change affect the outcome of the above experiment? How and why?

3.3) Include two screen shots covering the duration of your experiment: one for the bitrate of traffic going out of port 1 of *NPR1* and the other for the queue size at output port 1 of *NPR1*. Report the *ping* measurements of the empty-path RTTs as well as the overall amount of data received by *n1p1* from each sender (as reported by *iperf*).

4 Post-run Analysis

4.1) Analyze the experiment in the original configuration and derive expressions for the overall amount of data received by *n1p1* from each sender.

4.2) Analyze queuing at output port 1 of *NPR1* during the one-minute experiment. In your analysis, assume that each traffic flow on an overloaded link loses the same fraction of traffic. When and why does the queue start to grow? When and at what size does the queue stabilize? When and why does the queue start to drain?

4.3) Compute rates for the queue growth and queue drain and explain why the slopes are different.

5 Report

Prepare a report covering tasks 2.1 through 2.3, 3.1 through 3.3, and 4.1 through 4.3. A hard copy of the report must be submitted to the instructor on the following Monday before the class starts.