


## Project 1 – TCP Throughput Measurements

Assigned: Sept 10, 2011

Due: Sep 21, 2011 11:59pm



Topology for Project 1

For this project, we will use the *ns-3* simulator to measure the throughput of a single TCP connection through a bottleneck link, as a function of queue limits, segment sizes, and window sizes.

1. Log into the *jinx* cluster using the instructions found on the class web page.
2. Copy the *ns-3* simulator to your home directory (be sure to notice the period at the end of the command below):
 

```
cp /nethome/ECE6110/ns-3-ece-6110-gatech-09092011.tgz .
```
3. Compile the *ns-3* simulator following the instruction in the *ns-3* tutorial posted on the class web page.
4. Once *ns-3* has been properly compiled, create a simulation model of the topology shown in the figure above. It would be easiest if you called your new program `p1.cc` and put it in the `scratch/` directory.
5. To test your program, enter: `./waf --run scratch/p1`. Note, if you wish to use command line arguments with your program, you should wrap the run parameters in quotes: `./waf --run 'scratch/p1 --segSize=512 --queueSize=64000'`, as an example.
6. Use the TCP-Tahoe variation of TCP for this step.
7. Create a single TCP flow that connects to a server and sends 100,000,000 bytes. It might be easiest to use the `BulkSendApplication` and the `PacketSinkApplication` for this.
8. In the simulation, measure the *goodput* of a single TCP connection, using combinations of maximum receiver advertised window sizes ( 2000, 8000, 32000, and 64000 bytes), queue limits on the bottleneck link (2000, 8000, 32000, and 64000 bytes), and TCP Segment sizes (128, 256, and 512 bytes). This results in a total of 48 different combinations and 48 simulation runs. Here, the term *goodput* refers to the amount of data received at the packet sink application divided by the simulation time (resulting a value in units of bytes/sec).
9. Repeat the experiments, except use 10 *simultaneous* flows all sharing the same links. Each of the 10 flows should have the same advertised window size setting. Observe the relative fairness of the flows and comment on it in your report. Each of the 10 flows should start at a random time between 0 and 0.1 seconds.
10. Repeat both of the above experiments using TCP Reno instead of TCP Tahoe.

11. To run the experiments, use the *PBS-Torque* job scheduler that is installed on *jinx-login*. This is described in detail on the *jinx* web page linked from the class web page. You need 48 separate runs for each of the four experiment types (single flow and multiple flows, TCP Reno, TCP Reno) to measure the throughput under the conditions specified.
12. An example PBS scheduler script can be found at (and copied from) `/nethome/ECE6110/run-ns3`. This needs to be edited slightly to add your information.
13. Turn in your *ns-3* main program (via email to the instructor with cc: to the TA) used for the experiments, and all measured goodput results.
14. Turn in a two or three page report (hard copy and emailed soft copy) that describes your experiments and any conclusions you can draw.