

# Advanced Operating Systems:

## Lab 3 – TCP

### L41 Assignment

Prof. Robert N. M. Watson

2021-2022

Your lab report will analyse how network latency impacts TCP throughput, with a particular interest in its effects on congestion control. You will do this using the `tcp` socket mode of the IPC benchmark, using DUMMYNET to simulate various network latencies.

## 1 Approach

You will run a series of experiments using the IPC benchmark using the `tcp` IPC type, using DUMMYNET to simulate varying network latency. Configure the benchmark as follows:

- Use `-i tcp` to select the TCP benchmark mode
- Use `2proc` mode (as described in Lab 2)
- Hold the total I/O size (16M) constant
- Use verbose mode to report additional benchmark configuration data (`-v`)
- Use JSON machine-readable output mode (`-j`)
- Collect `getrusage()` information (`-g`)

## 2 Experimental questions: Latency and TCP bandwidth

Explore the following experimental questions, which consider only the TCP steady state (`ESTABLISHED`), and not the three-way handshake or connection close. For both questions, use a fixed 1MiB buffer (`-b 1048576`):

1. Explore how varying latency affects TCP bandwidth, considering two cases: one in which the socket-buffer is set explicitly (`-s`), and a second in which it auto-resizes. Plot latency (0ms, 5ms, .. 40ms in 5ms intervals) on the X axis, and effective bandwidth on the Y axis.

Characterize the two performance behaviours. Explain why socket-buffer auto-resizing helps, hurts, or fails to affect performance as latency varies, in comparison to the fixed buffer size.

2. Explore the behaviours of specific TCP connections, plotting time on the X axis and effective bandwidth on the Y axis. Consider four cases: two latencies (0ms and 30ms), each with and without automatic resizing of socket buffers.

Stack additional graphs showing the sender last received advertised window and congestion window on the same X axis, and use `snd_ssthresh` to label portions of the plot as being in slow start. For readability reasons, you may wish to generate two independent sets of plots, one for each latency, rather than overload a single set.

Characterise the performance behaviours. Document their differences, and explain how each of the two buffering strategies affect performance.

How do the results of these experiments contribute to supporting or rejecting our hypotheses? Be sure, in your lab report, to describe any apparent simulation or probe effects.

### **3 Plotting TCP time-bandwidth plots**

TCP time-bandwidth graphs plot time on a linear X axis, and bandwidth achieved by TCP on a linear or log Y axis. Bandwidth may be usefully calculated as the change in sequence number (i.e., bytes) over a window of time – e.g., a second. Care should be taken to handle wrapping in the 32-bit sequence space; for shorter measurements this might be accomplished by dropping traces from experimental runs in which sequence numbers wrap.

This graph type may benefit from overlaying of additional time-based data, such as specific annotation of trace events from the congestion-control implementation, such as packet-loss detection or a transition out of slow start. Rather than directly overlaying, which can be visually confusing, a better option may be to “stack” the graphs: place them on the same X axis (time), horizontally aligned but vertically stacked. Possible additional data points (and Y axes) might include advertised and congestion-window sizes in bytes.

### **4 Notes**

Graphs and tables should be used to illustrate your measurement results. Ensure that, for each question, you present not only results, but also a causal explanation of those results – i.e., why the behaviour in question occurs, not just that it does. For the purposes of performance graphs in this assignment, use achieved bandwidth, rather than total execution time, for the Y axis, in order to allow you to more directly visualise the effects of configuration changes on efficiency.