Advanced Operating Systems:
Lab 3 – TCP
L41 Assignment

Prof. Robert N. M. Watson

2022-2023

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 experience using the IPC benchmark using the `tcp` IPC type, using DUMMYNET to simulate varying network latency. Use DTrace and other tooling to explain the effects you see, with respect to the assignment hypothesis.

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. Plot latency (0ms, 5ms, .. 40ms in 5ms intervals) on the X axis, and effective bandwidth on the Y axis.
2. Measure each of the following dependent variable distributions for each latency point, and plot them using stacked plots aligned to the same X axis as the bandwidth plot:
   - Time from start of connection to its first exiting slow start, determined using `snd_ssthresh` (sender)
   - Last received advertised window and congestion window when slow start was first exited (sender)
   - Overall proportions of time spent in kernel, idle, and user time across the lifetime of a benchmark iteration (sender)
3. Using this data, and other data you may gather using DTrace, to explain why TCP performance varies as it does as latency varies. You might, for example, want to explore the performance of individual sample TCP connections at two different latencies (e.g., 0ms and 30ms) to better understand the effects you see.
4. Confirm or reject the assignment hypothesis, and explain why.

Note that connections may exit slow start more than once; be sure that you are capturing timing and other information only for the first exit.

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.