Lab2: Measuring Network Performance

COM-208: Computer Networks

The goal of this lab is to understand **the meaning of throughput and delay**, as well as **the nature of different delay components**. For this, you will use `ping` (which you discovered in the last lab) to **experimentally estimate** throughput and delay between your computer and various targets around the world. You will see that there are quantities that we cannot directly measure, yet we can reason about them if we understand what they are.

Setup

Download from Moodle `runping.sh` and `plot.sh`. If they are not already executable, make them, e.g., by running `chmod u+x runping.sh` and `chmod u+x plot.sh`. If you are doing the lab on your own computer, make sure you have the `gnuplot` application installed.

`runping.sh` is a shell script that takes as an argument a target (DNS name or IP address) and pings the target many times with different packet sizes. For example, `./runping.sh www.epfl.ch` executes

```
ping -s 22 -c 50 -i 1 www.epfl.ch > www.epfl.ch-p50
...
ping -s 1472 -c 50 -i 1 www.epfl.ch > www.epfl.ch-p1500
```

The first line pings www.epfl.ch 50 times, once per second, using 50-byte packets, and writes the output to a file called `www.epfl.ch-p50`. The last line pings www.epfl.ch 50 times, once per second, using 1500-byte packets, and writes the output to a file called `www.epfl.ch-p1500`.

`plot.sh` takes as an argument a list of file names (that were produced by `runping.sh`) and produces three new files: (1) `destination_delay.png` shows delay as a function of time, for different packet sizes (different colors correspond to different packet sizes). (2) `destination_scatter.png` shows delay as a function of packet size as a scatter plot. (3) `destination_avg.txt` contains the average (2nd column) and minimum (3rd column)

delay values for each packet size (1st column). For example, `./plot.sh www.epfl.ch-p*` produces the above three files for the target www.epfl.ch.

Measure and process

Run `runping.sh` with target www.epfl.ch, then with target www.yahoo.com, and then with target www.gov.kg. Given that you are sending one ping every second, these will take several seconds to complete.

When they finish, process the results by running:

```
./plot.sh www.epfl.ch-p*
./plot.sh www.yahoo.com-p*
./plot.sh www.gov.kg-p*
```

Understand your tool

We said earlier that `ping -s 22...` pings the target with 50-byte packets, while `ping -s 1472...` pings the target with 1500-byte packets.

- If the command specifies `-s 22` (or `-s 1472`), why might the resulting pings use 50-byte (or 1500-byte) packets, and not 22-byte (or 1472-byte) packets? To answer, start by reading what the `-s` parameter exactly means in the `ping` manual.
- What exactly does it mean that the target is pinged with "50-byte packets" or with "1500-byte packets"? Which part of the ping packet is 50 bytes? The application-layer message (header + data)? The entire packet (with transport-layer and network-layer and link-layer headers)? You can use Wireshark to answer, though that's not the only way.

Think about delay

A packet experiences different kinds of delay:

- The **transmission delay** experienced by a **packet on a link** is the amount of time it takes to push the bits of the packet onto the link, and it is equal to the packet length divided by the link transmission rate.
- The **propagation delay of a link** is the amount of time it takes for one bit to go from one end of the link to the other, and it is equal to the link length divided by the link propagation speed.

- The **queuing delay** experienced by a **packet at a switch** is the amount of time the packet sits inside a queue at the switch, waiting for other packets to be processed and transmitted. Queuing delay depends on the **other traffic** that traverses the same links and shares the same queues as the packet.
- The **processing delay** experienced by a **packet at a switch** is the amount of time it takes for the switch to process the packet (after it removes it from the queue and before starting to transmit it). Processing delay depends on the switch's processing capabilities and is typically independent of packet size, at least in most of the scenarios we will discuss in this class.

Look at the figures you have produced for each target and answer the following questions:

- Does delay change significantly over time? Whether you answer yes or no, explain why you think this is happening. E.g., if you answered that it does not change, consider each delay component and explain why it doesn't change over time. If you answered that it does change, which delay component(s) are the ones that change?
- Does delay change significantly with the size of the ping packets? Do bigger packets experience more or less delay?
- Is the delay-packet size dependence the same for all targets? Whether you answer yes or no, explain why you think this is happening.
- Suppose that, for a given target, there is no clear dependence pattern between delay and packet size, i.e., packet size does not affect delay significantly. What does that say about the path from your computer to that target?

Estimate delay components

Now we will explore the kind of "rough" estimations that network engineers often need to do.

Ping tells you the overall delay experienced by different packets of different sizes, but it does not (directly) tell you anything about individual delay components, like transmission delay or queuing delay. This is not surprising, because ping operates at the end-systems, it has no access inside the network to see, e.g., how long it takes to transmit each packet on each link, or how long each packet has to wait in a queue at each switch.

Still, if we know the overall delay experienced by *many* different packets and of *many* different sizes, then we *can* infer information about individual delay components.

Study the content of the file `www.yahoo.com_avg.txt`. Estimate the following components of the delay between your computer and www.yahoo.com:

- The propagation+processing delay.

- The transmission delay.
- The average queuing delay.

You don't have enough information to do any precise estimation. Do the best you can given the information you have. Think about the particular nature of each delay component: Should it be the same for all packets, independently of their sizes and when each packet was sent? Should it depend on any particular way on packet size?

Now study the content of the file `www.gov.kg_avg.txt` and estimate the same delay components between your computer and www.gov.kg. Observe the differences in the results relative to www.yahoo.com. What might explain these differences?

Estimate bottleneck throughput

The **bottleneck** link between a source and a destination is the link where data flows at the slowest rate. Estimate the **throughput** achieved by the bottleneck link between your computer and www.yahoo.com, *while transmitting a ping message*. Do the same estimation for target www.gov.kg.

Be careful: we are not asking for the *average* throughput achieved by the ping application (which is the total amount of ping-message data transferred from source to destination divided by the total transfer time). Rather, we are asking for the *instantaneous* throughput achieved *at the bottleneck link* while transmitting a ping message. To answer, you need to consider one particular delay component that you estimated above.