

**Experiment 1: “The basics – bandwidth and RTTs”**

The maximum speed at which bits can be sent is often known as the capacity (or “bandwidth”) and is often given in Mbps (Mega bits per second). One Byte consists of 8 bits.

If you share the Internet connection with other people, it’s your share of the capacity that’s being measured.

How fast is your connection to the Internet?  
Go to <http://www.speedtest.net/>  
Run the test and make a note of your upload and download speeds.

To find out how long it takes to send a piece of data, you need to multiply the size of the data by the speed:

**Data size \* speed = time to complete (in seconds)**

Consider a 3MB mobile phone photo: How much time will it take for you to upload this image over the capacity you measured?

The time it takes for a data packet to travel from your computer to an Internet server and back is called the “round trip-time” (or just “RTT”). It can be measured using a tool called “ping” that is present on most regular computers.

Measure the RTT between yourself and an Internet server using “ping”.

Open a terminal window:

Windows: <Windows key> then enter “cmd”

Mac: <Spotlight> then enter “Terminal”

Linux: you know how to open a terminal window :)

Type “ping nytimes.com”

The field showing “time=” displays the RTT in milliseconds (most times ping makes more than one measurement, and reports the results for each measurement it makes).

**Experiment 2: “The importance of RTTs”**

To get from one place in the Internet to another requires the signal to be sent along communications links (optical fibre, copper wire, wireless, satellite links, etc). Each link is connected by routers, sometimes known as a “hop”. Usually the data travels at (or very near to) the speed of light along each of the links.

In most cases though the data needs to be sent over multiple links using multiple routers. If data arrives at a router from a link and can’t be sent immediately on the next link, it is buffered (stored inside the router for a short period of time). This delays the data on its journey and accounts for the fact that you do not receive a response with “the speed of light”.

Go to <http://www.wolframalpha.com> and enter “distance between <your location> and New York”.

The section about “Direct travel times” will give you the direct line travel time for light in a fiber.

Open a terminal window on your computer and enter “ping metopera.org” The field labeled “time=” gives you the time it takes a data packet to travel to the server of the Metropolitan Opera in New York and back.

How does the “ping time” compare to the theoretical direct flight time?

Go to the following page:  
<https://wondernetwork.com/pings>

Search the list for a city near you. Press the city name in the list and you’ll see the ping time between your location and other places around the world.

Find a city in each continent, Africa, Asia, Australia, Europe, North America and South America, and note the RTT time.

Computers in the same continent often have ping times in the same range. What is the difference between the ping times you noted? Which continent has the highest ping from your current location? Why?

Sometimes we send a single big piece of data (such as a digital photograph). Each piece of data is sent across the Internet using a messages, each about 1500B – A 3 MB file therefore is sent using about 2000 messages.

However, we often use the Internet for a wide variety of uses, not just to send lots of data. For instance, most web pages consist not of one large file, but a collection of many smaller objects (pictures, text, buttons, etc). When we go to a web page, we need to see all these objects, and the network has to download each of these pieces of data – one after another, as a series of downloads.

Sometimes we send a single big piece of data (such as a digital photograph). Each piece of data is sent across the Internet using a messages, each about 1500B – A 3 MB file therefore is sent using about 2000 messages.

Let's assume that:

- 1) a webpage has 5 objects (e.g. some text and 4 pictures).
- 2) Each object is about 7500 bytes (sent as 5 data packets).
- 3) To transmit each object requires 4 messages passed back and forth (RTTs) across the Internet.

Using the highest and lowest of the ping times that you noted earlier, what would be the load time for this webpage? Let's assume you found a way of reducing the back-and-forth messages for each object so that the data can be sent in 3 messages instead of 5 messages. How much would you save in page load time for the two cases?

Recently, researchers have found ways of reducing the back and forth communication needed, effectively saving time for all of us, as described in this task, but there's still a lot of potential for improvements.

Your Internet router has memory for storing data that is waiting to be sent. This memory is called a "buffer". If the buffer is too big, it can create delay that adds to the time it takes to send a message, and therefore increases the measured ping time. This increase is more when your Internet connection is busy (and more messages need to be buffered). Internet engineers say the "latency" has been increased.

Go to the page <http://netalyzr.icsi.berkeley.edu/> or download the "Netalyzr" app on your computer/device. (You'll need Java installed)

Run the Netalyzr analysis for your network.

When the results show, scroll down to the part about "Network Access Link Properties". The "Network buffer measurements" section will tell you exactly how much time was added to the basic delay when your Internet connection becomes busy with traffic.

How much longer will it take to load the webpage when this extra delay is added?

Note: The capacity may be "too good" to measure the buffer on your home router if you have high-speed broadband.

Tip: You could try to use a mobile phone with Internet sharing to access the Internet. **Make sure you have an unlimited-use subscription if you do this.** Then redo the Netalyzr analysis.

There is a problem, sometimes routers need to have buffer memory to store messages for a short while because the network is busy at one moment in time, but if these buffers are used all the time they add very significant delay. Internet engineers are designing smart ways to control the way that routers use their buffers to control the additional delay that can be introduced. This is becoming known as "Active Queue Management".

## Experiment 3: “Webpages around the world”

Do the webpage latency test <http://tools.pingdom.com/fpt/>  
For instance using <http://www.metopera.org>

When the test ends, the statistics on the top of the page show how much data was downloaded. The diagram below shows the objects that were downloaded and how much time it took to receive them.

Calculate approximately how much time it would take to transfer this volume of data if it were instead sent as *one* large file using an Internet connection with a capacity of 2Mbps.

Was there a difference between the result you calculated and the actual load time from the test?  
If so, why?

You can choose the place from where the test is run using the “settings” menu, just below the box where you enter the URL of the website to use for measurements. Getting closer to the source will reduce the RTT for the test.

Do the same test from Amsterdam, Dallas, then New York.  
Note the difference in completion time.

Why does the difference in proximity account for so large difference in completion time?

Internet Engineers understand the benefits of being close to the web site. While it is not possible to move everyone to one place in the world (things would get very crowded with 8,000,000,000 people living near to the web site!), it is often possible to move the contents of web sites closer to people – popular web sites are now replicating their data to many places across the world, decreasing the latency, and improving the web experience.

Now you understand how Internet performance is impacted by delay you can:

- 1) Avoid excessive buffering in your home router (see fact sheet).
- 2) Appreciate the benefits of placing content near the consumer.
- 3) Know that its not “all about bandwidth”.

Researchers and industry are working to change the Internet so that both throughput and latency are served with excellent quality. Follow this quest on our webpages: [www.riteproject.eu](http://www.riteproject.eu)