Practical Sheet 7

Understanding the Internet

Aims

Section		Aim
1	Ping and latency	Understand the relationship between distance and delay
2	Traceroute	Explore the structure of the Internet a packets go from router to router.
3	Using DNS	Understand how names are converted to IP addresses.
4	Looking at packets	Get an idea of the variety of protocols used on the Internet.

Related topics

- **Topic 7.1** Principles of Communication
- **Topic 7.2** Internet Components

Preliminaries

This practical sheet uses some simple network tools to look at aspects of the network It is designed for use with the cloud2class.com Virtual Machine (VM) environment, but similar capabilities can be achieved on other systems, including other VMs and the Raspberry Pi. The tools used are:

Tool	cloud2class VM	Raspberry Pi	Windows?
ping	Available	Available	Available
traceroute	Available	Available	Should be available: name tracert (probably)
tcpdump	Available	Can be installed[1]	Not available easily
dig	Available	Can be installed	Not available easily

There are also a number of web-tools e.g. http://whois.domaintools.com/ In addition, to find information like IP addresses or Mac addresses, use:

- ifconfig (on Linux, Apple)
- ipconfig (on Windows probably)

You login to the VM using a web browser – please use Chrome or Firefox. You will be given a user name and password (both are the same for everyone).

WARNING. Unfortunately firewalls may mean that some of the practical steps described here do not work in all locations, no matter how the computer is set up. The cloud2class VM avoids these problems: provided your firewall allows access to cloud2class, everything else will work.

1 Ping and Latency

Ping is used to measure the round trip (there and back) latency from your machine to another. You can ping a URL (name)

```
ping www.eecs.qmul.ac.uk
```

or an IP address: `ping 138.37.95.211`

Here is an example:

```
light@node-10-6-1-198:~$ ping www.eecs.qmul.ac.uk
PING www.eecs.qmul.ac.uk (138.37.95.147) 56(84) bytes of data.
64 bytes from apricot.dcs.qmul.ac.uk (138.37.95.147): icmp_req=1 ttl=46 time=3.88 ms
64 bytes from apricot.dcs.qmul.ac.uk (138.37.95.147): icmp_req=2 ttl=46 time=3.41 ms
64 bytes from apricot.dcs.qmul.ac.uk (138.37.95.147): icmp_req=3 ttl=46 time=3.45 ms
64 bytes from apricot.dcs.qmul.ac.uk (138.37.95.147): icmp_req=4 ttl=46 time=3.31 ms
64 bytes from apricot.dcs.qmul.ac.uk (138.37.95.147): icmp_req=5 ttl=46 time=3.39 ms
64 bytes from apricot.dcs.qmul.ac.uk (138.37.95.147): icmp_req=6 ttl=46 time=3.53 ms
^C
--- www.eecs.qmul.ac.uk ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 5005ms
rtt min/avg/max/mdev = 3.311/3.497/3.886/0.200 ms
```

Exercise 1.1: Try the following:

- ping another VM (using the IP address)
- ping a webserver in UK
- ping a webserver in another country (universities often work)

Complete a table like this:

Domain Name	Location /Approx Distance (Km)	Response ?	Average Delay (ms)
www.bbc.co.uk	UK ? / <100 Km	Yes	
www.cmu.edu	East Coast USA/ 3000Km	Yes	
www.mit.edu			

Main points

- A server may be configured to ignore 'ping' messages.
- The latency (or delay) depends on distance (and how the busy the routers are).
- The delay varies: the distance sets a minimum but other delays are added.
- It is difficult to co-ordinate with another machine when the latency is high (relative to the bandwidth).

2 Traceroute

2.1 Using Traceroute

'traceroute' is a tool for determining the IP addresses of the network nodes (i.e. routers) between you are another host. In this section, you are to investigate its use. Traceroute is a command line tool.

- Traceroute can be used in two ways.

```
traceroute www.eecs.qmul.ac.uk
```

- The second version requires administrator privilege, so use 'sudo':

```
sudo traceroute -T www.eecs.qmul.ac.uk
```

Here is an example:

```
antelope@node-10-6-1-199:~$ sudo traceroute -T www.bbc.co.uk
traceroute to www.bbc.co.uk (212.58.246.94), 30 hops max, 60 byte packets
 1  10.6.1.193 (10.6.1.193)  1.296 ms  1.269 ms  1.233 ms
 2  rtr-peer-01.maid.v4.custdc.net (5.102.169.162)  2.286 ms  2.274 ms  2.257 ms
 3  rtr-23.maid.v4.custdc.net (109.74.255.173)  2.973 ms  2.963 ms  2.673 ms
 4  rtr-121.thn.v4.custdc.net (109.74.255.190)  6.352 ms  6.322 ms  6.312 ms
 5  rt-lonap-a.thdo.bbc.co.uk (5.57.80.90)  3.244 ms  3.191 ms  3.186 ms
 6  * * *
 7  ae0.er01.cwwtf.bbc.co.uk (132.185.254.93)  3.171 ms  2.831 ms  2.873 ms
 8  132.185.255.165 (132.185.255.165)  3.838 ms  3.785 ms  3.240 ms
 9  bbc-vip015.cwwtf.bbc.co.uk (212.58.246.94)  3.166 ms  3.595 ms  3.465 ms
```

Note that there is no information about hop 6.

Exercise 2.1: Try (*it does not always work*) tracing the route to the following domains:

www.eecs.qmul.ac.uk
www.qmul.ac.uk
qmplus.qmul.ac.uk
www.ucl.ac.uk

Count the number of hops and also answer the following questions:

1. What is the IP address of the QMUL EECS firewall (i.e. the router at the edge of EECS)?
2. Who runs the network that connects UK universities (such as UCL) to the Internet? Use google to find out about this organisation (an ISP for universities).
3. Who runs QMPlus (the VLE used at QMUL)?

2.2 Network Structure

Traceroute can be used to learn about the structure of a network (in terms of the number of routers and how they are linked).

Exercise 2.2: Select a University Computer Science department and find the URL of its web page (e.g. UCL, Imperial, Southampton, York, …) and also look for the URL of the same University's central web site. Use traceroute to find the path to the chosen web servers.

- List the machines (URL and IP address) on the route. For each node (i.e. router) find out what you can about the organisation that runs it.
- Is the Computer Science web server running on the same host as the University's central web server? If not, how does the route differ?

3 Using DNS

DNS is used all the time (including when you use traceroute above): whenever a URL is converted to an IP address.

We can use the dig tool to trace the delegation path to a domain:

```
> dig @f.root-servers.net <some domain> +trace
```

Exercise 3.1: Try this for an example such as www.nhs.net and answer the following questions about the output:

1. How many levels of delegation are there from the root server to the authoritative server for the domain you chose?

2. List the name servers that were used in the trace.
3. Is the trace the same each time it is run? (Can you think why not?)

Nowadays there is a lot of security information that adds to the complexity of this.

4 Looking at Packets

It is possible to look at the packet sent and received on an interface and to inspect their contents. Try the following command on the RPi (having installed tcpdump):

```
sudo tcpdump -v -c2 –ieth0
```

Here:

- '-v' is an option for verbose output
- -c2 means stop after capturing 2 packets; replace 2 with a larger number
- -ieth0 means capture on the interface eth0

However, you need to generate some traffic! The simple way to do this is to ask someone else to ping you. Here is an example output:

```
antelope@node-10-6-1-199:~$ sudo tcpdump -v -ieth0 -c2
tcpdump: listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
20:05:05.505902 IP (tos 0x0, ttl 64, id 17553, offset 0, flags [DF], proto ICMP (1), length 8 4)
    10.6.1.198 > 10.6.1.199: ICMP echo request, id 2090, seq 1, length 64
20:05:05.505951 IP (tos 0x0, ttl 64, id 30348, offset 0, flags [none], proto ICMP (1), length 84)
    10.6.1.199 > 10.6.1.198: ICMP echo reply, id 2090, seq 1, length 64
```

Exercise 4.1: Some other commands are shown below, together with a description.

Command	Explanation
`sudo tcpdump -v -c2 -ieth0 arp`	Capture 2 ARP packets: these are used for configuration.
`sudo tcpdump -v -c2 -ieth0 tcp`	Capture 2 TCP packets
`sudo tcpdump -v -c10 -A -ieth0 tcp port 80`	Capture 10 TCP packets used for HTTP (web page).
`sudo tcpdump -v -c5 -ieth0 udp port 53`	Capture 5 UDP packets used for DNS (domain name lookup).

Note that you need to generate some traffic. In general you have to work in pairs, both logged into the same VM. Do this by using ssh (see sheet 5). For example, to get a web page form the command line, type:

```
curl http://www.bcc.co.uk/news
```

5 Summary

- You cannot send a letter to someone with a secret address. In the same way, Internet machines must have visible addresses.
- You can learn about the Internet by playing around.
- Signals cannot travel faster than the speed of light: it takes time to transmit information over a long distance. This illustrates the difference between latency and bandwidth.
- There are many protocols used on the Internet for different purposes.