

COMP 123 — INTERNET SERVICES & PROTOCOLS

Kevin Jeffay

Spring 2004

Homework 3, January 21

Due: 2 pm, February 2

Investigations into the Structure and Performance of the Internet

In this assignment you will perform some simple experiments using tools called *ping* and *traceroute* in an attempt to understand the structure and performance of the Internet. *ping* and *traceroute* are network probing utilities that exist on both UNIX and Windows. *ping* is a program that sends a 1-packet IP datagram to a remote host and requests the host to return the datagram to the sender. The *ping* program measures the time it takes for the datagram to be delivered to the remote host and then be returned to the sender. This time is commonly referred to as the “round-trip time” between the sender and receiver. *traceroute* is a more complex tool. Briefly, it discovers and reports the names of the routers along the path between a source and a destination (it “traces the route”), as well as reports on the round-trip-times between the source and each router on the path. A brief tutorial on the use of each tool is included below.

The experiments you will perform will consist of a series of “ping trials” to a remote host. In each experiment you will first determine the route from your local computer to a remote host using *traceroute*. Once you have the route you are to ping each router along the path from the local machine to the remote host 60 times in one minute and record the average delay and loss rates seen. You are to repeat the ping experiments 5 times during a single weekday. The 5 time periods are early morning (7-9am), mid-morning (10am-12pm), mid-afternoon (2-5pm), evening (6-8pm), and night (10pm-6am). During each of these time periods you are to run 1 ping trial to each router on the path to your remote host. Once you have gathered your data, make a series of plots and tables that show the how the delay and loss rates change over time. In addition, write some text that describes and explains your data.

You are to perform three series of identical experiments. One experiment should be to a major university on the West coast (*e.g.*, UC Berkeley, UCLA, USC, University of Washington, *etc.*), one should be to your favorite (non-local) web site (to the server that hosts your favorite web site), and one to an institution in Europe or Asia. You can include more sites in your experiments if you like but you must include at least one of each of these three classes of sites.

Ultimately, you are to hand in a written report describing the experiments you performed, the data you gathered (in the form of plots or tables), and any interesting anomalies you observe or conclusions you are able to draw from the data.

A Quick Primer on *ping* and *traceroute*

Ping on UNIX

The syntax for the ping command on UNIX is:

```
/bin/ping [-c count] [-s packetsize] host
```

As per the UNIX convention, the parameter options in square brackets (“[]”) are optional. (There are lots of additional parameters to ping — read the on-line man page for details; here we just list the most useful ones.) The parameters are:

<code>-c count</code>	The number of times to ping the host (pinging once per second).
<code>-s packet_size</code>	The size of the ping packet to send (in bytes).
<code>host</code>	The host name (or IP address) of the host to ping.

For example, if we ping the main Web server for Duke (*www.duke.edu*) 5 times with 100 byte packets we get the following:

```
(classroom) 118> ping -c 5 -s 100 www.duke.edu
PING www.duke.edu (152.3.233.7) from 152.2.131.245 : 100(128) bytes of data.
108 bytes from farmer.acpub.duke.edu (152.3.233.7): icmp_seq=1 ttl=251 time=0.689 ms
108 bytes from farmer.acpub.duke.edu (152.3.233.7): icmp_seq=2 ttl=251 time=0.716 ms
108 bytes from farmer.acpub.duke.edu (152.3.233.7): icmp_seq=3 ttl=251 time=0.694 ms
108 bytes from farmer.acpub.duke.edu (152.3.233.7): icmp_seq=4 ttl=251 time=0.667 ms
108 bytes from farmer.acpub.duke.edu (152.3.233.7): icmp_seq=5 ttl=251 time=0.705 ms

--- www.duke.edu ping statistics ---
5 packets transmitted, 5 received, 0% loss, time 4039ms
rtt min/avg/max/mdev = 0.667/0.694/0.716/0.023 ms
```

First *ping* echos the name and IP address of the machine to be pinged, the IP address of the pinging machine, and the message size. *ping* then reports the results of each of the 5 pings. Note that the messages received are 108 bytes and not 100. The additional 8 bytes are because of protocol headers (which again will be explained later in the course). *ping* then reports the fully qualified name and IP address of the machine that responded to the ping, namely *farmer.acpub.duke.edu*. Note that the name of the machine responding, *farmer.acpub.duke.edu*, is different from what we specified on the command line. This is because when *ping* attempts to resolve the domain name *www.duke.edu*, it learns that *www.duke.edu* is a synonym for the machine *farmer.acpub.duke.edu* (we’ll study the problem of name assignment and name resolution later in the course). Next, *ping* reports a sequence number indicating which ping message this is in response to (“icmp_seq=#”). If a ping message is lost there will be a gap in the sequence numbers (*i.e.*, if ping *k* is lost, you won’t see a ping report with icmp_seq=*k*). The next field reports the time-to-live value of the response (“ttl=251”). This field can be ignored for this assignment. Finally *ping* reports how long it took the ping message to be delivered to the destination and for the response to return (“time=*x*”). After the 5 summary lines, it reports a summary of the exchange including the number of messages sent, received, and lost as well as the min, average, and maximum round-trip times.

Ping on a PC

The format of the command on the PC is slightly different:

```
ping [-n count] [-l packet_size] host
```

but it works the same way:

```
$ ping -n 5 -l 1000 www.duke.edu

Pinging farmer.acpub.duke.edu [152.3.233.7] with 1000 bytes of data:

Reply from 152.3.233.7: bytes=1000 time=3ms TTL=251
Reply from 152.3.233.7: bytes=1000 time=3ms TTL=251
Reply from 152.3.233.7: bytes=1000 time=3ms TTL=251
Reply from 152.3.233.7: bytes=1000 time=3ms TTL=251
Reply from 152.3.233.7: bytes=1000 time=3ms TTL=251
```

```
Ping statistics for 152.3.233.7:
  Packets: Sent = 5, Received = 5, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
  Minimum = 3ms, Maximum = 3ms, Average = 3ms
```

Once again there is a summary of the command request, though in a different format. After the replies are received, a summary is again reported with the same information as the UNIX version.

On the PC, running the ping command with only a host is the same as running “*ping -n 4 host*”; you get 4 summary lines.

Traceroute on UNIX

The *traceroute* command just takes a hostname (or IP address) as an argument:

```
/bin/traceroute host
```

traceroute probes routers along the path from the local machine to the host with a message that is similar to a ping. Each router will be “pinged” 3 times. *traceroute* reports the series of routers that are encountered along the path from source to destination and the statistics about the time it takes to reach each of these gateways. For example, a traceroute to *www.duke.edu* results in the following trace:

```
(classroom) 120> traceroute www.duke.edu
traceroute: Warning: www.duke.edu has multiple addresses; using 152.3.233.20
traceroute to www.duke.edu (152.3.233.20), 30 hops max, 38 byte packets
 1 ciscokid-cs.net.unc.edu (152.2.31.1)  0.319 ms  0.212 ms  0.200 ms
 2 gsr-12000.internet.unc.edu (128.109.36.254)  0.147 ms  0.122 ms  0.114 ms
 3 dukegsr-gw-to-ncren-oc48.ncni.net (128.109.52.3)  0.452 ms  0.408 ms  0.414 ms
 4 roti-internet-vlan200.netcom.duke.edu (152.16.51.97)  2.063 ms  1.702 ms  0.905 ms
 5 wells.acpub.duke.edu (152.3.233.20)  0.818 ms  0.683 ms  0.650 ms
```

This trace shows that there are 5 routers (“hops”) traversed when data is transmitted from the local UNIX host to *www.duke.edu*. The first two lines of output simply reports the machine to contact and some other information that is unimportant for now. The line beginning “1” is information about the first gateway reached. This is the router *ciscokid-cs.net.unc.edu*, with IP address *152.2.31.1*. *ciscokid* is the name of the UNC campus “egress router” — the router that connects UNC to the outside world. This is a machine located in Phillips Hall. The three time values listed on line 1 are the amount of time it took to reach *ciscokid* on each of the 3 attempts to probe the route.

The other routers along the way to Duke were *gsr-12000.internet.unc.edu*, also a UNC machine. Next was a hop through NCNI (the North Carolina Networking Initiative) at MCNC in the Research Triangle (*dukegsr-gw-to-ncren-oc48.ncni.net*), then to a machine on the Duke campus (*roti-internet-vlan200.netcom.duke.edu*) and finally to *wells.acpub.duke.edu*. Note that this last machine is different from the machine that we pinged and that this fact was indicated by *traceroute* when it discovered that the hostname *www.duke.edu* has multiple IP addresses associated with it. We’ll study all these details later in the course.

When you use *traceroute* it is highly likely that errors will be encountered. Common errors include lost probe messages and an inability to determine the name of a router or gateway. In the former case, if probe messages are lost or if the host was unreachable, asterisks will be printed in place of a time value (e.g., instead of printing three time values, you might see something like:

```
ciscokid-cs.net.unc.edu (152.2.31.1)  0.319 ms  *  0.200 ms
```

indicating that the second probe message to *ciscokid-cs* was lost). In the latter case, if *traceroute* is unable to determine the name of a router or gateway you will also see a series of “*’s” for the name of the host. This may indicate that either the probe is failing or that the network administrators do not

wish to report any statistics for that gateway. It is recommended that you study the on-line UNIX man page for *traceroute* to completely understand the program's output.

Traceroute on a PC

The *traceroute* command on the PC has a different name (to conform to the archaic MSDOS "8.3" filename format) but basically works the same way:

```
$ tracert www.duke.edu
```

```
Tracing route to farmer.acpub.duke.edu [152.3.233.7]
over a maximum of 30 hops:
```

```
  1    80 ms     1 ms    <10 ms  ciscokid-cs.net.unc.edu [152.2.31.1]
  2     1 ms     <10 ms   1 ms   gsr-12000.internet.unc.edu [128.109.36.254]
  3     1 ms     1 ms     1 ms   dukegsr-gw-to-ncren-oc48.ncni.net [128.109.52.3]
  4     1 ms     1 ms     1 ms   roti-internet-vlan200.netcom.duke.edu [152.16.51.97]
  5     1 ms     1 ms     1 ms   farmer.acpub.duke.edu [152.3.233.7]
```

```
Trace complete.
```

The results are generally the same, though formatted slightly differently. The summary lines simply list the time required for each probe first and then list the hostname and IP address. Note that the PC version doesn't report precise times, instead reporting times like 1 *ms*, less than 10 *ms*, or 80 *ms*.