Routing
Due Wednesday February 5, 1997
Note: This is also the date of your first exam.

For this homework you will use *udcps* or the Sun systems in Anderson 131 to explore the operation of routing in the Internet. You will use the *traceroute* commmand to determine the routes to several different network hosts. If you are not familiar with the Sun systems, this is a good time to experiment and learn about them. You will be writing network programs on these systems for future homework assignments.

The traceroute program attempts to trace the route an IP packet would follow to some Internet host by launching UDP probe packets with a small TTL (time to live). The program then listens for an ICMP "time exceeded" reply from a gateway. We start our probes with a TTL of one and increase by one until we get an ICMP "port unreachable" which means we got to the requested destination host.

To use the program type: *traceroute <destination>* at the shell prompt. For example, to determine the path to the White House World-Wide Web server we would type: *traceroute www.whitehouse.gov*

Note: The traceroute command is in the "/opt/bin" directory on the Sun systems and "/usr/sbin" on *udcps*.

Part 1: Data Collection

Use the traceroute command to determine the path to each of the following hosts.

1. *localhost* - back to the machine you are using
2. *sunspot.cps.udayton.edu* - the local Sun server
3. *saber.udayton.edu* - the campus DEC Alpha
4. *udecc.engr.udayton.edu* - the Engineering Sun server
5. *cis.ohio-state.edu* - the CIS department at Ohio State
6. *gatech.edu* - the Georgia Institute of Technology
7. *mit.edu* - the Massachusetts Institute of Technology, a Ga Tech wannabe
8. *berkeley.edu* - that left coast school
9. *spectrum.com.au* - a company in Australia
10. *www.inria.fr* - a company in France
11. *www.msn.com* - a company from another planet

Save the output from each of these tests to a file. You can do that by redirecting output with ">" or ">>". Print these files to turn in with your homework. Make sure you label the file so I know where each trace is going.

Part 2: Data Analysis

Each line of output from the traceroute command indicates the name and address of one hop along the way to the destination. The three numbers on the right of each line indicate the time in milliseconds for three different round-trip messages to the indicated router. By averaging these three numbers you can get an idea of the round-trip time between yourself and the remote router.

1. For each of the destinations listed, indicate the number of **network layer routers** crossed before reaching the destination. If there is no router involved, indicate that.

2. From the information gathered in part 1, draw a partial map of network connectivity in the Internet. Do this by matching the routers found in each of the traces and drawing the links and routers that are found. Label each network node with the name and address found in your traceroute output. For many of the hops you can identify city names where the routers are located. Try to use these to give some geographic realism to your network map. (e.g. put the St. Louis router in the center of the US) For locations you aren't sure about, just put them somewhere that might make sense.

   For some routers there will be no name given. This means that the name service was not able to find a name for this IP address.

3. In some of your traces you will find that the round trip time reported for a given hop is greater than the time reported for a more distant hop. Does this make sense? Should it be possible that it takes longer to go to a router than it does to go through a router to some node on the other side? **Come up with an explanation for this seemingly implausible data?**