Name: **AndrewID:**

15-441
Homework #2/Fall 2012

1. *dig* is a tool that, among other things, queries the Domain name System (DNS) for the IP address that corresponds to a DNS name. Pick some wired host on the CMU campus and use dig to find its IP. What is the hostname? The IP address?

2. The *WHOIS* database contains information about various aspects of domain name and AS registration. It can be queried from the command-line using *whois*. For example:
   a. whois -h radb.ra.net ***IPaddr***
   b. whois -h whois.arin.net AS***N***

   Lookup the host by IP address. What is the Autonomous System Number (ASN) associated with this host?

3. Now, lookup your ASN from question #2. What ISP does CMU use?

4. The program *traceroute* details the path that a trace packet follows to a specific destination. The path provided by *tracerout*e identifies the routers along the way, often using DNS names that describe the ISP and other and their role within the ISPs network.

   This problem asks you to consider the following hosts:

   a. www.washington.edu
   b. www.openafs.org
   c. www.cam.ac.uk

   This question asks you to do the following for each of the three hosts:
   a. Use *traceroute* and *whois* to determine the AS number (ASN) associated with each of the routers along each path
   b. Identify the name of the ISPs along the way
   c. Guesstimate if each of the identified ISPS is local, regional, or backbone

5. *Traceroute* several universities or major research labs. Now *traceroute* your favorite major retail organizations (or other purely commercial interests). What is the deal with *ucaid.edu*? Google if you need a hand.

6. In class we discussed fragmentation and reassembly as responsibilities of the network layer, but the MTU is a function of the link layer. Why doesn't the link layer just handle the fragmentation and reassemble itself?

7. We often think of an IP address as being a unique name for a computer. But, this is certainly not necessarily the case. Each IP addressable interface on a computer has a unique IP address. This is true if these interfaces are on the same network, or on different networks.

    a. Why can’t interfaces on the same network share an IP address?

    b. Why can’t interfaces on different networks, e.g. on a multi-homed host, share the same IP address?

8. As you saw in class, a routing loop might easily occur when routes are managed by the Routing Information Protocol (RIP), preventing proper routing due to the circulation of incorrect routing information. The symptom of such a routing loop is counting to infinity: while routing updates on an unreachable network are incorrectly replaced by the older routing information, the metric gradually increases as the information repeatedly gets passed from router to router.

    a. Create an example of this “count to infinity” problem using not more than five routers.
        a.i. Draw the initial topology of the network.
        a.ii. Annotate the change that causes the instability, e.g. label which connection will fail, etc.
        a.iii. Draw the initial state of the tables on these routers
        a.iv. Draw not more than three additional sets of tables that demonstrate the count-to-infinity problem, annotating the tables to explain the progression.

9. Please distinguish between I *split horizon* and *poison reverse*. Each is used to address the count-to-infinity problem – but do they have any relative drawbacks? If so, please explain.

10. Consider the following network with unit-cost links:

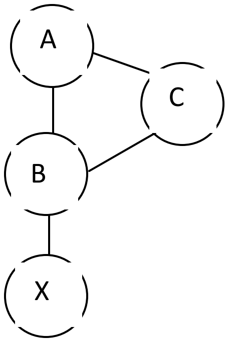

a. Imagine node X fails. Assuming A, B and C are using split horizon, could a routing loop be Formed (e.g. if B’s infinity packet to A gets lost)? Would it be different if they were also running poison reverse? Justify your answer by writing the states.

b. Suppose split horizon routers A and B somehow reach a state in which they forward traffic for a given destination X through each other. Describe how this situation would evolve with and without the use of poison reverse.

11. Suppose a hacker obtains control of all the BGP-speaking routers in several different Autonomous Systems (ASes). Our hacker has each AS “hijack” several IP blocks. That is, each AS under his or her control announces via BGP that it owns IP blocks for which it does not. For example, our hacker has AS (CMU) announce a one-hop path to the IP block 18.0.0.0/8 (MIT).

(a) Assuming that the AS graph still converges to a stable state, can this attack cause routing loops to form? Explain why or why not.

(b) Suppose the ASes under attack are identified. Can other ASes change their routing policies to ensure that their traffic still reaches the hijacked IP blocks? Explain.

(c) In response to this attack, suppose all ASes agree to check a central registry for IP block ownership before a path is considered valid. That is, whenever an AS receives a route to a prefix P, it checks that the last AS in the route actually owns P. For example, upon receiving a path to 18.0.0.0/8 (MIT), an AS will check that the last AS in the route is 3 (MIT). Can a hacker still hijack IP address blocks belonging to ASes he or she does not control? (i.e., can he or she cause traffic destined to those IP blocks to be routed to the ASes he controls?) Explain.