Homework 3, CSE 461 Fall 2009

Instructor: Wetherall, Out: 11/4, Due: 11/18

Part A – ISP Routing Programming Assignment

The goal of this part of the homework is for you to explore the differences in routes produced by different Internet routing designs. You will code up your designs as a program and run it in an emulated environment to see how well it performs. *We recommend that you do this part with a partner. We require at most two people per team and a different partner from earlier assignments. Ask Josh for help if you cannot find a partner.*

Internet routing and ISP policies

Internet routes across ISPs are selected with the BGP. A major factor in the routes that are ultimately selected is that ISPs can decide which routes they prefer to use rather than pick the shortest path. In this assignment, you will explore ISP policies within a much simplified BGP framework to understand how multiple ISPs and their policies impact Internet routes. Your job is to develop route-finding algorithms that select paths for either performance or economic cost and compare the routes they select. You will do this with the help of the programming framework that we provide.

ISPsim

In ISPsim, the Internet is made up of multiple ISPs. Each ISP has POPs (points of presence) in some of the top twenty cities in the world and a mesh of network links between these POPs. At each of its POPs, an ISP connects to a local customer that sends and receives traffic to all other customers in the Internet. At their POPs, ISPs also connect to the other ISPs that have a POP in the same city. You write code that selects routes between customers.

The goal of ISPsim is to let you explore how routes change depending on whether paths are selected to be good in terms of performance and/or economic cost. For performance, the paths taken between customers through the Internet should be fairly geographically direct. We do not really want to send packets via Hong Kong to get between Seattle and Vancouver. Directness is measured as miles along the network path compared to miles as the crow flys between the customers. That was the easy part.

For economics, the cost to the ISP of providing the paths should be low compared to the revenue it makes from customers and other ISPs. These costs and revenues depend on business arrangements that are proprietary and can be difficult to compare. We will use a simple model. ISPs charge their customers $1 for each route that sends or receives traffic to another other customers. ISPs charge $1 for each route that carries traffic between two other ISPs; so one ISP pays another ISP $1 for it to carry traffic to third ISP. ISPs charge each other nothing to hand over or take traffic that is sent to their own customers. It costs ISPs $d to carry packets across their own network, where $d is the normalized length of the path ranging from 0 to 1 for the longest possible path in the world.

ISPsim as you find it defines the classes to represent an Internet (a collection of ISPs and their POPs) and routes (destinations, POP path, and cost) and objects that give you a sample Internet, Internet0. When run, it will print out the sample Internet and trivial routes and their statistics. Your job is to extend ISPsim to compute different kinds of routes.

Your assignment

Your job is to develop code for two kinds of routes and run those routes over the sample Internet to explore how they are similar and how they differ.

The first routes you should develop are shortest path routes. Here, cost is the distance along the path. To do this, provide a distance vector style of algorithm in which every POP tells the POPs it is connected to (in different cities belonging to the same ISP and in the same city belonging to different ISPs) the best routes it knows, gathers routes advertised by these other POPs, updates its best routes, and iterates until it is converged. Note that the route structure you exchange has the full POP path of the route, not simply the next hop.

The second routes you should develop are lowest economic cost routes. By this we mean that each ISP as a matter of policy chooses for each destination the route that is lowest cost from its own point of view. This is different than the "cost" in the route, which can be ignored. To prevent loops, you must check that a POP receiving a route is not already in the POP path. Each ISP must also use the same policy; in reality ISPs may use different policies but must follow some conventions to ensure that the Internet is stable, and our simple convention is that everyone behaves with the same set of rules. One other restriction for both lowest cost and shortest path routes is that you must make deterministic choices, e.g., break ties with a consistent rule, so that the system converges.

Questions and Turn-in

Question 1: Turn in your shortest path route logic (as code). Describe your route selection rule and its rationale.

Question 2: Turn in your lowest cost route logic (as code). Describe your route selection rule and its rationale.

Question 3: Give a bar graph that shows the path lengths of routes, ordered shortest to longest, and compares shortest-path to lowest-cost routing.

Question 4: Give a bar graph that shows the economic cost for ISPs, in ISP order, and compares shortest-path to lowest-cost routing.

Question 5: Explain why shortest-path and lowest-cost routes can be different using an example route.

Part B – Problems

1. P&D4.40 and 4.45

2. P&D 5.3

3. Path MTU designs. Consider the following alternatives for how a router may handle a packet that are is too big:
    a. Fragment the packet; reassemble it at the destination.
    b. Return an ICMP message "Packet too big" to the source.
    c. Return an ICMP message "Packet too big. N bytes is the limit" to the source
    d. Truncate the packet to the maximum size and forward it.
    e. Return an ICMP message "Packet too big" to the destination.
    f. Return an ICMP message "Packet too big. N bytes is the limit" to the destination.

    Compare these designs in terms of delay and packets before content starts to reach the destination, delay and packets before the path is operating efficiently with large packets, and any other factors you argue matter. Provide a better/worse ranking of paths where possible. Realize that the path may include several routers at which the largest packet that can be sent on the next link changes.

4. Nagle and delayed ACKs. First, read about the Nagle algorithm and the Delayed ACK algorithm in your textbook and perhaps online. These two algorithms, designed to protect the network, can interact badly together to slow down transfers. Consider a small file of a few packets being transferred in one direction.
    a. Draw a time sequence diagram of the progression without Nagle and delayed ACKs.
    b. Now draw one with Nagle and delayed ACKs that shows how the transfer can be slowed down.
    c. Describe the interaction
    d. Describe how much it may slow the transfer and whether this is significant.