# Compsci 514: Computer Networks and Distributed Systems Homework 2

Instructor: Bruce Maggs

February 15, 2013

## OVERVIEW

This problem set has three questions, each with several parts. Answer them as clearly and concisely as possible. You may discuss ideas with others in the class, but your solutions must be your own. Do not look at anyone else's solutions or copy them from anywhere. (Please refer to the Duke University honor code). Turn in your solutions in on February 20, 2013 in class.

## P1: NETWORK OPERATOR FOR A DAY (WITH RCC)

To work on this problem, you will need the following resources:
The first routing table dump (or dumps, if you need them) on January 1, 2013 from the Internet2 backbone network. BGP table dumps are available at:
`http://ndb7.net.internet2.edu/bgp/RIBS/ATLA/2013/01/01/`
You should choose the first dump made on January 1, 2013 (rib.20130101.0113.gz).

1. To parse rib* files you may use a tool called zebra dump parser.

(a) Other than the sessions to private AS numbers, what are the ASes with the most number of eBGP sessions?

(b) At what routers does Microsoft have eBGP sessions to Internet2?
(Hint: You will first have to figure out Microsoft's AS number)

(c) Note that Microsoft is corporate, but Internet2 is supposedly a research and education network; why might Microsoft have eBGP sessions to Abilene?

(d) What prefixes that are advertised by Microsoft are reachable from Internet2? Which routing table did you look at to answer this question (and does it matter)?

2. Observe an output of running rcc verifier at:
`http://www.cs.duke.edu/courses/spring08/cps214/hw/ps1/rcc-html/`

(a) Click on IS-IS Errors and then on MTU Mismatch Checks. What is an MTU mismatch, and why could it cause a problem? The pair of interfaces in question start with ge-*, which typically stands for giga-bit ethernet. Which value is likely the correct value for the MTU?

(b) Under BGP Errors, click on Information Flow. These warnings indicate places where an import or export policy was configured in different ways on different routers for the same neighboring AS. What is a reasonable explanation for why anomalous import (i.e., different import policies on different neighboring routers) might be a reasonable thing for an operator to do?

(c) Under BGP Errors, click on iBGP Signaling. What is meant by an iBGP Signaling Partition, and why is it bad?

## P2: UNDERSTANDING BGP USING TABLE DUMPS

For this question, you will need to download the Routeviews routing table from
`http://www.cs.duke.edu/courses/spring08/cps214/hw/ps1/oix-full-snapshot-2008-01-20-1800.dat.bz2`
This file contains a Cisco BGP4 routing table snapshot, taken at Oregon Route Views (http://www.routeviews.org/) on January 20, 2008. (Beware: This is a text file that is 13MB, compressed. You should be able to analyze it without uncompressing it using, for example bzcat, grep, less, searching into the file - be patient when searching this is a really huge file.)
If you are curious about what other snapshots look like, you can find daily snapshots at `http://archive.routeviews.org/`

1. Find the routing table entry for Duke University network.

(a) What is the IP address of the best next hop from this router to Duke? How does this router know how to reach that next hop IP address?

(b) From the routing table file, what is the AS number for Duke?

(c) How many routes are there to get from this router to Duke?

(d) What is the best route to Duke? Why was this route selected as the best route?

(e) How many ASes must a packet traverse between the time it leaves the router and the time that it arrives at Duke?

(f) What are the AS numbers of all Dukes upstream providers? What ISP does the above AS correspond to? (Hint: You can discover this information using a whois query.)

(g) In paths where Duke University uses Time Warner Telecom (AS4323) as an upstream, the AS path ends with two instances of the same AS number. Why? What is the likely relationship between this AS number and Time Warner Telecom?

(h) Use `traceroute` to measure route from some machine at Duke to the router that took the snapshot. Please include the output of your `traceroute` with your problem set. Is the sequence of ASes from Duke to the router the same as the reverse route in the trace data? Why might the reverse path differ? (Please list reasons other than the fact that your `traceroute` was performed at a different time as the table snapshot!)

2. Look at the routing table entry for 12.108.254.0/24. This entry has several routes marked with a "d" for "damped". Give a short, one-to-two sentence explanation for (1) why routers damp routes and (2) why routers keep damped routes. To answer this question, you may want to look at RFC 2439.

3. Several of the IP prefixes in the table are formatted as w.x.y.z/m. The mask field, m, specifies the length of the network mask to use when matching input destination addresses to entries in the table.

(a) Write down the bit-wise operation to determine whether a destination address, $A_i$, matches a prefix $A$/m in the routing table. $A_i$ and $A$ are 32 bits each.

(b) Find the first "class C" CIDR address in the table (address prefix ≥ 192.0.0.0). How many class C networks does this address correspond to? What is the maximum number of routing table entries that this single CIDR address saves? Why is it that we can only infer the maximum, and not the actual, number of addresses that this CIDR address saves?

(c) In the table, there are examples of groups of prefixes that have the same advertised AS path, but show up as separate entries in the routing table.

- Provide an example of non-contiguous prefixes (and the corresponding AS path) for which this is true. Why might non-contiguous prefixes have the same AS path?
- Provide an example of contiguous prefixes (and the corresponding AS path) for which this is true. This practice is often called de-aggregation. Why might this be done?

4. RouteViews makes available table snapshots from 1997 to present. Suppose you had access to all of these snapshots, as well as some routing table snapshots from pre-CIDR. For each of the following pieces of information available in the table snapshot, what information might you be able to infer about the evolution of the Internet?

1. Only the destination addresses.
2. Only the lines marked *>.
3. Only the paths, with best next-hops marked.

# P3: UNDERSTANDING IS-IS USING PACKET TRACES

Obtain the IS-IS packet traces from the Abilene network for January 1, 2013. For example, the trace from the Atlanta router is located at
`http://ndb7.net.internet2.edu/isis/ATLA/2013/01/isisd.20130101.gz`
Five of the 9 Abilene backbone routers capture such traces. You will need all five IS-IS traces for this day to answer this question. To open the downloaded file, you will need to install Wireshark.

1. In the trace, list the different types of IS-IS messages that you see and the purpose of each message.

2. From the traces, what is the LSA refresh interval? What are the advantages of setting this value to a small value? What are the disadvantages?

3. Looking and accounting for failures: How many occurred on this day? What type?

4. Compute the average propagation time between each pair of routers (there are 5 routers). To compute the average propagation time use at least 10 values.