In the name of God

Sharif University of Technology
Department of Computer Engineering

CE 693: Advanced Computer Networks

Mehdi Kharrazi *Azar 20^{th}, 1391*

Homework 4

Please email your answers/report in **PDF format** to Mr. Baki "baki@ce.sharif.edu" and CC me at "kharrazi@sharif.edu". The HW file name should be **"Your Lastname-693-HW-4"**. It should be used as the subject of your email, too. In order for us not to miss your homework please follow the formatting. This homework is due by **Day 1st, 11:59 PM**. You are alsof

Part I: Measuring Internet topology with BGP Updates [1]

A network topology consists of the interconnection and arrangement of equipment (routers, links) that makes up the network. The complete topology of the Internet is unknown. Yet knowing the structure of the Internet is crucial in designing simulation/evaluation environments to test new protocols, in understanding the nature of how different networks interconnect, in searching for weak points and analyzing resilience of the Internet, and in provisioning for future growth patterns. To deal with this problem, researchers perform measurements to infer the structure of the Internet graph on various levels. In this problem we will infer the AS-level (ISP-level) topology of the Internet graph. In particular, for each AS in the Internet, we will determine which other ASes are connected to it. This is a simpler problem than determining how many times ASes peer with one another, or determining the router-level Internet topology, or inferring what routing policies are used at each router, etc.

The University of Oregon Route Views project (www.routeviews.org) maintains eBGP peering sessions with routers in a number of ASes. Route Views maintains historical records of BGP routing updates logged on each of these sessions. Analyzing these routing updates can provide insights into stability of the global routing system, distribution of addresses, and the structure of the Internet topology. Here, we will study the structure of the AS-level topology.

Initial Setup

Routing updates (UPDATES) and BGP table snapshots (RIBS) are located at http://archive.routeviews.org/bgpdata/. Download the first BGP table snapshot from November 1 2012. You can convert the snapshot to human-readable format by using libbgpdump at www.ris.ripe.net/source/bgpdump/libbgpdump-1.4.99.12.tar.gz

Analysis

1. Using the vantage point 67.17.82.114 plot the CDF of AS outdegree (the number of links adjacent to an AS), over all ASes in the Internet.

[1] Obtained from an earlier assignment given in CS598 Advanced Internetworking, Fall 2008 at UIUC by Matthew Caesar.

2. Using the vantage point 67.17.82.114 how many ASes are there in the Internet on November 1 2012? How fast is this number growing over time (explain your methodology for arriving at this number – you may have to look at RIBS on multiple dates).

3. Using the vantage point 67.17.82.114, how large are Internet routing tables (in terms of number of prefixes) on November 1 2012? How fast is this number growing?

4. Using the vantage point 67.17.82.114, how many links (inter-AS adjacencies) are there in the Internet?

5. Using all vantage points in the trace, how many links are there in the Internet? Using all vantage points, how many ASes are there in the Internet? Are these answers different from your answers above? Explain why your results are different, or why they are the same.

Submission:

1. A write-up that describe your methodology and observation with related graphs. Avoid printing one graph per page. Logical organization of content (text and graphs) is expected!

2. Submit your scripts or codes.

Part II: Measuring Egyptian route withdrawals [2]

On January 25, 2011, a popular uprising began in Egypt that would ultimately bring an end to the 29-year regime of Hosni Mubarak. On January 27, 2011, attempting to inhibit the Facebook- and Twitter-organized protests, the Egyptian government shut off essentially all Internet service to the country of 82 million people.

We want to use the Route Views data to figure out how long it took Egypt to leave the Internet. To do this, we will use an imperfect but simple approach: We will count how many Egyptian-related prefix withdrawals we have seen over time. Due to the dynamic nature of the Internet, there are continually announcements and withdrawals even under normal conditions. But well see one period of time with a very high rate of Egyptian withdrawals, and that period will correspond to Egypts disconnection from the Internet. We have written an incomplete parser for you (a program called simple_bgp_parse.c).This program expects to receive, on its standard input, the output of bgpdump. It scans through the BGP RIB entries or updates, and does two things. First, when it sees "interesting" entries, it remembers that the associated IP prefixes are interesting. Second, it keeps track of how many withdrawals of "interesting" prefixes it has seen, and prints out a running total. What is "interesting" is a matter of opinion. The default version of the program is quite dull and thinks nothing is interesting. As described below, you will need to decide how to pick out the "interesting" entries, in order to learn which prefixes are associated with Egypt.

1. From Route Views, download the London Internet Exchange (LINX) collectors Updates data (http://archive.routeviews.org/route-views.linx/bgpdata/), from near the time of Egypts departure from the net. That happened sometime between 21:00 and 23:00 UTC on January 27, 2011, so youll want to grab all the LINX Updates data at least in that interval. Process these files with bgpdump, and send the output of all that to the (unmodified) simple_bgp_parse program. It should output the total number of updates seen.

2. Modify simple bgp parse so that it thinks BGP RIB entries for advertisements that originated at Egyptian ASes are "interesting". (Hint 1: this only takes a few lines of code. Hint 2: the following Autonomous System (AS) numbers belong to Egyptian ISPs: 5536, 8452, 24835, 24863, and 36992. Hint 3: Think about what part of the BGP RIP entry information you can to figure out which advertisements were originated by Egyptian ASes.)

[2] Obtained from an earlier assignment given in CS598 Advanced Computer Networks, Fall 2012 at UIUC by Brighten Godfrey.

3. Now, run your modified simple_bgp_parse. This time, on standard input, feed it the output of bgpdump from the RIB file that you downloaded, concatenated with the output of bgpdump on the Update files. (The RIB output lets us learn which prefixes are interesting, and then we can count the occurences of interesting withdrawals in the updates.) The output should now be a list of pairs of numbers; read the comment near the end of simple_bgp_parse.c for a description. Using that data, draw a plot showing time on the x axis, and total number of Egyptian prefix withdrawals seen so far on the y axis.

Submission:

1. If everything worked, it should show a slowly increasing number of withdrawals seen, and then it should increase quickly for a period a "withdrawal storm", corresponding to the disconnection of Egypt from the Internet – and then the rate of withdrawals should slow down again. Based on your plot, how long did that high-rate "withdrawal storm" last?

2. Submit modified code.