Report on Programming Assignment 4: Analyzing BGP routing tables

Gaurang Naik
123079009
Collaborator: Aniruddh Rao

October 28, 2014

1 Understanding routing table data

1.1 Unique prefixes and ASs

The number of unique IP prefixes in the BGP routing table provided are 525483**.** This was calculated using a bash script which first isolated the field showing the AS path. The script then sorted all the prefixes and then extracted only the unique prefixes.

The number of unique ASs in the BGP routing table provided are 48119**.** This excludes the ASs mentioned within the {braces}. The unique ASs were also calculated using a bash script where a similar approach was followed is in the case of unique prefixes.

1.2 Announcements and withdrawals in the update file

In order to determe the prefixes that saw announcements and withdrawals, the 'updates-20140815-0000-ascii.txt' file was used.

The line in the file showing a withdrawal is as follows.

```
BGP4MP|1408060823|W|96.4.0.55|11686|182.160.108.0/24
```

The line in the file showing an announcement is as follows.

```
BGP4MP|1408060823|A|208.51.134.246|3549|112.215.16.0/24|3549 6453 24203|IGP|
208.51.134.246|0|2719|3549:2683 3549:31528|NAG|72.14.194.165|
```

In either case (withdrawal or announcement), the prefix being withdrawn or announced is specified in the sixth field (delimited by |). Thus, to identify the prefixes being announced or withdrawn, the line was first separated to give only the sixth field. Then, the prefixes were sorted and the unique prefixes were extracted. The first few lines of the output are as follows.

The number of prefixes seeing announcements and withdrawals are 17147, while the total number of prefixes are 525483.**Thus only** 0.032 **or** 3.2% **of the prefixes undergo withdrawals and announcements in the update file.**

```
1239 112.215.16.0/24
765 202.70.88.0/21
728 202.70.64.0/21
359 202.123.88.0/24
```

```
267 208.75.206.0/24
...
```

The second coulumn is the prefix and the first column indicates the number of times the prefix was announced or withdrawn. The prefix that saw the largest number of announcements and withdrawals was '112.215.16.0/24'. On closely examining the announcements and withdrawals of the prefix 112.215.16.0/24, it was observed that there were only 2 withdrawals, while all other entries were for announcements. A look at the announcements reveals that every time a withdrawal was made, the same set of announcements were made again. Also, for every profix, there were multiple announcements with all other parameters same, but with a different MED attribute. This indicates that, it is possible for every med attribute, the BGP router announces a path. But the BGP router at RouteView considers all the paths and stores only the path with lowest med attribute in its Base table.

1.3 AS path leading to IITB, MIT and AS path from IITB to MIT

The AS corresponding to IIT Bombay's prefix 103.21.127.0/24 is 132423. All AS paths that lead to IIT Bombay's AS are listed below.

```
11686 3356 6453 4755 132423
1221 4637 6453 4755 132423
1239 6453 4755 132423
1299 6453 4755 132423
13030 3549 6453 4755 132423
1668 6453 4755 132423
2152 3356 6453 4755 132423
22652 6453 4755 132423
2497 6453 4755 132423
286 6453 4755 132423
2914 6453 6453 4755 132423
293 6453 4755 132423
3130 1239 6453 4755 132423
3257 6453 4755 132423
3303 6453 4755 132423
3356 6453 4755 132423
3549 6453 4755 132423
3561 6453 4755 132423
37100 6453 4755 132423
3741 3356 6453 4755 132423
40191 3257 6453 4755 132423
5413 1299 6453 4755 132423
6539 577 6453 4755 132423
6762 6453 4755 132423
6939 1299 6453 4755 132423
701 6453 4755 132423
7018 6453 4755 132423
7660 4635 6453 4755 132423
8492 9002 6453 4755 132423
852 6453 4755 132423
40191 3257 6453 4755 132423
```

There are a total of 30 unique AS paths that lead to IIT Bombay's AS 132423.
The AS corresponding to MIT's prefix is 3. All AS paths that lead to MIT's AS are listed below.

```
11537 10578 3 3 3
11686 19782 11537 10578 3 3 3
1221 4637 174 3
1239 174 3
1299 3356 3
13030 11164 10578 3
1668 3356 3
2152 2153 11537 10578 3 3 3
22388 11537 10578 3 3 3
22652 3356 3
2497 3356 3
286 3356 3
2914 3356 3
293 3
3130 1239 3356 3
3130 2914 3356 3
3257 3356 3
3303 174 3
3356 3
3549 3356 3
3561 3356 3
37100 3356 3
3741 3356 3
40191 174 3
5413 3356 3
6539 577 174 3
6762 3356 3
6939 11164 10578 3
701 3356 3
7018 174 3
7660 22388 11537 10578 3 3 3
8492 20485 3356 3
852 174 3
```

There are a total of 33 unique AS paths that lead to MIT's AS 3.

It is mentioned in the question that the BGP router of RouteView peers with multiple BGP routers to obtain the routing table. Since, all the BGP routers mentioned in the routing table are peers of RouteView's BGP router, the only announcements made are that of customers of the respective ASs. Therefore every AS path shown above is valid and can be used to carry data.

Looking at the AS paths of IIT Bombay and MIT, there are multiple paths from IITB to MIT. One such path is, $132423 -> 4755 -> 6453 -> 293 -> 3$**.**

2 Inferring AS relationships

2.1 Obtain degree of the ASs within the routing table

As defined in Lixin Gao's paper, the degree of any AS is the number of neighbors of the AS. Based on the algorithm mentioned on page 7, figure 4, a bash script was used to determine the degree of all the ASs in the RouteView's routing table. The output of the code agrees to the sample output provided on the web page with a difference of 1 or 2 degree on the higher side.

The script file makes a two-dimensional array in which one dimension is the AS number of an AS (say AS1), and the other dimension is the neighboring ASs of the AS. In every entry of the BGP routing table, the script the neighbors of an AS (say AS2) as per the algorithm. It then checks, if AS2 is already present in the neighbor list of AS1. If not, it adds AS2 to the neighbor list of AS1. If, however, AS2 is already within the neighbor list of AS1, the script file moves on to the next AS. This is repeated for all the ASs in the routing table. At the end, degree of each AS is calculated as the number of elements in the neighbor list of the AS and printed.

The first few lines of the output are as follows.

```
174  4305
3356  4042
6939  3400
7018  2386
4323  1870
209  1547
701  1500
6461  1221
...
```

Note: The script file 'pa4_part2_degree.sh' provides updates for smaller tables fairly fast. However, for the file 'rib-20140815-ascii.txt', the script takes about an hour for the output to display. A sample output is stored in the file 'part2a.txt'.

2.2 Inferring customer-provider relationships on the IITB-MIT path

The path from IITB to MIT picked in Part 1.3 was the following. $132423- > 4755- > 6453- > 293- > 3$

Now, as mentioned in the assignment description, if you see an AS path "A B C D", and B is the AS with the largest degree on this path, then you will infer that A-B is a customer-provider link, and B-C and C-D are provider customer links.

In the above AS path, the AS that has highest degree is 6453, which has a degree of 619, followed by the AS 4755 which has a degree of 376. The AS number 293 has a degree 167, and the IITB and MIT AS have degrees of 5 and 11 respectively. Thus, we infer that the IITB is a customer of the AS 4755, which in turn is a customer of AS 6453. This is the uphill path, where the AS 6453 is connected to most ASs. Now follows the downhill path, where AS 293 is the customer of AS 6453 and MIT is the customer of AS 293.

From the WHOIS database, the AS number 6453 belongs to TATA COMMUNICATIONS (AMERICA) INC, while the AS 4755 belongs to TATACOMM-AS. The AS 293 belongs to ESnet in the US. Thus, packets forwarded from IITB to MIT first go to the TATACOMM-AS in Mumbai, and then to TATA COMMUNICATIONS (AMERICA) INC AS in the US. This is the AS with the highest degree in the above path. The packet is the forwarded to the ESnet AS in the US, which finally forwards packets to the MIT AS. The ASs along the path predicted as Providers by the algorithm. do sound like providers from the WHOIS database.