17-10-2007

On the Characteristics of BGP Routes

J. Clement

University of Wollongong, uow_clement@uow.edu.au

K. Chin

University of Wollongong, kwanwu@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/infopapers>

Part of the [Physical Sciences and Mathematics Commons](#)

Recommended Citation

Clement, J. and Chin, K.: On the Characteristics of BGP Routes 2007.
<https://ro.uow.edu.au/infopapers/658>

On the Characteristics of BGP Routes

Abstract

This paper studies the characteristics of Border Gateway Protocol (BGP) routes. In particular, it identifies autonomous systems (ASs) that appear frequently in path vectors. In addition, it analyzes path vector lengths. Finally, it investigates the prevalence of inflated paths, commonly used by ASs for traffic engineering purposes.

Disciplines

Physical Sciences and Mathematics

Publication Details

This conference paper was originally published as Clement, J, Chin, K, On the Characteristics of BGP Routes, International Symposium on Communications and Information Technologies ISCIT 2007, 17-19 Oct, 655-659.

On the Characteristics of BGP Routes

Julien Clément and Kwan-Wu Chin

*School of Electrical, Computer and Telecommunications Engineering
University of Wollongong
Northfields Ave, NSW, Australia 2522
{jyc157, kwanwu}@uow.edu.au*

Abstract—This paper studies the characteristics of Border Gateway Protocol (BGP) routes. In particular, it identifies autonomous systems (ASs) that appear frequently in path vectors. In addition, it analyzes path vector lengths. Finally, it investigates the prevalence of inflated paths, commonly used by ASs for traffic engineering purposes.

I. INTRODUCTION

BGP plays a key role on the Internet. It interconnects autonomous systems (ASs), each with a unique number (ASN), and facilitates the construction of routing tables. Each AS has one or more BGP speakers that exchange reachability information with other speakers located in neighboring or peer ASs. For example, a speaker may receive the following reachability information from peer AS-1384: [130.130.0.0/16 → {1384 38 383 211 666}], which indicates that AS-1384 has a route to the network 130.130.0.0/16, via the ASs: {1384, 38, 383, 211, 666}. Hence, by exchanging reachability information, each speaker is able to construct a routing table containing paths to various destinations on the Internet.

Many researchers have used BGP routing tables to investigate different aspects of the Internet. Their main motivation and that of this paper is to provide an up-to-date understanding of the growth and characteristics of the Internet. Govindan et al. [4] in 1995 were the first to measure the size of the global network using BGP. They found around 900 ASs, and 1,300 links for 30,000 prefixes. Then, in 2004, Zhang et al. [8] presented a detailed study of the Internet at the AS level using four data sources, namely RIPE, RouteViews, Looking Glasses and Routing Registries. In addition, they used route updates accumulated over time to construct a comprehensive topology. They witnessed 19,000 ASs and 60,000 links. Likewise, Lad et al. counted more than 20,000 ASs in 2006 [6]. Wang et al. [7] studied the consequences of having inflated paths on ASs' inbound traffic. They found path inflation to be a widely used method for inbound traffic engineering. In fact, in June 2004, 40% of the ASs on the Internet used such export policy. According to their study, 7% of paths in 1997 were inflated and this increased to 12% in 2004. Recently, in [1], we showed that the frequency distribution of ASs in BGP routes follows a Power-Law relationship, and the distribution of path vector lengths is characterized accurately by stable distributions. This paper extends [1] by delving into plausible reasons that influence an ISP/company's ranking. It studies inflated paths, and provides more results with regards to path vector lengths.

We first describe our research methodology in Section II. In particular, we explain the datasets containing reachability information collected from BGP speakers around the world, and how they are used in our investigations. Then, in Section III, we present our key findings, followed by a summary in Section IV.

II. RESEARCH METHODOLOGY

All datasets originate from two sources: RIPE¹ and RouteViews². RIPE is a major organization that collects BGP routing information bases (RIBs) and publishes them on the Internet. RIPE provides views of the Internet from fourteen locations around the world. Thereby, its datasets capture a significant number of BGP updates exchanged by ASs. Similarly, the RouteViews project provides BGP RIBs, but only from six locations. At each location, RIPE and RouteViews have a BGP collector that peers with a set of BGP speakers, and periodically dumps all route updates into a file, called a snapshot. In this paper, we use the snapshots dumped at midnight on January 1, 2007.

Each snapshot contains millions of records or route updates. Figure 1 shows an example record. The main fields of interest are PREFIX and AS_PATH. The record in Figure 1 indicates to a BGP speaker that the prefix/network 200.189.56.0/21 is reachable from a router with an address 195.66.224.83 located in AS-5511, and the route involves the following ASs: 5511, 17379, 6140 and 14282. Here, the last AS, i.e., 14282, is assumed to be the owner of the prefix 200.189.56.0/21. The first AS in the path, i.e., AS-5511, is the peer from which BGP updates are collected. Note, there can be more than one paths to a given prefix. In other words, there may be multiple records for the same prefix, each advertising a different path. We denote the path vector advertised in Figure 1 as an inflated path, since it contains duplicate ASNs; i.e., 14282 and 6140.

A small number of duplicate records exist in RIPE snapshots. To find them, we compare records according to the following fields; PREFIX, AS_PATH, NEXT_HOP and FROM. If all fields match, the record in question has a duplicate. In three snapshots from RIPE, namely RRC1, RRC03 and RRC12, there are respectively 15 duplicates amongst 2.8 million records, 250 duplicates out of 3.7 million records and 19 out of 1.8 million records. These numbers are very small

¹<http://www.ripe.net/projects/ris/rawdata.html>

²<http://www.routeviews.org>

```

TIME: 2007-1-1 11:00:26
TYPE: MSG_TABLE_DUMP/AFI_IP
VIEW: 0 SEQUENCE 16411
PREFIX: 200.189.56.0/21
STATUS: 1 ORIGINATED:Wed Dec 20 02:15:10:2006
FROM: 196.66.224.83 AS5511
AS_PATH: 5511 17379 6140 6140 14282 14282 14282 14282 14282 14282
14282 14282 14282 14282 14282 14282 14282 14282 14282 14282
14282 14282 14282 14282 14282 14282 14282 14282 14282 14282
14282 14282 14282 14282 14282
NEXT_HOP: 195.66.224.83
COMMUNITIES: 5511:540 5511:542 5511:999 17379:6140 17379:230

```

Fig. 1. An example record.

in comparison to the total number of records. Nevertheless, we remove them to avoid any ambiguities.

Table I shows key characteristics of the snapshots used in our investigations. There are 17.2 million and 15.9 million records in RouteViews and RIPE snapshots respectively. In the table, the total number of prefixes is less than 225,000. Likewise, the number of ASs is around 24,000. The AS with the highest degree is AS-701 for fourteen snapshots, AS-7018 for five snapshots and AS-3303 for one snapshot. While ASs like MCI (AS-701) and AT&T (AS-7018) have a high degree, the average degree is rather small. The average degree is 3.30 and 3.40 for RIPE and Routeviews respectively. Many ASs involved in the snapshots have a degree of only one. These ASs are local ISPs or very small networks. They represent a very large majority of ASs. This confirms the Power-Law relationship reported by Faloutsos et al. [3], which states that only a few ASs on the Internet have a high degree.

III. RESULTS

In the following discussions, we refer to snapshots by their name. Those from RouteViews are called, EQIX, OREGON, ISC, OIX, WIDE and LINX, whereas snapshots from RIPE are referred to as RRC0, RRC1, and RRC02 ... RRC15. Note, we omit RRC08, and RRC09 since there were no data for January 1, 2007.

A. Growth of the Internet

We first report on our investigations pertaining to the growth of the Internet; as viewed from the AS level. We count the total number of links and ASs appearing in both RIPE and RouteViews snapshots. The results obtained are then compared to the following previous works [4], [8], and [2].

In 2004, Zhang et al. [8] found around 19,000 ASs. We however witnessed 24,400 ASs. This means, in a little bit more than two years, the number of ASs has grown around 28%. Over a year, using the results from [2], which recorded around 21,000 ASs in January 2006, the Internet grew by 16%. Zhang et al. [8] found the presence of 60,300 links in October 2004. Since then, there has been a 4% rise in the number of links. However, this 4% rise is a conservative value, given that Zhang et al. [8] also used data from Looking Glasses³

³Looking glasses are services that allow users to gain access to a router's routing table. More than 75% of them do not peer with RIPE or RouteViews, hence they are an alternate source of BGP routes information.

and Route Registries⁴. Govindan et al. [4] found in June 1995 that the average degree is 2.99. In our investigation, we found the average degree to be 5.15, an increase of 72% in a little bit more than ten years.

B. AS Frequency in BGP Routes

In this section, we count the number of times or frequency an AS has appeared in routes from all snapshots. The aim here is to determine whether there are ASs that are crucial to Internet routing.

Table II shows the top ten ASs, ranked according to their frequency. These ASs appeared in more than 71% of all observed routes. In fact, the popularity of these ASs agrees with the Power-Law relationship reported in [1]. The highest ranked AS in terms of frequency is Level 3 Communications. It is witnessed more than five million times, and is 20% more frequent than Sprint. Table II also shows other characteristics, columns 4-9, for each AS. The goal of these columns is to determine whether they explain the position/ranking of an AS in Table II. For example, is an AS's ranking related to the number of IP addresses it owns? We explain these characteristics next.

1) *AS Degree*: On a per snapshot basis, an AS's degree does not have any impact on its ranking. The AS with the highest degree is MCI but it is only ranked fifth in frequency. The ninth and tenth ranked ASs are Tiscali and Teleglobe Inc, but they are only ranked 34th and 29th in terms of degree. As a consequence, there is no relation between an AS's degree and its popularity.

2) *Address Block Size*: The address block size corresponds to the number of addresses that an ISP owns. We like to point out that the address block size reported herein is an approximation given that companies/ISPs merge or may be bought by a different company, hence it is difficult to determine the exact number of addresses an ISP/company actually owns. From Table II, it can be seen that there is no correlation between the address block size an ISP owns to its ranking. Level 3 Communication and AT&T own more than 30 million addresses each. They both own two "/8" prefixes and their address block size is by far larger than any other ISPs in Table II. However, AT&T is only the seventh most frequent AS.

3) *Unique Prefixes Advertised*: The unique prefixes advertised correspond to the number of non-aggregatable addresses that an AS owns. As can be seen in Table II, the number of unique prefixes advertised is not linked to ASs ranking. For example, even though MCI owns the highest number of prefixes, it is only ranked fifth in terms of frequency.

4) *Average Hop Count*: The hop count refers to the average number of hops to reach an AS. If the hop count is low, an AS is more likely to provide transit. However, even if the overall hop count is low for every AS in Table II, no strong correlation can be drawn. The AS with the lowest average hop count is Tiscali but it is only ranked ninth.

⁴Route registries contain ASs peering information, and their import and export policies.

Dataset	Peer Location	N ⁰ Records	N ⁰ Prefixes	N ⁰ ASs	ASN Range	N ⁰ Links	Highest Degree	Average Degree
EQIX	Virginia, USA	923,466	203,952	24,030	2-65,509	41,045	1874 (AS701)	3.42
OREGON	Oregon, USA	8,876,550	221,542	24,328	2-65,535	50,664	2405 (AS701)	4.16
ISC	California, USA	1,888,877	207,061	24,186	2-65,500	43,711	2341 (AS701)	3.61
OIX	Eugene, USA	215,136	215,136	24,069	3-65,517	32,354	1840 (AS7018)	2.69
WIDE	Tokyo, Japan	616,050	207,344	24,068	3-65,124	32,860	1672 (AS701)	2.73
LINX	London, UK	4,672,238	207,406	24,085	2-65,509	45,999	2265 (AS701)	3.81
RRC0	Amsterdam, Holland	2,043,298	210,040	24,249	1-65,530	45,761	2344 (AS701)	3.77
RRC1	London, UK	2,822,990	201,741	24,112	1-65,509	46,629	2216 (AS701)	3.87
RRC02	Paris, France	5,364	5,281	806	71-65,531	840	81 (AS3303)	2.08
RRC03	Amsterdam, Holland	3,710,613	209,427	24,209	2-65,517	46,619	2215 (AS701)	3.88
RRC04	Geneva, Switzerland	745,712	207,250	24,145	2-65,509	40,910	1981 (AS701)	3.39
RRC05	Vienna, Austria	1,044,820	206,533	24,144	2-65,509	40,455	1786 (AS701)	3.35
RRC06	Otemachi, Japan	311,964	156,820	21,041	3-65,124	27,928	1456 (AS7018)	2.65
RRC07	Stockholm, Sweden	328,860	66,082	11,343	2-42,148	19,303	1162 (AS7018)	3.40
RRC10	Milan, Italy	257,483	118,769	18,698	2-65,500	27,808	1341 (AS7018)	2.97
RRC11	New York, USA	848,773	204,297	24,037	2-65,509	40,530	2313 (AS701)	3.37
RRC12	Frankfurt, Germany	1,810,207	209,262	24,220	2-65,509	44,772	2104 (AS701)	3.70
RRC13	Moscow, Russia	1,225,236	206,217	24,125	3-65,509	41,283	1930 (AS701)	3.42
RRC14	Palo Alto, USA	215,210	69,169	11,836	2-65,518	19,291	1366 (AS701)	3.26
RRC15	sao Paulo, Brazil	517,361	207,249	24,013	3-65,124	35,569	1664 (AS7018)	3.05

TABLE I
GENERAL CHARACTERISTICS OF THE DATASETS.

5) *Transit Percentage*: The transit percentage corresponds to the percentage of path an AS provides transit to. It is the number of non-Origin ASs divided by the total frequency of an AS. The more that an ISP is willing to provide transit to other ASs, the more likely it is a Tier-1 ISP. Again, this cannot be linked with an AS's ranking since smaller ISPs can still have a transit percentage of 99%. On the other hand, AT&T shows a transit percentage of only 91%.

6) *Market Capitalization*: Market capitalization corresponds to a company's value. It is calculated by multiplying the number of shares of a company by the value of its share at a given point in time. The last column of Table II shows the market capitalization of each company on June 13th, 2007. AT&T is by far the biggest company, followed by MCI and Sprint. Some companies like Global Crossing are relatively small compared to AT&T. From the table, we conclude that a company's market presence has no bearing on its ranking.

7) *A Possible Explanation*: Almost all of the ASs in Table II are located in the United States. More than 17,000 ASs are in the US [5], which accounts for 46% of the total AS allocation by IANA. No other country owns more than 2,000 ASs. Therefore, it is conceivable that American companies are more likely to have a high degree and frequency in BGP routes.

Unfortunately, the aforementioned explanation and hypotheses do not fully justify the popularity and the ranking of ASs outlined in Table II. Therefore, more hypotheses need to be tested. We leave this as future work.

C. Route Lengths

The next question of interest is the length of route vectors. In other words, the number of ASs in the AS_PATH field. Note, we remove all duplicate ASNs. Hence, AS_PATH: 12-23-34 has a length of three while AS_PATH: 12-23-34-34-34-34 also has a length of three. Tables III and IV show the path length recorded for each snapshot.

Snapshot	Path Length Range	Average Path Length	Variance
EQIX	1-12	3.63	0.85
OREGON	1-13	3.84	1.08
ISC	1-15	3.19	0.65
OIX	1-8	2.67	0.71
WIDE	1-10	4.08	0.96
LINX	1-12	3.77	0.85

TABLE III
AS PATH INFORMATION FROM ROUTEViews SNAPSHOTs

From Table III, the longest path is 15. Considering the size of the Internet, with more than 24,000 ASs, this is rather short. Recall that our analysis is at the AS level. If it was carried out at the router level, route lengths will certainly be longer than those reported here. From the table, we see that the longest path is very homogenous between snapshots. The longest path of six snapshots is 12 hops whereas for six other snapshots it is 13 hops.

The average path length is very small. It goes over five in only one snapshot. Nine snapshots have an average path length between four and five whereas ten of them have an average

ASN	Company	Frequency (in millions)	Degree	Address block size	Unique prefixes advertised (in thousands)	Average hop count	Transit %	Market Cap (in \$billion)
3356	Level 3 Comms.	5,495	1345 (5 th)	> 30,000	728	2.36	98.9	8.49
1239	Sprint	4,196	1727 (3 rd)	1,049	1173	2.43	96.9	62.50
3549	Global X	4,167	871 (6 th)	114	496	2.18	98.9	0.76
1299	Telianet	3,558	443 (25 th)	227	64	2.08	99.8	N/A
701	MCI	3,062	2406 (1 st)	2,180	6,641	2.59	95.0	125.08
2914	NTT America	3,028	551 (13 th)	481	276	1.80	98.8	31.67
7018	AT&T	2,766	2036 (2 nd)	> 30,000	2,249	2.58	91.2	240.95
174	Cogent	2,280	1628 (4 th)	74	2,689	2.09	93.1	1.27
3257	Tiscali	1,899	338 (34 th)	< 50	32	1.75	99.8	N/A
6453	VSNL Intl.	1,752	372 (29 th)	66	168	1.99	99.0	N/A

TABLE II
TOP 10 ASS IN FREQUENCY OUT OF ALL SNAPSHOTS AND THEIR CHARACTERISTICS.

Snapshot	Path Length Range	Average Path Length	Variance
RRC0	1-13	4.34	1.30
RRC1	1-12	3.89	0.94
RRC02	1-6	2.38	0.52
RRC03	1-13	4.00	1.06
RRC04	1-13	4.08	0.82
RRC05	1-12	4.02	0.89
RRC06	1-12	3.94	0.98
RRC07	1-13	4.28	1.28
RRC10	1-10	3.98	0.83
RRC11	1-12	4.01	1.07
RRC12	1-13	4.04	1.13
RRC13	1-14	5.03	1.68
RRC14	1-11	3.59	0.74
RRC15	1-13	4.57	0.96

TABLE IV
AS PATH INFORMATION FROM RIPE SNAPSHOTS

Fig. 2. Path Length versus Frequency.

path length shorter than four. Overall, the average path length is 3.87. It means that on average, it takes a packet less than four hops to reach its destination.

Figure 2 shows the distribution of path lengths for the OREGON snapshot. The average path length is 3.84. From the figure it can be seen that the distribution exhibits a long tail. In fact, it can be characterized by stable distributions [1].

D. Inflated Paths

Next, we investigate the prevalence of inflated paths. Table V and VI show the percentage of inflated paths, the longest inflated path, the biggest inflation factor and the average inflation factor for each snapshot. Here, the inflation factor is defined as the number of times an AS is present more than once in a path. For example, AS_PATH 1-2-4-4-4, has an inflation factor of two.

For OIX, 5.3% of its paths are inflated. RRC13 has the highest proportion of inflated paths with more than 25%. The average proportion of inflated paths are 12.3% and 14.0% for RouteViews and RIPE respectively. Overall, 13.1% of paths are inflated. In comparison, Wang et al. [7] reported that 12% of paths were inflated in 2004 and 7% in 1997.

From our investigations, we found that some paths are

repeated over thirty times! Figure 1 shows an example path vector from RRC1 that has a length of 41, and has an inflation factor of 37. Here, AS-14282 has added itself to the path more than thirty times before advertising the path to AS-6140. This means AS-14282 is discouraging AS-6140 from using the link 6140-14282. One reason is that 6140-14282 may be a backup link. Another explanation is that the link from 6140 to 14282 could be an expensive link, monetary wise. Another possibility is that the link may have a large delay, e.g., satellite links.

The average inflation factor for each snapshot is less than one. For nineteen snapshots, the inflation factor is less than 0.50. Overall, the average inflation factor is 0.32. These results imply that when a path is inflated, a small inflation factor is used most of the time. Cases of big inflation factors are rare.

The inflated path distribution for each snapshot is carried out to confirm that inflated paths in general have a rather low inflation factor. Figure 3 and 4 plot the cumulative distribution of path lengths for both OREGON and RRC03.

Snapshot	Percentage of inflated paths	Longest inflated path	Biggest inflation factor	Average inflation
EQIX	14.8%	39	35	0.41
OREGON	12.8%	41	37	0.33
ISC	13.2%	39	35	0.37
OIX	5.3%	31	26	0.12
WIDE	9.3%	33	26	0.21
LINX	18.8%	41	37	0.43

TABLE V
INFLATION INFORMATION FROM ROUTEVIEWS SNAPSHOTS.

Snapshot	Percentage of inflated paths	Longest inflated path	Biggest inflation factor	Average inflation
RRC0	13.3%	39	35	0.33
RRC1	14.8%	41	37	0.44
RRC02	21.7%	17	13	0.95
RRC03	14.0%	39	35	0.41
RRC04	12.5%	39	35	0.32
RRC05	13.1%	39	35	0.38
RRC06	9.9%	33	26	0.28
RRC07	10.0%	34	26	0.33
RRC10	12.4%	32	26	0.33
RRC11	13.9%	39	35	0.36
RRC12	13.8%	39	35	0.39
RRC14	10.8%	32	26	0.29
RRC15	10.8%	38	35	0.28

TABLE VI
INFLATION INFORMATION FROM RIPE SNAPSHOTS.

Other snapshots have similar shape, therefore are omitted from presentation. Figure 3 and 4 show that most inflated paths have an inflation factor of one. The higher the inflation factor, the less likely it is going to be present in a snapshot. In all snapshots, an inflation factor of ten can be considered as very unlikely; as a consequence, the percentage of these paths is minimal (less than 1% of all paths). As discussed earlier, there are paths with an inflation factor of more than thirty. However, they constitute only a small proportion of the total AS paths.

Fig. 3. Cumulative distribution for RRC03

Fig. 4. Cumulative distribution for OREGON

IV. SUMMARY

We have investigated the characteristics of BGP routes using snapshots dumped at midnight on January 1, 2007 by the RouteViews and RIPE collectors. Our key findings are as follows:

- 1) As of January 2007, there are more than 24,000 ASs. This corresponds to a 28% growth since October 2004.
- 2) The ASs/ISPs that appear most frequently in BGP routes are Level 3 Communications, Sprint and Global Crossing. MCI and AT&T have the highest degrees.
- 3) The average path length is 3.9. Paths over eight in length are very unlikely.
- 4) The proportion of inflated paths is significant. Out of the 32 million paths, 13% are inflated. However, their inflation factor is rather small. In fact, the average inflation factor is only 0.32.

V. ACKNOWLEDGMENT

The authors would like to thank the reviewers for their constructive feedback, and the folks at RouteViews and RIPE for collecting RIBs and making them publicly available.

REFERENCES

- [1] J. Clément and K.-W. Chin. Power law and stable distributions of BGP paths, June 2007. Tech. Rep., School of Electrical, Computer, and Telecommunications Engineering, University of Wollongong, Australia.
- [2] M. Eubanks. The number of active autonomous systems in the commodity internet. Web Page, Jan. 2007. <http://www.multicasttech.com/status/>.
- [3] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proceedings of ACM SIGCOMM*, Sept. 1999.
- [4] R. Govindan and A. Reddy. An analysis of internet inter-domain topology and route stability. In *IEEE INFOCOM*, Japan, 1997.
- [5] G. Huston. BGP routing table analysis reports. Web Page, Jan. 2007. <http://www.potaroo.net/tools/asn16/>.
- [6] M. Lad, D. Massey, and L. Zhang. Visualizing internet routing changes. *IEEE Transactions on Visualization and Computer Graphics: Special Issue on Visual Analytics*, Dec. 2006.
- [7] H. Wang, R. K. Chang, D. M. Chiu, and J. C. Lui. Characterizing the performance and stability issues of the AS path prepending method: Taxonomy, measurement, study and analysis. In *Proceedings of ACM SIGCOMM Asia*, Jan. 2005.
- [8] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the internet AS level topology. In *Proceedings of ACM SIGCOMM*, Jan. 2005.