A Study on BGP AS Path Characteristics

Xiaoliang Zhao[1], Beichuan Zhang[1], Dan Massey[1], Lixia Zhang [2]

***Abstract*— BGP, as a de-facto inter-domain routing protocol, plays an important role in the Internet. In this paper, we examine one of BGP major components, *AS path*, and its characteristics. Such study can lead to a deeper understanding of BGP behavior and benefit studies of BGP dynamics, operational practice, and the performance of applications which are sensitive to routing dynamics. By examining nearly one year's BGP data, we found for most of prefixes, each of them was primarily reachable via one single path in a time window of one month, which may indicate that the AS paths could be predictable in such a time window. For example, from a particular point of view, there were 83.7% of 116,544 prefixes which were reachable by a single path for at least 95% of reachable time in March 2002. Moreover, we found 28.5% prefixes were *continuously* reachable via one single path during the same month, which is evident that some prefixes are very stable from routing perspective. In addition, this paper studies an entropy-based measurement to measure the steadiness of a path.**

I. Introduction

Ideally, as a protocol, there would be a solid understanding of BGP's behavior, such as its stability, its response to faults, and its vulnerabilities to attacks. But in practice, the BGP infrastructure constitutes a large scale system and could exhibit complex behaviors under various conditions. Therefore, understanding BGP dynamic behavior continues to be an open challenge. In this work, we focus on one aspect of BGP behavior, the AS path characteristics. There are couple of reasons to motivate this study.

- *How BGP utilizes multiple paths?* Today, Internet has become a richly connected network, and increased multihoming keeps increasing such richness. It results in multiple paths between a source and a destination. However, it remains unclear that how BGP utilizes those multiple paths. What is the distribution of utilization time for paths between a pair of source and destination?
- *How frequent do AS paths change?* Frequent path changes not only increase routers' load and consume the bandwidth, but also impact the performance of end-to-end communications. A number of studies also show that frequent path changes may have complicated effects on routing performance [1] or network operations [2]. Study of path changes may help us better understand the inputs to the routing system.

We examined the AS paths viewed from one major ISP. During the 10 month study, we observed 236,395 prefixes, but only 116,544 were announced for at least 40% of the ten month study. In this study, we focus on the path behavior of these long-lasting prefixes because we are interested in long-term AS path behavior. For these prefixes, we made the following observations:

- In March 2002, 41.6% of studied prefixes were reachable via one path only. In other words, when a route to the prefix existed, it always consisted of this one path.
- In March 2002, 83.7% of studied prefixes advertised a single path for at least 95% of the total time when the prefix was advertised. In other words, when a route to the prefix existed, with more than 95% probability, that prefix relied on a single "primary path". Moreover, 97.912% of studided prefixes advertised two paths for at least 95% of the total time when the prefix was advertised.
- From March to December 2002, 18.5% of prefixes advertised only one path. 34.87% of prefixes advertised one path for at least 95% of the total time when the prefixes were advertised. 73.33% of prefixes were reachable via two paths for at least 95% of the total time when the prefixes were advertised. It also suggests that observations made on path stability may vary when we examine different time windows.
- When reachable, 28.5% prefixes were *continuously* reachable via one path during the whole month of March. 2002, which indicates there existed a set of prefixes which have very steady path. Meanwhile, we also noted a small number of prefixes which behaved abnormally unsteady, which may indicate some network problems for those prefixes.

In summary, this study reports that in a short time window such as one month, for most prefixes, it is highly likely that each prefix was reachable via *one* single path. The path which was primarily used is called "primary path" in this paper. For a small set of prefixes, we observed the presence of primary path for a very long time window such as ten months. We use *entropy* to measure a path's steadiness and the results show that in most cases, primary paths are reasonably steady. Switching to an alternative path occurred infrequently, and if it occurred, it would last for a brief time.

The paper is organized as follows. Section II talks about our methodology used for the data processing. Section III presents the results about primary path, its steadiness and how it was replaced. Section V concludes the paper.

II. Methodology

We analyzed BGP routing updates collected by Route-Views [3] from March to December of 2002. It should be noted

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No DABT63-00-C-1027. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the DARPA.

[1] Xiaoliang Zhao, Beichuan Zhang and Daniel Massey are with Information Sciences Institute, University of Southern California. E-mail: {xzhao,bzhang,masseyd}@isi.edu [2] Lixia Zhang is with Computer Science Department, University of California, Los Angels. E-mail: lixia@cs.ucla.edu

that BGP updates might be sent to the monitoring points via multi-hop BGP connections, which may cause measurement artifacts as discussed in [4]. We pre-process the update files to remove the updates that are generated due to session reset, resulting in a cleaned data set of BGP updates, for our analysis.

Within ten months, we observed totally 236,395 distinct prefixes. The number is almost doubled than the number of prefixes in a normal backbone BGP routing table. Some operational practices, like traffic engineering at BGP level, may introduce some short-lived prefixes into the global routing table. Router misconfigurations may inject some incorrect or bogus prefixes [1] into the routing table as well. Because those falsely announced prefixes, which are expected to be corrected shortly, are not of interest of this paper, we only consider those prefixes which were announced longer than 40% of the total time we examined. After the filtering, we obtained 116,544 distinct prefixes.

III. Primary Path and Its Steadiness

In this section, we present the results about the primary path, its steadiness and how it changes.

A. Primary Path

We are interested in the time fraction for which each path was announced. For each prefix, we calculated the time fraction for each path and found most of prefixes used one particular path more frequently than others.

Figure 1(a) shows that for March 2002, from Peer-A's point of view, 83.7% prefixes were reachable by a single path for at least 95% of total time. Similiar results are obtained for other months as well. Please note that the time fraction were counted over the total time for which a prefix was announced.

One explanation why most prefixes favor one single path could be because, as pointed out in [6], to choose best path, BGP uses the "static" metrics, such as AS path length, local preferences, and router id, etc. Unlike those dynamics metrics, such as link utilization rate, congestion information and routers' load, "static" metrics are not changed very often. Consequently, the best path computed from the "static" metrics will be the same one for most of time.

We also examined the data for a longer time scale. Figure 1(b) shows that for ten months periods, there are about 21,612 prefixes are reachable by one single path for the time for which the prefix was reachable during ten months. It might be because those prefixes are either directly connected with Peer-A (customers of Peer-A), or only one AS hop away from Peer-A (customers of Peer-A's peer). Therefore those prefixes do not have many alternative paths to choose. However, further study revealed that 11,704 prefixes were two AS hops away from Peer-A, and 2,185 prefixes were even longer than two AS hops away. It implies that, from one viewpoint, there exist a small set of prefixes which are heavily rely on one path. On the other hand, comparing with Figure 1a, one single path count for less percentage of total reachable time, which may indicate that paths do change over a long time period. Therefore, observations of path stability, such as the one in [7], may vary when observation window varies.

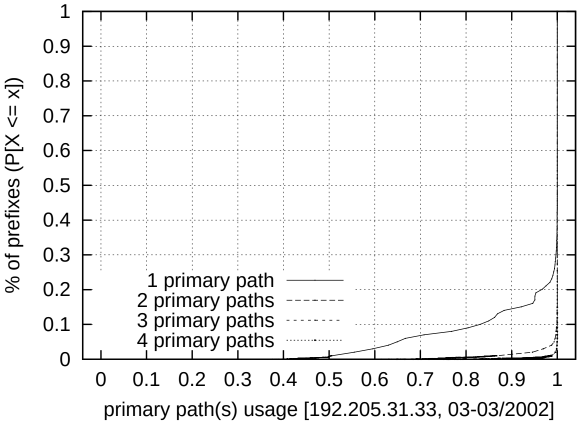

(a) Mar. 2002

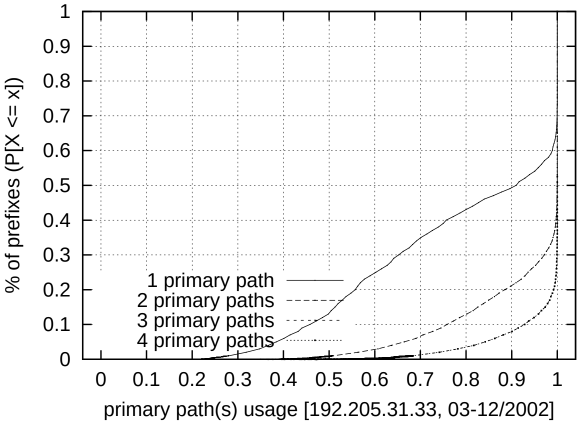

(b) Mar. to Dec., 2002

Fig. 1. Path usages.

Data suggest that for most prefixes, each of them was primarily reachable via one single path. In this paper, we called the path which was primarily used as *primary path*. In the following sections, we further study the steadiness and path change patterns specially for primary paths, and we only consider those prefixes which have a primary path. Empirically, we set a threshold as 90% of reachable time to judge if a prefix had a primary path. If a prefix was reachable by a single path for less than 90% of the total reachable time, we consider such prefix has no primary path. After the filtering, we finally got total 86,283 (74%) prefixes for Mar. 2002.

B. Primary path steadiness

A prefix may be reachable by its primary path for most of time, but may or may not be reachable by its primary path continuously. There could be two extremes, one is that the prefix was reachable by its primary path continuously for long

[1] Also known as 'bogon prefixes" [5]

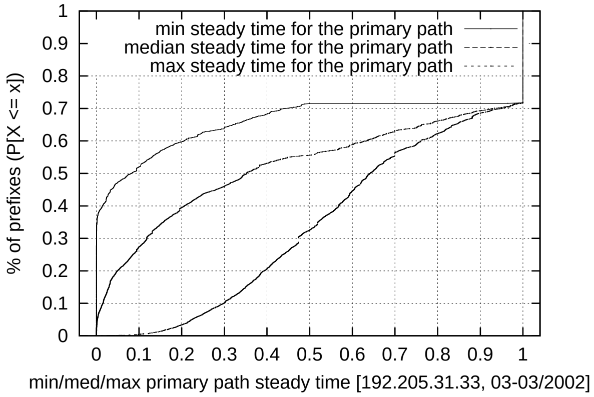

Fig. 2. Path steadiness.

time; another is its primary path frequently was replaced by other paths, but those paths were used for very short time period and then they were replaced by the primary path again. Figure 2 shows the minimal, mean and maximal continuous used time for each prefix's primary path. For about 24,176 prefixes, the minimal duration of the primary path is equal to the total time of the prefix was available. It suggests that there exist a small set of prefixes which was steadily reachable by their primary paths. For around 18 prefixes, the maximal duration of using primary paths are less than one day, which may imply those prefixes experienced the routing problems.

To further study the primary path steadiness, we define the path steadiness more formally. A path pp may be used for a while, then was replaced by another path (or was withdrawn), then was reused again. Therefore, it forms a series of time intervals for which p was continuously used. Let $I_p = (i_1, i_2, \ldots, i_{n_p})$ denotes such set of intervals for path p, $m_p \in I_p$ denotes the maximal value in I_p, n_p denotes the size of I_p, $T_p = \sum_{k=1}^{n_p} i_k$ denotes the total time p is used, we then adopt the concept of the *entropy*, denoted by $E(p)$, to measure the steadiness of path p. More precisely,

$$E(p) = -\sum_{k=1}^{n_p} \frac{i_k}{T_p} \ln(\frac{i_k}{T_p})$$

$E(p)$ has the following properties:

1) When $n_p = 1$, $E(p) = 0$. It corresponds to those very steady paths which never have been changed.
2) When $i_1 = i_2 = ...i_{n_p}$, $E(p) = \ln n_p$, and when $n_p \to \infty$, $E(p) \to \infty$. It corresponds the those very unsteady paths which are frequently replaced. Such paths will have very large entropy.
3) Given $E(p)$, we have $n_p \geq e^{E(p)}$ and $m_p \geq T_p e^{-E(p)}$. *Proof:* It has been proved that $E(p) \leq ln(n_p)$, easily we will have $n_p \geq e^{E(p)}$. According to the definition of $E(p)$, we have $-E(p) = \frac{m_p}{T_p}ln(\frac{m_p}{T_p}) + \sum_{k=1,i_k\neq m_p}^{n_p} \frac{i_k}{T_p}ln(\frac{i_k}{T_p})$. Since $\frac{i_k}{T_p} \leq \frac{m_p}{T_p}$, $\sum_{k=1,i_k\neq m_p}^{n_p} \frac{i_k}{T_p}ln(\frac{i_k}{T_p}) \leq \sum_{k=1,i_k\neq m_p}^{n_p} \frac{i_k}{T_p}ln(\frac{m_p}{T_p}) = ln(\frac{m_p}{T_p})(1 - \frac{m_p}{T_p})$. Then we will have $-E(p) \leq \frac{m_p}{T_p}ln(\frac{m_p}{T_p}) + ln(\frac{m_p}{T_p})(1 - \frac{m_p}{T_p}) = ln(\frac{m_p}{T_p})$, and $m_p \geq T_p e^{-E(p)}$.

Therefore, $E(p)$ provides hints about the path change frequency and the longest duration a path is continuously used. Also, we could draw the distribution of the $E(p)$ of primary paths for overall Internet prefixes, and use the distribution as a reference to compare if a primary path for a particular prefix is abnormally unsteady in a statistical sense. Moreover, similar to [8], we could use $E(p)$ as a statistical anomaly detection tool to detect abnormal behavior of a path.

Figure 3(a) shows the distribution of entropy of primary paths for Mar. 2002. We found that 24,564 (28.5%) prefixes have primary path with entropy $E(p) = 0$, which means those prefixes were continuously reachable by the primary path once the primary path was used. While at another extreme, we found few primary paths with $E(p) = 6.61$. For example, one prefix was reachable by the primary path for 30.1 days, but changed 1194 times, and the longest interval to continuously use that path is only 8.8 hours. The behavior of such prefixes are considered as abnormal, which may indicate some configuration errors or link flapping.

There were 75% primary paths with $E(p) \leq 1.277$, and because $\frac{m_p}{T_p} \geq e^{-E(p)} \geq e^{-1.277} = 0.279$, which means 75% examined primary paths are at least continuously used for 27.9% total used time. On the other hand, 25% primary paths with $E(p) > 1.277$, and because $n_p \geq e^{E(p)} > e^{1.277} = 3.59$, it means there are 25% examined primary paths being changed at least four times.

Using the similar approach, we can compute the entropy for a prefix to measure the stability of a prefix. Every time interval which identifies the time period of a stable state for a prefix is collected, and the entropy is computed over the set of the intervals. We plot the entropy distribution for all prefixes in Figure 3(a), and it shows a match between two curves, which may indicate correlation between two different type of entropy. Since Figure 3(a) plots the accumulative results, we verify the correlation by plotting the primary path entropy versus prefix entropy for each prefix in Figure 3(b). As can be seen, there is a strong correlation between prefix entropy and primary path entropy. It may imply that the steadiness of primary path in general determines the steadiness of the prefix.

C. *Path change patterns*

In this section, we study further about the path change patterns, i.e., how long for an alternative path replaced the primary path, and how many path changes occurred before the primary path was reused again. The path change pattern is of interest of some applications. For example, path change patterns may help us better understanding of inputs to BGP damping algorithm [1].

To study the path change pattern, we first assumed that initially a prefix was on the primary path. Then at some time, another path may replace the primary path, and we started to count how long the replacing path would be used until the primary path was restored. Such time interval will be counted as "leaving time". Figure 4(a) shows the minimum, median and maximum of leaving time of each prefix. From

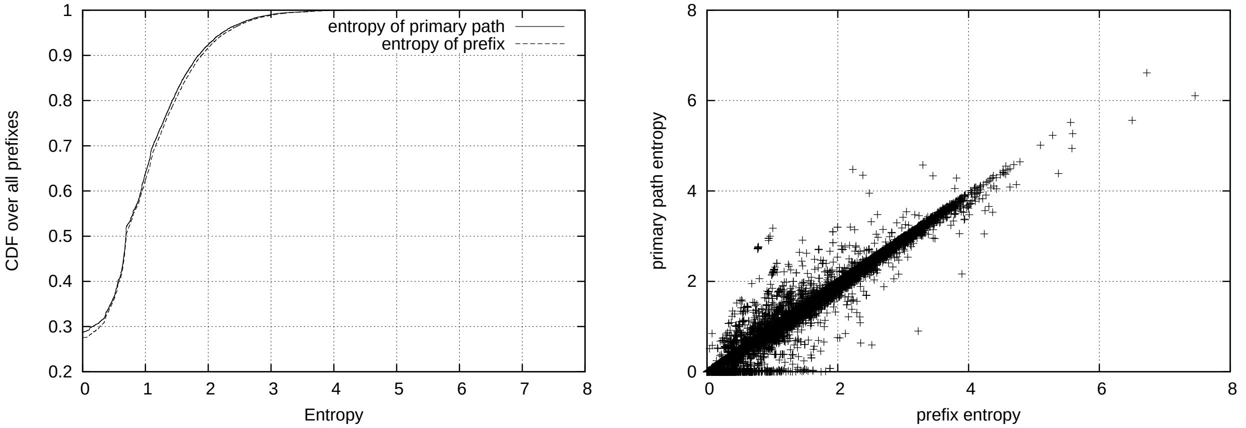

(a) Distribution of $E(p)$.

(b) Correlation between primary path entropy and prefix entropy.

Fig. 3. Entropy-based measurement of steadiness

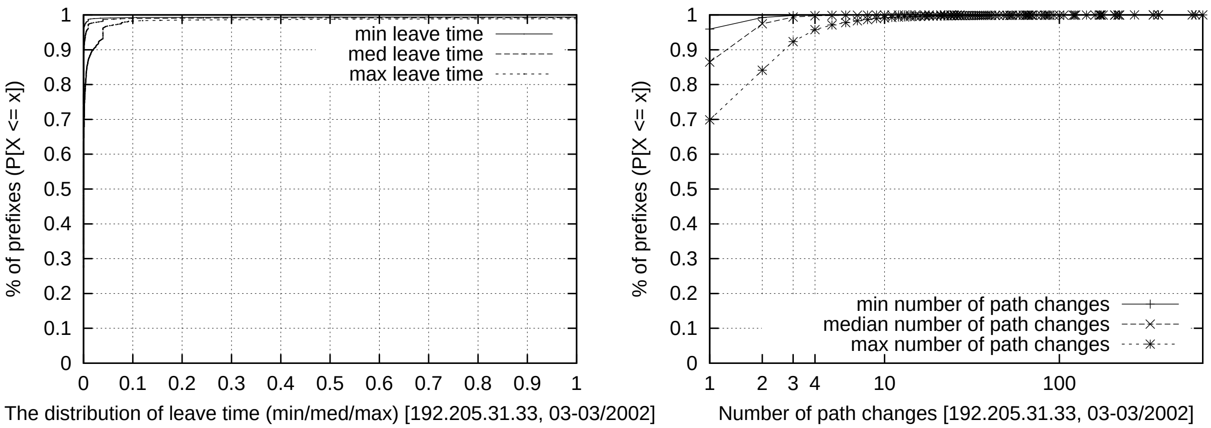

(a) Time for not using a primary path.

(b) Number of path changes before reusing the primary path.

Fig. 4. Path change patterns.

the graph, we can see that the leaving time usually were very short. As a matter of fact, if only the median of leaving time was concerned, 90% prefixes did not use an alternative path continuously for longer than ten minutes in Mar. 2002.

Figure 4(b) shows the number of path changes before reusing the primary path. Although 92.3% of prefixes explored no more than three paths before returning to the primary path [2], there were about 153 prefixes explored at least four paths before the primary paths was restored. In those cases, the prefixes might lose the connectivity for a while because the primary path might be dampened. Indeed, a closer look at Figure 4(b) revealed that when the number of path changes increased, the minimal and maximal value tended to be same, which means those prefixes only "leave" primary paths once, but experienced a burst path changes, which might trigger the route flap dampening.

In the summary, we could conclude that primary paths are frequently used, and most of them are very steady. Switching to an alternative path occurred infrequently, and if it occurred, it would be a brief time and only fewer paths being explored before restoring the primary path.

IV. Related work

The most relevant work [9] is done by SIganos et. al. Based on three years BGP data, they measured the prevalence and persistence of AS paths, which are equivalent to primary path and its steadiness in this paper. Althought the methodologies are slightly different, we share most of the results. In this work, we developed a new entropy-based metric, which is more effective to measure path steadiness.

Ramesh et. al. [6] have studied the Internet topology and routing stability in 1997. [6] looked at general Internet characteristics, such as topology, prefix availability, etc., and noted

[2] According to [1], at least four flaps will trigger flap dampening, thus in most cases, the primary path may not be dampened.

that most prefixes were reachable via one single primary path. However, [6] examined the early Internet back to 1995, which only contains 900 ASes and nearly 30K prefixes at that time. With the emergence of new technologies and the boost of economics, the Internet has changed dramatically, which contains more than 10,000 ASes and 100K prefixes today. We examined the most recent Internet, and focused on primary paths and their properties as well as their implications.

Vern Paxson [10] examined the end-to-end routing stability, and found at router level, one path was primarily used for the communication between a particular source and destination. However, the methodology used in [10] is quite different from ours. We are focusing on analysis of BGP traces to understand BGP behavior.

Rexford et. al. [7] studied the routing stability for "popular" prefixes, and found the paths to "popular" prefixes were in general quite stable. However, [7] only examined the a very small set of 31 prefixes. We examined almost all the prefixes to approximate to some routing property working for other networks as well.

V. Conclusion

Due to its large scale, understanding BGP behavior is difficult. This paper looked at one aspect of BGP behavior, AS path characteristics. We found for most of prefixes, each of them was primarily reachable via one single path in a time window of one month. Also, we report that the paths being primarily used often were "steady", i.e., the replacement of primary paths occurred infrequently and the duration for which the primary paths were continuously used were long in most cases. In addition, We studied the entropy-based metric to measure the steadiness of a path. Finally, we noted that if being replaced, primary paths would be reused quickly and often after few number of attempts of alternative paths. This study improves the current understanding of BGP behavior and benefits the related study such as BGP damping.

References

[1] George Varghese Zhuoqing Morley Mao, Ramesh Govindan and Randy Katz, "Route Flap Damping Exacerbates Internet Routing Convergence," in *Proceedings of the ACM SIGCOMM '02*, August 2002.

[2] N. Feamster, J. Borkenhagen, and J. Rexford, "Guidelines for Interdomain Traffic Engineering," in *ACM SIGCOMM Computer Communications Review*, Oct. 2003.

[3] "The Route Views Project," http://www.routeviews.org.

[4] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. Wu, and L. Zhang, "Observation and Analysis of BGP Behavior under Stress," in *Proceedings of the ACM SIGCOMM Internet Measurement Workshop 2002*, Nov. 2002.

[5] "Meet the Bogons," http://www.cymru.com/Documents/bogon-dd.html.

[6] Ramesh Govindan and Anoop Reddy, "An Analysis of Internet Inter-Domain Topology and Route Stability," in *Proceedings of the IEEE INFOCOM '97*, 1999.

[7] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "BGP Routing Stability of Popular Destinations," in *Proceedings of the ACM IMW 2002*, Oct. 2002.

[8] Diheng Qu, Felix S. Wu, and Feiyi Wang, "Statistical-based Intrusion Detection For OSPF Routing Protocol," in *In 6th IEEE International Conference on Network Protocols*, Oct. 1998.

[9] G. Siganos and M. Faloutsos, "BGP Routing Properties at a Large Time Scale," in *Global Internet Symposium*, Nov. 2002.

[10] Vern Paxson, "End-to-end routing behavior in the Internet," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, pp. 601–615, 1997.