

A Measurement Study on BGP AS Path Looping (BAPL) Behavior

Shenglin Zhang^{† ‡ §} Ying Liu^{† §} Dan Pei^{* ‡ §}

[†] Institute for Network Sciences and Cyberspace, Tsinghua University, Beijing 100084, China

[‡] Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China

[§] Tsinghua National Laboratory for Information Science and Technology, Tsinghua University, Beijing 100084, China

Email: slzhangsd@gmail.com, liuying@cernet.edu.cn, peidan@tsinghua.edu.cn

Abstract—As a path vector protocol, Border Gateway Protocol (BGP) messages contain the entire Autonomous System (AS) path to each destination for breaking arbitrary long AS path loops. However, after observing the global routing data from RouteViews, we find that BGP AS path looping (BAPL) behavior does occur and in fact can lead to multi-AS forwarding loops in both IPv4 and IPv6. The number and ratio of BAPLs in IPv4 and IPv6 for 1456 days on a daily basis are analyzed. Moreover, the distribution of BAPL duration and loop length in IPv4 and IPv6 are also studied. Some possible explanations for BAPLs are discussed in this paper. Private AS number leaking has contributed to 1.76% of BAPLs in IPv4 and 0.00027% in IPv6, and at least 2.85% of BAPLs in IPv4 were attributed to faulty configurations and malicious attacks. Valid explanations, including multinational companies, preventing particular AS from accepting routes, can also lead to BAPLs.

Index Terms—Forwarding Loops, BGP AS path, IPv6, RouteViews, Traceroute

I. INTRODUCTION

The Internet is made up of thousands of Autonomous Systems (ASes), each of which is a connected group of one or more IP prefixes that have a solely and clearly defined routing policy [18]. At least one intra-domain routing protocol is deployed in an AS to optimize routing within the domain, such as OSPF [19], RIP [20], and IS-IS [21]. The reachability information among ASes can be exchanged with the inter-domain protocol, Border Gateway Protocol (BGP) [16]. Each BGP route contains an AS-PATH attribute which lists the path of ASes used to reach the destination prefix.

BGP is supposed to eliminate path looping. When an AS receives a BGP routing update, it will check whether the AS-PATH attribute contains its own AS number. If so, it will discard this BGP routing message immediately to break the AS path loops. As described in RFC 4271 [16], “this information (the AS-PATH attribute) is sufficient to construct a graph of AS connectivity from which routing loops may be pruned”.

Despite BGP’s design intention of preventing AS path loops, previous research has shown the evidence of BGP AS path looping (BAPL) [4], [12], [22]. A BAPL occurs if there is a loop in the AS-PATH attribute. More precisely, suppose there are n ASes in the AS path, and the BGP AS path vector from the AS p_n to the destination AS p_1 is defined as $asp = (p_n, p_{n-1}, \dots, p_1)$. We call a BAPL happens if $p_i = p_j, j >$

$i, j - i \neq 1$. Although there were some previous works about BAPL, BAPL has not yet been systematically studied. In this paper, we try to conduct a systematic study on BAPL using real BGP and traceroute data, covering the following important aspects of BAPL.

The relationship between BAPL and forwarding looping. The BGP AS path denotes the list of ASes through which the BGP update messages propagate, while the forwarding AS path is the list of ASes that actually propagate the data packets.

Previous studies have shown that inter-domain forwarding loops exist in the Internet [5], [7]. In theory, BAPL can potentially cause forwarding loops. If a BAPL contributes to multi-AS forwarding loops, then analyzing the distribution of BAPL behavior and studying its possible causes will help us understand how to reduce loop-induced transmission delay and packet loss [11], [13], and to prevent attackers from interrupting the Internet [10]. However, whether BAPL can cause real forwarding loops in reality has not been studied, which is one of the aspects that we study in the paper. Our observation verifies that a small fraction of BAPLs (about 1%) can cause inter-domain forwarding loops.

The characteristics of BAPL. We observed that there were more than 8000 BAPL updates for IPv4 per day and more than 2000 for IPv6 on average. The majority (more than 91%) of the loops lasted shorter than one day, while non-trivial number of BAPL updates lasted longer than a month. Two-AS loops and three-AS loops dominated the loop length distribution.

Potential Causes of BAPL. We show that BAPL may occur for a few valid reasons, such as multinational companies and preventing particular AS from accepting routes. In addition, private AS number leaking contributed to 1.76% of BAPLs in IPv4 and 0.0027% in IPv6, and at least 2.85% of BAPLs in IPv4 could be attributed to malicious attacks or misconfigurations. BAPLs caused by invalid reasons (private AS number leaking, faulty configurations and intentional attacks) should be fixed by network operators.

The rest of this paper is organized as follows. Previous researches related to BAPL are summarized in Section II, and Section III describes the data set and methodology of this study. We discuss the relationship between BAPL behavior and forwarding looping in Section IV and present BAPL characteristics in Section V. The potential causes of BAPL behavior are discussed in Section VI. Section VII concludes

* Dan Pei is the correspondence author.

our work and discusses the future work.

II. RELATED WORK

Several researches have been carried out on routing loops, however, few of them focus on BAPL behavior, or the relationship between BAPL and forwarding looping.

Some researches have focused on *forwarding* AS path loops or *forwarding* routing loops. For example, Paxson has studied routing loops using end-to-end traceroute measurements collected in 1994 and 1995 [14]. Although this paper focused on persistent loops, it found a few transient loops and conjectured that such loops were caused by link failure information. Z. M. Mao et al. believed that some ASes did not broadcast their infrastructure addresses and others could announce the addresses of shared equipment at border points between ASes, which led to some forwarding AS path loops in traceroute [7]. J. Xia et al. presented a measurement study on persistent forwarding loops, and analyzed the possibility of flooding attacks that exploited persistent forwarding loops [5]. They performed extensive measurements to study persistent forwarding loops, and found that persistent loops across multiple ASes did exist in the Internet. Traceroute was also used for measurement in the paper, and 0.2% of routable addresses were found experiencing persistent forwarding loops. Nevertheless, loop detection just using end-to-end tools such as traceroute is error-prone and cannot successfully detect transient loops [11].

D. Pei et al. studied *transient* BGP path vector route looping behavior [8]. They analyzed the cause of transient BAPL behavior theoretically and explained how AS path loops would form and resolve as well as the duration. This paper believes that routing updates are slowed down by delays, owing to physical constraints and protocol mechanisms. And therefore, the inconsistent routing information on different nodes leads to AS path loops during convergence, which depends on the ability of each node to choose alternative path without loops. The Minimum Route Advertisement Interval (MRAI) is the main factor of the duration of transient AS path loops.

R. Mahajan et al. presented an example of BAPL: a key AS of Internet introduced BAPL intentionally to achieve some strategies, while this behavior was unnecessary for most operators of BGP routers [12]. Similarly, X. Shi et al. also introduced an instance for BAPL caused by intentional configuration: University of Washington and Georgia Institute of Technology carried out a rerouting experiment which applied to AS 47065 and led to BAPLs [1], [2], [4].

Different from previous work, this paper focuses on the distribution and causes of both *transient* and *persistent* BAPL behaviors. In order to avoid the disadvantage of measurement using only traceroute, we have carried out our measurement study on BAPL using both *RouteViews* [15] and *traceroute* [24].

III. DATA SETS AND METHODOLOGY

The BGP AS path, also called *signaling* AS path, is the list of ASes that propagate the BGP update messages, while the *forwarding* AS path denotes the list of ASes through which

data packets actually traverse. For example, as Figure 1 shows, AS0, which has a destination p , propagates the BGP update message to AS3 through AS1 and AS2, then (AS3, AS2, AS1, AS0) is AS3's BGP signaling AS path to destination p . On the other hand, AS3 forwards packets to destination p in AS0 along path of AS3, AS2, AS1, AS0, then (AS3, AS2, AS1, AS0) is AS3's forwarding AS path to destination p . However, these two types of AS paths are not always identical due to various reasons, such as route aggregation/filtering and forwarding anomalies [7], [9].

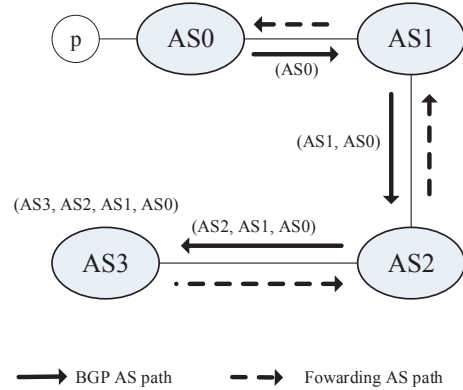


Fig. 1. An example of BGP AS path and forwarding AS path

To collect the forwarding AS paths, we employ the traceroute [24], which is widely used to observe routing problems and discover the underlying network topology. In traceroute, the interfaces on a forwarding path are identified and the round-trip time statistics for each hop along the way are reported. It is considered the only effective way to observe how packets pass through the Internet under the circumstance of no access to private routing data. For better understanding the connection between BAPL and forwarding looping behavior, we follow the methodology presented in [7] and measure the *signaling* AS path and *forwarding* AS path at the same time. When a *forwarding* AS loop is identical to the *signaling* AS path loop, we consider that the *forwarding* AS loop is attributed to BAPL.

Each BGP node announces its best paths to all destinations to its neighbors, and records the most recent paths received from all of its neighbors. BGP advertises the route to each destination only once, and sends subsequent updates only upon route changes. Consecutive updates for the same destination are spaced out by M seconds (default value 30) using a MRAI. When a current path to a destination is noticed to be no longer available, the BGP router will attempt to find an alternative path by checking all the saved paths it learned from its neighbors previously. If there is no alternative path, it will send an explicit path withdrawal message to its neighbors.

To obtain the signaling AS path used in this paper, we collect data from the publicly available Oregon RouteViews route-views4 collector [15], which gathers BGP data from its geographically distributed AS peers (sometimes also called *monitors*), for both IPv4 and IPv6.

The route-views4 collector dumps snapshots of the BGP routing table (RIB) for each of its peers every two hours in the Multi-threaded Routing Toolkit (MRT) [23] format. In addition, the collector receives BGP routing updates from its peers, and writes the collected BGP routing updates into files every 15 minutes in the MRT format [3]. BGP RIB and updates both contain attributes such as timestamp, peer IP, peer AS, prefix, AS-PATH, origin AS. Among these attributes, the AS-PATH attribute is the *signaling* AS path, and we use it to analyze BAPLs. The timestamp in the RIB is the time when the snapshot is dumped, while the timestamp in the update is the time when the update is received from a peer.

We collect the RIB data at 01/01/2010 00:00:00 and BGP update data from RouteViews in 1456 days from 01/01/2010 to 12/31/2013¹. Based on the RIB data and the update data, we obtained the routing table at anytime during the period. When a new update appears, a corresponding record will be added to the routing table. On the other hand, a record may be removed from the routing table as a result of a withdrawal or a different update.

IV. WILL BAPL LEAD TO FORWARDING LOOPS?

In general, a packet from the source traverses a sequence of routers to reach the destination. A packet experiences a forwarding loop if it traverses a set of routers more than once. Studies have shown that forwarding loops can cause packets in the loops with higher loss rate and longer delay. For other packets that traverse one or more links in the loop, they could have longer delay and higher jitters due to the resource consumption caused by the looping packets [11], [13]. Such a vulnerability can be exploited by attackers to overload the shared links for disrupting the Internet connectivity to some victim destination addresses or prefixes [10].

[5], [7] have shown multi-AS forwarding loops existed in the Internet. Will BAPL contribute to multi-AS forwarding loops? Suppose that BAPL may lead to inter-domain forwarding loops, then analyzing the distributions and explanations of BAPL behavior will help to prevent part of forwarding looping, thus reduce packets loss rate, hold back attackers from disrupting the Internet, and decrease link utilization and corresponding delay.

We conducted some case studies to analyze whether BAPLs observed by us can actually cause forwarding AS path loops or not. Among the observed BAPLs, we tried to find a RouterView peer AS who had a looking glass router which allowed us to run traceroute towards the destination prefix. For example, on 09/08/2013, with RouteViews [15], we observed a signaling AS path (AS1299, AS6453, AS577, AS7788, AS6407, AS7788) destined for prefix 64.26.148.0/24 in the RIB entry for the monitor 80.91.255.62 (from AS 1299), and this BAPL lasted more than a few days. The traceroute [24] resulted from 80.91.255.62 (which happened to be looking glass) to 64.26.148.28 (an IP address in the destination prefix)

¹Except the five-day period of 07/14/2011 to 07/18/2011 when no data was recorded by the route-views4 collector

TABLE I
AN EXAMPLE OF TRACES THAT CONTAINS FORWARDING LOOPS

Hop	Router address	AS number
1	213.155.133.147	1299
2	213.155.133.142	1299
3	213.155.130.51	1299
4	80.91.249.29	1299
5	213.155.131.139	1299
6	213.248.100.178	1299
7	63.243.128.42	6453
8	64.86.85.1	6453
9	216.6.87.9	6453
10	216.6.98.58	6453
11	64.86.85.1	6453
12	216.6.98.58	6453
13	67.69.218.3	577
14	209.217.64.37	7788
15	206.191.0.89	7788
16	67.230.128.70	6407
17	209.217.64.37	7788
18	206.191.0.89	7788
19	67.230.128.70	6407
20	209.217.64.37	7788
21	206.191.0.89	7788
22	67.230.128.70	6407
...

witnessed a forwarding loop as shown in Table I. Using the method introduced in [7], we converted the router-level forwarding path into forwarding AS path, which turned out to be identical to the signaling AS path. In particular, the forwarding AS loop was identical to the BGP AS path loop.

Moreover, we have repeated above experiments and found only 1% of the signaling loop accounted for forwarding loops. Note that this percentage might be biased because we have only sampled the signaling loops which we could use looking glass to run traceroute. Nevertheless, our results show that BAPL behavior indeed could cause inter-domain forwarding loops. This observation motivates us to carry out more in-depth researches on BAPL behavior in the rest of the paper.

V. MEASUREMENT RESULTS

Studies have shown that BAPLs exist in the Internet [1], [2], [4], [12], but the scale of BAPLs in IPv4 and IPv6 remains unclear. We used to believe that all of BAPLs are caused by misconfigurations, while [5] observed a great many persistent forwarding AS path loops which may be caused by persistent BAPLs. Moreover, we want to know the distribution among the duration of BAPLs and the explanations for persistent BAPLs. Because BAPL may lead to forwarding AS path looping, and the loop length is important for an attacker to amplify the traffic in the forwarding links, the distribution of BAPL loop length is also studied in this paper.

A. Total Number and Ratio of BAPLs

We call the BGP update with a BGP AS path loop a BAPL. With the daily BGP update data described in Section III, the number of BAPLs per day is counted.

Fig. 2 (a) shows the number of BAPLs, and the ratio of *the number for BAPLs to the number of all of the BGP updates* collected by RouteViews in IPv4 on a daily basis from

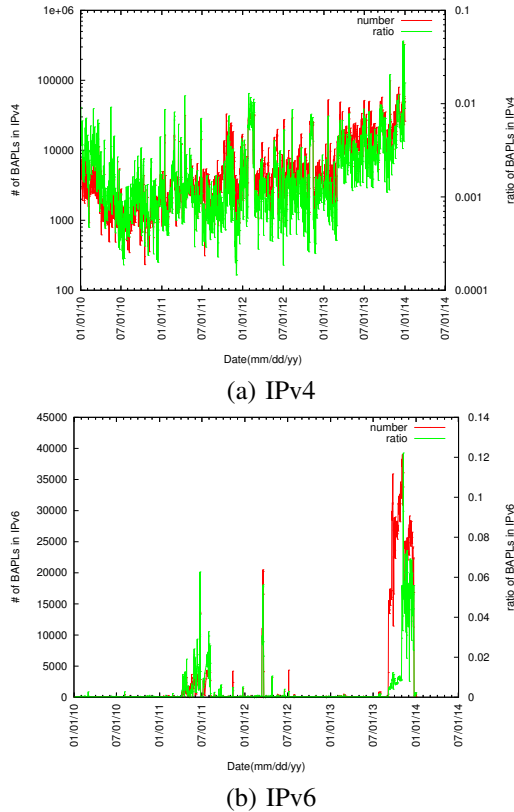


Fig. 2. The number and the ratio of BAPLs

01/01/2010 to 12/31/2013. Fig. 2 (b) shows the number and ratio of BAPLs in IPv6. Overall, 5973568 BAPLs have been observed in IPv4 and 1440104 in IPv6 during 1456 days.

The medians of the number and ratio of BAPLs for each year in IPv4 and IPv6 are listed in Table II. The median of the number for each year is the median number of the set of BAPLs numbers per day, and the median of the ratio for a certain year is the median ratio of the set of BAPLs ratios on a daily basis. In IPv4, the number of BAPLs increased dramatically from 2010 to 2013. Due to the explosion of global BGP routing table, the ratio of BAPLs kept stable in 2011 and 2012, and witnessed rapid growth in 2013. While in IPv6, the number of BAPLs increased in 2011, decreased in 2012, and stayed stable in 2013, so as the ratio.

The scale of deployment of IPv6 is much smaller than IPv4, and most of the facilities in IPv6 are deployed later than that in IPv4. Incidents like faulty configurations, malicious attacks, and other potential causes discussed in Section VI occur much less frequently in IPv6 than that in IPv4. As a result, as we can see from the above, the number and ratio of BAPL in IPv6 is much smaller than that in IPv4.

B. Duration of BAPLs

We also have studied the duration of BAPL with the BGP RIB data and updates data (defined in Section III). As described above, a BGP entry can be removed from the routing table due to a withdrawal or a new update. We define the duration of a BAPL as the time interval between

TABLE II
MEDIAN OF BAPLs PER YEAR

Year	Number of IPv4	Ratio of IPv4	Number of IPv6	Ratio of IPv6
2010	1580	1.03×10^{-3}	0	0
2011	2870.5	9.23×10^{-4}	21.5	7.65×10^{-5}
2012	4603	1.08×10^{-3}	14	5.01×10^{-5}
2013	15808	2.94×10^{-3}	21	5.07×10^{-5}

its announcement and its withdrawal or a new replacement announcement without the same AS path. For a BGP RIB entry, the period from 01/01/2010 00:00:00 to the time when it is withdrawn or replaced with a different update is the duration.

The persistency of the BAPLs is studied. Fig. 3 shows the Complementary Cumulative Distribution Functions (CCDFs) of the distribution of duration for BAPLs in IPv4 and IPv6. That is, 5473962 out of 5973568 (91.64%) BAPLs in IPv4 and 1438730 out of 1440104 (99.90%) BAPLs in IPv6 last shorter than one day. These short-lived BAPLs could be explained with configuration faults or malicious attacks. Excluding the BAPLs that last shorter than 1 day, the average duration is 13.11 days in IPv4 and 23.41 days in IPv6. Considering that many other short-lived BAPLs could also be attributed to faulty configurations or intentional attacks, we take the data set which contains only BAPLs lasting longer than 9 days into account (135426 BAPLs are involved in IPv4 and 464 in IPv6). The average duration for these BAPLs is 37.70 days in IPv4 and 60.97 days in IPv6. Table III lists the averages duration of BAPLs from different data sets. In IPv4, the longest duration is 1456 days, while it is 1102 days in IPv6.




Fig. 3. CCDFs of BAPLs duration

The results have astonished us, because we once expected that fault configuration was the only factor that led to BAPL behavior. Were it true, BAPL should last shorter than what is observed above. We will discuss this later in Section VI.

C. Loop Length of BAPLs

As discussed above, BAPLs may lead to multi-AS forwarding loops. AS path loops and forwarding paths may share one or more links to the destination prefixes or addresses. An attacker can use BAPLs to overload the shared links to

TABLE III
AVERAGES BAPLS DURATION

Measured data set	Averages in IPv4(days)	Averages in IPv6(days)
Longer than 0 day	9.91	4.72
Longer than 1 days	13.11	23.41
Longer than 9 days	37.70	60.97
Longer than 29 days	87.89	131.26
Longer than 89 days	224.76	322.11




Fig. 4. Loop length of BAPL

interrupt the connectivity with those reachable prefixes or addresses [10].

The length of an AS path loop is important for the traffic amplification in the links. When a packet enters an AS path loop, it is possible that the packet traverses the links in the loop several times before its TTL expires. Obviously, the shorter the loop length is, the more times the packet will spend to traverse the links in the loop. Since the BGP AS path vector from p_n to p_1 is $asp = (p_n, p_{n-1}, \dots, p_1)$, by definition, $p_i = p_j, j > i, j - i \neq$ for a BAPL, and the loop length of asp is $j - i$. Fig. 4 (a) shows the loop length distribution of BAPLs in IPv4 each year, and Fig. 4 (b) shows the distribution in IPv6. The number of BAPLs of AS path loop length l for a certain year denotes the number of different RIB entries or updates that contain looped AS path with the loop length l in the year. It is obvious that the bulk of BAPLs had 2-hop or 3-hop loops, for both IPv4 and IPv6, which makes it easy to amplify the amount of traffic remarkably to destination addresses in the links that appear in the loops.

As RFC 4271 [16] describes, BAPL should not occur in any case, but our observation shows the large scale of BAPLs in both IPv4 and IPv6. We have once believed that misconfiguration was the only explanation of BAPL, and all of BAPLs should be transient, while our observation presents that quite a number of BAPLs last longer than one day. Furthermore, most of BAPLs have a 2-hop or 3-hop loop, which is easily exploited by attackers for overloading the links. The frequency and duration of BAPLs that we observed is surprising.

VI. EXPLANATIONS OF BAPL

There are a few possible causes of BAPLs, such as transnational enterprises, private AS number leaking, preventing particular AS from accepting routes, faulty configuration, and intentional attacks.

A. Private AS Number Leaking

The Internet Assigned Numbers Authority (IANA) has reserved the AS numbers (65512 - 65535) for private use and private AS numbers should not be advertised on the global Internet [18]. However, we have observed a large number of AS paths which contain private AS numbers. As explained in [7], when a customer who uses a private AS number mistakenly leaks BGP routes learned from one upstream provider to another, an AS path containing private AS number may arise. Some private AS number leaking events even account for BAPLs, which is quite beyond our expectations. As described in the definition of *asp*, the context where BAPLs are caused by private AS number leaking can be described as $p_i = p_j, j > i, j - i \neq, \forall m \in (i, j), p_m \in [65512, 65535]$.

Specifically, in our observation, 105340 out of 5973568 (1.76%) BAPLs are definitely caused by AS number leaking in IPv4, but AS number leaking only account for 39 out of 1440104 (0.0027%) BAPLs in IPv6.




Fig. 5. Private AS number leaking

If a customer AS is requested to communicate with a single provider AS using BGP, it can use a private AS number. This should not happen unless the routing policy between the provider AS and the customer AS is not visible in the Internet.

As Fig. 5 shows an example of why private AS number leaking onto the Internet might lead to BAPL. ASp communicates with its single provider ASa using a private AS number for load balancing purpose. Normally, the private AS number should not be advertised to the Internet. When the update about a prefix in ASx propagates to ASa, ASa will forward it to ASp. ASp somehow (probably due to misconfiguration) propagates the path (ASp, ASa, ASx) back to ASa, and somehow (due to another misconfiguration) ASa just accepts the paths and

further propagates the path to ASy. As a result, the path (ASy, ASa, ASp, ASa, ASx) with private AS number and loop are leaked onto the Internet.

B. Multinational Companies




Fig. 6. Multinational Companies

Some multinational companies (e.g. AS number X) have exchange-points all over the world, and several exchange points may share the same AS number X , where operators configure their routers to accept routes whose AS-PATH attributes containing their own AS number X [17]. When BGP routing updates pass through exchange points with the same AS located in different countries, they may also go through one or more ASes among the exchange points. It appears as if the BGP updates loop in the AS path from the BGP perspective. For example, the NTT Communications Corporation [25] has exchange-points in Frankfurt, Tokyo, and some cities in the USA, which share the same AS number, 2914. When the prefix in ASx propagates the BGP updates to ASy, as Fig. 6 illustrates, the message will pass through the exchange point of NTT in Chicago and Frankfurt. If the BGP router of the exchange point in Frankfurt computes the degree of preference of the route based on preconfigured policy information, and does not discard the routing updates the AS-PATH attributes of which contain AS2914, a BAPL (ASy, AS2914, ASb, ..., ASa, AS2914, ASx) from the BGP perspective will occur.

C. Preventing Particular AS from Accepting Routes

Some BGP operators of AS X might prepend another AS's number Y so that Y will not pick up the routes from X . For example, the BGP operator of AS 3066 wanted to send routes to Sprint (AS1239) [26], but did not want the routes to be picked up by UUnet/Verizon Business (AS701) [27]. Then the path (AS3066, AS701) was prepended to the AS path, and the path (AS3066, AS701, AS3066) was sent to Sprint AS1239 [22], which accepted the route. As a result, a BAPL (AS1239, AS3066, AS701, AS3066) was propagated on the Internet. However, when the BGP routing updates containing the AS path vector (AS1239, AS3066, AS701, AS3066) arrived to AS701, AS701 would discard the message immediately, and no relative traffic towards AS30166 would traverse AS701. Note that the BAPL was artificially injected by the operator of AS 3066, which would not lead to any forwarding loop.

Similarly, on 08/18/2011, the University of Washington and the Georgia Institute of Technology conducted a rerouting experiment which applied to AS47065 [1], [2], [4]. In this experiment, a looped AS path (47065, x , 47065) for prefix

184.164.255.0/24 was announced, so that ASx could not accept this route later, and related traffic would not pass through ASx. Obviously, operator's such prepending configuration on BGP routers can lead to BAPLs, but these BAPLs do not account for any forwarding loop.

D. Faulty Configurations or Malicious Attacks

BAPLs can also arise when a BGP router incorrectly accept routing updates whose AS-PATH attributes contain the local AS number. This could occur due to configuration errors or even malicious attacks.

Argus [6] is an agile system to detect prefix hijacking and other anomalies which are caused by misconfigurations or malicious attacks, and starts to collect data since 06/01/2011. After cross checking the RouteViews data described in Section III and the data collected from Argus, we have found that in IPv4, at least 170036 out of 5973568 (2.85%) BAPLs were associated with prefix hijacking or other routing anomalies from 06/01/2011 to 12/31/2013. These prefix hijacking and routing anomalies can be attributed to faulty configurations or intentional attacks, which means that *at least* 2.85% of BAPLs were caused by misconfigurations or malicious attacks.

In summary, several valid explanations can contribute to BAPLs, such as multinational cooperation and preventing particular AS from accepting routes, while other BAPLs can be attributed to invalid reasons, such as misconfigurations and intentional attacks (contributed to at least 2.85% of BAPLs in IPv4), private AS number leaking (cause 1.76% of BAPLs in IPv4 and 0.00027% in IPv6), .

VII. CONCLUSION

The BAPLs studied in this paper can be helpful in understanding the operational behavior of BGP in both IPv4 and IPv6. As a motivation, at first, we try to explore the relationship between BAPL behavior and forwarding looping, but we find that only a small part (about 1%) of BAPLs can lead to forwarding loops. We have studied the global BGP routing data in 1456 days and analyzed the number and ratio of BAPLs in IPv4 and IPv6. In addition, the duration of BAPLs, and the distribution of loop length are also discussed in this paper. Different from our initial expectations, we find that nontrivial number of BAPLs lasts more than a month.

What's more, we have attributed BAPLs to various reasons. Private AS number leaking contributed to about 1.76% BAPLs in IPv4 and 0.00027% BAPLs in IPv6, and *at least* 2.85% of BAPLs in IPv4 are caused by malicious attacks or faulty configurations. Reasonable explanations, including multinational cooperations and preventing particular AS from accepting routes, also have contributed to BAPLs. BAPLs caused by invalid reasons like private AS number leaking, misconfigurations and intentional attacks should be fixed.

As future work, we plan to focus on the correlation between BAPL behaviors and other BGP anomalies. As we are surprised by the frequency that private AS numbers appear, we plan to explore reasons why private AS numbers are leaked to the Internet.

ACKNOWLEDGMENT

The work was supported by the National Natural Science Foundation of China under Grant No. 61161140454, the National Key Basic Research Program of China (973 program) under Grant 2013CB329105 and 2009CB320500.

REFERENCES

- [1] U. Javad, I. Cunha, D. R. Choffnes, E. Katz-Bassett, T. Anderson and A. Krishnamurthy, "PoiRoot: Investigating the Root Cause of Interdomain Path Changes". [C]. In *Proceedings of SIGCOMM 2013*, pp. 183-194, Hongkong, China, Aug, 2013.
- [2] E. Katz-Bassett, C. Scott, D.R. Choffnes, I. Cunha, V. Valancius, N. Feamster, H.V. Madhyastha, T.E. Anderson, and A. Krishnamurthy. "LIFEGUARD: practical repair of persistent route failures". [C]. In *Proceedings of SIGCOMM 2012*, pp. 42(4):395-406, Helsinki, Finland, Aug, 2012.
- [3] P. Cheng, X. Zhao, B. Zhang and L. Zhang, "Longitudinal Study of BGP Monitor Session Failures". [J]. *ACM SIGCOMM Computer Communication Review*, 40(2):34-42, 2010.
- [4] X. Shi, Y. Xiang, Z. Wang, X. Yin and J. Wu. "Detecting Prefix Hijackings in the Internet with Argus". [C]. In *Proceedings of the 2012 ACM Conference on Internet Measurement Conference (IMC)*, Boston, USA, Nov, 2012.
- [5] J. Xia, L. Gao, T. Fei, "A measurement study of persistent forwarding loops on the Internet". [J]. *Computer Networks*, 51:4780-4796, 2007.
- [6] Y. Xiang, Z. Wang, X. Yin, and J. Wu. "Argus: An accurate and agile system to detecting IP prefix hijacking". [C]. In *19th IEEE International Conference on Network Protocols (ICNP)*, Vancouver, Canada, Oct, 2011.
- [7] Z. M. Mao, J. Rexford, J. Wang and R. H. Katz, "Towards an accurate AS-level traceroute tool". [C]. In *Proceedings of SIGCOMM 2003*, pp. 365-378, Karlsruhe, Germany, Aug, 2003.
- [8] D. Pei, X. Zhao, D. Massey, and L. Zhang. "A Study of BGP Path Vector Route Looping Behavior". [C]. In *Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS)*, pp. 720-729, Tokyo, Japan, 2004.
- [9] T. G. Griffin, G. Wilfong, "On the correctness of IBGP configuration". [C]. In *Proceedings of SIGCOMM 2002*, pp. 17-29, Pittsburgh, PA, Aug, 2002.
- [10] J. Xia, L. Gao, T. Fei, "Flooding attacks by exploiting persistent forwarding loops". [C]. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, Berkeley, CA, USA, Oct, 2005.
- [11] U. Hengartner, S. Moon, R. Mortier and C. Diot. "Detection and analysis of routing loops in packet traces". [C]. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement*, Pittsburgh, PA, Aug, 2002.
- [12] R. Mahajan, D. Wetherall, and T. Anderson. "Understanding BGP Misconfiguration". [C]. In *Proceedings of SIGCOMM 2002*, pages 3-16, Pittsburgh, PA, Aug, 2002.
- [13] D. Pei, L. Wang, D. Massey, S. Wu and L. Zhang. "A Study of Packet Delivery Performance during Routing Convergence". [C]. In *Proceedings of 2003 International Conference on Dependable Systems and Networks (DSN)*, San Francisco, USA, Jun, 2003.
- [14] V. Paxson, "End-to-end routing behavior in the Internet". [J]. *IEEE/ACM Transactions on Networking*, 5(5): 610-615, Oct, 1997.
- [15] "University of Oregon Route Views Project". [Online]. Available: <http://www.routeviews.org/>
- [16] Y. Rekhter and T. Li, "RFC 4271: A Border Gateway Protocol 4 (BGP-4)", Jan. 2006. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc4271.txt>.
- [17] P. Marques, "RFC 6368 Internal BGP as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)", Sep. 2011. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc6368.txt>.
- [18] J. Hawkinson, "RFC 1930 Guidelines for creation, selection, and registration of an Autonomous System (AS)", Mar. 1996. [Online]. Available: <http://tools.ietf.org/html/rfc1930>.
- [19] J. Moy, "RFC 2328: OSPF Version 2". Apr. 1998. [Online]. Available: <http://www.ietf.org/rfc/rfc2328>.
- [20] G. Malkin, "RFC 2453: RIP Version 2". Nov. 1998. [Online]. Available: <http://www.ietf.org/rfc/rfc2453.txt>.
- [21] D. Oran, "RFC 1142: OSI IS-IS Intra-domain Routing Protocol". Feb. 1990. [Online]. Available: <http://tools.ietf.org/rfc/rfc1195.txt>.
- [22] J. Chandrasekar, "AS Path Loops in practice ?". NANOG mailing list, msg00255, Dec.8, 2003.
- [23] L. Blunk and M. Karir, "RFC 6396: Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format". Oct. 2011. [Online]. Available: <http://tools.ietf.org/html/rfc6396>.
- [24] "traceroute.org". [Online]. Available: <http://www.traceroute.org>
- [25] "NTT communications". [Online]. Available: <http://www.us.ntt.com/en/index.html>
- [26] "Sprint". [Online]. Available: <http://www.sprint.com/>
- [27] "Verizon". [Online]. Available: <http://www.verizonenterprise.com/>