

Understanding BGP Next-hop Diversity

Jaeyoung Choi*, Jong Han Park†, Pei-chun Cheng†, Dorian Kim‡, Lixia Zhang†

*Seoul National University, jychoi@mmlab.snu.ac.kr

†University of California, Los Angeles, {jpark,pccheng,lixia}@cs.ucla.edu

‡NTT Communications Inc., dorian@blackrose.org

Abstract—The Internet topological connectivity becomes denser over time. However the de facto routing protocol of the global Internet, BGP, lets each BGP router select and propagate only a *single* best path to each destination network. This leads to a common concern that the rich connectivity is not fully utilized and the lack of alternative paths can reduce a network’s robustness to failures as well as flexibility in traffic engineering, and can lead to slow adaptation to topological changes. Yet there have been few quantitative measurement studies on path diversity in today’s operational Internet. In this paper we use iBGP routing data collected from a Tier1 ISP, ISP_A , over a 2-year time period to quantify BGP next-hop diversity for all destinations. Our results show that ISP_A reaches the majority of prefixes through multiple next-hop routers. We use several case studies of prefixes with different diversity degrees to identify two major factors that impact the number of observed next-hops: the ISP’s path preference and the number of peering routers between large ISPs. This observation provides operational input to the current efforts on augmenting BGP to increase path diversity.

Index Terms—BGP; Next-hop Diversity; Tier1 ISP; Measurement

I. INTRODUCTION

Although a BGP router may learn multiple paths from its peers for a given destination, the BGP specification requires the router to select and propagate only one single best path. As the topological connectivity of the Internet grows denser over time [11], it becomes increasingly desirable to fully utilize multiple available paths. However, previous studies on BGP path diversity mostly focus on measuring the *AS level* path diversity and its impact [9, 3, 17, 5]. They do not fully address practical questions raised by ISPs, whose routine operations mainly deal with and are confined to the diversity of paths to reach their directly connected neighbor ASes, which we coin as *next-hop diversity* in this work.

Currently in IETF, the operation community expresses avid desire to increase the next-hop diversity, as it represents the opportunities in fast failure recovery, traffic engineering, and load balancing. As a result, several modifications to BGP have been proposed to allow BGP routers to propagate multiple paths for the same destination [12, 10, 7, 16].

Nevertheless, despite the promising effort on *adding* diversity, there has been little understanding on the more fundamental question: what is the *existing* next-hop diversity in the operation networks? Knowing the existing next-hop diversity is of significance as it can help us better understand the actual operational needs, and can shed light on important operational practices that influence the degree of next-hop diversity.

In this work, we define and measure the next-hop diversity as observed from all the *backbone routers* in a Tier1 ISP (ISP_A) for *all prefixes in the global routing table*. Our findings can be summarized as follows:

- The majority of prefixes can be reached via multiple next-hop routers: without any modifications to BGP, more than 88% and 78% of all prefixes can be reached via more than 2 and 5 next-hop routers respectively (Section III).
- The two major factors that impact the next-hop diversity are the ISP’s path preference and the number of peering routers with its neighbor ISPs. In addition, the prefixes with a very high next-hop diversity are mostly caused by the lack of geo-presence of ISP_A in some regions (Section IV).
- Although the overall next-hop diversity has not changed significantly over the past two years, the maximal next-hop diversity has shown noticeable increase due to the gradually increased number of routers inside the Tier1 ISP (Section V).

II. METHODOLOGY

In this section, we first present the high level network topology of ISP_A . Then, we describe details in measuring path diversities.

A. ISP_A ’s Topology and Next-hop Diversity

ISP_A is a Tier1 ISP who uses AS confederations [14]. In particular, ISP_A maintains one subAS ($subAS_1$ shown in Figure 1) as its backbone network, which consists of more than one hundred routers that are distributed globally across 14 countries and 3 continents. The routers in the backbone network connect to tens of other smaller subASes within ISP_A using confederation BGP sessions as well as routers of neighbor ASes using eBGP sessions as depicted in Figure 1. In this paper, we refer the routers in the backbone subAS as *backbone routers* (e.g., R_{A1} , R_{A2}).

Between ISP_A and its neighbor ASes, two routers maintaining an eBGP session are located in the same POP¹ (e.g., R_{A1} and R_{11} in POP1), and there are typically multiple eBGP sessions between ISP_A and its neighbor AS for redundancy. For example in Figure 1, ISP_A maintains two eBGP sessions with AS_1 , using R_{A1} - R_{11} and R_{A2} - R_{12} across two different POPs. This path redundancy is useful not only in avoiding a single point of failure, but also in faster failure recovery.

¹Point of Presence

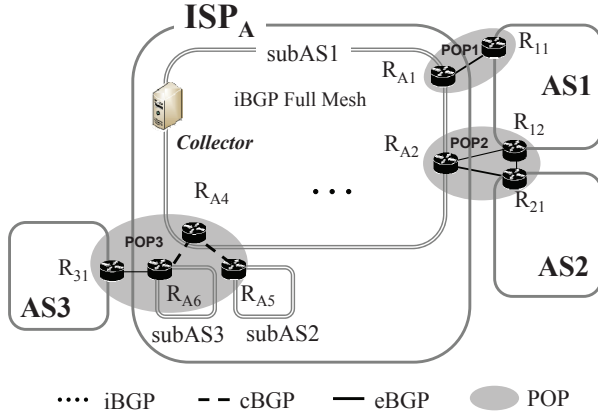


Fig. 1. High Level Topology of ISP_A

Assume that R_{A4} sees two BGP paths to reach a prefix in AS_1 , through R_{A1} and R_{A2} . When R_{A4} notices that R_{A1} or any internal path from R_{A4} to R_{A1} failed, it can immediately failover to the path through R_{A2} without waiting for BGP to converge. Even in the absence of failures, these redundant paths can be used to optimize network-wide requirements, such as load balancing and traffic engineering. If R_{A1} is overloaded (or R_{A2} is underutilized), the operator may be able to reduce the load (or increase the load in R_{A2}) by shifting the data traffic to R_{A2} . All of the above benefits are only possible when there exist multiple paths to a given destination.

For a given network destination (*i.e.*, prefix), BGP announcement messages can be received from multiple neighbor ASes, POPs, and routers, which leads to different granularity of path diversity. In this work, we define *next-hop ASes*, *next-hop POPs*, and *next-hop routers* as the distinct number of neighbor ASes, POPs, routers to reach a specific prefix. For example, in Figure 1, suppose that ISP_A receives the reachability information to prefix p through AS_1 . Then, the next-hop diversity would be 1 next-hop AS (AS_1), 2 next-hop POPs (POP_1 and POP_2), and 2 next-hop routers (R_{11} and R_{12}) respectively.

B. Measuring Path Diversity

ISP_A deploys a BGP data collector in the backbone subAS as depicted in Figure 1. The collector is an iBGP speaker and maintains iBGP peering sessions with all the other routers in in the backbone subAS to passively record all iBGP updates received. The collected iBGP update messages and the routing table snapshots of each peer are periodically stored to files in MRT [4] format. We used *bgpparser* [1] to convert MRT binary format to text (XML) format. Then, for each entry in routing table snapshots (RIBs) gathered from all iBGP peers, we extracted NEXT_HOP and AS_PATH attributes and counted how many unique next-hop routers (*i.e.*, next-hop router in NEXT_HOP attribute) along with their geographical locations (*i.e.*, geographical location of the next-hop router) and next-hop ASes (*i.e.*, the first AS appearing in AS_PATH attribute) are visible to the collector for each destination.

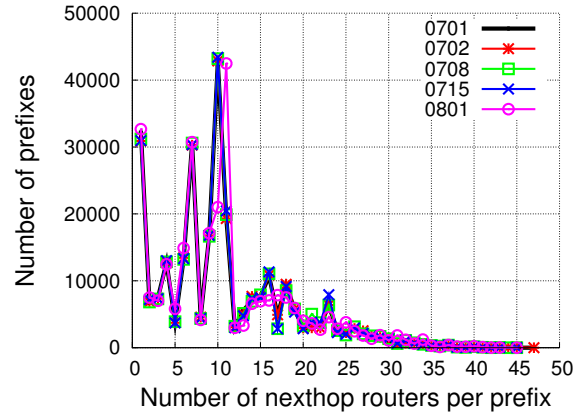


Fig. 2. Verifying Representativeness of Dataset

Note that ISP_A does not use *next-hop-self* option. Therefore, NEXT_HOP attribute contains the IP address of the router residing in the neighbor AS. Since all routers in the backbone are connected in a full-mesh, the number of visible next-hops observed by each backbone router should be the same as that observed by the collector.

In this measurement study we exclude two types of prefixes: internal prefixes and potential bogon prefixes. Internal prefixes are meant to be used only inside ISP_A . Since the goal of our measurement is to understand the path diversity of commonly visible prefixes to all ASes in the Internet, we filter out such internal prefixes. In addition, we exclude the prefixes longer than 24 bits because the BGP messages containing reachability information on these prefixes are generally filtered out by BGP routers, and can lead to inaccurate results.

To quantify the next-hop diversity, we use the routing table snapshots taken from all backbone routers on July 1st, 2009. To ensure that the snapshots are representative, we also measured next-hop diversity using routing table snapshots taken at different times. Figure 2 compares the number of unique next-hop routers from routing table snapshots on July 1st, 2009 with those from four other snapshots on July 2nd (one day after), 8th (one week), 15th (two weeks), and August 1st (one month) respectively. As depicted in Figure 2, the distribution of the number of next-hop routers for a given prefix from the five snapshots are very similar. In addition, we checked that the total number of prefixes in each snapshot and the set of unique neighbor ASes are roughly the same. Note that we also performed the same measurements described in this paper on all the dates shown in Figure 2, however due to space limit, in the following sections, we only present the results on July 1st.

III. QUANTIFYING NEXT-HOP DIVERSITY

In this section, we start our work by quantifying the next-hop diversity of all prefixes in the global routing table observed by ISP_A with three different granularities: next-hop ASes, POPs, and routers. Then, we focus on characterizing and analyzing the router level diversity. In the following sections,

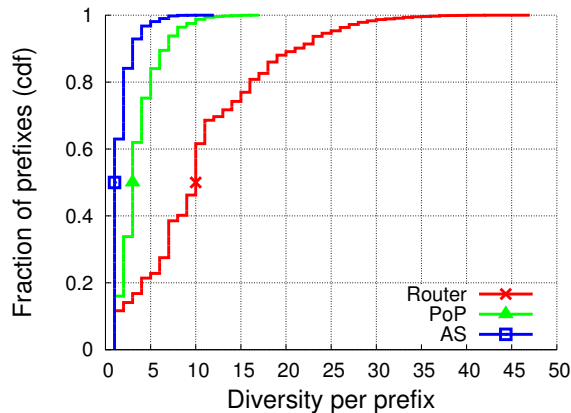


Fig. 3. Distribution of Next-hop Diversity

without further specification, we use *next-hop* and *next-hop diversity* to refer to the next-hop router and the router level diversity respectively.

A. Next-hop ASes

We first measure, for each prefix, how many next-hop (*i.e.*, neighboring) ASes can be used to reach a given network destination. Figure 3 shows the cumulative distribution (CDF) of the number of next-hop ASes to reach a prefix. For the total 276,712 prefixes, we observe that about 62% of all prefixes are reached via 1 neighbor, and almost all prefixes (about 96%) can be reached via less than or equal to 5 neighboring ASes.

Note that the number of next-hop ASes represents a gross diversity at the inter-domain routing level. For those prefixes that can only be reached through one neighbor, ISP_A must wait for BGP to explore and settle down on the routes via other neighbors (if there is any) when the particular neighbor AS fails. The prolonged convergence delay in this case can potentially degrade the performance in the data plane [18]. However, such number of next-hop ASes only describe an abstract reachability at the logical AS level. In a typical operation settings, two ASes often set up peering sessions at different geographical locations using multiple BGP routers as explained earlier.

B. Next-hop POPs and Routers

We further measure the number of available next-hop routers and their geographical locations (*i.e.*, POPs) to reach a given destination. Figure 3 shows the distribution of the number of observed next-hop routers and POPs to reach each destination prefix. We observe that even though 18% of prefixes can still be reached via only one POP from one neighboring AS, the majority of the prefixes can be reached via 2 to 5 POPs. Furthermore, given that there often exist multiple routers in a given POP, the next-hop diversity is further amplified and varies widely from 1 up to 47. Most of the prefixes (88%) have more than 2 next-hops, and around 47% of all prefixes have their next-hop diversity between 6 and 12. We also observe that there exists a small fraction of prefixes (1.6%) with a very high next-hop diversity (≥ 30).

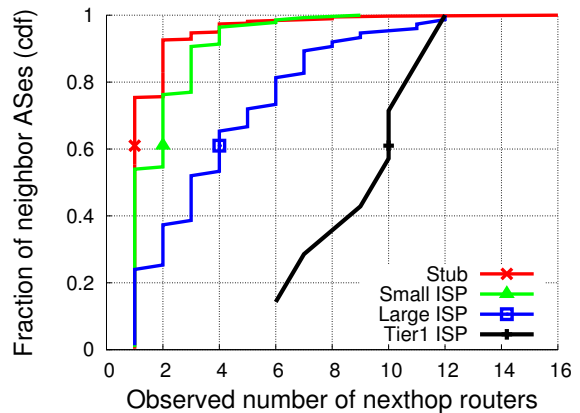


Fig. 4. Observed Connectivity of Different Neighbor Types

In Figure 3, we observe that prefixes with the same next-hop AS diversity can have different next-hop POP and router diversity. This indicates that the amount of visible next-hop diversity can depend on not only the number of neighbor ASes but the number of peering routers with neighbor ASes through which ISP_A reaches a given destination.

To show the number of peering routers differ across different neighbor ASes, we classified each neighbor AS as one of “Tier1”, “Large ISP”, “Small ISP”, and “Stub” based on the classification found in [19]. Then, we measured the number of peering routers for each of the types. Figure 4 shows that, in general larger neighbor ASes tend to have a higher number of routers peering with ISP_A . This tendency is reflected in next-hop diversity. For example, if two prefixes are reached via a Tier1 and a small ISP neighbor respectively, then based on Figure 4, the former prefix can have its next-hop diversity ranging from 6 to 12 while the diversity of the latter prefix can range from 1 to 9. Note that there exist few Stub ASes (*e.g.*, UltraDNS, Amazon, Akamai, etc) whose number of peering routers is exceptionally high. This is due to their specific business needs to provide global wide services, and connect to ISP_A with many routers at different locations globally. The high number of peering routers enables these ASes to increase next-hop diversity in ISP_A by simply announcing their prefixes over multiple peering routers. The existence of highly connected neighbors such as these large stub ASes, large ISPs, and Tier1 ISPs shown in Figure 4 suggests that, by utilizing available connectivity, there exist opportunities in ISP_A ’s current network to increase and exploit the existing path diversity.

IV. MAJOR FACTORS IMPACTING NEXT-HOP DIVERSITY

In this section, we take a closer look at representative cases of prefixes with the low, moderate, and high next-hop diversity to shed lights on the main factors that determine the amount of next-hop diversity for a given prefix.

A. ISP’s Path Preference

In our study, prefixes with lowest next-hop diversity have announced by a neighboring AS through a single next-hop router like prefix 201.133.104.0/24 shown in Figure 5(a).

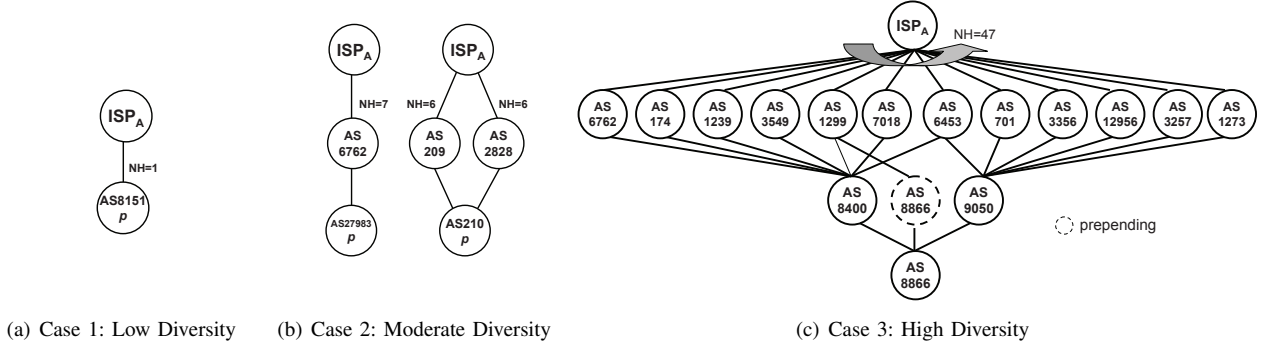


Fig. 5. Representative Cases of Prefixes with Low, Moderate, and High Next-hop Diversity

There can be two reasons: (i) there is only one path to reach this prefix, and/or (ii) BGP’s design choice to select and propagate only the best path to the neighbors prevents ISP_A from being able to see other alternative paths. By further investigating the update messages, we found that the main reason is the latter; when the best (and the only visible) path fails, we could observe that oftentimes many alternative paths, which were hidden previously due to the BGP path selection, got exposed during iBGP convergence process.

Suppose that in Figure 1, AS_2 announces its prefix p to its neighbors, and the reachability information arrives at ISP_A via 3 router level paths: $R_{21}-R_{A2}$, $R_{21}-R_{12}-R_{A2}$, and $R_{21}-R_{12}-R_{11}-R_{A1}-R_{A2}$. After receiving these 3 announcements from its peers, ISP_A will choose the best path based on BGP best selection criteria. In this particular example, ISP_A would choose the path $R_{21}-R_{A2}$ as the best path because of the shortest AS_PATH length, making the other two paths hidden at the receiving border routers (in R_{A1} and R_{A2}). As briefly explained earlier, this is due to the design of BGP to avoid routing loops. As the result, the next-hop diversity for this prefix p would be 1 despite the fact that there do exist other paths via AS_1 .

The above example is also observed in practice. Later in this section, we present a prefix announced by AS8866 in Figure 5(c) as the prefix with the very high next-hop diversity. An interesting observation on this case is that it was able to maintain the high diversity due to the usage of AS path prepending. If AS8866 did not use prepending on its shortest AS path (*i.e.*, AS1299-AS8866), this AS path would have the shortest AS path length (*i.e.*, 2) and will be preferred over other previously visible paths with their AS path length equal to 3 (*e.g.*, AS6762-AS8400-AS8866). As a result, the previously visible path with longer AS paths will be withdrawn and hidden.

The above examples show that how BGP path preference limits the next-hop diversity. However, as the path preference are configurable by design (*i.e.*, via tunable parameters such as $weight^2$, LOCAL_PREF, etc.), a network operator may be able to adjust path preference to achieve higher next-hop diversity

²supported by router vendors

while respecting the network’s routing policy.

B. Number of peering routers

In Figure 5(b), we present two representative cases in moderate diversity. We classify prefixes whose next-hop diversity is between 6 and 12 as moderate, and almost half of prefixes in our measurement belong here as Figure 3 shown. Prefix 190.103.225.0/24 announced by AS27983 is the first case. This prefix can be reached from ISP_A through AS6762, a large ISP. The number of next-hops between ISP_A and AS6762 were 7. Another representative case of a prefix with moderate next-hop diversity was prefix 204.113.217.0/24 announced by AS210. The AS path and next-hop diversity are 2 and 12 respectively. In both examples, the prefixes were reached through at least one neighbor AS which is a large ISP with at least 6 BGP peering sessions with ISP_A like AS1 in Figure 1. From these two cases, we can see that (i) the number of peering routers has an impact on the next-hop diversity; 190.103.225.0/24 announced by AS27983 has moderate diversity because its provider (AS6762) has 7 multiple peering routers with ISP_A . In addition, we also see that (ii) multi-homing helps increase path diversity; 204.113.217.0/24 announced by AS210 has 12 next-hop diversity in ISP_A by multi-homing with AS209 and AS2828.

C. Lack of Geographical Presence

Our last case study explores prefixes with very high degree of next-hop diversity. Figure 5(c) shows a prefix 83.228.80.0/23 announced by AS8866, a regional ISP. AS8866 multi-homes with two providers (AS8400 and AS9050) which connect to many Tier1 and large ISPs. By becoming a customer of these two highly connected providers, prefix 83.228.80.0/23 in AS8866 inherently becomes visible through highly diverse paths from the perspective of ISP_A .

In general, a common characteristic observed in prefixes with high degrees of next-hop diversity is that their origin ASes do not directly connect to ISP_A . From this observation, we hypothesized that the lack of geographical presence of ISP_A can be a factor that determined the set of prefixes with high next-hop diversity. In the regions that ISP_A does not provide connectivity, the origin ASes would connect to

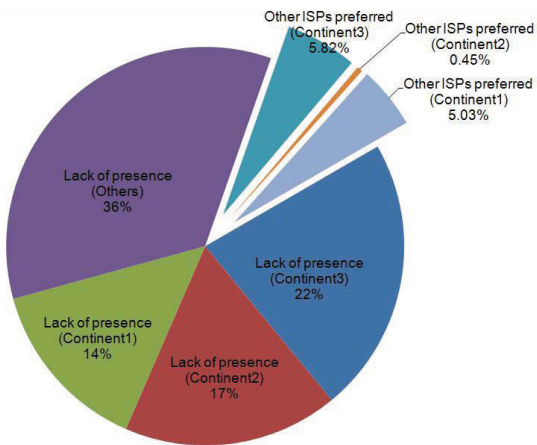


Fig. 6. Geographical Presence of ISP_A

other ISPs when they wish to connect to the Internet. If these local ISPs happen to multi-home with many large ISPs except ISP_A , then there will be many paths with equal AS_PATH length between the origin AS and ISP_A , which leads to the very high next-hop diversity.

To verify our hypothesis, we checked the prefix origination point of prefixes with very high next-hop diversity against the POPs covered by ISP_A . To find the location of prefix origination point, we used MaxMind GeoLite package [2] to map each prefix into a city. Then for these cities, we checked whether any POP of ISP_A is present. Figure 6 verifies our hypothesis; in 89% of prefixes with very high next-hop diversity, ISP_A did not have a presence.

This observation that some prefixes can have a very high diversity regardless of the ISP's intention can be an important input to the proposed BGP modifications [12, 10, 7], which increase the diversity for all prefixes. Our results suggest that more intelligent approaches could be used to utilize the router resources more efficiently by increasing diversity selectively for interested prefixes, rather than over-provisioning these high diversity prefixes altogether.

V. TRENDS OF NEXT-HOP DIVERSITY IN TIME

In this section, we seek to find out if there is a general trend of next-hop diversity changes over time. Due to the large amount of iBGP routing data and the processing loads, we sampled the next-hop diversity of the first day of each month from July 2007 to July 2009. In addition, to better capture the next-hop diversity change in time, we only consider the prefixes that continuously exist over the entire two-year measurement period, which leaves us total 220,432 prefixes.

Figure 7 depicts next-hop router diversity changes at 25, 50, 95, 99 percentile, and maximum in next-hop router diversity distribution curves at different times. For example, on July 2007 (the leftmost data points), the median, 99%, and maximum next-hop diversity were 8, 25 and 36. Figure 7 shows that over the last two years, the median value stayed almost the same, though we checked that the individual prefix does shift its diversity to some extent. We do not observe any significant pattern of changes. As the individual prefix's diversity is

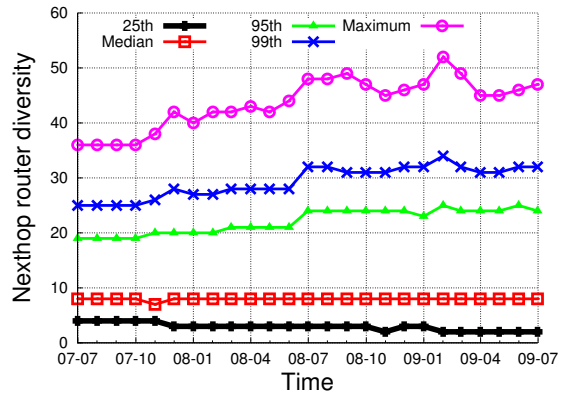


Fig. 7. Next-hop Diversity Change in Time

determined by a complex interaction between the topological and geographical location of the origin AS, the inter-domain routing path from the origin to ISP_A , the number of next-hop routers, and the BGP routing decisions, the path diversity changes in time with a seemingly unpredictable manner.

However, we also observed that the maximal next-hop diversity slowly increases in time, mainly due to the increased number of backbone routers inside ISP_A . The maximum, 99 percentile, and 95 percentile next-hop router diversity gradually increased in time, and 25 percentile value decreased slightly. After further investigation, we found that the increasing trend in maximum, 99 percentile, and 95 percentile is mainly due to the increased number of peering routers between ISP_A and its neighbors. Since July 2007, the number of backbone routers in ISP_A gradually increased up to 19 additional routers by the end of July 2009. This also confirms to the findings of case studies we made previously in Section IV.

VI. DISCUSSIONS OF RELATED WORKS

Prior works on path diversity fall into two classes: (i) quantifying existing path diversity and (ii) increasing path diversity.

Among prior works in the first class, Teixeira *et al.* [13] measured the IP level path diversity inside a Tier1 ISP (Sprint)'s backbone network, and showed that Sprint has significant *IP level* path diversity among their POPs. In contrast, we measure the *BGP level* exiting point diversity. On the other hand, [9, 3, 17, 5] measured path diversity at the AS granularity. While our work measures path diversity from the perspective of a Tier1 ISP, they mostly focused on multi-homing stub ASes because their common goal was to understand the impact of path diversity on data forwarding performance for a given multi-homing AS. Uhlig *et al.* [15] quantified path diversity in a Tier1 ISP which configured its network with route reflection [6]. Because they focused on the impact of route reflection on path diversity reduction, they used simulations to measure the path diversity inside an ISP using a small set of sampled prefixes. Our measurements confirm with their results in that many paths are not visible from the iBGP routers because of the local routing policy.

However, we measure next-hop diversity of production routers in a full-mesh network, and quantifies path diversity for all prefixes in the global routing table as well as the trend in overtime using the actual number of routers, which yield a more tangible and comprehensive understanding of path diversity and different impacting factors of path diversity in both general and corner cases.

The second class of prior works involve efforts to increase path diversity. Recently, the operator community starts to demand higher path diversity to accommodate the newly emerging applications [12, 7, 8]. This led to on-going efforts to increase path diversity by modifying the behavior of BGP. Walton *et al.* and Schrieck *et al.* [7, 8] propose a BGP capability, *Add-Path*, to distribute multiple paths for a given destination. This new extension increases the availability of additional paths, and can help reduce persistent route oscillations and route convergence within a network. While we deem it necessary to have a general way to exchange multiple paths between BGP routers, in this paper we showed that the majority of prefixes can be reached via more than one next-hop routers without changes to BGP. One interesting question is whether ISPs actually utilize the existing diversity before moving forward to increase it.

On the other hand, instead of modifying BGP, Raszuk *et al.* [12] proposed to deploy multiple BGP route reflectors planes, and each additional plane incrementally increases the number of alternative paths. The key idea is to configure each reflector such that the N_{th} reflector could select and distribute the N_{th} best path. This technique echoes our observations in this paper that changing BGP path preferences can greatly affect the diversity, but note that this technique might not be applicable to networks such as ISP_A , that does not use route reflection to organize its network. We hope that our measurement results can serve as valuable input on these efforts to decide whether such mechanisms to increase path diversity are necessary.

VII. SUMMARY AND FUTURE WORK

BGP has gone through many changes as it operates as the de-facto routing protocol in the Internet. Its original design required a BGP router to select and propagate only a single best path to its neighbors. This design choice is being reconsidered to increase path diversity. However, there has been little understanding on path diversity in the existing system, and the necessity and effectiveness of different proposals are not clear.

Using iBGP routing data collected from more than one hundred production backbone routers inside a Tier1 ISP, we show that there already exist opportunities in the existing network for the ISP to utilize its diversity by showing that the majority of prefixes could be reached through multiple next-hop routers.

Also, our analysis on the two impacting factors, the path preference and the number of peering routers, show that the ISP may further increase path diversity without any modification to BGP, by adjusting path preference values while

respecting the network's policy. Furthermore, we find that a very small number of prefixes maintain a high degree of diversity, and in most cases, they happen specifically and regardless of the ISP's intention, caused by the lack of geographical presence of ISP_A in the regions where origin ASes are located. Such observed overall next-hop diversity have not changed much over the past two years, but the maximal next-hop diversity slowly increases in time, mainly due to the increased number of backbone routers inside ISP_A .

In this paper, we take a first step in reporting the existing diversity and the causes inside one Tier1 ISP to help better understand the operational needs and practices that influence the degree of next-hop diversity. More data processing and analysis are necessary to investigate whether ISP_A utilizes and benefits from the diversity, and how such diversity interacts with the BGP dynamic behaviors. This is a subject for our future work.

ACKNOWLEDGMENT

This work was supported by NAP of Korea Research Council of Fundamental Science & Technology and by the MKE, Korea, under the ITRC support program supervised by the NIPA-2010-C1090-1011-0004. This work is also supported partially by the US NFS CNS-0721859.

REFERENCES

- [1] BGP Parser. <http://irl.cs.ucla.edu/bgpparser/>.
- [2] MaxMind - GeoIP. <http://www.maxmind.com/app/ip-location>.
- [3] Measuring Provider Path Diversity from Traceroute Data: work in progress. <http://www.caida.org/outreach/isma/0112/talks/krishna/index.pdf>.
- [4] MRT Routing Information Export Format. <http://www.ietf.org/internet-drafts/draft-ietf-grow-mrt-07.txt>.
- [5] P. D. Arjona Villicana, C. C. Constantinou, and A. S. Stepanenko. The Internet's Unexploited Path Diversity, 2009.
- [6] T. Bates, E. Chen, and R. Chandra. RFC 3345: BGP Route Reflection: An Alternative to Full Mesh Internal BGP.
- [7] A. R. D. Walton and E. Chen. Advertisement of Multiple Paths in BGP, March 2010.
- [8] V. V. den Schrieck and P. Francois. Analysis of Paths Selection Modes for Add-paths, July 2009.
- [9] J. Han, D. Watson, and F. Jahanian. An Experimental Study of Internet Path Diversity. In *IEEE Transactions on Dependable and Secure Computing*, October 2006.
- [10] P. Marques, R. Fernando, E. Chen, and P. Mohapatra. Advertisement of the Best External Route in BGP, February 2010.
- [11] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (in)Completeness of the Observed Internet AS-level Structure. In *IEEE/ACM Transactions on Networking*, 2010.
- [12] R. Raszuk, K. Patel, I. Kouvelas, R. Fernando, and D. McPherson. Distribution of Diverse BGP Paths, March 2010.
- [13] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker. In search of path diversity in ISP networks. *ACM Sigcomm conference on Internet measurement*, 2003.
- [14] P. Traina, D. McPherson, and J. Scudder. RFC 5065: Autonomous System Confederations for BGP, August 2007.
- [15] S. Uhlig and S. Tandel. Quantifying the BGP Routes Diversity Inside a Tier-1 Network. *Networking*, 3976, April 2006.
- [16] J. Uttaro, V. V. den Schrieck, P. Francois, R. Fragassi, A. Simpson, and P. Mohapatra. Best Practices for Advertisement of Multiple Paths in BGP, November 2010.
- [17] V. Vasudevan, D. G. Andersen, and H. Zhang. Understanding the AS-level Path Disjointness Provided by Multi-homing, 2007.
- [18] F. Wang, Z. M. Mao, L. G. Jia Wang, , and R. Bush. A Measurement Study on the Impact of Routing Events on End-to-End Internet Path Performance. In *Proceedings of ACM SIGCOMM 2006*, 2006.
- [19] B. Zhang, R. Liu, D. Massey, and L. Zhang. Internet Topology Project. <http://irl.cs.ucla.edu/topology/>.