A Comparative Study of Architectural Impact on BGP Next-hop Diversity

Jong Han Park*, Pei-chun Cheng[†], Shane Amante[‡], Dorian Kim[§], Danny McPherson[¶], Lixia Zhang[†]
*AT&T Labs, jpark@research.att.com
[†]University of California, Los Angeles, {pccheng,lixia}@cs.ucla.edu
[‡]Level-3 Communications Inc., Shane.Amante@Level3.com
[§]NTT Communications Inc., dorian@blackrose.org
[¶]Verisign Inc., dmcpherson@verisign.com

***Abstract*—Large ISPs have been growing rapidly in both the size and global connectivity. To scale with the sheer number of routers, many providers have replaced the flat full-mesh iBGP connectivity with a hierarchical architecture, using either Route-Reflection (RR) or AS confederation. Given that each intermediate iBGP router in the hierarchy selects and propagates only one best path per destination network, there is a common perception that, compared to full-mesh, a hierarchical iBGP connectivity is likely to lose sight of alternative paths to external destinations. To gauge the path diversity reduction in the operational networks, we performed a comparative study by using iBGP data collected from two global-scale ISPs, with full-mesh core and RR architecture respectively. Our results show that both ISPs suffer a significant reduction (up to 42%) in the overall path diversity. However the specifics of different iBGP architectures only made a minor impact (less than 2.9%) on this reduction. Rather, in both ISPs the majority of the alternative paths are eliminated by the first two criteria in BGP best path selection, *i.e.,* LOCAL_PREF and AS_PATH length.**

***Index Terms*—BGP; Next-hop Diversity; Measurement**

I. INTRODUCTION

BGP is the routing protocol used in the global Internet to exchange reachability information among autonomous systems (ASes). As the Internet grows in size and connectivity density over time [1], it is highly desirable to utilize multiple alternative paths to reach a network destination. However although a BGP router may learn multiple available paths by connecting to different neighbor routers, by design it can only select and propagate a single best path for each destination network. This leads to a concern in many large ISPs whose BGP routers are connected hierarchically to scale with the network size. Intuitively, if an AS deploys a hierarchical iBGP topology which results in reduced numbers of connectivity between its iBGP routers as compared to an AS with a full mesh iBGP connectivity, the AS may potentially miss many alternative paths to reach external destinations.

To answer the question of whether hierarchical iBGP topologies may have a negative impact of hiding alternative paths in operational networks, we performed a comparative study on BGP path diversity using iBGP routing data collected from two global-scale ISPs, referred to as ISP_{FM} and ISP_{RR} based on their internal *full-mesh* iBGP and *route reflection* iBGP connectivity, respectively. Our main findings can be summarized as follows.

- We show that, for each given destination network, the number of next-hop POPs and ASes varies widely in both ISPs. A significant fraction of prefixes (50.22% and 28.97% in ISP_{FM} and ISP_{RR}, respectively) do have a high path diversity with more than 10 next-hop POPs, mainly due to the topological connectivity between the origin AS and the two measured ISPs. On the other hand, a noticeable amount of prefixes (9.95% and 34.02%) are reached via a single next-hop POP only (Section IV).
- Our simulations using the collected iBGP data show that as much as 42% of alternative paths are eliminated in both ISPs, mainly by the first two criteria in the BGP best path selection. The specifics of different iBGP architectures have only a minor impact (less than 2.9%) in reducing the number of alternative paths, and the architecture-specific reduction can further be mitigated by a well-engineered iBGP placement and connectivity (Section V).

II. BACKGROUND

In this section, we provide a brief overview on BGP and describe BGP next-hop diversity.

A. Routing in the Internet

BGP runs between routers both of different ASes (eBGP) and inside a single AS (iBGP). In the iBGP case, routers use BGP to distribute externally learned routing information within the network. To avoid routing loops, iBGP requires that all iBGP routers within the same AS connect in a full-mesh, and that reachability information learned from one iBGP router must not be forwarded to another iBGP router. This full-mesh requirement results in a intractably high number of iBGP sessions (square of the number of iBGP routers inside an AS). To mitigate the scalability problem, two alternative architectures are proposed and used widely by large ISPs: AS confederations [2] and route reflection [3].

Regardless of which iBGP architecture is used, all BGP routers select only one best path for each destination prefix and propagate this selected path to neighbor routers. The best path selection process considers the following criteria in

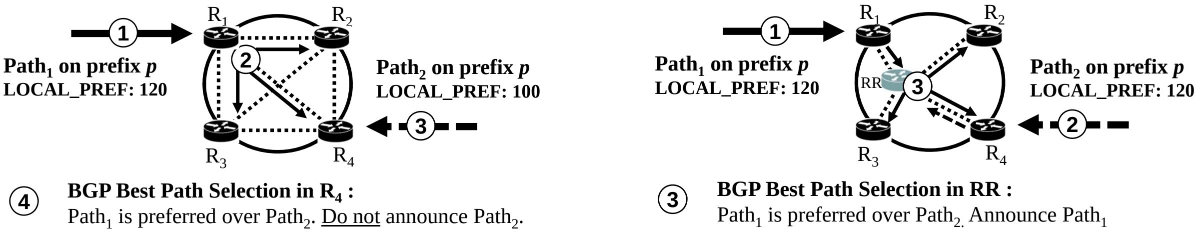

(a) Topology-independent Hidden Path

(b) Topology-dependent Hidden Path

Fig. 1. Hidden Path Phenomenon in iBGP

the order listed [4]: (1) highest LOCAL_PREF[1], (2) shortest AS_PATH length, (3) lowest ORIGIN, (4) lowest MED, (5) prefer path learned from eBGP session over path learned from iBGP session, (6) lowest IGP cost, and (7) lowest Router ID. The first four criteria examine BGP attributes whose values are *independent* from the router's location in the internal iBGP topology (*i.e.,* the preference of a path based on these four criteria would be the same regardless of the topological location of the router inside the AS). The last three criteria examine values that are topology-dependent and can result in different preference by routers of different topological locations and connectivity inside the AS.

B. iBGP Hidden Path

1) Hidden Paths at Border Routers: When an AS border router does not announce into the AS the learned, but less preferred external paths for a given destination, we say that these paths are hidden at the border routers. As a result, the less preferred paths are known only to the border router itself, and other iBGP routers do not have the visibility to these less preferred paths to reach a given external destination.

Figure 1(a) shows an example of an external path (due to lower LOCAL_PREF attribute value in this case) hidden at the border router in a full-mesh iBGP configuration. In this example, the less preferred path ($Path_2$) is not selected and known only by the border router (R_4) unless the current best path fails[2]. When $Path_1$ fails, no router except R_4 can switch immediately to use $Path_2$, until R_4 announces $Path_2$ to the rest routers. This inability to failover immediately to an available alternatively path has a negative impact on the data plane performance [5].

2) Hidden Paths due to iBGP Hierarchy: In addition, depending on the iBGP architecture, the number of paths learned by a router to reach a destination may differ. Figure 1(b) shows an example of a hidden path due to the route reflection iBGP configuration. That is, although all equally preferred external paths are announced into the route reflector by the border routers, the route reflector chooses only one best path based on its topology-dependent BGP best path selection criteria and propagates only the selected path to its clients, preventing other iBGP routers from learning the other alternative paths.

Because only the best paths are propagated from one side of sub-AS (or route reflector) boundary to the other side, the number of overall paths learned can be further reduced.

III. METHODOLOGY

We used iBGP data collected from ISP_{FM} and ISP_{RR}. In this section, we describe the high level network topology of the 2 ISPs, followed by data collection settings and how we measure the next-hop diversity.

A. A Brief Description of ISP_{FM}

ISP_{FM} is a global-scale ISP which uses a single AS number globally in the Internet. It has several hundreds of iBGP routers distributed across many countries in multiple continents, and uses AS confederations [2] to scale with its network size. Figure 2(a) depicts a simplified topology of ISP_{FM} at a high level, where *backbone sub-AS* represents the backbone network of this ISP, consisting of more than one hundred iBGP routers connected in a full-mesh (hence referred to as ISP_{FM}). ISP_{FM} deploys a BGP data collector which establishes an iBGP peering session with each of the iBGP routers in the backbone sub-AS to passively record all iBGP updates received.

B. A Brief Description of ISP_{RR}

ISP_{RR} is another global-scale ISP and uses one AS number globally in the Internet. It also has several hundreds of iBGP routers distributed across many countries in multiple continents. It deploys hierarchical route reflection architecture by recursively applying route reflection. Figure 2(b) depicts a simplified hierarchical route reflection system built by ISP_{RR}. The diamond-shape RRs at the top level represent continent level RRs; the square-shape RRs are at the 2^{nd} level of hierarchy, each represents a regional RR, and the 3^{rd} level circle-shape RRs represent POPs. A collector (an iBGP router) is configured as RR client to all route reflectors in the 2nd level route reflectors and passively record all iBGP updates received.

C. Quantifying Next-hop Diversity

Potentially, next-hop diversity can be measured at different granularity levels: Note that ISP_{FM} and ISP_{RR} interconnect with other *ASes*, at *peering points* in different cities, further

---

[1] In BGP, ISPs use LOCAL_PREF attribute value to indicate the policy preference on each path. Typically, a path through a customer is preferred over that of peer. This is true in both ISP_{FM} and ISP_{RR}.

[2] In the case that R_4 learns $Path_2$ first, it will explicitly withdraw the path after learning about the more preferred path ($Path_1$).

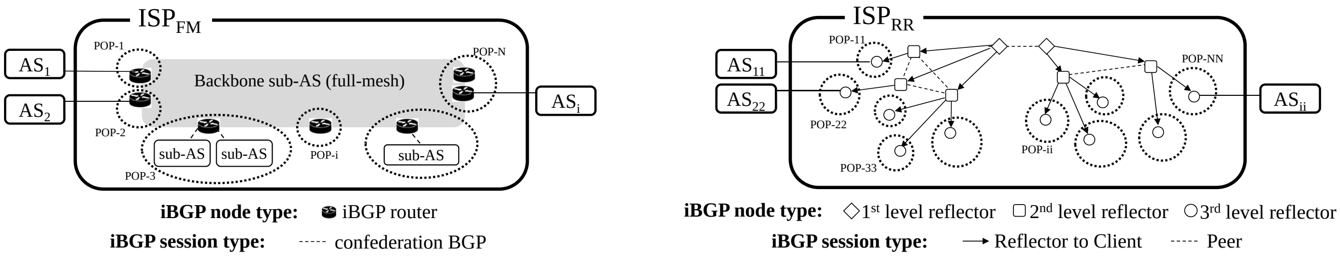

(a) Simplified Topology of ISP_{FM}

(b) Simplified Topology of ISP_{RR}

Fig. 2. High Level iBGP Topology of Two ISPs

through multiple next-hop neighbor *routers*. One can measure the number of *next-hop ASes*, *next-hop POPs (Point of Presence)*, and *next-hop routers* respectively. In this work, we focus only on the *next-hop POP and AS diversity*, due to the limitation that ISP_{FM} and ISP_{RR} configure differently the *next-hop-self option* at the border routers, which prevent a fair comparison of the router level diversity. [3]

From ISP_{FM} and ISP_{RR}, we gathered routing table snapshots (RIBs) from all backbone iBGP routers. We first exclude two types of prefixes from this measurement study: internal prefixes and prefixes with their length shorter than 8 or greater than 24 (3.5% and 10.7% of overall prefixes in ISP_{FM} and ISP_{RR}). Then, from each RIB entry, we extracted NEXT_HOP and AS_PATH attributes to measure how many distinct next-hop POPs and ASes are visible collectively in the view of the backbone routers for a given destination.

IV. BGP NEXT-HOP DIVERSITY

A. Next-hop Diversity in ISP_{FM}

We start by measuring next-hop diversity in ISP_{FM}. In this paper, we present our measurement results based on the routing table snapshots taken on June 3rd, 2010 for clarity. To ensure that the snapshots are representative, we performed the same measurements using routing tables taken on each day during one week of June 3rd to 9th and on every 1st day of each month from January to May in 2010. We verified that the distributions of next-hop POP and AS diversity are very similar across all the samples. In addition, we checked that the total number of prefix entries and the set of unique POPs and neighbor ASes are roughly the same.

Figure 3(a) shows the distributions of next-hop POP and AS diversity of 307,212 prefixes. We observe in Figure 3(a) that a significant number of prefixes can be reached via more than 10 next-hop POPs and 2 ASes (50.22% and 37.7% respectively). However, the number of prefixes with multiple next-hop ASes is quite small; only 3.78% of all prefixes can be reached via more than 4 next-hop ASes. The number of POPs to reach a given prefix is generally higher than the number of neighbor ASes, indicating that ISP_{FM} peers with some of its neighbor ASes in multiple sites. In our previous study [6] which studied the nature of BGP next-hop diversity in one ISP in detail, we found that the high diversity of these prefixes in ISP_{FM} is mostly (more than 89%) due to the AS-level topological connectivity as viewed by ISP_{FM}; we showed that an origin AS which is located in the regions that ISP_{FM} is not present cannot directly connect to ISP_{FM}, and has to connect to ISP_{FM} via other (regional) ISPs. Given the dense connectivity of today's Internet, the increased topological distance translates to a high number of equally preferred paths inside ISP_{FM}. We also observe that there exist two large groups of prefixes sharing the same degree of POP diversity. About 15% and 14% of prefixes have their POP diversity equal to 14 and 9 respectively. We further investigate why these prefixes have the same degree of next-hop diversity, and find that the paths to reach these prefixes are learned from a handful of large neighbors and thus share particular next-hop AS and POPs. Lastly, we observe that more than 9.95% and 31.46% of prefixes can only be reached via 1 POP and AS.

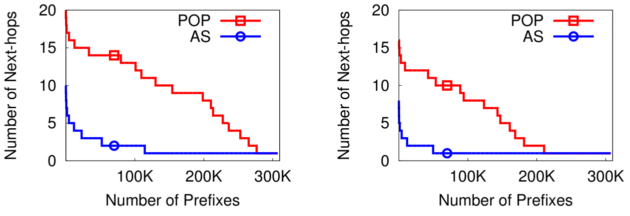

(a) ISP_{FM}

(b) ISP_{RR}

Fig. 3. Next-hop POP and AS Diversity

B. Next-hop Diversity in ISP_{RR}

We measure next-hop diversity in ISP_{RR} and compare the results with the next-hop diversity in ISP_{FM}. Figure 3(b) shows the distributions of next-hop POP and AS diversity of the same 307,212 prefixes. The difference in the number of total prefix between the two ISPs mainly comes from the different announcements made by the neighboring ASes.

We make a number of common observations compared to ISP_{FM}. First, there is a significant number of prefixes that can be reached with more than 10 next-hop POPs and 2 ASes (28.97% and 16.29%). The number of prefixes with multiple next-hop ASes is small as in the case of ISP_{FM};

---

[3] ISP_{FM} does NOT use next-hop-self option, meaning that the number of measured next-hop router indicates the routers belonging to the neighboring ASes; On the contrary, ISP_{RR} uses next-hop-self option, meaning that number of next-hop router indicates the routers belonging to ISP_{RR} itself.

only 1.3% of all prefixes can be reached via more than 4 next-hop ASes. Overall, next-hop POP diversity is relatively higher than next-hop AS diversity, indicating that ISP_{RR} also peers with its neighbor ASes in multiple POPs. We find that the highest degree of diversity in ISP_{RR} is mostly related to how the origin ASes connect to ISP_{RR}. We identified the top 8,881 prefixes with the highest degree of next-hop diversity inside ISP_{RR}, announced by 1,336 unique origin ASes. Then, we used MaxMind GeoLite package [7] to map each prefix into a city. Finally for these mapped cities, we checked whether any POP of ISP_{RR} is present. We found that all 1,336 (100%) origin ASes that announced the prefixes with the highest degree of diversity do not directly connect to the two ISPs and that more than 91% of these origin ASes are located in regions that ISP_{RR} is physically absent. There are a few groups of prefixes sharing the same degree of POP diversity (*e.g.,* POP diversity equal to 12 and 8), representing the prefixes that use the same next-hop POP and next-hop AS. Lastly, a considerable and relatively larger number of prefixes can be reached via only one neighbor POP and AS; 34.02% and 84.42% of all prefixes have both their next-hop POP and AS diversity equal to 1.

Although both ISPs are classified as global-scale large ISPs, there is a noticeable difference in the next-hop diversity distribution. First, we observe that the maximum number of next-hop POP and AS is different, potentially caused by the difference in their external connectivity. More importantly, we observe that the overall number of ISP_{RR}'s next-hop POPs and ASes to reach a given prefix is relatively lower, compared to ISP_{FM}. For example in ISP_{FM}, there are 9.95% and 31.46% of all prefixes with 1 next-hop POP and AS respectively. However in ISP_{RR}, we observe that relatively more prefixes (34.02% and 84.42%) have only 1 next-hop POP and AS respectively.

V. IMPACTING FACTORS ON NEXT-HOP DIVERSITY

In this section, we further investigate different impacting factors on path diversity by examining the iBGP updates collected from the two ISPs for 6-month time period from January 2010 to June 2010. More specifically, we focus on understanding the following 3 factors and their impact on the overall next-hop diversity: (1) *external connectivity*, (2) *Topology-independent hidden paths*, and (3) *Topology-dependent hidden paths*.

A. External Connectivity

As we have seen in the previous section, next-hop POP and AS diversity of a prefix can potentially be upper-bounded by the external BGP connectivity of the ISP with its neighbor ASes. The most accurate approach to obtain the exact amount of external connectivity between an AS and its neighbor ASes is to examine the configurations of all the border routers of the AS. However, this requires access to all border routers in each of the two measured ISPs, which we did not have at the time of our measurement. To get around this obstacle, we estimate ISP_{FM} and ISP_{RR}'s external connectivity from observing the routing dynamics in the iBGP data collected. More specifically, we examine the iBGP updates to estimate the external connectivity by recording the next-hops exposed by the prefixes that have route flaps (*i.e.,* completely lost and restored the reachability) at least once during the time period. Oliveira *et al.* [8] show that the most number of paths are exposed during such events from the monitored location to the destination.

One challenge in estimating the external connectivity by observing the routing dynamics is to determine the measurement time duration. If the time duration is too long, the routing changes can include the permanent topology changes in the Internet [1]. On the other hand, if the time duration is too short, we will not observe the prefixes which are inactive during the observation period, and the number of observed prefixes can be too small. To capture as many prefixes as possible without including the permanent topology changes, we decided to look at multiple short time durations of one week that do not overlap over a longer period of time; to estimate the external connectivity, we use the iBGP data collected over 6 months during the 1st week January, February, March, April, May, and June in 2010. Overall, we identified 88,236 prefixes announced by 12,727 origin ASes (38% of all ASes: 10 Tier-1s, 1,346 Transits, and 11,371 Stubs) which are approximately 1/3 of all prefixes and origin ASes in the global routing table. Additionally, we checked that the prefixes and their origin ASes cover various AS types, topological locations, and the overall next-hop diversity. Although we did not capture all prefixes and ASes in the global routing table, our goal in this paper is to compare the relative difference of the external connectivity of the two ISPs, rather than precisely estimating all external connectivity for a given ISP. For this purpose, we believe that the total number of identified prefixes and origin ASes is sufficient.

In each of above 6 independent measurement periods, the percentage in diversity reduction varies slightly. However, we essentially make the same observation across the multiple independent measurements, and the generality of our conclusion does not change. Therefore in this paper, we present the results on one week from June 3rd to June 9th in 2010 as the representative result for clarity. The number of identified prefixes during this week is 24,244 (about 7% of all prefixes), announced by 4,457 unique origin ASes (13.59% of all ASes: 5 Tier-1s, 648 Transits, and 3804 Stubs).

The blue lines (labeled *PathExplored* marked with filled square) in Figure 4 and 5 show the number of next-hop POPs and ASes based on the estimated external connectivity for the identified prefixes in ISP_{FM} and ISP_{RR}. The distributions of the estimated external connectivity between the 2 ISPs reveal that there is no significant discrepancy, and therefore, we concluded that the external connectivity is not the dominating cause for the discrepancy observed in Figure 3.

B. Topology-independent Hidden Path

Given that the distribution of external connectivity of the 2 ISPs is similar, we measure the amount of topology-

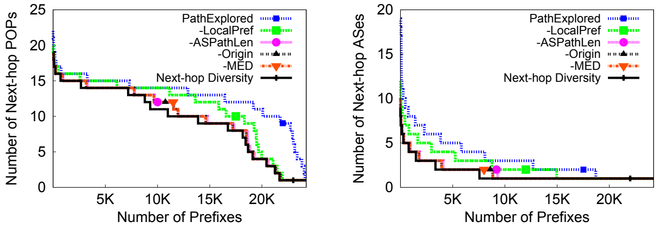

(a) Next-hop POP Diversity (b) Next-hop AS Diversity

Fig. 4. Next-hop Diversity Reduction in ISP_{FM}

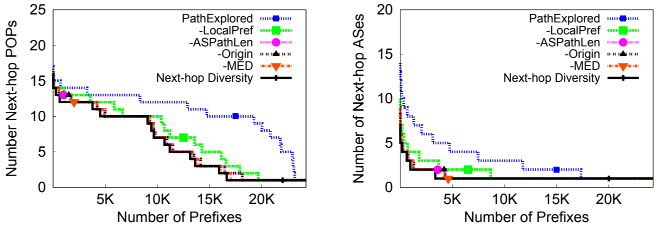

(a) Next-hop POP Diversity (b) Next-hop AS Diversity

Fig. 5. Next-hop Diversity Reduction in ISP_{RR}

independent hidden path, which happens regardless of the iBGP architecture or router topology as described earlier in Section II-B. To quantify the reduced next-hop diversity due to iBGP hidden path, we simulate the BGP best path selection algorithm with the first 4 topology-independent criteria, and count how many external paths remain equally preferred by routers inside the ISP after each criterion. The number of such remaining paths represents the paths that would be announced by the border routers into the AS after each of the first four BGP best path selection criteria.

1) ISP_{FM}: Figure 4 summarizes our simulation results for ISP_{FM}. In Figure 4(a) and Figure 4(b), each green (marked with a square), pink (marked with a circle), dotted black (marked with a triangle), orange (marked with an upside-down triangle) colored lines show the remaining next-hop POP and AS diversity respectively after each step of the first 4 best path selection criteria in ISP_{FM}. For example in Figure 4(a), our estimated external connectivity (*i.e.,* blue line marked with a square) indicates that there are only 0.4% of prefixes initially with their next-hop POP diversity equal to 1. After considering the 1st criterion (LOCAL_PREF comparison), the green line (labeled *-LocalPref*) shows that more prefixes (7.36%) have the next-hop POP diversity equal to 1. This means, among multiple external paths to reach a given prefix, only *one* path stands out due to its higher LOCAL_PREF value, making other (less preferred) paths hidden inside the border routers.

Overall, the first 2 criteria contribute most to the next-hop diversity reduction. After the 1st criterion (LOCAL_PREF comparison), about 10% of overall next-hop POP diversity is reduced. Then additional 12% next-hop POP diversity reduction happened after the 2nd criterion (AS_PATH length comparison).

2) ISP_{RR}: Figure 5 summarizes our simulation results for ISP_{RR}. As in the case of ISP_{FM}, the first 2 criteria of the best path selection are identified as the dominating factors that reduce next-hop diversity. However, the amount of reduction caused by each of the 2 criteria is quite different. In case of ISP_{RR}, the 1st criterion (LOCAL_PREF) had the most impact on next-hop diversity reduction (of about 29%), and is the main reason why the 2 ISPs have such discrepancy in the measured next-hop diversity in Figure 3. Our results reveal that although ISP_{RR} has a comparable amount of external connectivity compared to ISP_{FM}, relatively less number of paths are equally preferred after examining LOCAL_PREF attribute value and the subsequent topology-independent criteria.

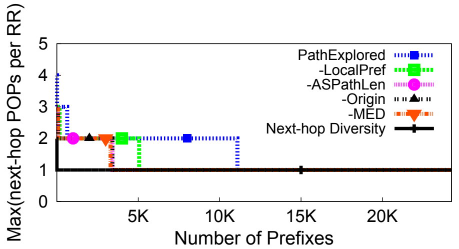

Fig. 6. Maximum Number of Next-hop POPs per RR

C. Topology-dependent Hidden Path

The iBGP hidden path due to the first 4 topology-independent criteria of the best path selection happens regardless of the iBGP topology. This implies that even in the full-mesh topology, the remaining next-hop diversity after the 4th criterion is the upper-bound, and that further reduction caused by the topology-dependent criteria represents the cost of moving away from the full-mesh topology.

Thus, we define the difference between measured diversity as (*i.e.,* black line labeled *Next-hop Diversity*) and the diversity after the 4th criterion of best path selection (orange line labeled *-MED*) as the amount of diversity reduced due to topology hierarchy and connectivity between the border routers and the backbone routers.

1) The Impact of iBGP Topology and Next-hop Diversity Reduction: In both Figure 4 and 5, we observe that the difference between the solid orange line (labeled *-MED* marked with an upside-down triangle) and the solid black line (labeled *Next-hop Diversity* marked with a short vertical line) is relatively small. This indicates that the overall reduction due to the topology-dependent factors across all simulated prefixes is small; even with ISP_{RR}'s multi-level hierarchical route reflection architecture and its topology, there is only up to 2.9% reduction.

2) iBGP Topology Design and Next-hop Diversity: In route reflection architecture, path diversity reduction happens essentially by deploying a relatively smaller number of route reflectors, compared to the available number of paths per route reflector. Given that ISP_{RR} has only a minor reduction in overall next-hop diversity, we further verify that in ISP_{RR} the number of route reflectors in the backbone routing infrastructure roughly match the number of available next-hop POPs. We first calculated the number of distinct next-hop POPs observed by each route reflector before and after considering the first 4 BGP best path selection criteria, and then chose the maximum number across all route reflectors. For example, if the number of observed next-hop POPs by two route reflectors are 2 and 5 respectively, the maximum number of next-hop POPs per route reflector (as shown in Figure 6) is max(2,5) = 5. If this number is equal to 1 for a given prefix, it implies that there is sufficient number of route reflectors in the network to preserve the observed next-hop POP diversity for that prefix.

Figure 6 summarizes our results. First, the number of maximum next-hop POP is 1 for the majority (more than 54%) of the prefixes. This indicates that the route reflectors are sufficiently placed for these prefixes in terms of their next-hop POP diversity density per a given route reflector. For the prefixes with the available next-hop POP greater than 1, there is a noticeable decrease in the maximum number of observed next-hop POPs per route reflector after considering the first 4 topology independent BGP best path selection criteria; for more than 32% of simulated prefixes, the number of next-hop POP decreased to 1. This result suggests again that in the current iBGP operation where the major amount of path diversity are hidden at the AS borders regardless of iBGP topology, it is practical to use more scalable iBGP architectures (with careful design) without much sacrifice in the overall path diversity.

VI. DISCUSSIONS AND FUTURE WORK

Although several recent measurement and analysis studies have addressed the issues of BGP path diversity [6], [9], [10] by proposing BGP modifications to support multiple paths [11]–[17], there is a lack of general understanding on the degree of BGP path diversity in today's operational networks, their impacting factors, and in particular whether or not the different iBGP architectures such as route reflection have significant impacts on reducing the path diversity at the AS-level.

Our comparative measurement study based on the iBGP data collected from two large ISPs with different iBGP architectures quantifies and compares the degree of path diversity in these two ISPs. Our results reveal the most influential factors on path diversity reduction. We show that, although there is a significant overall path diversity reduction in iBGP, the reduction caused by the specifics of iBGP architecture inside ISP_{RR} is not substantial as commonly perceived. There are two main reasons. First, topology-independent criteria are high in the order of BGP best path decision process and contribute significantly to the overall reduction. Second, a well-engineered iBGP topology mitigates the topology-dependent reduction as described in Section V.

Although the overall alternative path reductions is mainly due to the topology-independent factors as our results showed, there was a noticeable difference in the amount of reduction due to LOCAL_PREF attribute in BGP best path selection. We conjecture that this difference can be explained by the economical factors such as the access-circuit prices, transit prices, SLA's and peering policies which are affected by the different geographical regions that the two ISPs serve and leave the detailed analysis and verification as our future work. We would also like to study other ISPs with various iBGP topologies to verify whether the same observations can hold true. Lastly, we focused on understanding the static path diversity in different iBGP architectures in the absence of failures in this work. It remains as an open question how different iBGP architectures may impact BGP convergence in the presence of topological changes, which is the subject of our ongoing effort.

REFERENCES

[1] 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.

[2] P. Traina, D. McPherson, and J. Scudder, "Autonomous System Confederations for BGP." RFC 5065 (Draft Standard), Aug. 2007.

[3] T. Bates, E. Chen, and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)." RFC 4456 (Draft Standard), Apr. 2006.

[4] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)." RFC 4271 (Draft Standard), Jan. 2006.

[5] 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.

[6] J. Choi, J. H. Park, P. chun Cheng, D. Kim, and L. Zhang, "Understanding BGP Next-hop Diversity," in *14th IEEE Global Internet Symposium*, 2011.

[7] "MaxMind - GeoIP." http://www.maxmind.com/app/ip-location.

[8] R. Oliveira, B. Zhang, D. Pei, R. Izhak-Ratzin, and L. Zhang, "Quantifying Path Exploration in the Internet," in *ACM SIGCOMM Internet Measurement Conference (IMC)*, 2006.

[9] S. Uhlig and S. Tandel, "Quantifying the BGP Routes Diversity Inside a Tier-1 Network," *IFIP Networking 2006*, vol. 3976, April 2006.

[10] W. Mühlbauer, S. Uhlig, A. Feldmann, O. Maennel, B. Quoitin, and B. Fu, "Impact of Routing Parameters on Route Diversity and Path Inflation," *Computer Networks*, 2010.

[11] R. Raszuk, R. Fernando, K. Patel, D. McPherson, and K. Kumaki, "Distribution of Diverse BGP Paths." http://tools.ietf.org/html/draft-ietf-grow-diverse-bgp-path-dist-03, January 2011. (Internet Draft).

[12] P. Marques, R. Fernando, E. Chen, and P. Mohapatra, "Advertisement of the Best External Route in BGP." http://tools.ietf.org/html/draft-ietf-idr-best-external-03, March 2011. (Internet Draft).

[13] D. Walton, A. Retana, E. Chen, and J. Scudder, "Advertisement of Multiple Paths in BGP." http://www.ietf.org/id/draft-ietf-idr-add-paths-04.txt, August 2010. (Internet Draft).

[14] "BGP Multipath." http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094431.shtml#bgpmpath. Cisco online documentation, [Online].

[15] V. Van Den Schrieck and P. Francois, "Analysis of Paths Selection Modes for Add-paths." http://tools.ietf.org/html/draft-vvds-add-paths-analysis-00, July 2009. (Internet Draft).

[16] R. Raszuk, "To Add-Paths or not to Add-Paths." http://www.nanog.org/meetings/nanog48/presentations/Tuesday/Raszuk_To_AddPaths_N48.pdf, February 2010. [Online].

[17] V. Van Den Schrieck and P. Francois and O. Bonaventure, "BGP Add-Paths: The Scaling/Performance Tradeoffs," *IEEE Journal on Selected Areas on Communication*, vol. 28, no. 3, pp. 1299 – 1307, 2010.