Is Your IP Address Prefix Well-Served by Internet Routing?

Yibo Wang and Jun Li, University of Oregon
{*wangyibo, lijun*}*@cs.uoregon.edu*

ABSTRACT

In today's Internet, users or network operators cannot easily find out how reachable their IP address prefixes are from the rest of the Internet, nor do they know what kind of packet delivery quality they receive from Internet routing. Instead of setting up probing sites from multiple locations on the Internet and probing every prefix, we leverage the fact that there are continuous collections of BGP updates concerning the reachability of every IP address prefix. We introduce a novel data mining method and cluster IP address prefixes according to their BGP attributes, and verify if each cluster corresponds to a different level of packet delivery performance. Users can classify their prefix to see to which cluster the prefix belongs to.

1. INTRODUCTION

IP address prefixes on the Internet receive packets with different levels of "luck." Some can be reached by most hosts in the blink of an eye, some may seem faraway from the rest of the Internet; some can always be reached with very little packet loss, some often face bad luck in receiving packets. End-users or network operators can sometimes "feel" the poor packet delivery quality when running a VoIP or web browsing application, but they have little idea whether their address prefix is simply experiencing some temporary reachability or stability snafu, or they have chosen the wrong ISP to provide them connectivity to the Internet. In receiving the packet delivery service from the Internet routing as a whole, there are yet to be tools that a user can use to query a variety of questions related to learn the connectivity quality of his address block.

To provide such a tool, one could set up many probing sites, at vantage points throughout the Internet, and have each of them probe the address prefix in question. The probing results can then be collected to analyze the delay and loss in reaching the address prefix. Unfortunately, doing so is not scalable and incurs a daunting setup and measurement overhead if one must target every address prefix on the Internet. If the service is intended to be available at *any time*, the measurement also has to be done continuously.

The goal of Internet routing is for every router to figure out how to reach any IP address on the Internet; in particular, in running the Border Gateway Protocol (BGP), every BGP router exchanges messages with its peers on how to reach every IP address prefix. Noticing there have been efforts in collecting BGP routing updates [1, 4], can we study the pattern of BGP updates related to any specific address prefix to deduce whether BGP routers throughout the Internet do a good job in setting up "good" routes from everywhere toward the address prefix?

We verify this possibility. We cluster a large set of randomly selected IP address prefixes according to certain BGP features, and study if every cluster corresponds to a different level of packet delivery performance in general. If so, users can judge whether routes toward their address prefix provide good reachability or stability: They can classify an IP address prefix and see which cluster the address prefix would map to. This method is probably only a coarse-grained indicator regarding a prefix's performance in receiving packets, but it would offer a scalable and convenient approach to an important, yet-to-be-addressed problem.

2. CLUSTERING IP ADDRESS PREFIXES

The tool we used for clustering IP address prefixes is based on the classical `k-medoids` algorithm [3]. The clustering process starts with a root node (the root cluster) that contains all randomly selected IP address prefixes, (we omit the details of how we randomly select them), and finally splits them into multiple leaf clusters, where each cluster contains a set of IP address prefixes without overlapping. We verify that the pre-

*CoNEXT'*07, December 10-13, 2007, New York, NY, U.S.A.

fixes in every leaf cluster share similar packet delivery performance—if they do, we say a clustering process is successful. otherwise, it is unsuccessful.

In running the clustering process, we need to decide whether to further split a node or treat that node as a leaf cluster. Specifically, we calculate the similarity degree of prefixes within a cluster in terms of their BGP attributes, and compare the similarity degree with a threshold. While we are only trying an empirical threshold value for now, we have yet to systematically investigate what threshold values to use.

We run a separate clustering process for every different metric of packet delivery performance. In addition to loss rate and delay, we also study the following metrics: *delta-loss-rate* that indicates the variation degree of the packet loss rate toward a prefix, *delta-delay* that indicates the changes of packet delays toward a prefix, and *route-change* that shows how often the routes toward a prefix are changing. While the delay and loss metrics show the reachability toward a prefix, we use the other three to capture the stability in reaching a prefix.

For every IP address prefix, we devise a matrix structure per prefix as the basic element to represent the prefix. Every row of the matrix corresponds to BGP data for that prefix during a time interval. We use BGP data from Route Views [1] and RIPE [4]. Every column is a specific BGP attribute (or feature). Here, we carefully test different possible combinations of BGP attributes to see if at least one of them would lead to a successful clustering. Our BGP data source may not provide us a thorough view of the whole Internet, but at this preliminary stage, this should not be a problem for us.

3. EVALUATING CLUSTERED PREFIXES

We need to verify whether we can use BGP routing data to deduce the packet delivery performance for an IP address prefix. More specifically, we must verify we can check to which cluster an IP address prefix belongs in order to know the packet delivery performance towards that address prefix. Our methodology is as follows: we compare the packet delivery performance similarity degree of prefixes within a leaf cluster node (i.e., inner-degree) with that of prefixes from different leaf cluster nodes (i.e., inter-degree). If the inner-degree is higher than the inter-degree, our verification is then successful.

We use packet delivery performance data from two measurement projects: Skitter [2] and PingER [5]. They provide delay, loss rate, and traceroute information from multiple vantage points throughout the Internet to IP addresses in different IP address prefixes. For every address prefix in question, we can evaluate the packet delivery performance in terms of a given metric from those vantage points to several addresses in that prefix.

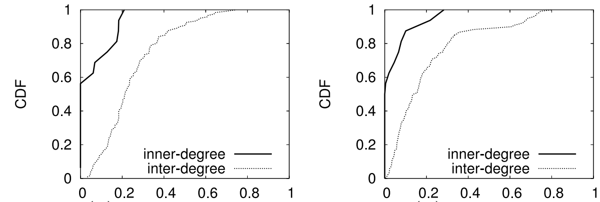

(a) Loss rate (b) RTT

Figure 1: CDF of inner-degree and inter-degree of leaf clusters.

We calculate the inner-degree of a cluster regarding a specific packet delivery performance metric as follows: We find out the value in terms of that metric for each IP address prefix in a cluster, calculate the difference for every pair of prefixes in the cluster, calculate the average of them, and finally use the normalized average of all leaf clusters as the inner-degree of a specific metric. We similarly calculate the inter-degree between leaf clusters.

We have obtained some preliminary results. Figure 1(a) and Figure 1(b) show an example of our verification results for a loss-oriented and delay-oriented clustering. In terms of loss rate, Figure 1(a) shows that almost 100% of the leaf clusters have an inner-degree smaller than 0.2, while only 40% of the leaf cluster pairs have an inter-degree of 0.2 or less. Similarly, as shown in Figure 1(b), the inner-degree is clearly different from the inter-degree for the delay.

4. SUMMARY

This research attempts to run data mining against BGP data to cluster IP address prefixes, and see if prefixes clustered together will share similar packet delivery performance in terms of a specific metric. If such a clustering can be done, one then can classify prefixes, and even if the classification is coarse-grained, it would provide a scalable and convenient tool for users to query the packet delivery performance toward their address prefix.

5. REFERENCES

[1] University of Oregon Route Views Project. `http://www.routeviews.org/`.
[2] CAIDA. Caida skitter internet topology and performance measurement tool. `http://www.caida.org/tools/measurement/skitter/`.
[3] J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers, 2001.
[4] RIPE NCC. RIPE routing information service raw data. `http://data.ris.ripe.net/`.
[5] SLAC. Pinger: Ping end-to-end reporting. `http://www-iepm.slac.stanford.edu/pinger/`.