

Detecting and Assessing the Hybrid IPv4/IPv6 AS Relationships

Vasileios Giotsas
University College London
v.giotsas@cs.ucl.ac.uk

Shi Zhou
University College London
s.zhou@cs.ucl.ac.uk

ABSTRACT

The business relationships between the Autonomous Systems (ASes) play a central role in the BGP routing. The existing relationship inference algorithms are profoundly based on the valley-free rule and generalize their inference heuristics for both the IPv4 and IPv6 planes, introducing unavoidable inference artifacts. To discover and analyze the Type-of-Relationship (ToR) properties of the IPv6 topology we mine the BGP Communities attribute which provides an unexploited wealth of reliable relationship information. We obtain the actual relationships for 72% of the IPv6 AS links that are visible in the RouteViews and RIPE RIS repositories. Our results show that as many as 13% of AS links that serve both IPv4 and IPv6 traffic have different relationships depending on the IP version. Such relationships are characterized as hybrid. We observe that links with hybrid relationships are present in a large number of IPv6 AS paths. Furthermore, an unusually large portion of IPv6 AS paths violate the valley-free rule, indicating that the global reachability in the IPv6 Internet requires the relaxation of the valley-free rule. Our work highlights the importance of correctly inferring the AS relationships and the need to appreciate the distinct characteristics of IPv6 routing policies.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design; C.2.3 [Computer Communication Networks]: Network Operations—*Network Management*

General Terms

Measurement, Algorithms

Keywords

Internet, Autonomous Systems, inter-domain routing, BGP, AS relationship, topology, IPv6, inference algorithms

1. INTRODUCTION

The business relationships between the Autonomous Systems (ASes) play a central role in the Internet BGP decision process. Knowledge on the AS relationships is essential for

measuring and analyzing the Internet inter-domain properties and operational trends. AS relationships are coarsely categorized as the transit relationship, i.e. provider-to-customer (p2c) or customer-to-provider (c2p); and the peering relationship, i.e. peer-to-peer (p2p). In a transit relationship the customer pays the provider to transit its traffic to the rest of the Internet. In a peering relationship the peers freely exchange traffic but only between themselves or their customers.

AS relationships are usually treated as confidential. A number of algorithms [1, 2, 4] have been proposed to infer the Type-of-Relationship (ToR) between ASes, utilizing a variety of heuristics. Their common ground is the valley-free property [1], according to which an AS path is valid only if it has one of the following formats: (1) $n \times c2p + m \times p2c$, or (2) $n \times c2p + p2p + m \times p2c$, where $n, m \geq 0$.

The existing ToR algorithms analyze the IPv4 and IPv6 AS links using exactly the same principles. However, the AS links carrying IPv6 traffic may follow unconventional BGP routing policies, including relaxed peering requirements and even free IPv6 transit. These distinct policies may result in AS links with different relationship type between the IPv4 and the IPv6 Internet. Such relationships are called *hybrid* IPv4/IPv6 relationships and cannot be captured by the existing ToR algorithms. Hence, measurement artifacts are unavoidable under the current ToR inference approaches.

To rigorously analyze the IPv6 AS relationships and detect the hybrid relationships we rely on the BGP Communities relationship information [3], and the Local Preference (LocPrf) attribute. We utilize the metric of “customer tree” [2] to assess the impact of hybrid links on the IPv6 routing structure.

2. DATA COLLECTION METHODOLOGY

The Communities attribute tags a BGP advertisement with additional information. Communities values are not standardized but the Internet Routing Registries (IRR) provide abundant information on their interpretation. Only Communities that describe relationship types and traffic-engineering requests are relevant to our analysis (the latter to facilitate the interpretation of LocPrf values).

LocPrf expresses the degree of preference that an AS has to a certain route. Usually LocPrf are ordered as follows: $LocPrf_{customer} > LocPrf_{peer} > LocPrf_{provider}$. However, LocPrf can be adjusted to any value for traffic engineering purposes.

The Communities relationships function as the Rosetta Stone for the LocPrf values. A LocPrf value is assigned to a

Copyright is held by the author/owner(s).
SIGCOMM'11, August 15–19, 2011, Toronto, Ontario, Canada.
ACM 978-1-4503-0797-0/11/08.




Figure 1: The change in the customer tree when the link 1–2 is (a) p2c or (b) p2p. In (a) AS1 can reach all the nodes through p2c links, while in (b) it can reach only AS3 through a p2c link.

relationship type only if we can validate it from the collected Communities. The traffic engineering Communities allow us to detect and filter-out LocPrf values used for non-standard preference tuning of specific paths.

3. ANALYSIS AND RESULTS

We collect daily BGP data from RouteViews and RIPE RIS repositories to extract the IPv6 AS paths and AS links. In August 2010, there are 346,649 IPv6 AS paths and 10,535 IPv6 AS links. 7,618 IPv6 AS links are also visible in the IPv4 topology. From the Community and the Local Preference attributes we are able to extract the actual AS relationship for 72% (7,651) of the all IPv6 links and for 81% (6,160) of the IPv4/IPv6 links. We have a number of interesting observations.

Firstly, 779 (or 13%) of the IPv4/IPv6 links have hybrid AS relationships. 67% of such hybrid links have a peering relationship for IPv4 and a transit relationship for IPv6; the rest are p2p for IPv6 and p2c for IPv4, except a single case where the two ASes have a p2c for IPv4 and a c2p for IPv6.

Secondly, the hybrid links usually happen among tier-1 or tier-2 ASes with large numbers of connections. As a result the hybrid links have a high visibility in IPv6 AS paths. More than 28% of the IPv6 paths contain at least one IPv4/IPv6 link with hybrid AS relationships.

Thirdly, 13% of the IPv6 paths do not follow the valley-free rule. We call them the *valley* paths. The large number of IPv6 valley paths is a major reason underlying the inference errors of the existing ToR algorithms. Our analysis indicates that 16% of the valley paths are due to the relaxation of the valley-free rule in order to expand the reachability of IPv6 prefixes. The IPv6 topology is partitioned in terms of valley-free routing¹ and the relaxation of the valley-free rule is necessary in some cases to maintain IPv6 reachability.

4. DISCUSSION

The substantial existence of hybrid IPv4/IPv6 links suggests that the IPv4 and IPv6 Internet topologies should be studied separately. This is consistent with our recent study on the evolution of IPv4 and IPv6 AS topologies [5].

¹An example is the peering dispute between two transit-free ASes in the IPv6 plane, AS6939 and AS174, as described in <http://mailman.nanog.org/pipermail/nanog/2009-October/014017.html>




Figure 2: The change of the average shortest path and the diameter of the IPv6 AS customer trees as we gradually correct the misinferred relationship of the 20 hybrid AS relationships with the highest visibility in the IPv6 AS paths.

The ToR-annotated AS topology is very sensitive to misinference of AS relationship. An expressive metric to assess the impact of misinferred relationships is the “customer tree” of an AS (root), which contains all the ASes that the root can reach through p2c links. Figure 1 shows an example where the change of relationship between two ASes results in two different topologies. We find that when we replace the IPv6 AS relationships that were misinferred in [4] with the correct relationships inferred from the BGP Communities, the average length and the longest length (diameter) of the shortest valley-free AS paths of the union of the IPv6 customer trees are reduced from 3.8 to 2.23, and from 11 to 7 hops, respectively (Figure 2).

In summary, our results reveal substantial differences between the IPv4 and IPv6 relationships, including a significant number of IPv6 paths that are not valley-free, sometimes in exchange for better reachability. The IPv6 topology should be studied separately using new models that capture its distinct characteristics. Our future work will validate and expand our inference results, investigate the reasons for hybrid links and revisit the valley-free rule. Our dataset and more details are available at the project’s website [6].

5. REFERENCES

- [1] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans. Netw.*, 9(6):733–745, 2001.
- [2] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy, and G. Riley. AS relationships: inference and validation. *SIGCOMM Comput. Commun. Rev.*, 37(1):29–40, 2007.
- [3] J. Xia and L. Gao. On the evaluation of as relationship inferences. In *Global Telecommunications Conference, 2004. GLOBECOM '04. IEEE*, 2004.
- [4] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (in)completeness of the observed internet as-level structure. *IEEE/ACM Trans. Netw.*, 18:109–122, February 2011.
- [5] G.Q. Zhang, B. Quoitin, and S. Zhou. Phase changes in the evolution of the IPv4 and IPv6 AS-Level Internet topologies. *COMPUT COMMUN*, 34(5):649–657, 2011.
- [6] <http://www.cs.ucl.ac.uk/BAB/>