

IPv6 AS Relationships, Cliques, and Congruence

Vasileios Giotsas, Matthew Luckie^(✉), Bradley Huffaker, and Kc Claffy

CAIDA, UC San Diego, La Jolla, USA
{vgiotsas,mjl,bradley,kc}@caida.org

Abstract. There is increasing evidence that IPv6 deployment is maturing as a response to the exhaustion of unallocated IPv4 address blocks, leading to gradual convergence of the IPv4 and IPv6 topologies in terms of structure and routing paths. However, the lack of a fully-connected transit-free clique in IPv6, as well as a different economic evolution than IPv4, implies that existing IPv4 AS relationship algorithms will not accurately infer relationships between autonomous systems in IPv6, encumbering our ability to model and understand IPv6 AS topology evolution. We modify CAIDA’s IPv4 relationship inference algorithm to accurately infer IPv6 relationships using publicly available BGP data. We validate 24.9% of our 41,589 c2p and p2p inferences for July 2014 to have a 99.3% and 94.5% PPV, respectively. Using these inferred relationships, we analyze the BGP-observed IPv4 and IPv6 AS topologies, and find that ASes are converging toward the same relationship types in IPv4 and IPv6, but disparities remain due to differences in the transit-free clique and the influence of Hurricane Electric in IPv6.

1 Introduction

Depletion of the unallocated IPv4 address pool increases the pressure for widespread adoption of IPv6. IPv6 deployment has long been characterized as largely experimental, dominated by research and education networks [7, 16]. However, recent studies suggest that the IPv6 network is maturing, reflected in increasing similarity of the IPv6 and IPv4 networks in terms of topological structure, routing dynamics and AS path congruity [9]. Czyz *et al.* also found the IPv6 traffic mix (set of applications using IPv6) in 2013 much more similar to the IPv4 traffic mix than in the past [8].

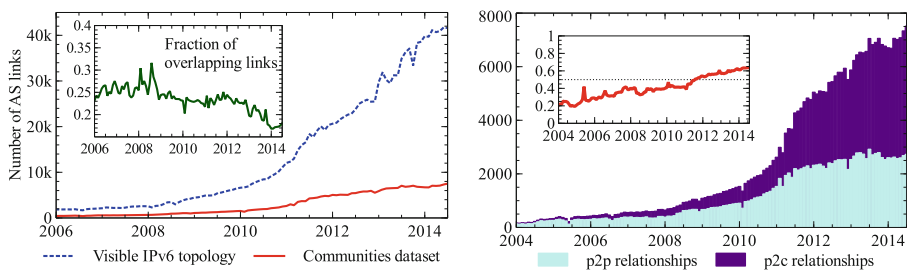
Despite these signals of convergence, noticeable differences remain between IPv4 and IPv6 routing relationships. In August 2010, Giotsas *et al.* found disparity in IPv4 and IPv6 AS relationships as inferred from BGP communities and local preference values [13]. Dhamdhere *et al.* showed that while 40–50% of dual-stacked paths observed in public BGP data were identical in 2012, if the ASes followed the same routing policies in IPv4 and IPv6, then 60–70% of paths could have been congruent [9]. They also found significant deviation between the most prominent ASes (those that appeared most frequently in AS paths) in IPv4 and IPv6, with the most prominent AS in the IPv6 topology (Hurricane Electric) appearing in a much larger fraction of IPv6 AS paths than the most prominent AS in the IPv4 topology (Level 3) appeared in IPv4 AS paths.

Because IPv6 deployment did not build on the existing IPv4 network, the IPv6 topology evolved in parallel, and not all assumptions relied upon by IPv4 AS relationship inference algorithms hold in IPv6. Inferring AS relationships is more challenging in IPv6 than in IPv4 for two reasons. First, given its still low deployment and different economics compared to IPv4 [9], IPv6 business policies are less rigorously enforced, leading to more policy violations [14] which impede the accuracy of relationship inference heuristics. Second, the IPv6 graph is not fully connected due to peering disputes between large transit-free providers [17, 23]. These challenges have discouraged both research [21] and commercial [22] efforts from inferring IPv6 AS relationships.

We make the following contributions. First, we adapt our IPv4 AS relationship algorithm [18] to accurately infer IPv6 AS relationships by accounting for IPv6-specific realities: in particular, the IPv6 AS topology is still not fully connected due to peering disputes [17]. We use our algorithm to infer AS relationships for January 2004 to July 2014 and publicly release our inferences. Second, we evaluate our algorithm’s accuracy by validating 10,357 (24.9%) of our 41,589 inferences using three sources of validation data, and find our provider-customer and peer-peer inferences have a 99.3% and 94.5% positive predictive value (PPV), respectively, in July 2014. We publicly release our validation data, which we derive quarterly between 2004 and 2014. Finally, we use our inferences to understand the growing congruity between IPv4 and IPv6 AS topologies. We show that despite growing congruity between the graphs, IPv6 AS relationships have evolved differently from those in IPv4. Disparate dual stack relationships are decreasing, from 15% in January 2006 to 5% in 2014, consistent with previous findings of growing similarity between IPv4 and IPv6 [8, 9]. However, Hurricane Electric (HE) is the main contributor of disparate relationships, and over 50% of their dual stack relationships differed between IPv4 and IPv6 in July 2014.

2 Background on Inferring as Relationships

AS relationships are often abstracted into three conventional classes [10]. In a provider-customer (p2c) relationship, a customer AS buys transit from a better connected AS to expand its reachability. In a peer-peer (p2p) relationship, two ASes provide access to their own and their customers’ networks. In a sibling-sibling (s2s) relationship, two ASes under common ownership may provide mutual transit to each other. ASes that can reach every network in the routing system without purchasing transit are known as Tier-1 ASes. Tier-1 ASes maintain p2p links between each other to ensure their global reachability, forming a clique that serves as the backbone of inter-domain routing. AS relationships translate into BGP routing policies that determine the economics of traffic exchange [11]. Accurate knowledge of AS relationships is thus essential to understanding not only inter-domain routing but also Internet economics [18]. Unfortunately, AS relationships are often treated as proprietary by ISPs and controlled by non-disclosure agreements, leading researchers to build algorithms



(a) Communities data compared to visible IPv6 topology. The inset graph shows the fraction of links in common.

(b) The composition of the communities data by relationship type. The inset graph shows the fraction of p2c relationships.

Fig. 1. Summary of the communities validation dataset over time. For July 2014 the dataset includes 7,514 relationships that cover 18.1% of the visible topology, 64% of which are p2c relationships and the rest p2p.

that heuristically infer AS relationships using publicly available BGP routing data. We recently developed an algorithm for inferring IPv4 AS relationships; we validated 34.6% of 126,082 p2c and p2p inferences for April 2012 to have a 99.6% and 98.7% PPV, respectively [18]. Our approach began by inferring a Tier-1 clique, applied heuristics to infer p2c links based primarily on how neighbors were observed to export routes, and inferred the remainder to be p2p. Section 4 describes how we modified this algorithm to infer AS relationships in the IPv6 topology graph.

Our IPv6 AS relationship algorithm infers conventional p2c and p2p relationships and does not infer complex AS relationships by design. We have developed and validated an algorithm to infer hybrid and partial transit relationships in IPv4 [12]. That algorithm uses conventional AS relationship inferences as input, and it is possible to apply the same heuristics to the output of our conventional IPv6 algorithm to infer complex IPv6 AS relationships.

3 Data

3.1 BGP Paths

We extracted AS paths from every vantage point providing BGP data to Route Views (RV) [4] and RIPE RIS [3] by downloading one RIB file per day between the 1st and 5th of every month between January 2004 and August 2014 and extracting AS paths that announced reachability to IPv6 prefixes.

3.2 Validation Data

For validation, we used three sources of IPv6 AS relationship data: BGP communities, RPSLng, and local preference (LocPref). We had access to BGP community data every month, quarterly RPSLng dumps, and three LocPref collections.

BGP communities are an optional transitive attribute that operators use to annotate routes [6]. The meaning of communities values are not standardized and each operator defines their own community values and meanings. We compiled a dictionary of community values and corresponding meanings that encode relationship types by mining WHOIS records and websites where operators document their specific use of community values; we also used historical documentation of communities values in archived WHOIS records and the Wayback web archive service [15] to obtain a dictionary for each April from 2004 to 2014. We assembled monthly validation datasets by applying the dictionary to corresponding public BGP data; the composition of this set of validation data over time is summarized in Fig. 1. For April 2014, our dictionary included 1,560 communities values defined by 284 ASes, and we used the dictionary to obtain validation data for 7,514 IPv6 links for the July 2014 IPv6 AS topology.

RPSLng is the Routing Policy Specification Language next generation [5], which network operators can use to store routing policies in public databases. The largest source of such data is RIPE’s WHOIS database; many European IXPs require operators to register routing policies with RIPE NCC. An import rule specifies which route announcements to accept from neighbors, and an export rule specifies what routes to advertise to neighbors. The special rule ANY is used to import/export all routes from/to a neighbor, and indicates a customer/provider relationship. Using RIPE’s WHOIS database from July 2014, we extracted 739 c2p relationships with the following method: if X has a rule that imports ANY from Y, and Y has a rule that exports ANY to X, we infer X is a customer of Y. Because RIPE NCC no longer provides the changed dates in their WHOIS dumps, we were unable to filter by freshness and used all records.

Despite the many links in our communities and RPSLng datasets, they include less than 2% of the IPv6 links observed in public BGP data for Hurricane Electric (HE), the most prominent AS in the IPv6 graph [9]. To extend our validation dataset to include HE’s relationships we use the local preference (**LocPref**) attribute, which does not directly encode relationship information but often reflects it [11]. LocPref is a number that expresses the level of preference an AS gives a route if multiple routes are available for the same prefix. LocPref values are also non-standardized, but many ASes assign the highest value to their customers and the lowest to their providers, which maximizes transit revenue. We collected LocPref values for HE’s neighbors by querying its public route server in July 2014, and we used two older datasets from [13]. Figure 2 summarizes the collected LocPref values for HE’s IPv6 neighbors; with few exceptions (22/2325 neighbors, marked with red crosses) HE assigns a single LocPref value to all prefixes received by each neighbor. Where HE assigned multiple values for different prefixes received from the same neighbor, we chose the value assigned to the most prefixes, since altering LocPref values is not typical behavior. When comparing HE’s LocPref values in IPv4 to inferred IPv4 relationships, we found a consistent mapping between LocPref 140 and HE’s customers (2591/2593) and LocPref 100 and HE’s peers (601/603). This mapping is HE-specific and not valid for every AS.

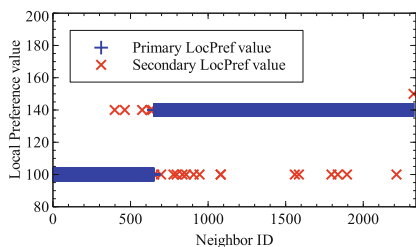


Fig. 2. Summary of HE’s LocPref validation dataset. In all three snapshots the LocPref values are concentrated around 100 and 140.




Fig. 3. Summary of agreement across validation data sources (first number inside intersections is number of overlapping relationships that agree). Communities and RPSLnG data have the largest agreement, over 98 %.

Algorithm 1. is a summary of our IMC 2013 IPv4 AS relationship inference algorithm. The bold lines were updated to accommodate IPv6.

Require: AS paths, Allocated ASNs, IXP ASes

- 1: Discard or sanitize paths with artifacts
 - 2: Sort ASes in decreasing order of computed transit degree, then node degree
 - 3: **Infer clique at top of AS topology (updated)**
 - 4: Discard poisoned paths
 - 5: Infer c2p relationships top-down using above ranking
 - 6: Infer c2p relationships from VPs inferred not to be announcing provider routes
 - 7: Infer c2p relationships for ASes where customer transit degree exceeds provider’s
 - 8: Infer customers for ASes with no providers
 - 9: **Infer c2p relationships between stub and clique ASes (removed)**
 - 10: Infer c2p relationships where adjacent links have no relationship inferred
 - 11: Infer remaining links represent p2p relationships
-

Figure 3 shows the overlap between the BGP communities, RPSLnG, and LocPref validation data sources for July 2014. The BGP communities and RPSLnG data had the largest overlap, and were consistent 98 % of the time.

4 Inference Methodology

4.1 Overview of Existing IPv4 Algorithm

Our IPv6 AS relationship algorithm is based on our IPv4 algorithm [18], with adjustments to account for differences in the routing ecosystems [9, 13]. In particular, the IPv6 AS topology lacks a fully connected clique that serves as the transit backbone because of a long-standing peering dispute between Cogent and Hurricane Electric [17].

Algorithm 1 summarizes our IPv4 AS relationship inference algorithm (details in [18]), highlighting the two steps we changed to accurately infer IPv6 relationships. First, we sanitize the input data by removing paths with artifacts, i.e., loops,

reserved ASes, and IXPs (step 1). We use the resulting AS paths to compute the node and *transit degree* (the number of unique neighbors that appear on either side of an AS in adjacent BGP links) of each AS, and produce an initial rank order (step 2). We then infer the clique of ASes at the top of the hierarchy (step 3). After filtering out some poisoned paths (step 4), we apply heuristics to identify c2p links (steps 5–10). In step 5, we infer c2p relationships top-down using the ranking from step 2, inferring an AS X is a customer of Y if Y exports routes received from X to peers or providers; this step infers 90% of all the c2p relationships we infer. In step 6, we infer c2p relationships from VPs we find announcing no provider routes, which we define as VPs that provide paths to fewer than 2.5% of the ASes. In step 7, we infer c2p relationships for ASes where the customer has a larger transit degree than its provider, to infer c2p relationships for links skipped in step 5. In step 8, we infer customers for provider-less non-clique ASes, which were also skipped in steps 5 and 7 because those steps require a non-clique AS to have a provider in order to infer customer relationships. In step 9, we infer that stub ASes are customers of clique ASes even if we do not observe the clique AS exporting the customer’s route to other peers; in IPv4 stub networks are unlikely to meet the peering requirements of clique members, and are most likely customers. In step 10, we resolve relationships where we observe triplets with adjacent unclassified links. Finally, we classify all remaining unclassified links as p2p.

4.2 Inferring the IPv6 Clique

Our inference algorithm follows a top-down approach starting from the clique members, to avoid relationship cycles and errors caused by stub ASes with high peering visibility. Inferring the IPv4 clique is relatively straightforward, given the maturity of the IPv4 network. In contrast, the IPv6 transit market is still in its early stages, making it more difficult to determine clique ASes. Because the accuracy of the inferred clique impacts the overall accuracy of the inferred relationships, we first focus on challenges of inferring the IPv6 clique.

To infer the IPv4 clique, our algorithm from [18] first sorted ASes by decreasing transit degree and then applied the Bron/Kerbosch algorithm to find the clique involving the first ten ASes that has the largest transit degree sum. We label these first ten ASes as the *seed* ASes because inferences for other ASes descend from this initial set. For each remaining AS, we added the AS to the clique if we observed a link with every other clique AS, and the AS did not appear to receive transit from one clique member to reach a second clique member. This approach works well for inferring the IPv4 clique because the largest transit degree ASes have restrictive peering policies, maintaining a peering clique with only selected transit-free ASes. For April 2014, the largest IPv4 transit degree AS was Level3 which maintains a restrictive peering policy. In contrast, the largest IPv6 transit degree AS is HE, which has an open peering policy, and is part of large peering meshes with ASes that are not transit-free; calculating the clique starting from the ASes with the largest transit degrees returns incorrect cliques in IPv6. Furthermore, because the IPv6 network is still early in its evolution, the IPv6 network is more dynamic than the IPv4 network, making

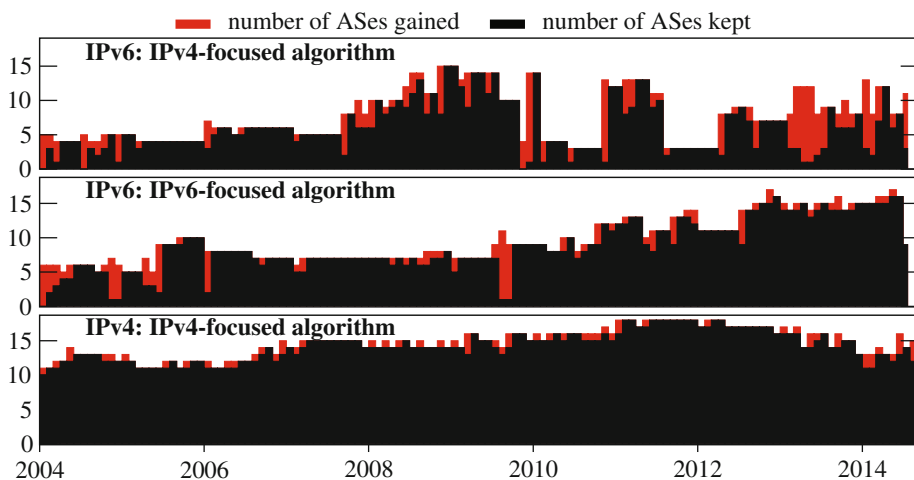


Fig. 4. By improving the way in which we infer the IPv6 transit-free clique, we reduce the average number of ASes that are added or removed between temporally adjacent cliques from 3.4 to 1.8, bringing the IPv6 clique’s stability closer to the average of 1.5 seen in IPv4.

transit degree alone an unreliable metric. Figure 4 illustrates the highly dynamic clique membership that results when applying our IPv4-focused algorithm to the BGP-observed IPv6 AS topology over the last decade; on average, 3.4 ASes changed between temporally adjacent cliques. We found that at least 11 of the 20 ASes most frequently inferred to be in IPv6 cliques had at least one transit provider in our validation data, contradicting the notion of the transit-free clique.

We therefore modified step 3 of Algorithm 1 to consider an AS’s *peering policy* and *reachability* in addition to the AS’s transit degree. An AS with an *open* peering policy will peer with other ASes, with few or no conditions; a *selective* policy requires conditions on traffic volume and symmetry; and a *restrictive* policy limits peering to as few networks as necessary. The peering policy of an AS expresses an important and relatively stable property of the AS, but is not easily inferred from the topology alone because most peering links are invisible in public BGP data [20]. We used the self-reported peering policy data in PeeringDB [2]; for networks with PeeringDB entries but without a registered peering policy, we assumed a restrictive policy, which operators tend not to disclose [19]. We required the seed ASes to follow a restrictive or selective policy; we did not select ASes with open peering policies as seed ASes even if they had the largest transit degree. In addition, we reduced the initial seed set to three ASes for years before 2007, and to five ASes for 2007 and onwards, based on the accuracy of inferences derived from these seed values. As with the IPv4 method, after we find a clique involving the seed ASes, we add other ASes to the clique whose addition




Fig. 5. The top 20 ASes most frequently inferred to be part of the IPv6 clique according to the improved inference algorithm that uses three metrics: peering policy, reachability degree and transit degree. We exclude another 18 ASes inferred less often to be in the clique. The improved algorithm yields a more stable inferred clique, which only include transit-free ASes for most snapshots.

does not result in triplets of consecutive clique members in the BGP-observed paths, implying one of the ASes in the triplet is receiving transit.

The use of a reachability metric is required because some transit-free ASes are partitioned from each other due to peering disputes [17,23]. The use of a partitioned AS as a seed can yield an incomplete clique. To minimize the chance of using a partitioned AS, we required that seed ASes provide direct BGP feeds to RV or RIS and announce routes to at least 90% of the BGP-visible IPv6 address space. Additionally, if an AS misses just one link from being part of the clique, we considered it a clique member to account for the reality of the currently partitioned IPv6 Internet [17,23], provided that the AS does not receive transit from one clique member to reach a second clique member, i.e. could not be transit-free. As with the IPv4 algorithm, if there are multiple cliques we select the clique with the largest transit degree sum. Note that some ASes previously used different AS numbers in IPv4 and IPv6, most notably Sprint (IPv4 AS1239, IPv6 AS6175) and Verizon (IPv4 AS701, IPv6 AS12702). Both ASes eventually used a single ASN for both IPv4 and IPv6, but when they were transitioning to a single ASN (i.e., the IPv4 ASN) they used both ASNs in the IPv6 AS topology. During the period when they used both ASNs, we merged the IPv6 AS links for both ASNs for these two organizations to capture their full connectivity during the period they were shifting all of their neighbors to their primary ASN.

Figure 5 shows the IPv6 clique inferred using the improved clique inference method. The improved method infers more stable IPv6 cliques that are composed of transit-free ASes with the number of ASes entering or leaving the clique reducing from an average 3.4 ASes with the IPv4-focused method to 1.8 with our IPv6-focused method. This improvement brings the edit distance between temporally adjacent cliques much closer to IPv4's average of 1.5 ASes.

4.3 Inferring Clique-Stub Relationships

After we infer the clique, we apply the rest of the steps in Algorithm 1 without modification until step 9, which infers stub ASes to be customers of clique ASes irrespective of whether we observed a clique AS exporting a stub AS as a customer. In IPv4, this step avoids misinferring backup transit links as p2p, and relies on the fact that no clique members have an open peering policy. Establishing backup transit links is a popular strategy for IPv4 ASes that need to ensure reliable connectivity in the face of failures, but backup transit links appear to be less critical in IPv6 given the low levels of traffic [8] and small size of the topology. Therefore, we skip step 9 of the algorithm.

4.4 Validation

We evaluated the positive predictive value (PPV) of our improved algorithm, defined as the proportion of inferences of a particular type that were correct. Figure 6 shows the PPV over time according to our validation datasets described in Sect. 3.2. Our algorithm achieves high PPV throughout the period of inferences (January 2004–July 2014), with PPV for both p2c and p2p inferences consistently above 96% after spring 2009 for the communities and after fall 2012 for the RPSLNg data. The PPV for inferences validated using the LocPref data is over 96% for the three points in time

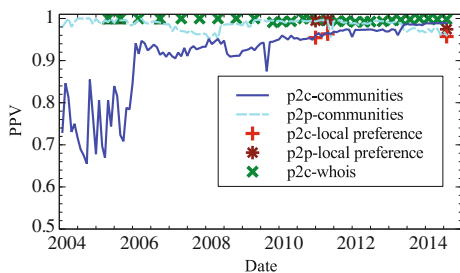


Fig. 6. Validation of our inferences over time using the three validation datasets described in Sect. 3.2. Validation results involving the BGP communities and local preference datasets are in strong agreement despite involving different ASes.

where we have LocPref data. The diversity of validation data sources and high PPV values strengthens our confidence in the suitability of our algorithm and the accuracy of our inferences. Figure 6 shows that both p2c and p2p relationships are inferred with high PPV, except for before 2006 when p2c inferences have a PPV of less than 80% for many BGP snapshots. However, our validation dataset (and the IPv6 AS topology) is considerably smaller (Fig. 1) prior to 2006.

5 Analysis

We compare IPv6 and IPv4 routing relationships starting from 2006 to avoid artifacts from inference errors on very small early topologies. *Congruent relationships* refer to dual-stack AS links with the same relationship type in IPv4 and IPv6; *disparate relationships* are dual-stack links where the relationship differs from IPv4 to IPv6. The fraction of disparate relationships decreases linearly over time (Fig. 7a), from 15% in 2006 to 5% in 2014, because congruent relationships increase in number faster than disparate ones, suggesting convergence.

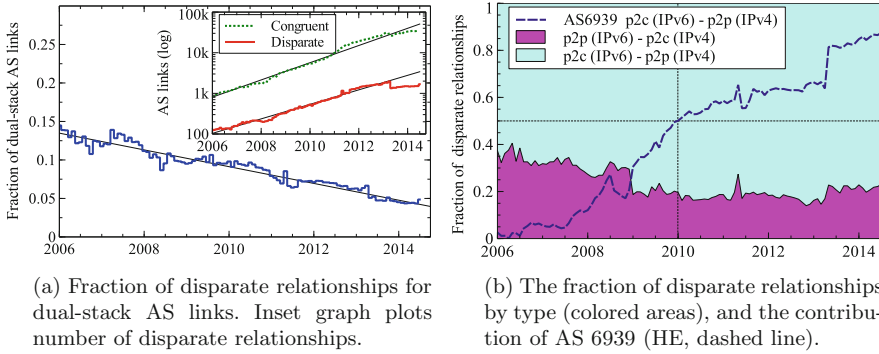


Fig. 7. The fraction of disparate relationships decreases over time to about 5% in July 2014, showing convergence between the IPv4 and IPv6 topologies. Most disparate relationships after 2010 are due to HE’s free IPv6 transit service.

Figure 7b shows the fraction of disparate relationships by relationship type. Most inferred disparate relationships are p2c in IPv6 and p2p in IPv4, and the remaining disparate relationships are p2p in IPv6 and p2c in IPv4. HE (AS6939) contributes over 50% of the disparate $IPv6_{p2c}/IPv4_{p2p}$ relationships after 2010, peaking in July 2014 when it contributed 87% of observed disparate relationships. These observations are consistent with the behavior of the IPv6 tunnel broker service provided by HE, which allows free transit, so that IPv4 peers can be IPv6 customers without cost [1]. This strategy allows HE to acquire many IPv6 (free-transit) customers compared to IPv4, illustrated by comparing HE’s IPv4 and IPv6 customer cones. The customer cone is defined as the ASes that an AS can reach by following a customer link (the AS’s customers, customers of those customers, and so on) and is a metric of influence of an AS in the transit market. Figure 8 compares the relative size of the customer cones between IPv4 and IPv6 over the last 9 years for the 9 largest providers as of July 2014. HE is the only AS with a significantly larger customer cone in IPv6 than in IPv4 (over




Fig. 8. The fraction of ASes in the customer cones of the largest 9 providers. The left plot compares the customer cones between IPv6 and IPv4 for July 2014. The right plot shows their growth over time. HE (AS6939) has a larger customer cone in IPv6 than in IPv4; most other ASes have smaller, but still growing customer cones (except for AS3549, Global Crossing, which merged with Level3 in 2012).

50 % of IPv6 ASes). However, the relative sizes of customer cones of the largest providers have an increasing trend, in contrast with the trend observed in the IPv4 topology [18]. The only exception is Global Crossing (AS3549), which was acquired by Level3 (AS3356) in 2012.

6 Conclusion

The low level of IPv6 deployment has hindered efforts to accurately infer IPv6 AS relationships. We tackled this challenge by modifying CAIDA’s IPv4 relationship inference algorithm, which required a focus on the correct inference of the IPv6 clique. The clique is a crucial component of AS topology, but with fundamental disparities between IPv4 and IPv6, including the extreme peering openness of some IPv6 ASes, and long-lived peering disputes among transit-free IPv6 networks. To overcome these obstacles, we used two new metrics to help filter out topological inconsistencies in IPv6: peering policy and BGP-observed reachability.

We validated ten years of our algorithm’s inferences against three data sources: BGP communities, RPSLNg, and local preference values, which covered 25 % of the BGP-observed topology for July 2014. Our inferences achieved an overall Positive Predictive Value of at least 96 % for each dataset since 2009, with increasing accuracy over time. We found that dual-stack relationships are increasingly congruent, as disparate relationships decreased from 15 % in 2006 to 5 % in 2014, while the number of nodes and links increased by a factor of 14.5 times and 22 times, respectively. Notably, disparate relationships are now dominated by a single AS, Hurricane Electric, whose long-standing offer of free IPv6 transit has enabled it to become the dominant transit-free provider in IPv6, with the largest customer cone in the IPv6 topology, despite not even being a transit-free network in IPv4.

Our validation and inference data is available at <http://www.caida.org/publications/papers/2015/asrank6/>

Acknowledgements. The work was supported by U.S. NSF grant CNS-1111449, DHS S&T Cyber Security Division (DHS S&T/CSD) BAA 11-02 and SPAWAR Systems Center Pacific via contract number N66001-12-C-0130, and by DRDC pursuant to an Agreement between the U.S. and Canadian governments for Cooperation in Science and Technology for Critical Infrastructure Protection and Border Security. The work represents the position of the authors and not of NSF, DHS, or DRDC.

References

1. Hurricane electric IPv6 tunnel broker. <https://tunnelbroker.net/>
2. PeeringDB. <http://www.peeringdb.com>
3. RIPE Routing Information Service (RIS). <http://www.ripe.net/ris>
4. Route Views. <http://www.routeviews.org/>
5. Blunk, Y., Damas, J., Parent, F., Robachevsky, A.: Routing policy specification language next generation (RPSLNg). RFC 4012, March 2005

6. Chandra, R., Traina, P., Li, T.: BGP communities attribute. RFC 1997. August 1996
7. Colitti, L., Gunderson, S.H., Kline, E., Refice, T.: Evaluating IPv6 adoption in the internet. In: Krishnamurthy, A., Plattner, B. (eds.) PAM 2010. LNCS, vol. 6032, pp. 141–150. Springer, Heidelberg (2010)
8. Czyz, J., Allman, M., Zhang, J., Iekel-Johnson, S., Osterweil, E., Bailey, M.: Measuring IPv6 adoption. In: SIGCOMM, pp. 87–98 (2014)
9. Dhamdhere, A., Luckie, M., Huffaker, B., claffy, k., Elmokashfi, A., Aben, E.: Measuring the deployment of IPv6: topology, routing and performance. In: IMC, pp. 537–550 (2012)
10. Gao, L.: On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans. Netw.* **9**(6), 733–745 (2001)
11. Gill, P., Schapira, M., Goldberg, S.: A survey of interdomain routing policies. *CCR* **44**(1), 28–34 (2013)
12. Giotsas, V., Luckie, M., Huffaker, B., claffy, k.: Inferring complex AS relationships. In: IMC, November 2014
13. Giotsas, V., Zhou, S.: Detecting and assessing the hybrid IPv4/IPv6 AS relationships. In: SIGCOMM Poster, pp. 424–425 (2011)
14. Giotsas, V., Zhou, S.: Valley-free violation in Internet routing - analysis based on BGP community data. In: IEEE ICC 2012 CQRM, pp. 1208–1212, June 2012
15. Internet Archive: Wayback Machine. <http://archive.org/web/>
16. Karpilovsky, E., Gerber, A., Pei, D., Rexford, J., Shaikh, A.: Quantifying the extent of IPv6 deployment. In: Moon, S.B., Teixeira, R., Uhlig, S. (eds.) PAM 2009. LNCS, vol. 5448, pp. 13–22. Springer, Heidelberg (2009)
17. Leber, M.: IPv6 internet broken, cogent/telia/hurricane not peering, NANOG. <http://mailman.nanog.org/pipermail/nanog/2009-October/014017.html>
18. Luckie, M., Huffaker, B., claffy, K., Dhamdhere, A., Giotsas, V.: AS relationships, customer cones, and validation. In: IMC 2013, pp. 243–256 (2013)
19. Norton, W.: Restrictive peering policy. DrPeering International. <http://drpeering.net/FAQ/What-is-a-restrictive-peer.html>
20. Oliveira, R., Pei, D., Willinger, W., Zhang, B., Zhang, L.: The (in)completeness of the observed Internet AS-level structure. *IEEE/ACM Trans. Netw.* **18**(1), 109–122 (2010)
21. UCLA: Internet AS-level topology archive. <http://irl.cs.ucla.edu/topology>
22. Zmijewski, E.: A Bakers Dozen, 2013 Edition. Renesys Blog. <http://www.renesys.com/2014/01/bakers-dozen-2013-edition/>
23. Zmijewski, E.: World IPv6 Day. Renesys Blog. <http://www.renesys.com/2011/06/world-ipv6-day/>