Inferring Complex AS Relationships

Vasileios Giotsas
CAIDA / UC San Diego
vgiotsas@caida.org

Matthew Luckie
CAIDA / UC San Diego
mjl@caida.org

Bradley Huffaker
CAIDA / UC San Diego
bradley@caida.org

kc claffy
CAIDA / UC San Diego
kc@caida.org

ABSTRACT

The traditional approach of modeling relationships between ASes abstracts relationship types into three broad categories: transit, peering, and sibling. More complicated configurations exist, and understanding them may advance our knowledge of Internet economics and improve models of routing. We use BGP, traceroute, and geolocation data to extend CAIDA's AS relationship inference algorithm to infer two types of complex relationships: hybrid relationships, where two ASes have different relationships at different interconnection points, and partial transit relationships, which restrict the scope of a customer relationship to the provider's peers and customers. Using this new algorithm, we find 4.5% of the 90,272 provider-customer relationships observed in March 2014 were complex, including 1,071 hybrid relationships and 2,955 partial-transit relationships. Because most peering relationships are invisible, we believe these numbers are lower bounds. We used feedback from operators, and relationships encoded in BGP communities and RPSL, to validate 20% and 6.9% of our partial transit and hybrid inferences, respectively, and found our inferences have 92.9% and 97.0% positive predictive values. Hybrid relationships are not only established between large transit providers; in 57% of the inferred hybrid transit/peering relationships the customer had a customer cone of fewer than 5 ASes.

Categories and Subject Descriptors

C.2.5 [**Local and Wide-Area Networks**]: Internet; C.2.1 [**Network Architecture and Design**]: Network topology

Keywords

AS relationships; complex routing policies

1. INTRODUCTION

Analysis and modeling of the Internet's AS topology requires accurate knowledge of AS relationships, which reflect often private negotiations between ASes. The details of

these relationships are not comprehensively published anywhere, and many of them are subject to non-disclosure agreements, motivating many researchers to develop algorithms to heuristically infer AS relationships based on AS path information available in public BGP data [7–10, 12, 15, 22–24, 30, 32, 33]. To simplify the development of inference heuristics, researchers traditionally abstracted relationships into three (and sometimes only the first two) classes: provider-customer (p2c), peering (p2p) and sibling (s2s). In a p2c or transit relationship, a customer buys access to achieve global reachability. In a p2p relationship, two ASes share access to their networks and their customers' networks. In an s2s relationship, two ASes under common ownership may provide mutual transit to each other.

This oversimplification ignores more complex relationships and may introduce artifacts into the study of inter-domain routing [31], such as spurious relationship cycles [11], artificial policy violations [24, 30], and generally inaccurate AS path prediction [9, 23]. These problems have inspired the development of a relationship-agnostic AS topology model as an alternative [25, 26]. While BGP is a prefix-based routing protocol and AS relationships are a course-grained abstraction, operators themselves express their policies in terms of AS relationships, and understanding classes of relationships can help researchers predict routing decisions.

In this paper we propose an algorithm to infer the two most common types of complex AS relationships: *hybrid relationships*, where two ASes have different relationships at different interconnection points [28] (e.g. p2c in one location and p2p elsewhere), and *partial transit relationships*, which restrict the scope of a p2c relationship to the provider's peers and customers (but not providers) [10, 13, 22, 28, 29]. We can define these complex relationships as special cases of the traditional p2c and p2p types, which allows us to leverage CAIDA's relationship inference algorithm instead of designing an entirely new algorithm.

Section 2 provides an overview of Internet routing and related work. Section 3 describes the BGP, traceroute, and geolocation data we used in this study, and presents our methodology to infer complex AS relationships. Section 4 presents our inferences: 1,071 hybrid relationships and 2,955 partial transit relationships, used mostly by ASes with a small customer cone. Our data suggests that IXPs facilitate the establishment of hybrid relationships by stub ASes at diverse geographical locations. Using feedback from operators and relationships encoded in BGP communities and RPSL, we found that our hybrid and partial transit inferences had positive predictive values of 92.9% and 97.0%, respectively.

Figure 1: **Hybrid relationship between ASes** A **and** B **at different points of presence (PoPs). In** PoP_1, A **is a customer of** B, **while in** PoP_2 **they are peers.** B's **export policy depends on the relationship between the ASes at each** PoP**; prefix** p_1 **is announced to all neighbors of B because A receives transit at** PoP_1, **while** p_2 **is only announced to customers of B because A receives peering at** PoP_2.
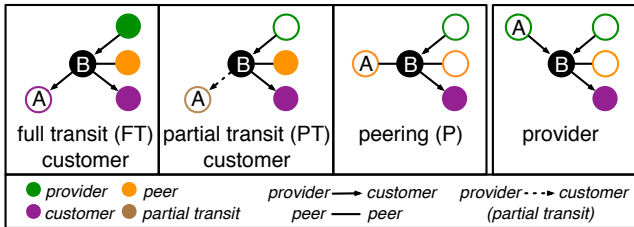

Figure 2: **Partial transit compared with traditional relationship types.** B's **export policy to** A **changes based on** B's **relationship to** A, **with routes from the filled AS nodes exported to** A **in each scenario. In a partial transit relationship with** A, B **exports only customers and peers to** A.

2. BACKGROUND

Interdomain routing is a collaborative effort among ASes, which interconnect and exchange routing information using the BGP protocol. Many ASes negotiate contractual agreements that impose technical restrictions on traffic exchange. The export policies that derive from relationships between ASes combined with traffic engineering largely define routing at the AS level. Unfortunately, operators are reluctant to publish their AS relationships which are often confidential. For more than a decade, researchers have studied how to infer the type of business relationships between networks through the analysis of AS-level paths [7,10,12,15,22,23,32].

Gao's seminal work [15] proposed a classification into three abstract relationship types; p2c, p2p, and s2s. Although these classifications capture the majority of AS interconnections, more complicated relationships exist. Norton defines two additional relationship types: *hybrid* and *partial transit* [28,29]. Hybrid relationships arise when two ASes agree to different relationship types at different inter-connection points of presence (PoPs); typically, the ASes will be peers in particular regions, but one AS will be a transit customer of the other AS elsewhere. Figure 1 illustrates the routing that results in such a relationship; B may advertise prefix p_1

Type	Definition
Transit	An AS offers another AS access to its peers, customers, and providers.
Peering	Two ASes offer access to their customer routes to each other.
Partial transit	An AS offers another AS transit to its peers and customers, but not providers.
Hybrid	An AS offers another AS a combination of full transit, partial transit, or peering according to interconnection location.

Table 1: **Relationship types that we inferred. This paper focuses on methods to infer hybrid and partial transit relationships.**

to all of its neighbors, but B may only advertise p_2 to its customers. In some cases, B may only advertise p_2 to customers in the region served by PoP_2. Partial transit relationships arise when a provider sells access to its customers and peers but not to its providers. Figure 2 compares the routing differences between the traditional provider, peer, and customer relationships and the partial transit relationship; in the traditional customer relationship, A receives routes from B's customers, peers, and providers, but in the partial transit case A receives just customer and peer routes. Table 1 summarizes the four relationship types.

Using only conventional transit and peering relationships to study the AS topology is problematic [31]. For example, in 2005 Mao *et al.* concluded that complex relationships may contribute to the inability to accurately predict AS-level paths [23]; in 2013 Deng *et al.* [9] reached a similar conclusion. Mühlbauer *et al.* proposed a relationship-agnostic AS topology model to capture complex BGP policies that cannot be modeled based on the conventional relationship abstraction [25, 26].

A survey of operators by Dimitropoulos *et al.* in 2007 [10] confirmed the existence of relationships that vary across peering points and prefixes. In 2008, Faratin *et al.* suggested that AS relationships will become increasingly complex because the growth of Content Distribution Networks (CDNs) and eyeball networks change the perceptions of symmetry in traffic delivery costs [13]. Yet, there is little work that tackles the complicated challenges of inferring complex relationships, and little validation data available. In 2013, Neudorfer *et al.* proposed a method to infer complex relationships by examining policy violations (unconventional export policies) using the conventional transit and peering definitions at different Points-of-Presence (PoPs) [27]. However, they were only able to manually identify a single hybrid relationship, illustrating the difficulties involved in the inference of complex relationships.

In 2013, we presented and validated a new AS relationship inference algorithm that also only inferred conventional transit and peering relationships [22]. Our top-down approach began by inferring a Tier-1 clique, applied heuristics to infer c2p links based primarily on how neighbors were observed to export routes, and inferred the remainder to be p2p. That algorithm inferred hybrid and partial transit relationships as conventional transit relationships because at least some routes from the neighbor AS were exported to peers or providers. To account for the effect of hybrid relationships, we presented a new algorithm for inferring the *customer cone* – the set of ASes that can be reached by

an AS using its customer links. Previous customer cone algorithms used a recursive model, inferring all customers of downstream customers to be in an AS's customer cone; complex relationships and traffic engineering behavior can distort such inferences. We addressed this limitation with the *provider/peer observed customer cone*, where we computed the customer cone of an AS using routes we observed from providers and peers of the AS. Because B will not export routes received from the peering portion of a hybrid relationship with customer C to provider A, we did not include these routes when computing the customer cone of A. Similarly, because B will not export routes from a partial transit customer C to provider A, we did not include these routes when computing the customer cone of A.

3. INFERENCE METHODOLOGY

Figure 3 shows a high-level overview of our process to infer complex relationships. First, we infer conventional transit and peering relationships and customer cones using our IMC 2013 algorithm [22] and publicly available BGP data. Because our IMC 2013 algorithm infers hybrid and partial transit relationships as conventional transit relationships, our new algorithm uses BGP data to infer per-prefix export policy for all conventional transit relationships, resulting in sets of partial transit and candidate hybrid relationships. Then, we use traceroute and geolocation data to infer if export policies of candidate hybrid relationships differ by geographic location. Our algorithm then classifies candidate hybrid relationships that have geographic characteristics as hybrid, with specific export policy classifications per geographic location.

3.1 Data

BGP data: We obtained routing information from BGP table snapshots collected by Route Views (RV) [5] and RIPE RIS (RIS) [4], which record BGP routing information by peering with globally distributed ASes. For each RV and RIS collector we downloaded one RIB file every five days for March 2014; we retained only paths that were observed in all files to reduce false positive complex relationship inferences. We extracted three attributes from each BGP route: the AS path, the prefix, and the communities string. We discarded AS paths with symptoms of misconfiguration or poisoning, e.g., loops, unassigned ASes, and non-adjacent Tier-1 ASes [22]. We also discarded prefixes longer than /24, which many ASes do not route. We used path and prefix attributes to identify partial transit and candidate hybrid relationships. We used available communities strings as a source of geographically-tagged interconnection points and to validate our complex relationship inferences.

Traceroute data: To derive the geographical footprint of AS relationships, we processed interface-level paths derived from two sources of traceroute data to infer the PoP-level AS connectivity. First, we used CAIDA's IPv4 Routed /24 Topology Dataset gathered by the Archipelago (Ark) [34] measurement infrastructure in March 2014 [19]. As of March 2014, Ark supports continual Paris traceroute measurements from 94 monitors in 84 different ASes distributed in 39 different countries, probing to one random IP address in each /24 of the entire routed IPv4 space approximately every three days. Second, we coordinated probing from thousands of public traceroute servers to expand coverage beyond what these existing Ark measurements capture. We developed
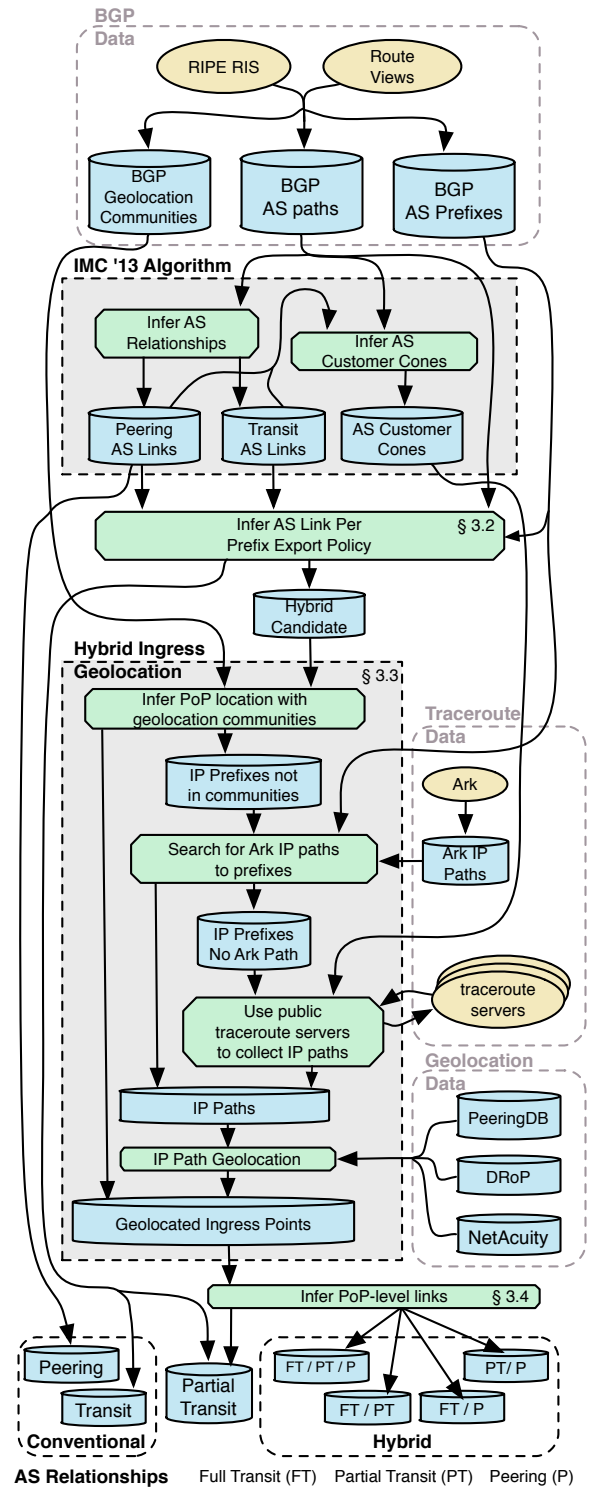

Figure 3: Overview of our process to infer complex AS relationships. We begin by inferring conventional peering and transit relationships and customer cones using AS paths from public BGP data. We infer partial transit and candidate hybrid relationships using paths and prefixes from BGP and conventional relationships. We use traceroute and geolocation data to infer hybrid relationships, and classify the arrangement by policy and location.

an overlay interface to interact with 2,509 public traceroute servers among 507 different ASes in 77 different countries. To avoid being blocked by traceroute servers for too frequent probing, we limited probing through this interface to one query every 10 seconds per server. We used customer cone data to inform target selection so that we would traceroute paths toward destinations most likely to cross a candidate hybrid relationship and therefore help with geolocation.

Geolocation data: To infer if export policies differ by geographic location, we used four sources of data, in order of preference: published BGP community information; PeeringDB's reverse DNS scan on IXP peering prefixes; strings found in DNS hostnames; and a commercial database.

Some ASes use communities to record the ingress location where a route entered their network. Community values and their corresponding meanings are not standardized, but many operators document their usage of communities in IRR records or network support pages, allowing us to compile a dictionary of values and their corresponding meanings. In total, we interpreted 1,533 values from 117 ASes that encoded the geographic location of AS interconnection points.

We used three data sources to map addresses observed in traceroute to geographic locations: PeeringDB's reverse DNS lookups of IXP prefixes [3], CAIDA's DNS-based Router Positioning (DRoP) [17], and Digital Envoy's NetAcuity [2]. We used PeeringDB's list to map addresses to IXP locations. We used DRoP to interpret DNS strings for geographic hints such as airport codes, city names, and CLLI codes (a geographic naming convention developed by the telecommunication industry [1]). If neither PeeringDB or DRoP reveals a location, we used NetAcuity, a commercial geolocation provider, to map the IP address to a geographic location. NetAcuity is more accurate at router positioning than other commercial geolocation databases, though it is optimized to infer edge host location (servers and end users) rather than router infrastructure IP addresses [18].

3.2 Inferring Prefix Export Policies

Hybrid relationships result in different export policies for different sets of prefixes across the same AS relationship. We therefore divide export policies into three types of increasingly restrictive policy: full transit (FT), partial transit (PT), and peering (P). To illustrate, in figure 2, FT is the only policy where the provider exports prefixes from its providers (in addition to peers and customers) to the customer. PT is more restrictive because the provider does not export provider routes, and P is the most restrictive because the provider exports only customer routes. We therefore annotate a PoP where two ASes interconnect with the most restrictive policy we observe in BGP for all prefixes we observe with traceroute to cross the link at that PoP.

Our algorithm begins by inferring the most restrictive export policy we observe in BGP for *link-prefix tuples*: for each link our IMC 2013 algorithm inferred to have a conventional transit relationship, we tag individual prefixes that propagate in BGP through the link. Figure 4 illustrates the approach; for a transit relationship between A and B where our IMC 2013 algorithm inferred A is a customer of B, we tag each prefix p_x that B exports to other neighbors it received from A. If we observe a triplet where B exports p_x to a provider, we tag the A-B p_x tuple with FT. If we do not observe an FT-triplet, but do observe a triplet where B exports p_x to a peer, we tag the tuple with PT. If we do
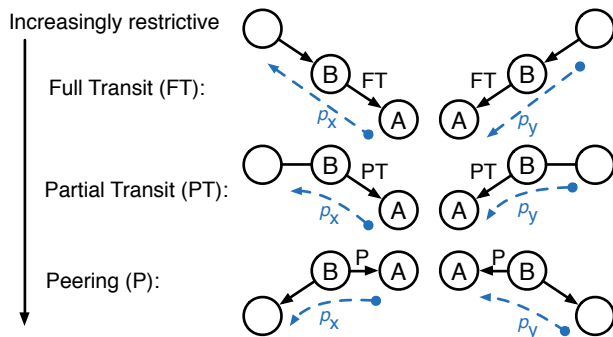

Figure 4: Tagging a link-prefix tuple between ASes A and B where our IMC 2013 algorithm infers A to be a customer of B. p_x is a prefix from A exported by B; p_y is exported by B to A. We tag a link-prefix tuple with the most restrictive export policy that could explain the observed routing for the prefix: FT if exported by B to (from) its providers; PT if exported by B to (from) its peers and customers but no further; P if exported by B to (from) its customers but no further.

observe either FT or PT triplets, we tag the tuple with P. We use a similar process with each prefix p_y that B exports to A, and it is rare to BGP-observe B export the same prefix both to and from A, i.e., $p_x \neq p_y$ for any A-B link.

If all tuples for a link are tagged FT, the relationship remains conventional provider-customer, and if all tuples are tagged PT, we classify the relationship as PT. However, if a given link has a mixture of FT, PT, or P tuples, we tag it as a candidate hybrid. At the end of this process, we infer 2,955 partial transit and 6,682 candidate hybrid relationships.

3.3 Candidate Hybrid Exploration

The presence of different export policies for different prefixes for the same AS link is an indication of a hybrid relationship, but it may also reflect traffic engineering practices that restrict the scope of route advertisements. To distinguish between hybrid relationships and traffic engineering, our technique relies on the ability to geolocate the ingress point of each link-prefix tuple. Depending on the combination of prefixes tagged as FT, PT, and P, a candidate hybrid relationship may fall into one of four categories; FT + PT + P, FT + PT, FT + P, or PT + P. We attempt to group prefixes with the same tag into mutually exclusive PoPs such that each PoP has a consistent BGP-observed export policy. For example, if all FT-tagged prefixes can be geolocated to paths that cross PoP_1 and all the P prefixes to PoP_2, then we infer the hybrid relationship to be FT + P. Conversely, if a prefix tagged as P crosses PoP_2, but so do prefixes tagged as FT or PT, then PoP_2 cannot represent a peering PoP, and we infer PoP_2 as FT or PT.

Therefore, given a candidate hybrid relationship with N_P prefixes tagged P, N_{PT} prefixes tagged PT, and N_{FT} prefixes tagged FT, we begin by inferring all interconnection PoPs for the P prefixes. We then try to rule out the PoPs as being peering PoPs by testing as many PT prefixes as necessary to find the same set of PoPs with a PT export policy; i.e., an AS cannot have a P policy at a PoP if it also exports PT prefixes at the PoP. Finally, we test as many FT prefixes as necessary to find the same set of PoPs as found for the P and PT cases.

We used a combination of BGP, traceroute, and geolocation data to infer PoPs for candidate hybrid relationships. For each candidate hybrid relationship's prefixes we searched BGP paths for a community that stores a geographic hint, which is our most trusted source of data. For each remaining prefix we tried to obtain a traceroute path toward the prefix that crosses the link. If a link is not observed with traceroute in any path toward a tagged prefix, we cannot infer the relationship as hybrid because we do not have adequate data to make a robust classification. In this case, the inferred relationship remains conventional transit to avoid misinterpreting traffic engineering as a complex relationship.

3.3.1 Obtaining traceroute paths

To minimize the workload we impose on public traceroute servers, we begin by using traceroute paths collected in existing Ark traces. For the set of prefixes and candidate hybrid relationships for which we find no usable Ark trace, we coordinate a distributed traceroute campaign using public traceroute servers. We traceroute only to prefixes that B exports from A, rather than the inverse, because in some cases B may export a complete routing table to A; i.e., in figure 4 we use p_x but not p_y. For each candidate hybrid link, we divide the available traceroute servers into two sets, depending on their relationship to the provider of the candidate hybrid relationship. The *full-visibility* set includes traceroute servers under the customer cone of the provider of the candidate relationship, which may be able to reach any of the targeted prefixes through the candidate hybrid relationship. The *limited-visibility* set includes traceroute servers in ASes that peer with the provider of the candidate relationship, or in ASes under whose customer cone the provider belongs. We consider traceroutes from these servers able to reach prefixes tagged either FT or PT through the candidate hybrid relationship, but not prefixes tagged P, since a peer should not export prefixes obtained from a peer to another peer. In figure 1, a traceroute server inside AS E would belong to the full-visibility set for the B-A relationship, a server inside AS C or AS D would belong to the limited-visibility set, and a server in a provider of D would not be used.

After classifying traceroute servers, we randomly select an IP address from each target prefix and traceroute it with the servers. Note that having a traceroute server in the full or limited-visibility groups does not imply that a traceroute will cross the targeted link, only that it could. Capturing possible peering PoP-level connectivity for a candidate hybrid relationship requires at least one traceroute server of full-visibility reaching the target prefixes through the link. For 97% of the candidate hybrid relationships we have at least one traceroute monitor of full-visibility, and for 85% of the of the links we have more than 10 full-visibility monitors. For 462 of our 6,682 candidate hybrid relationships, we could not identify ingress points due to lack of full-visibility traceroute servers and geographic communities.

3.3.2 Geolocation of Hybrid Links

If a traceroute toward a selected prefix crosses the candidate hybrid relationship, we use the path to geolocate the PoP. Identifying the interdomain link is difficult, since it is not clear who assigned the prefix used to establish routing. Consider the observed IP path a^0 a^1 b^0 in figure 5: changing the interdomain link from a^0–a^1 to a^1–b^0 has no effect on the IP path. Although it is not clear whether the interdomain
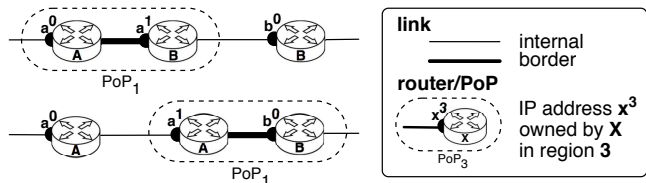

Figure 5: Regardless of which two routers form the interdomain link, interface a^1 is always part of the pair forming the interdomain link and so will be located in the same region.

Geolocation	Number	Percent
Communities	61,198	34.4%
PeeringDB	6,228	3.5%
DRoP	56,111	31.5%
Netacuity	54,590	30.6%

Table 2: The number of link-prefix tuples which were geolocated by each system, in the order they were geolocated.

link was just before or after the AS changed in the path, it does not matter for geolocation. The IP address just before the change will be on one of the two routers connected to the interdomain link and so will be in the same location as the target link. Therefore, we geolocate the IP address just before the AS changed, first using PeeringDB's reverse DNS scan, then by DRoP's inference given the hostname string, and then with Netacuity.

3.3.3 Classification of Hybrid Links

We use the BGP, traceroute, and geolocation data we collected and the logic in section 3.3 to infer which candidate hybrid links represent hybrid links. Table 2 shows which method geolocated each link-prefix tuple. For hybrid relationships where we find PoPs with a consistent export policy, we annotate the geolocated PoPs with specific FT, PT, and P classifications.

4. RESULTS

Of the 90,272 p2c relationships we inferred for March 2014, we inferred 4,026 (4.5%) of these relationships to be complex, including 1,071 hybrid AS relationships and 2,955 partial-transit relationships. The inferred hybrid relationships include 969 FT/P relationships, 72 PT/P relationships, 30 FT/PT relationships, and 4 FT/PT/P relationships. We found that hybrid relationships are not only established between large transit providers; in 57% of the hybrid transit/peering (T/P) relationships the customer had a customer cone size of less than 5 ASes. However, other metrics reflect network size differently than customer cone. For instance, we inferred 21 hybrid relationships that involve Akamai (AS 20940); although Akamai has no customers it is reported to be one of the top 10 networks in terms of inter-domain traffic volume [21]. For 34.6% of the hybrid T/P relationships the customers have traffic levels of at least 100 Gbps, according to traffic volumes self-reported in PeeringDB. The corresponding percentage for FT (customer relationships) in the whole BGP-observed graph is only 13.3%.

For 68.5% of the inferred hybrid T/P relationships, the peering link crossed a European IXP. Some ASes with open or selective peering policies are establishing peering relation-

	Hybrid			Partial		
	TP	FP	FN	TP	FP	FN
Direct report	33	2	1	2	0	0
Communities	124	10	4	158	5	0
RPSL	45	-	-	38	-	-
Validated	214 / 1,071: 20.0%			203 / 2,955: 6.9%		
Confirmed	202 / 1,071: 18.9%			198 / 2,955: 6.7%		
PPV	157 / 169: 92.9%			160 / 165: 97.0%		

Table 3: Validation results, showing true positives (TP), false positives (FP), and false negatives (FN). We validate 20.0% of our hybrid inferences and 6.9% of our partial transit inferences and found our inferences have 92.9% and 97.0% positive predictive values (PPV = TP/TP+FP), respectively.

ships with their customers over IXPs, sometimes unknowingly, as revealed by our discussions with some operators.

Partial transit relationships are also more frequently observed among European ASes. 88% of the inferred partial transit relationships were established in European PoPs. We believe that the dense peering connectivity among European ASes [6,16] makes partial transit more attractive, since partial transit providers can offer discounted access to a large fraction of their routing table through partial transit only. For instance, Init7 (AS13030), a large European provider, reports that 60% of their routing table is accessible through peering links [20] as of September 2014.

Despite the many complex relationships we inferred, and positive validation, we are limited by the well-known AS topology incompleteness problem [31], as our algorithm can only infer complex relationships where the peering component is revealed in public BGP data.

5. VALIDATION

We used three sources of validation data: direct e-mail feedback, BGP communities that signal relationship type, and relationship types expressed in different Routing Policy Specification Language (RPSL) objects.

Direct feedback: We obtained feedback from seven operators (of twelve contacted) that previously contributed AS relationship corrections through CAIDA's web interface [22]. We sent each operator our inferences for their ASes and asked them to specify if they were correct, and asked them if they were involved in other hybrid or partial-transit relationships not included in our inferences.

BGP communities: We compiled a dictionary of 1,502 communities defined by 281 ASes, which we used to extract a set of 40,820 relationships for March 2014, as explained in [22]. Relationship communities enable the implementation of sophisticated routing policies, so operators have a strong incentive to configure community values with correct relationship annotations [14]. We considered five types of relationship communities: customer, partial-customer, peer, partial-provider, and provider. We used the partial-customer and partial-provider communities to obtain validation data for partial transit relationships. We captured hybrid relationships when we observed that an AS tagged different inbound prefixes from the same neighbor with different community values depending on the ingress PoP. To mitigate transient misconfigurations we required that the same communities be observed in all BGP snapshots we collected.

RPSL: We used RPSL objects to evaluate only true positives, since RPSL objects are commonly expressed at a high level and in most cases do not encode complex relationships.

Table 3 summarizes our validation results. Overall, we were able to confirm 202 hybrid relationships (18.9% of the total hybrid inferences), and 198 partial-transit relationships (6.7% of the total partial-transit inferences). For the validation datasets that allowed us to test both true and false positives (directly reported and BGP communities) we had the correct inference for 157/169 (92.9%) of the inferred hybrid relationships, and for 160/165 (97.0%) of inferred the partial transit relationships. We failed to infer five hybrid relationships present in our validation dataset. For four, both the transit and the peering PoP were located in the same city, and our city-level geolocation granularity was too coarse to identify the different PoPs.

Via operator feedback we found that three of the inferred hybrid relationships were the result of misconfigurations, where the intended relationship was provider-customer. One operator followed up with a request for a semi-live system to alert operators of hybrid peering relationships, since they may be unintentional.

6. CONCLUSION

We presented a new algorithm to infer the two most common types of complex AS relationships: hybrid and partial transit using BGP, traceroute, and geolocation data. We inferred 1,071 hybrid and 2,955 partial-transit relationships for March 2014. We validated our inferences against direct feedback from operators, BGP communities, and RPSL data and found our hybrid and partial transit inferences have 92.9% and 97.0% positive predictive values, respectively. We have published our complex relationship inferences and validation data derived from BGP communities and RPSL at http://www.caida.org/publications/papers/2014/complex

We believe this is the first partly validated attempt to infer complex AS relationships. Our results reveal that complex relationships are more prevalent in the periphery of the AS topology than previously thought, while 61% of hybrid peering and 88% of partial transit relationships were inferred between European ASes that leverage Europe's extensive IXP ecosystem. Generally, unconventional relationships can have arbitrary complexity that may not be expressible in terms of relationship types, which is why we focused on the two most common types of complex relationships according to operator feedback and the existing literature, which are hybrid and partial transit. In the future we plan to expand our active measurements by integrating traceroute probes from RIPE Atlas, and we will continue to interact with AS operators for input and feedback to improve our inferences.

7. REFERENCES

[1] Common Language Location Identifier (CLLI) code. http://en.wikipedia.org/wiki/CLLI_code.

[2] NetAcuity. http://www.digitalelement.com/our_technology/our_technology.html.

[3] PeeringDB. http://www.peeringdb.com.

[4] RIPE RIS. http://www.ripe.net/ris.

[5] RouteViews. http://www.routeviews.org/.

[6] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a large European IXP. In *ACM SIGCOMM*, pages 163–174, Aug. 2012.

[7] G. D. Battista, T. Erlebach, A. Hall, M. Patrignani, M. Pizzonia, and T. Schank. Computing the types of the relationships between autonomous systems. *IEEE/ACM Transactions on Networking*, 15(2):267–280, Apr. 2007.

[8] H. Chang, S. Jamin, and W. Willinger. To peer or not to peer: Modeling the evolution of the Internet's AS-level topology. In *INFOCOM*, pages 1–12, Apr. 2006.

[9] W. Deng, W. Muhlbauer, Y. Yang, P. Zhu, X. Lu, and B. Plattner. Shedding light on the use of AS relationships for path inference. *Communications and Networks*, 14(3):336–345, June 2013.

[10] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy, and G. Riley. AS relationships: inference and validation. *ACM SIGCOMM CCR*, 37(1):29–40, Jan. 2007.

[11] X. Dimitropoulos, M. A. Serrano, and D. Krioukov. On cycles in AS relationships. *ACM SIGCOMM CCR*, 38(3):102–104, July 2008.

[12] T. Erlebach, A. Hall, and T. Schank. Classifying customer-provider relationships in the Internet. In *IASTED International Conference on Communications and Computer Networks*, Nov. 2002.

[13] P. Faratin, D. Clark, S. Bauer, W. Lehr, P. Gilmore, and A. Berger. The growing complexity of Internet interconnection. *Communications and Strategies*, (72):51–71, Dec. 2008.

[14] K. Foster. Application of BGP communities. *The Internet Protocol Journal*, 6(2):2–9, June 2003.

[15] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, Dec. 2001.

[16] V. Giotsas, S. Zhou, M. Luckie, and k claffy. Inferring multilateral peering. In *ACM CoNEXT*, pages 247–258, Dec. 2013.

[17] B. Huffaker, M. Fomenkov, and k. claffy. DRoP: DNS-based router positioning. *ACM SIGCOMM CCR*, 44(3):6–13, July 2014.

[18] B. Huffaker, M. Fomenkov, and k. claffy. Geocompare: A comparison of public and commercial geolocation databases. Technical report, CAIDA, May 2011. http://www.caida.org/publications/papers/2011/geocompare-tr/.

[19] Y. Hyun, B. Huffaker, D. Andersen, M. Luckie, and K. C. Claffy. The IPv4 Routed /24 Topology Dataset, 2014. http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml.

[20] INIT7. Init7 partial transit. http://www.init7.net/en/wholesale/partial-transit/.

[21] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. In *ACM SIGCOMM*, pages 75–86, Aug. 2010.

[22] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and k claffy. AS relationships, customer cones, and validation. In *ACM SIGCOMM IMC*, pages 243–256, Oct. 2013.

[23] Z. M. Mao, L. Qiu, J. Wang, and Y. Zhang. On AS-level path inference. In *ACM SIGMETRICS*, pages 339–349, June 2005.

[24] R. Mazloum, M.-O. Buob, J. AugÃİ, B. Baynat, D. Rossi, and T. Friedman. Violation of interdomain routing assumptions. In *PAM*, pages 173–182, Apr. 2014.

[25] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig. Building an AS-topology model that captures route diversity. In *ACM SIGCOMM*, pages 195–206, Aug. 2006.

[26] W. Mühlbauer, S. Uhlig, B. Fu, M. Meulle, and O. Maennel. In search for an appropriate granularity to model routing policies. In *ACM SIGCOMM*, pages 145–156, Aug. 2007.

[27] L. Neudorfer, Y. Shavitt, and N. Zilberman. Improving AS relationship inference using PoPs. In *IEEE INFOCOM Workshops*, pages 459–464, Apr. 2013.

[28] W. Norton. Dual Transit/Peering. http://drpeering.net/white-papers/Art-Of-Peering-The-Peering-Playbook.html#4.

[29] W. Norton. Partial Transit (Regional). http://drpeering.net/white-papers/Art-Of-Peering-The-Peering-Playbook.html#7.

[30] S. Qiu, P. McDaniel, and F. Monrose. Toward valley-free inter-domain routing. In *IEEE ICC*, pages 2009–2016, June 2007.

[31] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush. 10 lessons from 10 years of measuring and modeling the Internet's autonomous systems. *IEEE Journal on Selected Areas in Communications*, 29(9):1810–1821, Oct. 2011.

[32] U. Weinsberg, Y. Shavitt, and E. Shir. Near-deterministic inference of AS relationships. In *IEEE ConTEL*, pages 191–198, June 2009.

[33] H. Yihua, M. Faloutsos, S. Krishnamurthy, and M. Chrobak. Obtaining provably legitimate Internet topologies. *IEEE/ACM Transactions on Networking*, 20(1):271–284, Feb. 2012.

[34] Young Hyun. CAIDA: Archipelago Measurement Infrastructure. http://caida.org/projects/ark.