

Computing Customer Cones of Peering Networks

Jinu Susan Varghese
Iowa State University
jsusan@iastate.edu

Lu Ruan
Iowa State University
ruanlu@iastate.edu

ABSTRACT

We present a method to compute the customer cones of peering networks using PCH data. Our method computes location dependent customer cones (LDCCs) for networks that are present at more than one IXP instead of computing a single customer cone for each network. We use our method to compute 5753 LDCCs for 3290 IXP participants. Our preliminary analysis of the LDCCs reveals that IXP participants often have different customer cones at different locations.

CCS Concepts

•Networks → Topology analysis and generation; Public Internet;

Keywords

Customer cones; IXPs

1. INTRODUCTION

The autonomous systems (ASes) in the Internet engage in two types of business relationships to exchange traffic: provider-to-customer (p2c) and peer-to-peer (p2p). In a p2c relationship, one AS (the customer) pays the other AS (the provider) for transiting traffic to the rest of the Internet. In a p2p relationship, two ASes exchange traffic between themselves and their customers on a settlement-free basis. Several studies have developed heuristic algorithms to infer AS relationships using BGP data [9, 15, 13, 6, 12, 8, 11]. A recent work by Varghese and Ruan [14] presented a machine learning approach to inferring AS relationships (i.e., edge types) for AS graphs derived from any data source. Dimitropoulos et al. [8] defined the customer cone of an AS as the set of ASes that can be reached from it following only p2c links and used inferred AS relationships to compute AS customer cones. Luckie et al. [11] presented a new method of inferring customer cones (i.e., provider/peer observed cone) and used

inferred customer cones to show the flattening of the Internet topology and the financial consolidation of the Internet transit industry. CAIDA's AS rank project uses inferred AS relationships to rank ASes based on their customer cones and publishes AS ranking at [1]. Customer cones have also been used to select traceroute servers for identifying PoPs for hybrid AS relationships [10] and assess market concentration in the European Internet upstream market [7].

In this paper we present a new method of computing AS customer cones using BGP data archived by PCH [3], which operates route collectors at more than 100 IXPs around the world. Our method computes the customer cones of IXP participants that announce their customer routes to the PCH collectors. Unlike the existing method, our method does not make use of AS relationship inference algorithms which could produce inaccurate inferences. In addition, our method computes *location dependent* customer cones for ASes that are present at more than one IXP instead of computing a single customer cone for each AS.

2. CUSTOMER CONE COMPUTATION

2.1 PCH Data

At an IXP where a PCH collector is present, some IXP participants peer with the collector and send BGP route announcements and BGP updates to it. Such IXP participants are referred to as Vantage Points (VPs). There are two types of VPs: full VP (FVP) and partial VP (PVP). A FVP announces all its routes to the collector while a PVP announces only its customer routes to the collector. With a few exceptions, all PCH VPs are PVPs.

We download the routing table snapshots of all PCH collectors on 3/1/2016. There are 3290 PVPs and 2 FVPs in our dataset. We compute location dependent customer cones (LDCCs) for all PVPs using the method described below.

2.2 Method

We first extract (prefix, AS path) pairs from each snapshot in our dataset. We discard a (prefix, AS path) pair if the AS path contains a loop or an invalid ASN¹. Next we extract the view of each PVP from each snapshot, where the view of a PVP X contains all the (prefix, AS path) pairs in which X is the first ASN in the AS path.

Consider the snapshot of a PCH collector at IXP I . After we use the above procedure to compute the views of all

¹An ASN is invalid if it is unallocated or reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ANRW '16, July 16 2016, Berlin, Germany

© 2016 ACM. ISBN 978-1-4503-4443-2/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2959424.2959428>

PVPs at I , we compute the customer cone of each PVP X at I as follows. Let $\langle X \ A \ B \ C \rangle$ be an AS path in X 's view. We know A is a customer of X because X only announces customer routes to the collector. This implies that B is a customer of A and C is a customer of B because an AS generally only exports customer routes to a provider (i.e., provider routes and peer routes are not exported to a provider). So A , B , and C all belong to the customer cone of X . Thus, we compute X 's customer cone at I to be the union of all ASes that appear in X 's view at I minus X itself.

The above method works if all ASes follow the policy of not exporting provider routes or peer routes to providers. In this case, all AS paths are valley-free. However, we note that some ASes violate the policy. For example, the snapshot of the PCH collector at VIX in Vienna, Austria contains AS path $\langle 6663 \ 25145 \ 6762 \ 12874 \ 196753 \ 34281 \rangle$. Our algorithm will put all ASNs after 6663 in the customer cone of 6663. However, 6762 cannot be in the customer cone of 6663 because it is a Tier 1 network. Thus the ASNs after 6762 cannot belong to the customer cone of 6663. The problem is that the AS path is not valley-free. Specifically, 25145 exported the provider route $\langle 6762 \ 12874 \ 196753 \ 34281 \rangle$ to its provider 6663, which leads to a valley path. We should discard valley paths so that they are not used to derive customer cones. We use a simple approach to detect valley paths: if an AS path contains a Tier 1 AS², then it is a valley path and is discarded.

3. RESULTS

We compute a total of 5753 location dependent customer cones (LDCCs) for 3290 PVPs. For each PVP, we also compute a combined customer cone (CCC) which is the union of all its LDCCs.

Comparison with CAIDA Customer Cones. CAIDA uses inferred AS relationships to compute AS customer cones and publishes the customer cone datasets at [2]. We download CAIDA's customer cone dataset for 3/1/2016 and compare it with the CCCs computed by our method. There are 3265 common ASes in our dataset and CAIDA dataset. Our cone and CAIDA cone have the same size for 2008 (61.5%) ASes. Our cone is larger than CAIDA cone for 201 (6.2%) ASes and smaller than CAIDA cone for 1056 (32.3%) ASes.

We rank the 3265 common ASes based on cone size (i.e., number of ASes in cone). Table 1 shows the top 10 ASes based on our cones. The number in parentheses next to an ASN indicates the rank of the AS based on CAIDA cones. Our ranking is mostly consistent with CAIDA ranking. The only difference is that CAIDA ranks AS1273 before AS3491. We look up the 10 networks in PeeringDB [4] and find that all of them are NSPs.

Location Dependency of Customer Cones. There are 1158 ASes that peer with PCH at more than one IXP and thus have more than 1 LDCC. Out of the 1158 ASes, 512 (44.2%) ASes have nonidentical LDCC sizes and 681 (58.8%) ASes have nonidentical number of prefixes in their LDCCs. Thus, it is common for peering networks to have different customer cones and different set of prefixes at different locations.

We look up the 1158 ASes in PeeringDB and find the

²We manually identify 19 Tier-1 ASes according to Wikipedia [5] and CAIDA's inferred clique [2].

Table 1: Top 10 Networks Based on Our Cones

ASN	Our Cone Size	CAIDA Cone Size
AS3257 (1)	19,256	18,886
AS6762 (2)	12,379	14,319
AS6939 (3)	10,771	9,501
AS3491 (5)	7,825	4,561
AS1273 (4)	5,322	5,805
AS6461 (6)	4,330	4,415
AS9002 (7)	3,471	3,656
AS20485 (8)	2,883	3,153
AS12389 (9)	2,589	2,815
AS4323 (10)	2,265	2,288

business types for 956 of them; 92.5% of these ASes are of type NSP, Cable/DSL/ISP, and Content. For each of the three business types, Table 2 shows the number of ASes and the percentage of ASes that have nonidentical cone sizes and nonidentical number of prefixes at different locations. We observe that NSPs tend to have different cones and different prefixes at different locations. Access providers and content providers tend to have identical cones but different prefixes at different locations. ASes of all three business types are more likely to have location dependent prefixes than location dependent cones.

Table 2: Location Dependency of Customer Cones for Three Business Types

Business Type	Count	Nonident. Cone	Nonident. Prefix
NSP	386	66.8%	76.2%
Cable/DSL/ISP	305	39.7%	54.7%
Content	193	29.5%	58.5%

We examine the LDCCs of three big content providers: Apple (AS714), Akamai (AS20940), and Google (AS15169). We find that Apple announces 119-140 prefixes at 6 IXPs in North America, 19-23 prefixes at 5 IXPs in Europe, and 43 prefixes at one IXP in Asia Pacific. This is consistent with a note in Apple's PeeringDB entry stating "We are announcing +100 routes in the US now, fewer in Japan, Singapore, Hong Kong, Sydney and Europe." Akamai also announces different number of prefixes at different locations: the number varies from 1 to 147 across 29 IXPs. On the other hand, Google announces the same number of prefixes at 24 out of the 28 IXPs where it peers with PCH.

4. CONCLUSION AND FUTURE WORK

We present a method to compute location dependent customer cones (LDCCs) of peering networks using PCH data. Our preliminary results show that peering networks often have different customer cones and different prefixes at different locations. In our future work, we will perform an in-depth analysis of the LDCCs of different types of peering networks to gain insight on the location dependent nature of AS routing policies. We will create a web site to publish LDCCs of peering networks on a monthly basis. We believe LDCC data is of value to network operators as well as researchers. It can aid network operators in selecting peers at different IXPs. Researchers can use LDCC data to study the evolution of the Internet transit ecosystem in different regions of the world, discover correlations between AS peering behaviors and their customer cone properties, and investigate how peering and depeering events cause changes in the customer cones of ASes and vice versa.

5. REFERENCES

- [1] *CAIDA AS Rank Project*. <http://as-rank.caida.org>.
- [2] *CAIDA's AS Relationship Inference Dataset*. <http://data.caida.org/datasets/as-relationships/serial-1/>.
- [3] *Packet Clearing House*. <https://www.pch.net/>.
- [4] *PeeringDB*. <https://www.peeringdb.com>.
- [5] *Tier-1 network*. http://en.wikipedia.org/wiki/Tier_1_network.
- [6] G. Di Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between autonomous systems. In *IEEE INFOCOM*, volume 1, pages 156–165, 2003.
- [7] A. D'Ignazio and E. Giovannetti. Antitrust analysis for the Internet upstream market: A border gateway protocol approach. *Journal of Competition Law and Economics*, 2(1):43–69, 2006.
- [8] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy, and G. Riley. AS relationships: inference and validation. *ACM SIGCOMM Computer Communication Review*, 37(1):29–40, 2007.
- [9] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, 2001.
- [10] V. Giotsas, M. Luckie, B. Huffaker, et al. Inferring complex AS relationships. In *The Internet Measurement Conference*, pages 23–30, 2014.
- [11] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and k. claffy. AS relationships, customer cones, and validation. In *ACM IMC*, pages 243–256, 2013.
- [12] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (in)completeness of the observed Internet AS-level structure. *IEEE/ACM Transactions on Networking*, 18(1):109–122, 2010.
- [13] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *IEEE INFOCOM*, volume 2, pages 618–627, 2002.
- [14] J. Varghese and L. Ruan. A machine learning approach to edge type inference in Internet AS graphs. In *8th IEEE International Workshop on Network Science for Communication Networks (INFOCOM Workshop)*, pages 725–730, April 2016.
- [15] J. Xia and L. Gao. On the evaluation of AS relationship inferences. In *IEEE Global Telecommunications Conference*, volume 3, pages 1373–1377, 2004.