Obscure Giants: Detecting the Provider-Free ASes

Syed Hasan and Sergey Gorinsky

Institute IMDEA Networks
Avenida del Mar Mediterraneo, 22
Leganes, Madrid, 28918, Spain
syed.anwar@imdea.org, sergey.gorinsky@imdea.org

ABSTRACT

We study the detection of the provider-free AS set (PFS), i.e., the set of those Autonomous Systems (ASes) that reach the entire Internet without paying anyone for the traffic delivery. Using trustworthy but non-verifiable sources for sanity checks, we derive the PFS from public datasets of inter-AS economic relationships. Whereas a straightforward method for extracting the PFS performs poorly because the datasets are noisy, we develop a more sophisticated Temporal Cone (TC) algorithm that relies on topological statistics and exploits the temporal diversity of the datasets. The evaluation shows that our TC algorithm detects the PFS with a high accuracy.

1. INTRODUCTION

Economic relationships between ASes (Autonomous Systems) matter for Internet routing. For example, it is financially more attractive for an AS to route traffic through a peering link than a transit connection of the AS to its provider. Despite a trend towards flattening [1], the Internet routing ecosystem is essentially hierarchical [1–4]. A vast majority of ASes are relatively small and route traffic either as customers of transit links or by peering with local ASes of a similar stature. There exists only a handful of *provider-free ASes* that reach the entire Internet without paying anyone for the traffic delivery. These global provider-free ASes peer with each other and sell IP (Internet Protocol) [5] transit to their numerous customers.

Definite answers to Internet economic questions remain elusive. First, ASes typically do not disclose their business agreements. Without direct access to the agreements, researchers infer the inter-AS economic relationships from BGP (Border Gateway Protocol) [6] route advertisements or actual IP forwarding routes. Such inferences are imperfect, e.g., a router misconfiguration can trigger an inference of an invalid relationship. Second, the inference algorithms are heuristic and can cause additional deviations from the reality. Third, the economic relationships are dynamic: while it takes time to collect a comprehensive set of measurements, changes in the relationships can decrease the inference accuracy.

In this paper, we explore the detection of the *PFS, the provider-free AS set*. As the transit core of the routing ecosystem, the provider-free ASes clearly play an important role in the Internet. Besides, humans prefer to think in discrete categories, and designating (or not designating) an autonomous system as provider-free can have tangible marketplace implications. Furthermore, some algorithms use the PFS as a basis for inferring the economic relationships for all ASes [4, 7]. Despite their importance and large individual sizes, the provider-free ASes are difficult to discern because the difference between them and some non-provider-free ASes is small. For example, if an AS pays for less than 1% of its interdomain traffic, and these disqualifying payments are for paid peering that is subject to a non-disclosure agreement, the non-provider-free status of this AS can be very obscure to outsiders.

To the best of our knowledge, this is the first work focusing on the detection of the provider-free ASes. The two main contributions of this paper are in deriving:

- *PFS insights from non-verifiable sources*, such as Wikipedia [8], that do not disclose their data and methods for determining the PFS; whereas these sources are mostly trustworthy, we carefully filter out occasional spurious answers;
- *TC (Temporal Cone) algorithm that detects the PFS based on public datasets of inter-AS economic relationships*; while the noisy datasets lead to poor results with straightforward inference, the TC algorithm utilizes topological statistics and temporal diversity to compute the PFS accurately.

We structure the rest of the paper as follows. Section 2 reports the PFS insights from the non-verifiable sources. Section 3 describes the public datasets in our study. Section 4 considers a straightforward PFS detection method. After analyzing the failures of this straightforward method, Section 5 develops the more sophisticated TC algorithm. Section 6 evaluates the TC algorithm. Section 7 comments on related work. Section 8 concludes the paper by summing up its contributions.

2. NON-VERIFIABLE SOURCES

While the obscure inter-AS economic relationships do not reveal the ground truth about the PFS, a number of non-verifiable sources offer insights into this set. We consider three such non-verifiable sources: Wikipedia, Renesys, and Hurricane Electric.

Wikipedia maintains an article about provider-free ASes and refers to them as tier-1 networks [8]. While the development of our TC algorithm relies on public datasets collected in 2009, our primary interest is in the Wikipedia perspectives throughout that year. According to Wikipedia, the PFS consisted of 8 members on 1/1/2009: AT&T, Global Crossing, Level 3, NTT, Qwest, Sprint, Verizon, and Savvis [9]. The article has seen frequent revisions and expanded its PFS with Telia on 28/1/2009 [10]. The addition of Tata on 25/3/2009 resulted in the following PFS [11]:

$W_1 = \{$AT&T, Global Crossing, Level 3, NTT, Qwest, Sprint, Verizon, Savvis, Telia, Tata$\}$.

Except for few incidents in June and October when spurious modifications disappeared shortly after being made, the PFS preserved this 10-member composition until the end of 2009. In 2010 and 2011, Wikipedia continued the trend of the PFS expansion and typically recognized Tinet as the 11th member of the PFS, e.g., in the 10/2/2011 revision [12]:

$W_2 = \{$AT&T, Global Crossing, Level 3, NTT, Qwest, Sprint, Verizon, Savvis, Telia, Tata, Tinet$\}$.

Whereas Wikipedia is an online encyclopedia that anyone may edit, some short-lived revisions of this particular article certainly distorted the reality [13]. Nevertheless, experts think that on the whole the Wikipedia perspective reflects the PFS accurately [4].

Renesys is a private company that sells Internet business information. In 1/2009, Renesys announced a 12-member set of commercial default-free ASes [14], i.e., ASes that can route traffic to any Internet destination without relying on a default route. Default-free ASes are either provider-free or reaching the entire Internet by paying for peering but not for transit. The Renesys set subsumes W_1 and includes two more ASes: XO and AboveNet. Interestingly, the Wikipedia article explicitly stated in all its revisions that XO and AboveNet were not provider-free due to paid peering [9–13]. Thus, the Renesys perspective is consistent with limiting the PFS to W_1.

Hurricane Electric is an Internet service provider that offers an online tool for ranking the peers of an autonomous system [15]. We apply this tool to all ASes in W_1. For each such AS, all the other ASes in W_1 are among highly ranked peers of this AS. Thus, the ASes of W_1 do form a close-knit peering community as expected for provider-free ASes.

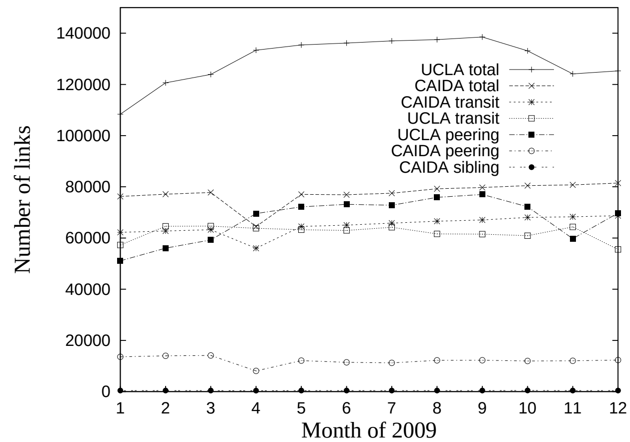

Figure 1: Inter-AS economic relationships in the UCLA and CAIDA datasets during 2009.

Based on the above considerations, our paper subsequently treats W_1 as the primary PFS answer from the non-verifiable sources for 2009.

3. PUBLIC DATASETS

The PFS insights in Section 2 came from the non-verifiable sources that did not disclose their data and methods. The rest of our study explores datasets from two public sources: *UCLA* (University of California, Los Angeles) [16] and *CAIDA* (Cooperative Association for Internet Data Analysis) [17]. Both UCLA and CAIDA convert BGP routing measurements into datasets that characterize the economic relationships between Internet ASes. UCLA classifies inter-AS links as transit or peering. CAIDA uses an additional category for sibling relationships: a sibling link connects two ASes belonging to the same Internet service provider.

While the UCLA datasets are available starting from 10/2008, CAIDA reports its datasets infrequently for 2009 and only thrice after 2009. During the development of our PFS detection algorithm, we would like to have similar time series for the two sources. Hence, our Sections 4 and 5 focus on the 12 months of 2009. Guided by the CAIDA dataset availability and trying to be as close as possible to the 20th day of the month, we select the following days for both sources: 22/1, 20/2, 11/3, 29/4, 20/5, 15/6, 20/7, 30/8, 20/9, 20/10, 20/11, and 15/12. June is the only exception: because the number of links in the UCLA 15/6 dataset is extremely low, we use 16/6 instead for UCLA.

Figure 1 depicts the inter-AS economic relationships in the UCLA and CAIDA datasets during 2009. For either source, the total number of links tends to grow with time, and the few down-and-up swings are most likely due to imperfect measurements rather than actual fluctuations in the number of economic relationships. The

Month of 2009	1	2	3	4	5	6	7	8	9	10	11	12
UCLA	8 (1)	6 (0)	7 (0)	17 (9)	16 (9)	17 (9)	15 (9)	19 (9)	19 (9)	16 (10)	17 (9)	18 (9)
CAIDA	23 (6)	26 (6)	26 (6)	29 (7)	30 (7)	27 (7)	28 (7)	29 (7)	25 (6)	26 (6)	27 (6)	27 (6)

Table 1: Size of the PFS according to the straightforward method for the UCLA and CAIDA datasets and (in parentheses) number of ASes from W_1 in this PFS.

sibling relationships in the CAIDA datasets constitute a negligible fraction of the overall link population. While the number of peering links is much higher for UCLA than for CAIDA, the number of transit links is rather similar for the two sources. The transit-link profiles are mostly consistent but do have some aberrations such as the dip for CAIDA in 4/2009. The numbers of transit links for the two sources remain most stable and close to each other between 5/2009 and 7/2009.

When evaluating our TC algorithm in Section 6, we utilize the UCLA datasets for all 32 months of their availability from 10/2008 to 5/2011. We select the 20th day for all 20 additional months except the last one, for which we use the latest available day of 10/5/2011.

4. STRAIGHTFORWARD INFERENCE

Given a dataset of inter-AS economic relationships, one might hope to infer the PFS using the following *straightforward method: compose the PFS from all such ASes in the dataset that have no transit provider*. We apply this straightforward method to the UCLA and CAIDA datasets of Section 3. Table 1 sums up the generally disappointing results for all 12 months of 2009. Throughout the year, the straightforward method includes into its PFS at least 6 (and up to 23) non-W_1 ASes and excludes at least 1 AS (and up to all 10 ASes) of W_1. For the UCLA and CAIDA datasets from 6/2009 (when the numbers of transit links for the two sources remain most stable and close to each other), the PFS contains respectively 17 and 27 ASes, with respectively 9 and 7 of these ASes belonging to W_1.

For the UCLA 6/2009 dataset, the straightforward method excludes Tata from the PFS because NTT and GIT Telecom (a Cypriot AS) are transit providers for this missing member of W_1 according to the dataset. Among the 8 non-W_1 members of the PFS, Sunkist Growers (a not-for-profit cooperative of citrus growers in California and Arizona), Open Peering Initiative (a public peering IXP in Amsterdam), and Siemens seem highly unlikely to be genuine provider-free ASes. These 3 ASes do have providers in the CAIDA dataset from the same month.

For the CAIDA 6/2009 dataset, the straightforward method omits NTT, Savvis, and Tata from the PFS because these 3 members of W_1 have transit providers. Specifically, NTT has 3 providers: Verizon, Telia, and Easynet. Savvis has 5 providers: Telia, Tata, Tinet, XO, and Deutsche Telekom. Although Tata is a transit provider for Savvis, the straightforward method does not recognize Tata as a provider-free AS either: Tata appears as a customer of NTT, Telia, and Tinet. On the other hand, the PFS of the straightforward method includes 20 non-W_1 ASes such as the University of Texas System, NASA, and New Zealand Research Network, which do have providers in the UCLA 6/2009 dataset.

5. ALGORITHM DEVELOPMENT

Section 4 demonstrates that the straightforward inference yields disappointing PFS results with respect to both false positives and false negatives. Two factors undermine the straightforward method. First, while the UCLA and CAIDA datasets do not classify the inter-AS links fully and correctly, even a single error in the input dataset can mislead the straightforward method. For instance, the method can exclude a genuine provider-free AS from the PFS because the dataset mistakenly reports a provider for this AS; also, the method can wrongly include an AS into the PFS because the dataset misses the transit link between this AS and its only provider. Second, the straightforward method implicitly assumes that having no provider implies the ability to reach the entire Internet. In reality, some ASes in the Internet ecosystem do not strive for the universal reachability. For example, the main goal of an IXP (Internet eXchange Point) [1, 18] is to serve as a peering infrastructure that enables other ASes to exchange their local traffic. The straightforward method can incorrectly classify an IXP as a provider-free AS.

This section develops a more sophisticated algorithm for detecting the PFS. We tackle the challenge of the noisy data by means of topological statistics and temporal diversity. Each of these two elements contributes to making the algorithm more robust than the straightforward method. We examine three topological parameters:

- *customer count* of an AS is the number of direct customers of the AS;
- *customer cone* of an AS is the total number of direct and indirect customers of the AS, i.e., all customers reachable from the AS through a sequence of provider-to-customer transit links [19];
- *peering count* of an AS is the number of peering links of the AS.

Because the customer count and cone of a provider-free AS are both very large, misclassifying a single rela-

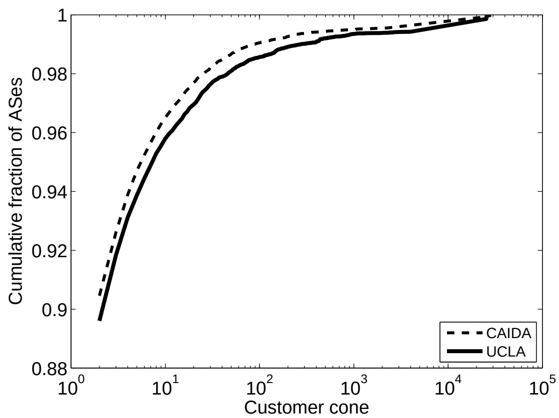

Figure 2: 6/2009 distributions of the AS customer cones.

Rank	AS name (AS number)	Customer cone	In W_1?
1	Sprint (1239)	28478	✓(1)
2	Level3 (3356)	28168	✓(2)
3	NTT (2914)	27650	✓(3)
4	AT&T (7018)	27613	✓(4)
5	Global Crossing (3549)	27236	✓(5)
6	Verizon (701)	27121	✓(6)
7	Telia (1299)	26833	✓(7)
8	Qwest (209)	26764	✓(8)
9	Deutsche Telekom (3320)	26263	–
10	Ipercast (34763)	26127	–
11	Savvis (3561)	26082	✓(9)
12	GIT Telecom (38925)	26015	–
13	Tata (6453)	26014	✓(10)

Table 2: UCLA customer-cone ranks of ASes for 6/2009.

tionship of the AS has a relatively small impact on these two statistical metrics. Hence, either customer count or customer cone represents a promising basis for robust PFS detection. The customer count is easier to compute but the correlation with the provider-free status is intuitively stronger for the customer cone. For example, a non-provider-free AS that is a direct customer of a PFS member can itself have a large number of direct customers. Due to multihoming [20] which is common throughout the Internet ecosystem, the customer cones of two ASes can overlap. We compute the customer cone of each AS using a recursive algorithm that takes the overlaps of the customer cones into account.

While the PFS members peer with each other, searching for close-knit peering communities is another potential approach to detecting the PFS. However, our preliminary analyses for the peering count and other peering-based parameters did not yield encouraging results. Consequently, the rest of our paper considers the customer-based parameters only.

Figure 2 plots the distributions of the AS customer cones in the UCLA and CAIDA datasets for 6/2009 and shows that only a tiny fraction of all ASes have a really large customer cone. Table 2 zooms in on the tail of the UCLA 6/2009 distribution. The tail covers set W_1 quite tightly: all 10 members of W_1 appear among the top 13 ASes ranked by the customer cone; this is an improvement over the straightforward method which includes only 9 members of W_1 into its 17-member PFS for 6/2009.

Figure 3 tracks the UCLA customer-cone ranks of all ASes in W_1 throughout 2009. The ranks remain close to the top 10 with few exceptions such as three dramatic dips for Tata. While Figure 3 corroborates the promising potential of the customer-cone statistics for PFS detection, the results also suggest that our algorithm needs additional features for overcoming the noise in the datasets.

Figure 4 depicts the CAIDA customer-cone ranks of all ASes in W_1 during 2009. In agreement with Table 1, the customer-cone results in Figures 3 and 4 imply that the UCLA datasets are less noisy and thus more suitable for PFS detection than the CAIDA datasets.

Figures 5 presents the 2009 customer-count ranks of all ASes in W_1 for the UCLA datasets. In comparison to the customer-cone ranks in Figure 3, the customer-count ranks are less effective in capturing W_1. Therefore, we choose the customer cone as the basis for our PFS detection algorithm.

To detect the PFS, our algorithm has to size this set. While Section 2 indicates that the PFS grows over time, no static size seems suitable. Instead, our algorithm sizes the PFS to a fraction of the overall AS population:

$$S_m = \lfloor k \cdot P_m \rfloor \qquad (1)$$

where S_m denotes the PFS size at time m, P_m represents the total number of Internet ASes at time m, and k is a fixed factor. Guided by Section 2 that sizes the PFS to 10 ASes in 2009, we recommend the k value of 0.00032, which means that one in about 3000 Internet ASes is provider-free.

With the PFS size determined, the algorithm still needs to identify the ASes of the set. We utilize the temporal dimension of the datasets to tackle the noise remaining in the customer-cone statistics. Our intuition is that the membership of an AS in the PFS is relatively stable. While a new AS can join the PFS and subsequently lose the provider-free status again, such transitions are infrequent, caused by rare mergers/acquisitions and guarded against by long-term business contracts. Therefore, to decide whether an AS is provider-free for month m, our algorithm looks w months back and ahead from month m and includes the AS into the PFS for month m only if the AS belongs to the set according to the customer-cone ranks for at least n out of these $2w + 1$ months. For an input with M months in the time series, our algorithm outputs the PFS for each month except for the first w and

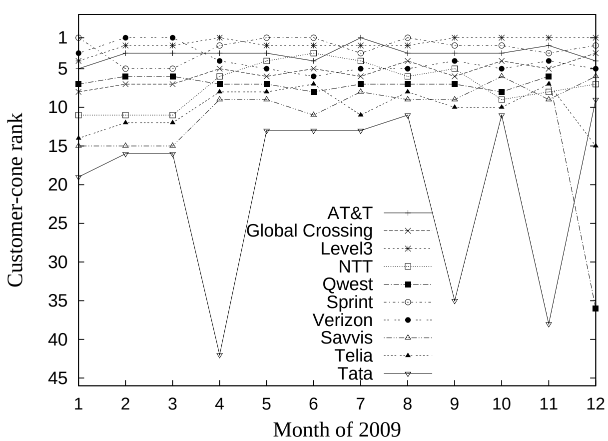

Figure 3: UCLA customer-cone ranks of the ASes in W_1.

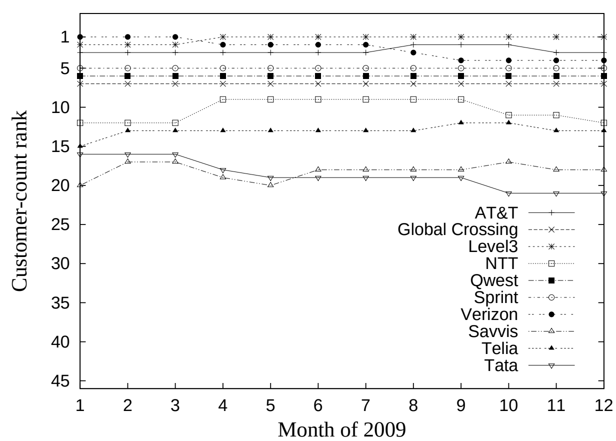

Figure 5: UCLA customer-count ranks of the ASes in W_1.

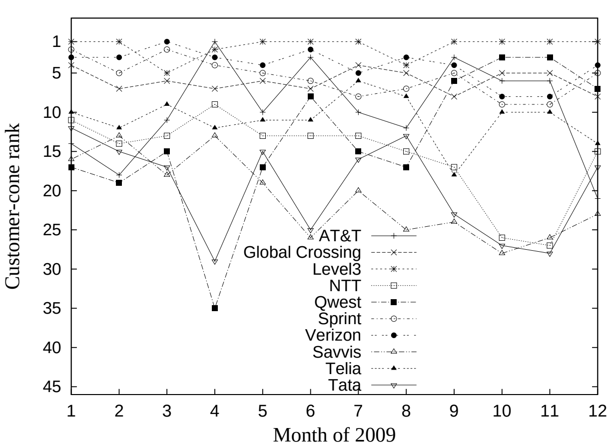

Figure 4: CAIDA customer-cone ranks of the ASes in W_1.

last w months, i.e., the algorithm computes the PFS for the $M - 2w$ middle months.

While one-year contracts between ASes are common, we recommend $w = 5$ and $n = 7$ as default settings for the w and n parameters of the algorithm, i.e., inclusion of an AS into the PFS requires from the customer-cone ranks to endorse the AS for at least 7 out of 11 months. These settings enable our algorithm to recognize a genuine one-year PFS membership in spite of multiple months of erroneous disqualifications by the customer-cone ranks; these settings also allow the algorithm to exclude a non-provider-free AS from the PFS despite multiple months of mistaken customer-cone endorsements.

We refer to the developed PFS detection algorithm as TC (Temporal Cone). Table 3 explains the notation used in Figure 6 that describes our TC algorithm in detail.

6. EVALUATION

According to Sections 3 through 5, the datasets from CAIDA are noisier and less available than the UCLA datasets. To evaluate the developed TC algorithm, this section relies on the UCLA datasets for the 32 months from 10/2008 to 5/2011. Following the recommendations in Section 5, we set the PFS sizing factor, number of the lookback/lookahead months, and PFS membership threshold to $k = 0.00032$, $w = 5$, and $n = 7$ respectively.

During its first iterative stage, the TC algorithm determines the AS customer-cone ranks and PFS sizes for all 32 months. Figure 7 plots the customer-cone ranks of the ASes in set W_2, i.e., all W_1 members and Tinet which became the 11th member of the PFS according to Wikipedia after 2009. All 11 members of W_2 consistently appear among the top 11 ASes ranked by the customer cone in 2010 and 2011, indicating a higher accuracy of the more recent UCLA datasets.

According to the TC calculations of the PFS size, the PFS contains 9 ASes from 10/2008 to 1/2009, 10 ASes from 2/2009 to 12/2009, and 11 ASes from 1/2010 to 5/2011. This expansion is also consistent with the PFS insights from the trustworthy but non-verifiable sources.

With 5 months to look back and ahead, the TC algorithm executes its second stage to compute the PFS for the 22 middle months from 3/2009 to 12/2010. Among the 10 months of 2009 (when the PFS size is 10), the PFS perfectly matches W_1 for 2 months and omits only Tata for other 6 months. The worst performance is in 4/2009 when the computed PFS misses 4 provider-free ASes. For all 12 months of 2010 (when the PFS size is equal to 11), the PFS matches W_2 exactly.

Table 4 sums up the performance of the TC algorithm. A quick comparison of these results with Table 1 reveals that the TC algorithm detects the PFS significantly better than the straightforward method.

Notation	Semantics
m or i	month
M	number of months in the time series
C_m	list of the Internet ASes ordered by their customer-cone ranks for month m
L_m	ordered list of PFS candidates for month m
S_m	size of the PFS for month m
w	number of the lookback/lookahead months
F_m	PFS for month m
a	AS
b_a	counter of months when AS a belongs to the PFS as per the customer-cone rankings
$r_{a,m}$	rank of a in L_m
n	PFS membership threshold

Table 3: Notation for our algorithm in Figure 6.

```
for m = 1, ..., M
    compute C_m;
    L_m ← C_m;
    calculate S_m according to Equation 1;
for m = M − w, ..., w + 1
    F_m ← ∅;
    a ← first AS in L_m;
    while |F_m| < S_m and a ≠ null
        b_a ← 0;
        for i = m − w, ..., m + w
            if r_{a,i} ≤ S_i
                then b_a ← b_a + 1;
        if b_a ≥ n
            then F_m ← F_m ∪ {a}
            else remove a from L_m; r_{a,m} ← ∞;
        a ← next AS in L_m
```

Figure 6: TC (Temporal Cone) algorithm for the PFS detection.

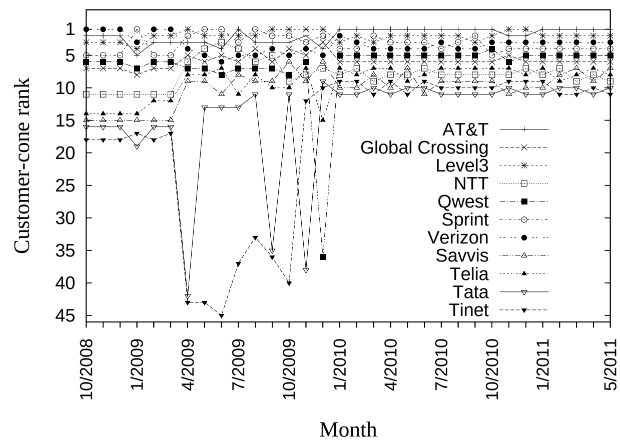

Figure 7: UCLA customer-cone ranks of the ASes in W_2 (i.e., the W_1 members and Tinet) from 10/2008 to 5/2011.

Year	2009				2010
Month	3, 5-7, 9, 11	4	8, 10	12	1-12
UCLA	10 (9)	10 (6)	10 (10)	10 (7)	11 (11)

Table 4: Size of the PFS according to the TC algorithm for the UCLA datasets and (in parentheses) number of ASes in this PFS that match the Wikipedia insights (W_1 for 2009 and W_2 for 2010).

7. RELATED WORK

The TC algorithm derives the PFS from inter-AS economic relationships. Since the pioneering work by Lixin Gao [2], the problem of inter-AS relationship inference has seen a variety of new heuristic solutions [4,7,19,21–26].

While our paper is the first to focus on detecting the PFS, other researchers used the PFS as an input to their inter-AS relationship inference algorithms [4, 7]. The PFS also served as a basis for studies of backbone networks and resilience of routing to failures [27, 28].

The PFS derivation from public inter-AS relationship datasets is challenging because missing or misclassified links make the datasets noisy. Addressing the problem of hidden links [29–33] has a potential for making the results of our TC algorithm even better.

While the TC algorithm exploits the temporal diversity of the inter-AS relationship datasets, prior works explored the temporal dimension for studying other problems such as network graph evolution [34, 35].

8. CONCLUSION

In this paper, we studied the detection of the PFS. Using Wikipedia and other trustworthy non-verifiable sources for sanity checks, we developed the TC algorithm that derives the PFS from public datasets of inter-AS economic relationships. The TC algorithm sizes the PFS to a fraction of the overall AS population, relies on the AS customer-cone ranks, and leverages the temporal diversity of the datasets. In comparison to the straightforward method for extracting the PFS, our TC algorithm detects the PFS significantly better.

9. REFERENCES

[1] A. Dhamdhere, C. Dovrolis, "The Internet is Flat: Modeling the Transition from a Hierarchy to a Peering Mesh", Proceedings of ACM CoNEXT 2010.

[2] L. Gao, "On Inferring Autonomous System Relationships in the Internet", IEEE/ACM Transactions on Networking, 2001.

[3] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, D. Rubenstein, "On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers", IEEE/ACM Transactions on Networking, accepted for publication, 2011.

[4] E. Gregori, A. Improta, L. Lenzini, L. Rossi, L. Sani, "BGP and Inter-AS Economic Relationships", Proceedings of IFIP Networking 2011.

[5] Information Sciences Institute, University of Southern California, "Internet Protocol DARPA Internet Program Protocol Specification", RFC 791, 1981.

[6] Y. Rekhter, T. Li, "A Border Gateway Protocol (BGP-4)", RFC 1771, 1995.

[7] J. Xia and L. Gao, "On the Evaluation of AS Relationship Inferences", Proceedings of IEEE Globecom 2004.

[8] Wikipedia, "Tier 1 Network", `en.wikipedia.org/wiki/Tier_1_network`.

[9] Wikipedia, "Tier 1 Network", 1/1/2009 revision, `en.wikipedia.org/w/index.php?&oldid=261328396`.

[10] Wikipedia, “Tier 1 Network”, 28/1/2009 revision, en.wikipedia.org/w/index.php?&oldid=267026426.
[11] Wikipedia, “Tier 1 Network”, 25/3/2009 revision, en.wikipedia.org/w/index.php?&oldid=279646779.
[12] Wikipedia, “Tier 1 Network”, 10/2/2011 revision, en.wikipedia.org/w/index.php?&oldid=413097463.
[13] Wikipedia, “Tier 1 Network”, 5/6/2009 revision, en.wikipedia.org/w/index.php?&oldid=294566542.
[14] M. A. Brown, C. Hepner, A. C. Popescu, “Internet Captivity and the De-peering Menace”, NANOG 45, www.renesys.com/tech/presentations/pdf/nanog-45-Internet-Peering.pdf.
[15] Hurricane Electric, “Hurricane Electric BGP Toolkit”, bgp.he.net.
[16] University of California, Los Angeles, “Internet Topology Collection”, irl.cs.ucla.edu/topology.
[17] Cooperative Association for Internet Data Analysis, “AS Relationships”, www.caida.org/data/active/as-relationships.
[18] B. Augustin, B. Krishnamurthy, W. Willinger, “IXPs: Mapped?”, Proceedings of ACM SIGCOMM 2009.
[19] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, K. C. Claffy, and G. Riley, “AS Relationships: Inference and Validation”, ACM SIGCOMM CCR, 2007.
[20] A. Akella, B. Maggs, S. Seshan, A. Shaikh, R. Sitaraman, “A Measurement-based Analysis of Multihoming”, Proceedings of ACM SIGCOMM 2003.
[21] Z. Ge, D. R. Figueiredo, S. Jaiwal, L. Gao. “On the Hierarchical Structure of the Logical Internet Graph”, Proceedings of ITCOM 2001.
[22] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, “Characterizing the Internet Hierarchy from Multiple Vantage Points”, Proceedings of IEEE INFOCOM 2002.
[23] Z. M. Mao, L. Qiu, J. Wang, and Y. Zhang. “On AS-level Path Inference”, Proceedings of ACM SIGMETRICS 2005.
[24] G. Battista, M. Patrignani, and M. Pizzonia, “Computing the Types of the Relationships Between Autonomous Systems”, Proceedings of IEEE INFOCOM 2003.
[25] H. Asai and H. Esaki. “Estimating AS relationships for Application-layer Traffic optimization”, Proceedings of ETM 2010.
[26] H. Asai, H. Esaki, and T. Momose, “A Solution Approach for AS Relationships-aware Overlay Routing”, IETF Internet draft(Informational), Expires: June 18, 2011. http://tools.ietf.org/html/draft-asai-cross-domain-overlay-01.
[27] J. Wu, Y. Zhang, Z. M. Mao, and K. G. Shin. “Internet routing Resilience to Failures: Analysis and Implications”, Proceedings of ACM CoNEXT 2007.
[28] R. Mahajan, M. Zhang, L. Poole, and V. Pai, “Uncovering Performance Differences among Backbone ISPs with Netdiff”, Proceedings of NSDI 2008.
[29] R. Oliveira, D. Pei, W. Willinger, B. Zhang, L. Zhang, “The (in)Completeness of the Observed Internet AS-level Structure”, IEEE/ACM Transactions on Networking, 2010.
[30] M. Roughan, S. J. Tuke, and O. Maennel. “Bigfoot, Sasquatch, the Yeti and other Missing Links: What We don’t know about the AS graph”, Proceedings of ACM SIGCOMM IMC 2008.
[31] K. Chen, D. R. Choffnes, R. Potharaju, Y. Chen, F. E. Bustamante, D. Pei, and Y. Zhao, “Where the Sidewalk Ends: Extending the Internet AS graph using Traceroutes from P2P users”, Proceedings of ACM CoNEXT 2009.
[32] B. Zhang, R. Liu, D. Massey, and L. Zhang, “Collecting the Internet AS-level topology”, ACM SIGCOMM CCR, 2005.
[33] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy, “A Systematic Framework for Unearthing the Missing Links: Measurements and Impact”, Proceedings of NSDI 2007.
[34] A. Dhamdhere and C. Dovrolis, “Ten Years in the Evolution of the Internet Ecosystem”, Proceedings of ACM SIGCOMM IMC 2008.
[35] J. Leskovec, J. Kleinberg, and C. Faloutsos, “Graphs over Time: Densification laws, Shrinking diameters and Possible explanations”, Proceedings of ACM SIGKDD 2005.