Stable and Practical AS Relationship Inference with ProbLink

Paper #311

Abstract

Knowledge of the business relationships between Autonomous Systems (ASes) is essential to understanding the behavior of the Internet routing system. Despite significant progress in the development of sophisticated relationship inference algorithms, the resulting datasets are impractical for many critical real-world applications, cannot offer adequate predictability in the configuration of routing policies, and suffer from inference oscillations. To achieve more practical and stable relationship inferences we first illuminate the root causes of the contradictions between these shortcomings and the near-perfect validation results of `AS-Rank`, the state-of-the-art relationship inference algorithm. Using a "naive" inference approach as a benchmark, we find that the available validation datasets over-represent AS links with easier inference requirements. We identify which types of links are harder to infer, and we develop appropriate validation subsets to enable more representative evaluation.

We then develop a probabilistic algorithm, `ProbLink`, to overcome the inference barriers for hard links, such as non-valley-free routing, limited visibility, and non-conventional peering practices. To this end, we identify key interconnection features that provide stochastically informative and highly predictive relationship inference signals. Compared to `AS-Rank`, our approach reduces the error rate for all links by $1.6\times$, and importantly, by up to $6.1\times$ for different types of *hard* links. We demonstrate the practical significance of our improvements by evaluating their impact on three applications. Compared to the current state-of-the-art, `ProbLink` increases the precision and recall of route leak detection by $4.1\times$ and $3.4\times$ respectively, reveals 27% more complex relationships, and increases the precision of predicting the impact of selective advertisements by 34%.

1 Introduction

The Internet, often referred to as a "network of networks", is composed of more than 60,000 of Autonomous Systems (ASes), each independently operated and managed, which cooperate via the Border Gateway Protocol (BGP) to exchange routing information with each other and obtain global reachability. Inter-AS connectivity involves contractual relationships that determine the economics and various technical aspects of traffic exchange. For over 15 years, researchers have studied the problem of inferring the different types of relationships between ASes from publicly available BGP routing data. Understanding of these inter-AS relationships is useful for identifying misconfigurations on BGP routers, troubleshooting pathological routing behaviors such as dilated routes or blackholes, predicting how Internet policies evolve over time, and inferring interdomain congestion, among several other contexts [28, 26, 9]. However, it has been found that the available relationship inferences perform poorly in many critical applications [34, 33, 5, 35, 32, 44].

In this paper, we revisit the AS relationship inference problem and study why the state-of-the-art inference techniques are not sufficient, despite extensive validation that indicates an error rate of just 1%. In particular, we consider the sophisticated `AS-Rank` technique [30] that synthesizes eleven deterministic heuristics. As a first step in assessing the performance of this technique, we create a benchmark comparison in the form of a simple three-step AS relationship inference algorithm, called `CoreToLeaf`, which utilizes just the valley-free assumption on paths through a core set of transit-free ASes. In spite of its simplicity, `CoreToLeaf` achieves accuracy that is almost as high as that of `AS-Rank`. At the same time, we conduct qualitative evaluations of both `CoreToLeaf` an `AS-Rank` with practical applications — detection of route leaks and the analysis of selective advertisements — which reveal that both algorithms achieve surprisingly poor accuracy for such use cases.

These insights allow us to identify that most of the links in the validation datasets have associated relationships that are relatively easy to infer, while it is necessary to achieve higher accuracy for harder to infer links to enable practical applications. We construct different subsets of the validation dataset that contain AS links that are considered *hard*, and find that both the `CoreToLeaf` and `AS-Rank` techniques have substantially lower accuracy on these validation subsets, indicating that there are distinct classes of links for which the current algorithms are not sufficiently accurate.

We next examine some of the challenges in developing a more accurate AS relationship inference algorithm. We observe first that the attributes of a link (and that of the paths that traverse the link) that can be potentially used by an AS relationship inference algorithm are noisy and only have a weak correlation with the link's relationship type. Second, a large number of links appear in paths that likely violate the valley-free assumption made by existing algorithms. Third, existing algorithms are sensitive to the locations of the vantage points and the order in which the link relationships are inferred. An AS relationship inference technique has to address the above challenges if it is to achieve higher accuracy on the various *hard* link benchmark datasets.

We develop a *probabilistic AS relationship inference* algo-

rithm, `ProbLink`, to address the above issues. In particular, it provides a framework that allows for easy integration of many noisy but useful attributes into the relationship inference algorithm. We identify a set of link attributes that take into account not only observed paths but also information gleaned from the fact that certain paths are not observed. It allows for links to appear in paths that violate the valley-free property but attributes a lower probability to such occurrences. Finally, the algorithm uses an iterative algorithm that repeatedly infers the link types until they reach a fixed point in terms of the inferences as well as the underlying statistical distributions of link attributes. We evaluate this algorithm and show that it achieves significantly lower error rates when compared to prior techniques such as `AS-Rank`. Our work makes the following contributions:

- We design a strikingly simple algorithm, `CoreToLeaf`, that uses just the valley-free assumption, yet achieves inference accuracy comparable to the more complicated `AS-Rank`. We use `CoreToLeaf` to identify five categories of *hard* links existing in the Internet and show that the current validation datasets are skewed toward *easy* links.

- Despite the seemingly high accuracy of the `AS-Rank` algorithm, we show that its accuracy is still insufficient for several real-world applications.

- We design a robust clique inference mechanism, a key set of route features and compose these features to devise a probabilistic AS relationship inference algorithm, `ProbLink`, which reduces the error rate of `AS-Rank` overall by $1.7\times$, and achieves 1.8-$6.1\times$ better error rate for various categories of *hard* links.

- We show that even the small improvement in overall accuracy brought by `ProbLink` has a significant impact when applied to real-world applications. Compared to the current state of the art, `ProbLink` increases the precision and recall of route leak detection by $4.1\times$ and $3.4\times$ respectively, reveals 27% more complex relationships, and increases the precision of predicting the impact of selective advertisements by 34%.

- We will publicly release the code for `ProbLink` and its periodically inferred relationships.

2 Background and Related Work

The Border Gateway Protocol (BGP) is the mechanism used by ASes to exchange reachability information. A BGP AS path is a sequence of ASes denoting the routing path that the first AS in the path prefers to reach a destination prefix. The last AS in the path is referred to as the "origin AS" of the prefix. Each AS uses a complex decision process to select the most preferred path toward each destination prefix [7]. A BGP route collector infrastructure is operated by Routeviews [4] and RIPE NCC [3], which consist of routers that peer with ASes that volunteer to provide their BGP feeds for research or operational reasons. A route collector is config-

ured to obtain the *best paths* from ASes it peers with, i.e., an AS peering with a route collector exports its best path towards each destination prefix to the collector. Route collectors typically peer with several ASes, and thus obtain the best paths from those ASes to each destination prefix.

AS relationships fall into two broad categories: customer-provider (c2p) and settlement-free peering (p2p). In a c2p relationship, the customer AS pays the provider AS for reachability to the rest of the Internet. In a p2p relationship, two networks agree to exchange traffic destined to their respective networks or their customers without an associated fee. In practice, AS relationships can span the spectrum between c2p and p2p. Such cases, of *hybrid or complex relationships*, can occur when two ASes have different contractual agreements in different geographical regions [20]. Also, sibling relationships happen between ASes under the same organization, which can exchange traffic without any cost and routing restrictions. The *customer cone* of an AS *X* is the set of ASes that *X* can reach using only p2c links. The size of the customer cone can be perceived as a metric of the market power of an AS. A clique of *Tier-1 ASes* at the top of the Internet AS hierarchy are "transit-free", meaning that they have routes to all the other networks on the Internet through customer or peering links without the need to pay for transit.

2.1 AS Relationship Inference Techniques

Beginning with the seminal work by Gao [16], most algorithms for AS relationship inference are based on the assumption that valid BGP paths are *valley-free*, i.e., a path consists of zero or more c2p links, followed by zero or one peering link, followed by zero or more p2c links. This assumption intends to capture the economic incentives that, at least partially, drives traffic exchange between ASes . An AS advertising routes learned from a peer or provider to another peer or provider would be offering "free transit" which would be against its financial interests.

Another assumption made by Gao and other subsequent researchers [42, 10, 11, 47] is that the providers usually have a higher *node degree* (i.e., the number of ASes to which an AS node directly connects to) than customers, while peers usually have similar degrees. *node degree* is the number of neighbors an AS directly connects to, irrespective of whether the neighbors are providers, peers, or customers to the AS; It has been exhibited that node degree is significantly biased by the topology incompleteness problem, namely the fact that the available topological data reveal only a subset of the complete Internet topology due to their limitation to observe a large fraction of peer-to-peer AS links [46].

The state-of-the-art AS relationship inference technique, called the "`AS-Rank`" algorithm [30], made three generally accepted assumptions: 1) there is a clique of large transit providers at the top of the hierarchy, 2) most customers purchase transit in order to be globally reachable, and 3) there are no cycles of p2c links. The `AS-Rank` algorithm takes 11

intricate steps to label each link as customer-provider (abbreviated as c2p or p2c depending on the directionality of the relationship) or peer-to-peer (p2p). The complete steps of the `AS-Rank` algorithm is available in Appendix A.

It is worth observing a few properties of the `AS-Rank` algorithm. First, it uses the *valley-free* assumption and the *transit degree* attribute as the primary pieces of information in determining the labels. The *transit degree*, expresses the number of ASes that appear on either side of an AS in adjacent links in BGP paths, and it does not count neighbors for which the given AS does not transit traffic. Transit connectivity is easily observable by Route Collectors (with the exception of backup or partial-transit links) [22], therefore it provides a better metric to describe an AS's prominence than node degree. Second, the algorithm visits ASes and links in a specific order, uses the *transit degree* information in certain cases (step 5) and not in others (step 7), and does not have a mechanism for resolving conflicting information in cases where the *valley-free* property is violated.

3 Input Datasets

3.1 BGP Paths

We collect BGP paths from RouteViews [4] and RIPE RIS [3]. In September 2018, both projects operate 22 collectors, which in total connect with more than 1,000 vantage points (VP) worldwide. Each RouteViews and RIPE RIS collector dumps a snapshot of their Adj-RIB-out tables every 2 hours and every 8 hours respectively. For the purpose of evaluating the various algorithms over longitudinal data (as discussed in §4 and §7), we consider snapshots of BGP paths on the first day of April, August, and December (i.e., every four months) since 2006.

After collecting BGP paths, we parse them to remove duplicated ASes which result from BGP path *prepending*. We also filter out paths with AS loops, i.e., when an ASN appears more than once and is separated by at least one different ASN. We also sanitize the BGP paths by removing paths containing reserved ASes [39]. Loops and reserved ASes showing up in a path are artifacts of route poisoning [6, 27].

3.2 Sibling Relationships

We use CAIDA's AS-to-organization mapping dataset [8], which is derived from WHOIS data, to identify sibling links. This dataset provides quarterly information starting from 2009. We infer links between ASes that are operated by the same organization as sibling relationships.

3.3 IXP List

ASes often establish p2p relationships over shared switching fabric provided by IXPs. To facilitate dense peering connectivity, IXPs provide BGP Route Servers over which ASes establish many-to-many (multilateral) interconnections. To enable layer-3 connectivity Route Servers typically have their own ASN, but according to best practices it should be

Date (MM/DD/YYYY)	# links in validation set	# links in total	Percentage
04/01/2012	7,833	117,872	6.6%
04/01/2013	11,644	133,459	8.7%
04/01/2014	44,875	159,678	28.1%
04/01/2015	47,036	176,791	26.6%
04/01/2016	52,931	204,309	25.9%
04/01/2017	56,326	213,441	26.4%

Table 1: Size of the validation dataset vs. all links observed from all VPs.

filtered-out from the AS path since the Route Server does not participate in routing decision process [25]. However, for debugging reasons some IXP members append the Route Server ASN in the BGP path. We sanitize BGP paths to remove Route Server ASNs since essentially the peering links are between the IXP members, and not between the IXP and ASes. To collect a list of AS Numbers (ASNs) used by IXP Route Servers, we query PeeringDB [2] for networks of type "Route Server" and we extract the ASN. We augment this list using the Euro-IX IXP Service Matrix [1] to extract the Peering LAN ASN and Route Server ASN for each IXP. There are 172 IXP ASes in this list on 12/01/2017.

3.4 Validation Dataset

AS operators frequently encode the relationship type with their neighbors directly in their prefix advertisements using BGP Communities, an optional transitive BGP attribute used to attach metadata on BGP paths. While the use of Communities is not standardized, many ASes publicly document the meaning of their BGP communities on websites and in IRR databases, enabling us to assemble a dictionary of BGP communities that denote relationship type. We used a dictionary of 1286 community values from 224 different ASes to construct a set of relationships from BGP data starting quarterly from April 2006 to April 2017. Similarly to [30], we treat this dataset as "best-effort" validation to evaluate existing inference techniques and our proposed approaches.

Table 1 shows the size of this validation dataset over the past 6 years. The coverage of our validation dataset increases from 6.6% to over 26% of the observed links between 2014 – 2016, due to the increasing popularity of BGP Communities and the deployment of additional VPs that allow more communities to propagate to BGP collectors. We believe this high coverage of our validation dataset is sufficient to evaluate different AS relationship inference techniques. As prior work has pointed out, links involving *Tier-1* ASes and VP ASes are over-represented because public data on BGP communities mostly comes from large ASes [30], while communities from non-VP ASes may be stripped out during the propagation of BGP routes. However, unlike prior work, we take these biases into consideration during our evaluation.

4 Establishing a Benchmark for Hard Links

As explained in the previous section, the evaluation dataset is extensive but biased toward specific types of links. It is im-

portant to understand if the links over-represented in the validation dataset are easier to infer correctly, compared to the under-represented links, which may skew the overall evaluation results. To this end we develop `CoreToLeaf`, a very simple algorithm that allows us to understand which links are easy to infer. `CoreToLeaf` uses only the *valley-free* assumption and the list of Tier-1 ASes to infer relationships. We show that the inference accuracy of this algorithm is almost as high as that of the more sophisticated `AS-Rank` algorithm elaborated in §2.1, while the accuracy of both algorithms suffer for certain categories of links. Our findings reveal that indeed certain types of under-represented links in the evaluation dataset are harder to infer, possibly inflating the overall accuracy of past work. We address this issue by constructing distinct validation sub-datasets as benchmarks for hard links.

4.1 The `CoreToLeaf` Algorithm

`CoreToLeaf` starts by inferring a clique of Tier-1 ASes using the same inference method as `AS-Rank`, which is available in Appendix B. For each path that traverses a Tier-1, we skip the first link after the Tier-1 and label all succeeding links as p2c. For example, if AS_2 is a clique member in a BGP path "$AS_1, AS_2, AS_3, AS_4, AS_5, AS_6$", we infer links $<AS_4 - AS_5>$ and $<AS_5 - AS_6>$ as p2c. We skip inferring the relationship for $<AS_2 - AS_3>$ because it could either be a p2c or a p2p, but all subsequent links need to be p2c assuming that the path is *valley-free*. (Note that if AS_1 is a clique member, we would have labeled $<AS_2 - AS_3>$ as a p2c link.) Finally, we label all remaining unclassified links as p2p.

In the second step of labeling p2c links, a link could be labeled more than once if it shows up in multiple paths. In some cases, a link could be labeled as a p2c in some path and as a c2p when traversing a different path. We label this link as a "conflict" link when we encounter such an inconsistency.

Note that `CoreToLeaf` does not take into account degree or transit degree information nor does it use paths that do not go through Tier-1s. The rationale behind `CoreToLeaf` is simply that there is greater certainty that it is customer routes that are being transitively exposed to Tier-1s and that there is less likelihood of paths being exported to Tier-1s due to complex peering mechanisms.

4.2 Evaluation

We evaluate this extremely simple algorithm against our validation dataset on 04/01/2017, which contains 23,528 p2p links and 32,798 p2c links (26.4% of the visible topology). Table 2 compares the *precision* (true positives / (true positives + false positives)) and *recall* (true positives / (true positives + false negatives)) of `CoreToLeaf` and `AS-Rank`. Surprisingly, `CoreToLeaf` achieves high precision and recall for both p2c and p2p links (comparable to `AS-Rank`), with higher *precision* on p2c relationships (98.9% compared to 97.8%), and a small fraction of links labeled as 'conflict'.

The 1.1% mistakenly inferred p2c links and the links

Algorithm	p2c		p2p		Conflict
	Precision (%)	Recall (%)	Precision (%)	Recall (%)	(%)
`CoreToLeaf`	98.9	95.8	95.0	98.8	0.12
`AS-Rank`	97.8	97.5	98.8	98.9	0

Table 2: Precision and recall of `CoreToLeaf` and `AS-Rank`.

which `CoreToLeaf` labels as 'conflict' by are due to *valley-free* violation, which we quantify later in §5.2. Since the second step of labeling p2c links uses just paths through Tier-1s, it fails to capture 4.2% (95.8% *recall*) of the actual p2c links. Consequently, these links are inferred as p2p in the third step and result in a 5.0% error rate for links labeled as p2p.

The accuracy of `CoreToLeaf` and `AS-Rank` seem quite high, but they perform poorly when applied to real-world applications. Route leaks constitute a type of prevalent routing incident which can cause significant disruptions to Internet routing [41, 21]. In §8, we describe how we can use inferred relationships to detect route leaks and evaluate the effectiveness of `AS-Rank` inferences. Only 19% of the route leaks detected using `AS-Rank` were real route leaks, and almost 80% of the real leaks were missed. We observe relatively poor performance for two more applications we tested, discussed in detail in §8, which illustrates that a better AS relationship algorithm is needed for real-world applications.

4.3 Identifying Hard Links

The surprisingly high accuracy obtained by `CoreToLeaf` has many implications. First, it indicates that simple techniques might suffice for inferring the types of many of the links in the validation dataset. Second, it underscores the need for more comprehensive validation datasets, that would represent better AS links beyond those associated with Tier-1 and VP ASes. Third, in the absence of more comprehensive validation datasets, one way to make progress on improving and evaluating AS relationship inference algorithms is to identify specific types of links for which the current algorithms do not work well. We therefore now attempt to extract collections of *hard* links from the overall validation dataset based on the inference performance of `CoreToLeaf` and `AS-Rank`.

We feed a large set of features of every link in the validation dataset along with each link labeled as "inferred correctly" and "inferred incorrectly" by `CoreToLeaf` and `AS-Rank` into the *gradient boost decision tree* [15] , and calculate the feature importance. Supporting feature importance calculation and results are in Appendix C. We extract five categories of "hard" links suggested by the feature importance analysis and the `CoreToLeaf` algorithm.

1) Links with max node degrees smaller than 100. The feature importance analysis shows that links whose endpoint ASes both have small node degrees are hard to be inferred correctly by `CoreToLeaf` and `AS-Rank`.

2) Links observed by more than 50 but less than 100 VPs. Driven by the feature importance results, links observed by the range of 50 to 100 VPs are hard to infer correctly. The

Category	CoreToLeaf (%)	AS-Rank (%)	Fraction all links	Fraction validation
Max node degree <100	13.7	8.6	16.1%	1.7%
Observed by 50-100 VPs	4.7	9.3	9.9%	8.1%
Non-VP & Non-Tier1	5.3	9.0	24.2%	11.6%
Unlabeled Stub-clique	95.5	33.4	0.3%	0.1%
Conflict	100.0	8.1	0.24%	0.16%

Table 3: Error rates of `CoreToLeaf` and `AS-Rank` on *hard* links on 04/01/2016. The fraction of each category of *hard* links that is in overall links vs. in the validation dataset shows that *hard* links are underrepresented in the validation dataset.

reason is that p2p links are often observed by few VPs, and transit links are often observe by many VPs, so the range in the middle is hard to distinguish these link types.

3) Non-VP and non-Tier1 links. In general, a link that is directly connected to a VP or a Tier-1 is likely to appear in many BGP paths, and the AS inference algorithm is likely to have access to more information regarding the link. Moreover, most of our validation dataset are links that are connected to a VP or Tier-1 AS, so we want to specifically analyze the performance of inference algorithms on the "underrepresented" links in our validation dataset.

4) Unlabeled stub-clique links in `CoreToLeaf`. A stub AS connects with only one other AS, through which it gains access to the entire Internet. A stub-clique link is a link whose one endpoint is a stub AS and the other endpoint is in the Tier-1 clique. In other words, the clique member is the only AS to which the stub AS connects. These links typically have very high transit degree difference, which is an important feature as shown by the feature importance analysis.

In `CoreToLeaf`, a stub-clique link $<X, Y>$ (where X is a stub AS, and Y is a clique AS) is inferred as a c2p *iff* there is a path containing an AS triplet "Z, Y, X" where Z is also a clique AS. We call the set of stub-clique links which are not inferred as c2p in the second step of `CoreToLeaf` (thus they are then inferred as p2p) "unlabeled stub-clique links".

In step 9 of `AS-Rank`, stub-clique links are classified as c2p by default based on the assumption that stub networks are extremely unlikely to meet the peering requirements of clique members. We believe this is an unjustified assumption and with the trend of "Internet flattening", peering relationships between high-tier ASes and low-tier ASes are becoming more prevalent [17].

5) Conflicts in `CoreToLeaf`. Recall that `CoreToLeaf` labels some links as "conflicts". These links appear to behave as p2c on some paths and c2p on others, and the main reason for this is violation of *valley-free* routing. We believe that this set of links is difficult to analyze because the two endpoints are likely to have unconventional routing policies.

Table 3 shows the error rates of inferences made by `CoreToLeaf` and `AS-Rank` on each category of hard links on 04/01/2016. We observe both algorithms yield more errors than their inferences on normal links, especially on unlabeled stub-clique links. Furthermore, the fraction of every category of the hard links in the validation dataset is less than

that in the overall links, especially for the "Max node degree < 100" category. This indicates that the validation dataset is skewed to *easy* links. In addition to the entire validation dataset, we will use these more specific datasets for evaluating the AS inference algorithms in the subsequent sections.

5 Challenges With AS Relationship Inference

In this section, we identify three main challenges with AS relationship inference, and describe how they hamper existing inference techniques. This analysis helps inform the design of a probabilistic algorithm for AS relationship inference.
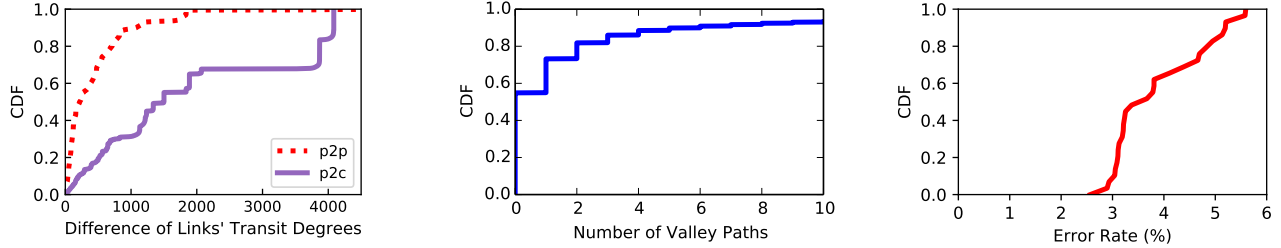
5.1 Degree Inversion

An AS inference algorithm can use any observed attribute associated with a link, its two endpoint ASes, AS links, and end-to-end paths that traverse the link in order to determine the link type. However, most attributes have stochastic information value, as we will illustrate below for AS degree.

Many existing techniques for inferring AS relationships make three assumptions: highest-degree ASes sit on the top of the routing hierarchy; peering ASes have similar degrees; and providers have larger degree than customers [16, 30, 12].

Over the past four years, the top two nodes with the largest transit and node degrees (as observed through BGP feeds from available VPs) have consistently been *AS*6939 (Hurricane Electric) and *AS*174 (Cogent Communications). However, both of these ASes are not Tier-1 ASes [45], so the assumption that the ASes with the highest degrees sit on top of the routing hierarchy is not universally valid. This fact influences the accuracy of some inference approaches since a key step in these approaches is identifying a clique of Tier-1 ASes at the top of the hierarchy [16, 30].

Figure 1a plots a CDF of the absolute transit degree differences of different link types. Transit degrees of ASes are computed from BGP paths observed on 04/01/2017, and link types are derived from the validation dataset on 04/01/2017 as described in §3.4. The validation dataset includes 55,016 links, 30,859 of which are p2c links and 24,157 are p2p links. We see that even though p2c links usually have larger degree differences than p2p links, over 14% of the p2p links have absolute transit degree differences larger than 1000, making many p2p links indistinguishable from p2c links in terms of transit/node degree difference.

According to this observation, the existence of substantial differences in node/transit degrees between peering ASes is common. This phenomenon is explained in part by the fact that, during recent years, large content providers such as Google, Akamai, and Microsoft, which usually have high degrees, are more willing to peer with large numbers of lower tier ASes to get free and more efficient traffic exchange [43, 23]. This trend is referred to as the "flattening" of the Internet [17], and it significantly influences the AS relationship inference techniques that differentiate peers from providers or customers based on transit/node degree differ-

(a) CDF of absolute transit degree difference.

(b) CDF of the number of paths that violate the valley-free property traversing each link.

(c) CDF of `AS-Rank`'s error rates on 30 consecutive 1-day BGP snapshots from April 1, 2016 to April 30, 2016.

Figure 1: Analysis of transit degree difference, valley-free violations, and error rates of `AS-Rank`.

ences, or rank ASes in decreasing order by degrees and label links based on the order in which ASes are considered (as is the case with `AS-Rank`).

5.2 Violation of Valley-Free Property

Next, we study the prevalence of *valley-free* violations, which is the culprit behind mistakenly inferred p2c links and 'conflict' links in `CoreToLeaf`.

3% of the BGP paths violate *valley-freeness* in the `AS-Rank` inference on 04/01/2012. This snapshot is incidentally one where the `AS-Rank` algorithm achieves the highest inference accuracy. We also find this level of *valley-free* violations is persistent over the various snapshots in our study.

Figure 1b shows a CDF of the number of paths that violate the valley-free property for the links in BGP paths on 04/01/2012. 47% of the links in the AS topology are traversed by paths that violate the valley-free property. This statistic is consistent with prior work that analyzes the prevalence of valley-free violations in Internet routing and characterize it is a result of the deliberate BGP policies of ASes that use unconventional economic models [19].

The existence of these violations has certain implications for AS relationship inference. First, a robust inference algorithm has to take into account the structure of all paths traversing a given link. Second, it might have to update the inference made for a given link after inferring the types of neighboring links. It is worth noting that the `AS-Rank` algorithm does not conform to the above guidelines and is therefore prone to erroneous inferences for links with paths that violate the valley-free property.

5.3 Current Techniques are Sensitive to VP and Snapshot Selection

We observe high variation in accuracy when applying the `AS-Rank` algorithm to consecutive snapshots of BGP paths. Figure 1c plots a CDF of `AS-Rank`'s error rates($1 - accuracy$) on 30 consecutive 1-day BGP snapshots in April, 2016. As shown in Appendix D.1, AS-Ranks accuracy is also quite sensitive to the VP selections.

The reason for the `AS-Rank` algorithm's sensitivity to snapshot and VP selections lies in the first step of its infer-

ence algorithm that identifies the Tier-1 clique and the subsequent steps that labels links in a particular order starting with the Tier-1 ASes. `AS-Rank` first finds the maximum clique from the AS-links involving the largest ten ASes by transit degree, then visits the rest of the ASes top-to-bottom, and adds an AS to the clique if it connects with all the members in the current clique. It then labels p2c links using path segments that radiate from the Tier-1 clique. Errors that creep into the clique determination step have a significant impact on the order in which AS links are analyzed and labeled. See Appendix D.2 for detailed discussions.

6 Probabilistic AS Relationship Inference

In this section, we present a new AS relationship inference algorithm, `ProbLink`, that is designed to address the challenges discussed above. First, `ProbLink` is a probabilistic algorithm that enables the use of link attributes with stochastic information value. Second, in determining a link's type, `ProbLink` simultaneously takes into account all information regarding the links and the paths that traverse it, and provides a framework for integrating conflicting information (e.g., paths that violate the valley-free property). Third, `ProbLink` does not prescribe a specific order in which ASes and links are considered, but rather continually updates the link type inferences and iterates till it reaches a fixed point in terms of the underlying stochastic distributions.

Crucially, our algorithm provides a framework for integrating various link attributes that might help infer a link's type. We therefore first design a set of link features or attributes that provide noisy but still informative signals regarding the AS relationships. In particular, we design features that capture routing behavior in terms of both observed and unobserved routes as well as integrate information regarding a link's endpoints. We then describe how we use these features to build a probabilistic inference model.

6.1 Overview

Our algorithm starts with an initial classification of links based on the inference result of `CoreToLeaf`, so at the beginning, each link has deterministic relationship probabilities. More concretely, if `CoreToLeaf` labels L as a p2p link,

6

we will convert it to $P(\mathsf{L} = p2p) = 1.0$, $P(\mathsf{L} = p2c) = 0.0$, $P(\mathsf{L} = c2p) = 0.0$ and provide that as the input to our algorithm. Note that `ProbLink` is essentially a meta-inference algorithm that can be bootstrapped by outcomes of any algorithm. Its performance is independent on the bootstrapping algorithm we choose, which we evaluate later in §7.1.

For each feature, `ProbLink` computes the conditional probability distribution based on observed data and the initial set of relationship types attributed to links. In each iteration, we update the probabilities of each link's types ($P(\mathsf{L} = p2p)$, $P(\mathsf{L} = p2c)$, $P(\mathsf{L} = c2p)$) by running our probabilistic algorithm described in §6.4, and recompute the distributions of features using the updated probability values of each link. We repeat this process until convergence, i.e., the percentage of links that change labels between each iteration drops below a small threshold.

6.2 Clique Inference

We attempt to infer the ASes that are at the top of the hierarchy, namely Tier-1 ASes, because it is used in a feature by `ProbLink`. Tier-1 ASes should have the largest *customer cones* [30], so estimating the customer cones is the core of doing clique detection.

First, we find top N ASes in terms of transit degree, denoted as D.[1] These ASes are either Tier-1 or Tier-2 ASes because of the large amount of neighbours they provide transit with. Then we estimate the customer cone size of each AS in the graph by determining the average number of destination ASes (last hops) for which an element of D uses this AS as part of a route. This is an effective way of estimating the customer cone size because irrespective of whether a node d in D is a Tier-1 or a Tier-2 AS, it will transit through Tier-1 ASes to reach their customer cones.

Second, we find the maximal clique C with largest estimated customer cone size sum in D. Then, we test every other AS in order by estimated customer cone size to add members to C. An AS is added to C if it has links with every other AS in C. If there are three consecutive members (X-Y-Z) in C showing up in paths, disconnect the edge between X and Z even though X and Z are connected in some paths, because no AS path should have three consecutive clique ASes.

6.3 Feature Design

An AS link can be characterized by the following three attributes: (A) The structure of paths that use the link; (B) The structure of paths that *do not* use the link; (C) Properties of the ASes on each side of the link. We carefully design six features that correspond to these three types of attributes.

Triplet feature (Type A). The triplet feature considers link triplets that appear in paths and attributes probabilistic values for the relationships of the first and the last links given the relationship of the middle link. Suppose three consecutive links "L1 - L - L2" show up in a BGP path, where L1, L,
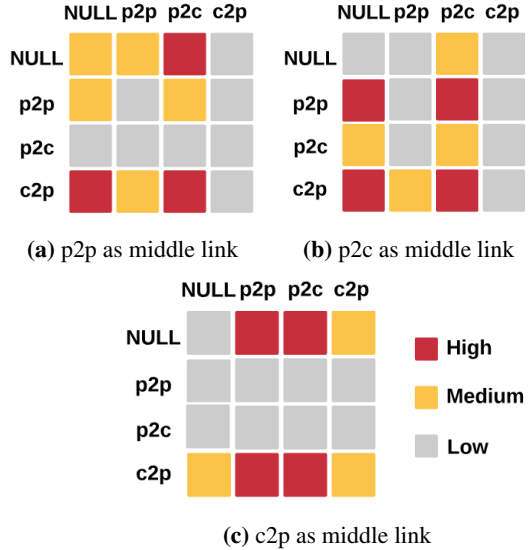
---
[1]Any value of N between 10–40 does not affect the final result.

(a) p2p as middle link   **(b)** p2c as middle link

(c) c2p as middle link

Figure 2: Conditional probability distribution for the triplet feature describing $P(previous, next \mid middle)$. Probability values in the ranges of > 0.1, $[0.01, 0.1]$, and < 0.01, are categorized as high, medium, and low respectively in the figure.

L2 are three links (AS pairs). "L1 - L - L2" is called a link triplet. We break down each BGP path in a snapshot into link triplets, and, for the first and the last links in each path, we insert a " NULL" link in front of and behind it. For example: a BGP path "8793 6939 1103 198499" is decomposed into 3 link triplets: "NULL - <8793, 6939> - <6939, 1103>", "<8793, 6939> - <6939, 1103> - <1103, 198499>", and "<6939, 1103> - <1103, 198499> - NULL". As a consequence, each link in the BGP paths appears as a middle link in at least one link triplet. We take into account sibling relationships (in §3.2) by skipping sibling links, i.e. treating a sibling link as a single AS, when constructing triplets.

The goal of the triplet feature is to model *valley-freeness* in a probabilistic way. For each middle link type, we compute the probability of the link type of its adjacent previous link and next link. If we put the link type of the previous link along the y-axis and the link type of the next link along the x-axis, we get a matrix view as shown in Figure 2, which is computed from `CoreToLeaf` initial labels. Each cube in the matrix represents a probability that is categorized as high, medium, and low depending on the it is in the range of > 0.1, $[0.01, 0.1]$, or < 0.01. For example, we can see from Figure 2a that when the middle link is of type p2p, the previous and next link are most likely to be <NULL, p2c>, <c2p, NULL> and <c2p, p2c>, but its previous link is very unlikely to be p2c no matter what its next link's type is.

Non-path feature (Type B). In addition to observed routes, unobserved routes also provide some information regarding AS relationships. The *non-path feature* describes the probability of how many adjacent p2p or p2c links a link has, but none of them appear as a previous link with this link on any of the paths. This feature is designed to capture the
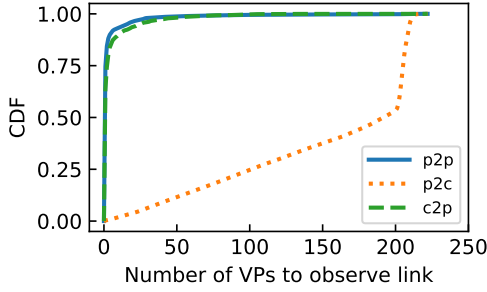
Figure 3: The visibility of each link type derived from `CoreToLeaf` inference results on 04/01/2017 BGP paths.

property that a link is unlikely to be a p2c link if it has lots of adjacent p2p/p2c links and none of them appear as a previous link on any of the paths containing the link.

Similar to the triplet feature, the non-path feature also models *valley-freeness* in a probabilistic way. The non-path feature is not necessarily applicable to all links. When a link does follow a p2p or p2c link or if a link does not have any adjacent p2p or p2c links, the non-path feature does not play a role in inferring the link's type.

Distance to clique feature (Type C). The *distance to clique* feature captures the fact that high-tier ASes are closer in distance (AS hops) to clique ASes than low-tier ASes, and that ASes in the same tier are likely to be peers, while high-tier ASes tend to be providers of low-tier ASes.

We first create an undirected graph by adding AS links as edges, and then compute the shortest path from each AS towards each clique member using Dijkstra's algorithm. For each AS, we compute its average distance to each member in the clique set, round it to a multiple of 0.1, and denote this value as dist(AS). We represent each link $<AS_1, AS_2>$ in the graph by a distance to clique tuple "dist(AS1), dist(AS2)".

Vantage point feature (Type C). The number of VPs observing a link also suggests the link type. The vantage point feature captures the probability of a certain number of VPs with at least one path traversing a link for each link type. This feature naturally folds in the following intuition: p2c links are more likely to be seen by more VPs compared to p2p and c2p links. This feature considers path directions. For example, let's consider a link L ($<AS1, AS2>$) where AS1 is the provider of AS2. `ProbLink` computes probabilities separately for both directions by counting how many paths traverse L in the direction of $<AS1, AS2>$, and how many paths traverse L in the direction of $<AS2, AS1>$

To evaluate the informational value of this feature, we analyze the number of VPs that observe a given link and correlate that with the link's type computed by `CoreToLeaf`. Figure 3 shows the CDF of the number of VPs that observe a link for each link type. We observe that 93% of p2p links and 90% of c2p links are observed by ≤ 10 VPs, while 98% p2c links are seen by more than 10 VPs.

Co-located IXP and co-located private peering facil-

ity feature (Type C). The *co-located IXP* and *co-located peering facility facility* information is extracted from PeeringDB [2]. These features are based on the intuition that the more IXPs or facilities two ASes are co-located in, the more likely they are peering with each other. Based on the validation data, 90% of transit links do not have any co-located IXPs or facilities, while more than 70% p2p links have at least one co-located IXP or facility.

6.4 Inference Algorithm

We begin by reviewing the Naive Bayes classifier. Given a link type variable C (which can be p2p, p2c, c2p) and a feature vector f_1 through f_n, *Bayes' theorem* states the following relationship:

$$P(C \mid f_1, ..., f_n) = \frac{P(C, f_1, ..., f_n)}{P(f_1, ..., f_n)} \qquad (1)$$

By assuming that each feature f_i is conditionally independent of every other feature:

$$P(f_i \mid C, f_1, ..., f_{i-1}, f_{i+1}, ..., f_n) = P(f_i \mid C) \qquad (2)$$

Using the *chain rule* to rewrite the numerator of Eq. 1:

$$P(C, f_1, ..., f_n) = P(C) \prod_{i=1}^{n} P(f_i \mid C) \qquad (3)$$

So,

$$P(C \mid f_1, ..., f_n) = \frac{P(C) \prod_{i=1}^{n} P(f_i \mid C)}{P(f_1, ..., f_n)} \qquad (4)$$

Since the denominator $P(f_1, ..., f_n)$ does not depend on the class C, the naive bayes classifier assigns a link being a type $\hat{C}$ by the following function:

$$\hat{C} = \arg\max_{C} P(C) \prod_{i=1}^{n} P(f_i \mid C) \qquad (5)$$

The inputs to `ProbLink` are BGP paths, link triplets extracted from these BGP paths, and initial relationship labels for each link as inferred by a bootstrapping algorithm. Algorithm 1 shows the pseudocode of `ProbLink`.

First, the algorithm calculates probabilities for each feature, conditional on the link type C (C in {p2p, p2c, c2p}) by accumulating probability values (line 2-3 in Algorithm 1). The parameter α is a smoothing parameter, which prevents a feature with examples in only one class from forcing the probability estimate to be 0 or 1. In our implementation, we use *Laplace (Add-1) Smoothing* [31] which sets the smoothing parameter to 1. The algorithm then assigns probability that link L is of each type by the prior probability distribution $P(C)$, which is the proportion of each link type in the data (line 5-8 in Algorithm 1). Then it goes through each feature and multiplies the probability that link L is of each type by the conditional probability of the feature given each link type (line 9-12). In the end, the final probability of the

Algorithm 1: `ProbLink`: probabilistic AS relationship inference algorithm based on Naive Bayes

---

Input : 1) BGP paths $\rightarrow$ link triplets
      2) Initial AS relationships R
      3) Feature vector f = [Triplet, Non-path, Distance to clique, VP, Co-located IXP, Co-located facility]

Output: Inferred probabilities of each link being p2p, p2c, c2p

/* Loop until convergence   */

1 **while** R $-$ *last*(R) $> \varepsilon$ **do**

    /* Compute conditional distribution of each feature   */

2   **foreach** *feature f_i in feature vector* f **do**

3        $P(f_i \mid C) = \frac{P(f_i, C)}{P(C)} = \frac{N(f_i, C) + \alpha}{N(C) + \alpha d}$

4   **foreach** *link* L **do**

5        $all \leftarrow N(p2p) + N(p2c) + N(c2p)$

6        $P(L = p2p) \leftarrow P(p2p) = \frac{N(p2p)}{N(all)}$

7        $P(L = p2c) \leftarrow P(p2c) = \frac{N(p2c)}{N(all)}$

8        $P(L = c2p) \leftarrow P(c2p) = \frac{N(c2p)}{N(all)}$

9   **foreach** *feature f_i in feature vector* f **do**

10           $P(L = p2p) *= P(f_i \mid p2p)$

11           $P(L = p2c) *= P(f_i \mid p2c)$

12           $P(L = c2p) *= P(f_i \mid c2p)$

13        $sum = P(L = p2p) + P(L = p2c) + P(L = c2p)$

14        $P(L = p2p) \leftarrow P(L = p2p)/sum$

15        $P(L = p2c) \leftarrow P(L = p2c)/sum$

16        $P(L = c2p) \leftarrow P(L = c2p)/sum$

        /* Update link's type   */

17   R $= \arg\max_C P(L = C)$

---

link L of being each type is calculated by the fraction of each type's probability over the sum of probabilities of all possible link types (lines 14-16 in Algorithm 1). Then we update L's type by picking the link type with the largest probability (line 17). We repeat this process of link type inference and updating probability distributions of features until convergence, i.e., the percentage of links that change labels between each iteration drops below a small threshold.

7 Evaluation

We now evaluate our probabilistic inference algorithm, `ProbLink`, from three aspects and show that:

- `ProbLink` consistently achieves low error rates across many years, reducing the average error rate of `AS-Rank` for all links by $1.7\times$, and achieving 1.8-$6.1\times$ better error rates for the different categories of *hard* links.

- `ProbLink` is not dependent on its bootstrapping algorithm and stable to the snapshot and VP selection.

- Each feature in our algorithm is meaningful but none of them by itself is sufficient to make good inferences.
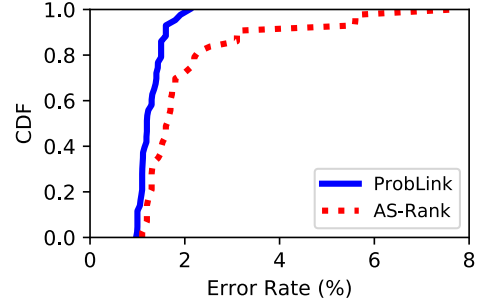

Figure 4: CDF of error rates of `ProbLink` and `AS-Rank` on the snapshots of BGP paths in the past 6 years.
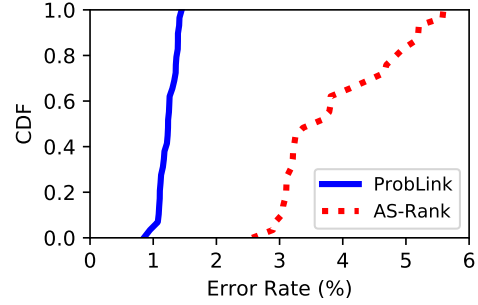

Figure 5: CDF of error rates of `ProbLink` and `AS-Rank` on 30 consecutive 1-day snapshots from April 1, 2016 to April 10, 2016

7.1 Accuracy

To evaluate the accuracy of `ProbLink`, we assemble daily snapshots of BGP paths on the first five days of April, August, and December (i.e., every four months) over the past 6 years. We apply our algorithm against these snapshots, and compare it with the `AS-Rank` algorithm over this time period. We translate the probability values of link types produced by our algorithm to one of the three link types that has the highest probability value among the three possible types.

Figure 4 compares the error rates of inferences made by `ProbLink` and those made by `AS-Rank`. Our probabilistic inference algorithm consistently yields a low error rate smaller than 2%, reducing the average error rate of `AS-Rank` for all links from 2.1% to 1.2%.

Figure 5 shows a comparison between error rates of `ProbLink` and `AS-Rank` on 30 consecutive snapshots of BGP paths during April 2016. The max and average error rates across these days for `ProbLink` are 1.4% and 1.2%, while the error rate of `AS-Rank` ranges from 2.6% to 5.6%, with an average error rate of 3.9%. `ProbLink` is not sensitive to the specific set of paths used in a snapshot and that it achieves uniformly low error rates in spite of clique inference inaccuracies that adversely impact `AS-Rank`.

Figure 6 plots the CDFs of error rates of inferences made by `ProbLink` and `AS-Rank` on the four categories of *hard* links identified in §4.3. Not only does our algorithm yield much smaller error rates, but also has less variation than `AS-Rank`. Table 4 lists the average error rate of our algo-
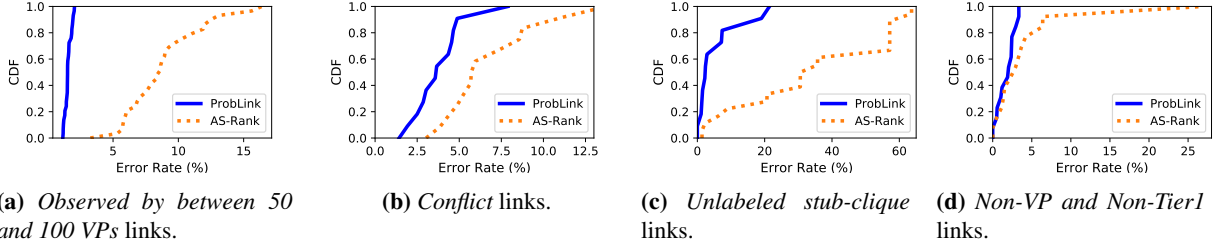
(a) *Observed by between 50 and 100 VPs* links.

(b) *Conflict* links.

(c) *Unlabeled stub-clique* links.

(d) *Non-VP and Non-Tier1* links.

Figure 6: CDFs of error rates of `ProbLink` vs. `CoreToLeaf` on *hard* links in a period of 30 days in April 2016.

Category	AS-Rank(%)	ProbLink(%)
Observed by 50-100 VPs	8.8	1.5
Non-VP & non-Tier1	4.4	1.7
Unlabeled Stub-clique	33.6	5.5
Conflict	6.8	3.8

Table 4: Average error rates on *hard* links. `ProbLink` achieves 5.9×, 2.6×, 6.1×, and 1.8× better error rate for the links observed by between 50 and 100 VPs, non-VP & non-Tier1 links, unlabeled stub-clique, and conflict than `AS-Rank` respectively.

rithm and `AS-Rank` for links *observed by 50 to 100 VPs*, *non-VP and non-Tier1* links, *stub-clique* links, and *conflict* links. Our probabilistic algorithm reduces the error rate on the four categories by a factor of 5.9, 2.6, 6.1, and 1.8 respectively compared to `AS-Rank`.

`ProbLink` is not dependent on the initial labels provided by the bootstrapping algorithm. Bootstrapping with `CoreToLeaf` and `AS-Rank` only results in 0.15% overall accuracy difference on average.

7.2 Feature Accuracy Breakdown

Eliminating any of the features used by `ProbLink` results in lower accuracy (data for this is omitted due to space constraints), but we now consider each feature in isolation. In order to show how effective each feature is, we run a version of `ProbLink` using each feature only, and calculate its error rates on the set of links inferred correctly/incorrectly by `ProbLink` with all features turned on against the 04/01/2017 snapshot of BGP paths (see Table 5).

The distance to Tier-1 feature is the most inaccurate feature, but if it is removed, the accuracy of `ProbLink` decreases. For the links which `ProbLink` infers incorrectly, all features have high error rates. This suggests that each feature adds some value, but none of them is perfect. So combining more features results in higher inference accuracy.

8 Practical Applications

8.1 Route Leak Detection

Route leaks are a class of common routing incidents which can cause large Internet service disruptions [41]. They are caused by violations of the policies among the ASes involved. For instance, on November 5, 2012, a Google peer Moratel (*AS*23947) improperly advertised Google routes to its provider, causing Moratel's providers to select the leaked

routes as the preferred ones destined to Google. As Moratel could not handle such large traffic volumes, Google's services went offline in parts of Asia for half an hour [21].

A conventional method for detecting route leaks is through checking valley-free violations in BGP paths. Mauch built a routing leak detection system based on this intuition by searching for valley paths containing three or more major networks with known relationships [24].

In the same spirit, we build a route leak detection system by detecting valley-free violations in paths based on the link relationship inference results of `ProbLink`. It is worth noting that a large fraction (more than 50%) of the valley-free violations do not result from route leaks but intended policies from ASes that are research/educational or IXP ASes [19]. Such ASes often establish a special type of AS relationship called indirect peering, where an AS functions as an intermediate link between two other ASes who wish to peer but through intermediate ASes in between. So we ignore a path if it contains research/educational or IXP ASes when detecting route leaks.

To evaluate the performance of `ProbLink` and other AS relationship inference techniques on route leak detection, we use only those links for which we have validation data in BGP paths. For example, suppose a path has link relationships "* * p2p * c2p", where * is unknown in the validation dataset. Even though a relationship inference algorithm should have predictions on the unknown links, we just detect route leaks by its predictions on the two known links in order to compare against the validation dataset. Figure 7 compares the precision and recall of `ProbLink`, `AS-Rank`, and `CoreToLeaf`, against the real route leaks implied by the validation dataset across 10 days in April 2016. The average precision for `ProbLink`, `AS-Rank`, and `CoreToLeaf` is 81.1%, 19.8%, and 8.1% respectively; and the average recall for `ProbLink`, `AS-Rank`, and `CoreToLeaf` is 76.2%, 22.1%, and 5.6% respectively.

8.2 Inference of Complex Relationships

AS relationships may be more complex than the traditional p2c/p2p model. Such complex agreements may take the form of a *hybrid* relationship with different relationship type for different Points-of-Presence (PoPs), or a *partial-transit* relationship, in which a provider offers transit only toward its peers and customers, but not its providers, or restricts transit

Feature	Links inferred correctly			Links inferred incorrectly			All links
	# correct	# links	Error rate	# correct	# links	Error rate	Error rate
Triplet	44,853	45,639	1.7%	222	459	51.6%	2.2%
Non-path	21,382	21,383	0.01%	169	242	30.2%	0.3%
Distance to Clique	42,432	45,639	7.0%	270	459	41.2%	7.4%
VP	45,346	45,639	0.6%	108	459	76.5%	1.4%

Table 5: Accuracy of each feature on the links inferred correctly/incorrectly by `ProbLink`. The Co-located IXP/facility features are omitted because they are used to distinguish transit and p2p relationships, and attribute equal probabilities to p2c and c2p links.
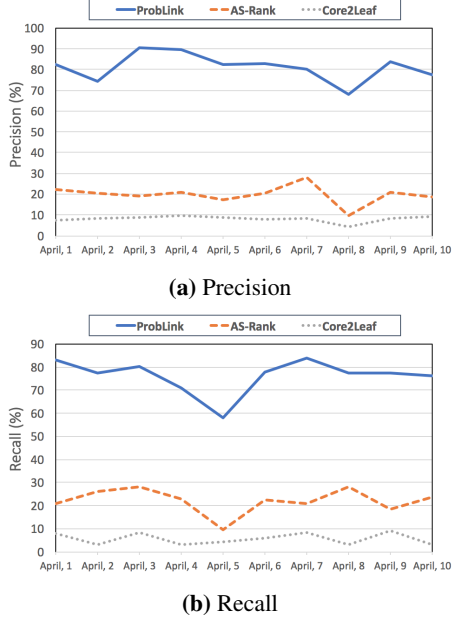
(a) Precision

(b) Recall

Figure 7: Evaluation of route leak detection across 10 days.

in a specific geographic region [11, 14]. The state-of-the-art algorithm for inferring complex relationships (`CR` algorithm) takes as input a set of conventional relationships and iteratively refines them by combining active traceroute measurements with geolocation data to discover the PoP-level propagation patterns of inter-domain paths [18]. Due to the high measurement cost in terms of traceroute queries required to infer complex relationships, `CR` utilizes customer cones to optimize the allocation of queries to traceroute probes and maximize the discovery of hybrid relationships within limited querying budgets used by platforms such as RIPE Atlas [40]. Therefore, the quality of the p2p/p2c relationships can affect the precision, accuracy and coverage of complex relationship inference.
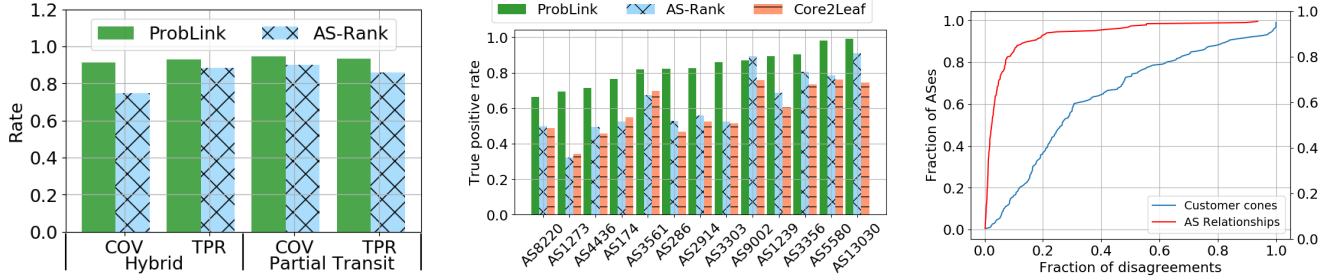
To test the performance of `ProbLink` for complex relationship inference, we implemented the `CR` algorithm and executed two 2-day measurement campaigns over the RIPE Atlas platform, on 2018/09/06 and 2018/09/08 using `ProbLink` and `AS-Rank` respectively. For each measurement round we allocated the maximum permissible number of measurement credits, which resulted in $125,529$ traceroute queries from $7,870$ Atlas probes. `CR+ProbLink` inferred $1,308$ hybrid relationships and $3,163$ partial transit links, while `CR+AS-Rank` inferred $1,029$ hybrid relationships and $3,009$

partial transit links. We evaluated these inferences against our validation dataset, which includes 346 hybrid links and 402 partial-transit links. As shown in Fig. 8a, combining `CR` with `ProbLink` not only improves the True Positive Rate (TPR) of the algorithm both for hybrid and partial transit relationships, but importantly, we significantly expand its coverage (COV) by capturing 91% of the hybrid and 95% of the partial transit relationships, compared to 76% and 90% respectively for `CR+AS-Rank`. Overall, `CR+ProbLink` discovers 27% more hybrid relationships than `CR+AS-Rank`.

8.3 Predicting the Impact of Selective Advertisements

The ability to predict the impact of traffic engineering policies on the active BGP paths can be valuable to network operators, as it would limit the need for trial-and-error experimentation, allow the configuration of more predictable and stable routing policies, and minimize the risk of propagating unintended routes [34]. However, past works have shown that the existing AS relationship datasets have poor predictive capabilities, making them impractical for such purposes. In this section, we evaluate the impact of `ProbLink`'s improved relationship inferences in predicting the outcome of a selective advertisement. Selective advertisement is a popular traffic engineering technique used by AS operators to achieve traffic load balancing, by advertising certain routes only to a subset of their inter-domain neighbors [36].

To predict the impact of selective advertisement "in the wild", we first need to explicitly capture the activation and the scope of such policies. We detect selectively advertised prefixes by utilizing route redistribution BGP Communities, which are increasingly utilized to implement selective prefix advertisement [37, 13]. In particular, many providers define an array of Community values that can be set by their customers, to allow them to control whether the provider should propagate or not a route to a specific peer or group of peers. For instance, if *AS*9002 (RETN) receives a prefix advertisement from a customer annotated with the BGP Community `9002:65535`, then RETN will propagate this route only to its customers, but not its peers or providers [38]. Redistribution Communities can further limit the scope of the prefix advertisement by determining a location for which the redistribution policy will be applied. For instance, when the Community `286:49` is applied on a prefix, *AS*286 (KPN) will not advertise this prefix to its US peers [29]. By parsing WHOIS records and NOC websites, we compile a dictionary of Community values that define one of the following types

(a) Comparison of coverage (COV) and True Positive Rate (TPR) between `ProbLink` and `AS-Rank` in complex relationship inference.

(b) True positive rate per relationship inference algorithm for the prediction of the impact of selective advertisement policies.

(c) Disagreements in AS relationships and ASes in the customer cone of the 200 ASes with the largest customer cones.

Figure 8: Evaluation of complex relationships inference and the prediction of path changes due to selective prefix advertisements.

of selective route redistribution:

- Do not announce route to neighbors of type *R*.
- Do not announce route to neighbors of type *R* at *L*.
- Announce route only to neighbors of type *R*.
- Announce route only to neighbors of type *R* at *L*.

R indicates relationship type (customer, provider, peer) and *L* indicates a city-level or country-level location identifier. In total, we extracted 644 Community values from 152 ASes.

After compiling our Communities dictionary, we monitor the BGP messages of the corresponding ASes to capture BGP Updates annotated with one of the redistribution Communities. Lets assume we observe a BGP Update for a destination prefix d annotated with a BGP Community C, which instructs AS_C to propagate p *only* to its neighbors of type *R*. We calculate which ASes will have to change their paths as follows: We first parse the BGP paths right before C was applied on the prefix d, and we collect all the paths P_{ALL} to d that traversed AS_C. Then, based on the inferred relationships, we find the paths $P_{R'} \subseteq P_{ALL}$ in which AS_C advertises the route toward d to a neighbor with relationship type $R' \neq R$. Since the Community C allows the prefix announcement only to neighbors with relationship type R, we infer that the paths $P_{R'}$ will be withdrawn, and the corresponding ASes in these paths will choose a different path. When C also defines a geographic scope for the prefix advertisement in addition to the relationship type, we use the techniques described in [18] to map the city-level location of AS interconnections, and calculate the affected paths in a similar manner. We validate our inferences by observing the withdrawn paths after C was applied on the path. We consider as false positive any AS $a \in P_{R'}$ that did not withdraw its path 15 minutes after we observed the BGP Update with C. We do not consider false negatives, as an AS may change its path to d for different reasons, and this change may simply coincide with the application of the Community C on the same prefix.

Figure 8b shows our validation results after executing the above experiment for the first week of April 2016. During that period we found 480 prefixes tagged with redistribution communities defined by 13 ASes. Overall, 83%

of `ProbLink`'s predictions were correct, compared to 62% for `AS-Rank` and 59% for `CoreToLeaf`. `ProbLink` outperformed `AS-Rank` for every AS except $AS9002$, and in some cases (e.g. $AS1273$) the true positive rate was 2x higher compared to the other algorithms. These results are surprising given that less than 4% of the relationship inferences differ between `ProbLink` and `AS-Rank`. To understand the significant improvement achieved by `ProbLink`, we investigate the impact of the relationship disagreements between the two algorithms on the customer cones obtained using the *Provider/Peer Observed* methodology proposed in [30]. We focus on the ASes with at least 100 ASes in their downstream path. For each of these ASes we calculate the fraction of their relationships and the fraction of their customer cones that disagree between `ProbLink` and `AS-Rank`. As shown in Figure 8c, while less than 10% of the ASes had more than 20% relationship mismatches, over 60% of the ASes had at least 20% difference in their customer cones. This finding highlights the fact that even a few incorrect relationship inferences can lead to significant differences in properties of the resulting downstream paths and substantial deviations in the predictive capabilities of `ProbLink` and `AS-Rank`.

9   Conclusion

We revisit the AS relationship problem and inference techniques. We first develop a simple inference algorithm that achieves accuracy comparable to that of the state-of-the-art inference technique, `AS-Rank`, indicating that the types of most links in validation datasets are easily inferred. We then construct different subsets of the validation dataset that might be considered *hard* and use these as benchmarks for evaluating improvements in AS relationship inference. Further, we observe that many of the features that can be used by inference techniques are of a stochastic nature, so we present a probabilistic AS relationship inference algorithm that provides a framework for easy integration of many noisy but useful attributes into the relationship inference algorithm. We show that this probabilistic algorithm is more accurate and less sensitive to the locations of vantage points and BGP paths compared to the state-of-the-art algorithms.

References

[1] IXP Service Matrix. https://www.euro-ix.net/en/tools/ixp-service-matrix/.

[2] PeeringDB. https://www.peeringdb.com/.

[3] RIPE (RIS). http://www.ripe.net/ris/.

[4] U. Oregon Route Views Project. http://www.routeviews.org/.

[5] R. Anwar, H. Niaz, D. Choffnes, Í. Cunha, P. Gill, and E. Katz-Bassett. Investigating interdomain routing policies in the wild. In *Proceedings of the 2015 Internet Measurement Conference*, pages 71–77. ACM, 2015.

[6] R. Bush, O. Maennel, M. Roughan, and S. Uhlig. Internet optometry: assessing the broken glasses in internet reachability. IMC '09.

[7] M. Caesar and J. Rexford. Bgp routing policies in isp networks. *IEEE network*, 19(6):5–11, 2005.

[8] CAIDA. Inferred AS to organization mapping dataset. http://www.caida.org/data/as-organizations/.

[9] A. Dhamdhere, D. D. Clark, A. Gamero-Garrido, M. Luckie, R. K. P. Mok, G. Akiwate, K. Gogia, V. Bajpai, A. C. Snoeren, and K. Claffy. Inferring persistent interdomain congestion. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, SIGCOMM '18, pages 1–15, New York, NY, USA, 2018. ACM.

[10] G. Di Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between autonomous systems. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 1, pages 156–165. IEEE, 2003.

[11] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, G. Riley, et al. AS Relationships: Inference and Validation. SIGCOMM '07, 2007.

[12] X. Dimitropoulos, D. Krioukov, B. Huffaker, G. Riley, et al. Inferring as relationships: Dead end or lively beginning? In *International Workshop on Experimental and Efficient Algorithms*, pages 113–125. Springer, 2005.

[13] B. Donnet and O. Bonaventure. On bgp communities. *ACM SIGCOMM Computer Communication Review*, 38(2):55–59, 2008.

[14] P. Faratin, D. D. Clark, S. Bauer, W. Lehr, P. W. Gilmore, and A. Berger. The growing complexity of internet interconnection. *Communications & Strategies*, (72):51, 2008.

[15] J. H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.

[16] L. Gao. On inferring autonomous system relationships in the internet. IEEE Press, 2001.

[17] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening internet topology: Natural evolution, unsightly barnacles or contrived collapse? In *International Conference on Passive and Active Network Measurement*, pages 1–10. Springer, 2008.

[18] V. Giotsas, M. Luckie, B. Huffaker, et al. Inferring complex as relationships. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 23–30. ACM, 2014.

[19] V. Giotsas and S. Zhou. Valley-free violation in internet routinganalysis based on bgp community data. In *Communications (ICC), 2012 IEEE International Conference on*, pages 1193–1197. IEEE, 2012.

[20] V. Giotsas, S. Zhou, M. Luckie, et al. Inferring multilateral peering. In *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*, pages 247–258. ACM, 2013.

[21] S. Goldberg. Why is it taking so long to secure internet routing? *Communications of the ACM*, 57(10):56–63, 2014.

[22] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. Lord of the links: a framework for discovering missing links in the internet topology. *IEEE/ACM Transactions on Networking (ToN)*, 17(2):391–404, 2009.

[23] G. Inc. Google peering policy. https://peering.google.com.

[24] M. Jared. BGP Routing Leak Detection System. https://puck.nether.net/bgp/leakinfo.cgi.

[25] E. Jasinska, N. Hilliard, R. Raszuk, and B. N. RFC 7947: Internet Exchange BGP Route Server. https://tools.ietf.org/html/rfc7947.

[26] M. S. Kang and V. D. Gligor. Routing bottlenecks in the internet: Causes, exploits, and countermeasures. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 321–333. ACM, 2014.

[27] E. Katz-Bassett, D. R. Choffnes, Í. Cunha, C. Scott, T. Anderson, and A. Krishnamurthy. Machiavellian routing: improving internet availability with bgp poisoning. In *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, page 11. ACM, 2011.

[28] M. Konte, R. Perdisci, and N. Feamster. Aswatch: An as reputation system to expose bulletproof hosting ases. In *ACM SIGCOMM Computer Communication Review*, volume 45, pages 625–638. ACM, 2015.

[29] KPN NOC. BGP Communities For AS286. https://as286.net/AS286-communities.html, September 2018.

[30] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, et al. AS Relationships, Customer Cones, and Validation. IMC '13.

[31] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008.

[32] R. Mazloum, M.-O. Buob, J. Auge, B. Baynat, D. Rossi, and T. Friedman. Violation of interdomain routing assumptions. In *International Conference on Passive and Active Network Measurement*, pages 173–182. Springer, 2014.

[33] A. Mitseva, A. Panchenko, and T. Engel. The state of affairs in bgp security: A survey of attacks and defenses. *Computer Communications*, 124:45 – 60, 2018.

[34] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig. Building an as-topology model that captures route diversity. SIGCOMM '06, pages 195–206, 2006.

[35] R. Nithyanand, O. Starov, P. Gill, A. Zair, and M. Schapira. Measuring and mitigating as-level adversaries against tor. In *23rd Annual Network and Distributed System Security Symposium, NDSS 2016, San Diego, California, USA, February 21-24, 2016*, 2016.

[36] B. Quoitin, C. Pelsser, L. Swinnen, O. Bonaventure, and S. Uhlig. Interdomain traffic engineering with bgp. *IEEE Communications Magazine*, 41(5):122–128, May 2003.

[37] B. Quoitin, S. Tandel, S. Uhlig, and O. Bonaventure. Interdomain traffic engineering with redistribution communities. *Computer Communications*, 27(4):355–363, 2004.

[38] RETN NOC. BGP Communities For AS9002. http://retn.net/support/bgp-communities/, September 2018.

[39] RFC. RFC 6996: Autonomous System (AS) Reservation for Private Use. https://tools.ietf.org/html/rfc6996.

[40] RIPE Atlas. User-Defined Measurements - Rate Limits. https://atlas.ripe.net/docs/udm/#rate-limits.

[41] K. Sriram, D. Montgomery, D. McPherson, E. Osterweil, and B. Dickson. Problem Definition and Classification of BGP Route Leaks. https://tools.ietf.org/html/rfc7908.

[42] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the internet hierarchy from multiple vantage points. IEEE, 2002.

[43] A. Technologies. Akamai Network partnerships. https://www.akamai.com/us/en/products/network-operator/akamai-network-partnerships.jsp.

[44] J. S. Varghese and L. Ruan. Computing customer cones of peering networks. In *Proceedings of the 2016 Applied Networking Research Workshop*, pages 35–37. ACM, 2016.

[45] Wikipedia. List of Tier-2 ases. https://en.wikipedia.org/wiki/Tier_2_network.

[46] W. Willinger and M. Roughan. Internet topology research redux. *ACM SIGCOMM eBook: Recent Advances in Networking*, 2013.

[47] J. Xia and L. Gao. On the evaluation of as relationship inferences [internet reachability/traffic flow applications]. In *Global Telecommunications Conference, 2004. GLOBECOM'04. IEEE*, volume 3, pages 1373–1377. IEEE, 2004.

A AS-Rank Algorithm

The `AS-Rank` algorithm takes the following 11 steps to infer the relationship of each link.

1. Discard or sanitize paths with artifacts.
2. Sort ASes in decreasing order of computed transit degree, then node degree.
3. Infer a transit-free clique (i.e., Tier-1) ASes at top of AS hierarchy and label the links between every pair of ASes in the clique as p2p links.
4. Discard poisoned paths.
5. Visit ASes in order of the ranking in (2), and label a link as c2p if its previous link in a BGP path is composed of two clique members, or if its previous link in a BGP path is already labeled as c2p.
6. Infer c2p relationships from VPs inferred to be announcing no provider routes.

7. Infer c2p relationships for ASes where the customer has a larger transit degree.
8. Infer customers for ASes with no providers.
9. Infer c2p relationships between stub ASes and clique ASes.
10. Infer c2p relationships where adjacent links have no relationship inferred.
11. Infer all other links left as p2p.

B AS-Rank Clique Inference

The clique inference algorithm in `AS-Rank` works as the following:

1. Find the top 10 ASes by transit degree.
2. If there are three consecutive members (X-Y-Z) in the top 10 ASes showing up in paths, and there are more than 5 ASes downstream from X Y Z (to make sure that the paths containing three consecutive members are not poisoned), disconnect the edge between X and Z even though X and Z are connected in some paths.
3. Find the largest clique in terms of transit degree sum among the top 10 ASes, denoted as C.
4. Visit the rest ASes top to down by transit degree, add an AS Z to C if Z has links with all members in C.
5. Similar to Step 2: If there are three consecutive members (X-Y-Z) in C showing up in paths, and there are more than 5 ASes downstream from X Y Z, disconnect the edge between X and Z.
6. Find the largest clique in C in terms of transit degree sum as the final inferred clique.

C Feature Importance Computation

Gradient boosting [15] is a widely used machine learning technique for solving classification problems. In gradient boosting, *GBDT* (Gradient boosting decision trees) produces a prediction model in the form of an ensemble of multiple decision trees. It is straightforward to retrieve importance scores for features when constructing *GBDT*. An importance score (F score) describes the number of times a feature is used to split the data across all trees. The more a feature is used to make key decisions with decision trees, the higher its importance score.

To decide what features can distinguish *hard* links from *easy* links in the Internet, we first split the validation dataset into two halves. The set of links which `CoreToLeaf` or `AS-Rank` infers incorrectly are labeled as "hard", while those which are inferred correctly are labeled with "easy". Then, we feed the features listed in Table 6 of links along with their labels into the *GBDT* and calculate the importance score corresponding to each feature.

Figure 9 plots the importance score of each feature divided by the sum of all features' scores. We can tell the features f1, f4, f7, f8, f9 are the most important ones, so we translate them into the various categories of features to characterize "hard" links in §4.3.
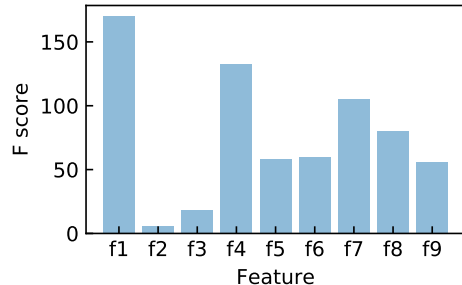

Figure 9: Feature importance scores provided by gradient boosting tree.

Feature label	Meaning
f1	Number of VPs which observe a link
f2	Max distance to Tier-1
f3	Min distance to Tier-1
f4	Max node degree
f5	Min node degree
f6	Node degree difference
f7	Max transit degree
f8	Min transit degree
f9	Transit degree difference

Table 6: Features fed into GBDT

D AS-Rank Sensitivity Analysis

D.1 Sensitivity to VP Selection

Each vantage point provides its own view of the Internet AS-level topology and the flow of traffic from the VP to rest of the Internet. VPs are located in different places, belong to different tiers, and they themselves have different import and export policies.

Even though the number of VPs have been growing over time, VPs are free to join or leave the set of public collectors, so the selection of VPs we have access to is arbitrary, biased, and under flux. A good AS relationship inference algorithm should not be sensitive to the selection of these VPs.
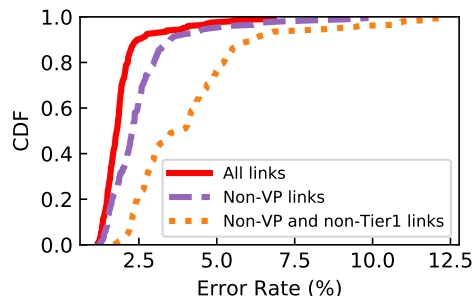

Figure 10: CDF of `AS-Rank`'s error rates on paths seen from 200 different half-VP sets.

We run the `AS-Rank` algorithm repeatedly, 200 times, against the 04/01/2017 BGP snapshot. For each of these 200 executions, we choose a random VP subset consisting

of half of all available VPs (which we denote as V below) and give as input to the AS-Rank algorithm only the BGP paths visible to the VPs in that subset. Figure 10 plots CDFs of AS-Rank's error rates using paths from these VP subsets. In the plot, we examine all links in the ground-truth dataset, links in the ground-truth dataset except links that directly connect with V (i.e., non-VP links), and the links in the ground-truth dataset except V's links as well as Tier-1 links (i.e., non-VP and non-Tier1 links). The inference error rates on overall links range from 1.2% to 6.9%, and the error rates on *non-VP and non-Tier1* links range from 1.8% to 12.3%. AS-Rank's accuracy is thus quite sensitive to the VP selections, especially for links which are relatively difficult to infer (i.e., not VP or Tier1 links).

D.2 Root Cause Analysis of AS-Rank Sensitivity

To illustrate the issue described in §5.3, let's consider two VP sets, V1 and V2, drawn from our 200 executions. AS-Rank's inference accuracy from V1 is low, while its inference accuracy from V2 is high. The largest ten ASes differ for different sets of VPs because the transit degrees of ASes are determined by paths observed by the VPs. For example, the 9th largest AS (AS2914) observed by V2 is the 12th largest AS observed by V1, so AS2914, which is a real Tier-1 AS, shows up in the clique chosen from the top 10 ASes using V2's paths, but it does not show up in the clique chosen using V1's paths.

For V1, the AS-Rank algorithm determines the maximum clique with the largest transit degree (from the top 10 ASes) to be "AS3356, AS6939, AS8220, AS9002, AS43531". AS43531 is not considered for V2 due to a relatively lower measurement of its transit degree, and it is not a higher-tier AS in reality. This affects the subsequent expansion of the clique, wherein ASes are considered in order by degree and added to the clique if they have connections to all members of the clique. So, in V1's execution, all added members are required to have a direct link with AS43531, and AS1764, AS8767, AS12389, AS12552, AS20485, AS25091, AS33891, AS43531, AS57724 are therefore all added to the clique, even though they are all low-tier ASes.

In a nutshell, the clique inference of AS-Rank algorithm is sensitive to the top 10 largest ASes ranked by transit degrees, which are determined by the selection of VPs and the selection of snapshots. Further, the clique membership determines the order in which links are analyzed by AS-Rank, impacts the computation of *customer cones* for each clique member (i.e., the set of ASes that a clique AS can reach using p2c links), and impacts the overall accuracy of the algorithm.