

On the Evaluation of AS Relationship Inferences

Jianhong Xia

Department of Electrical and Computer Engineering
University of Massachusetts
Amherst, MA 01002
jxia@ecs.umass.edu

Lixin Gao

Department of Electrical and Computer Engineering
University of Massachusetts
Amherst, MA 01002
lgao@ecs.umass.edu

Abstract—The commercial contractual relationship among Autonomous Systems (ASs) is important for understanding the reachability and traffic flow in the Internet. Several heuristic algorithms have been proposed for inferring AS relationships by extracting information from Border Gateway Protocol (BGP) routing tables. However, there is no evaluation of the AS relationships inferred from existing algorithms. In this paper, we present the methodology for obtaining partial AS relationships, and use them to evaluate these algorithms. Our result shows that although existing algorithms achieve good overall accuracy, the accuracy on the peer-to-peer relationships is quite poor. Furthermore, we propose an algorithm for inferring AS relationships by taking advantage of this partial information. We show that our proposed algorithm outperforms existing algorithms for inferring peer-to-peer relationships by 1.8 ~ 3.8 times in accuracy. In addition, the overall accuracy of our algorithm is higher than that of existing algorithms.

Keywords— Border Gateway Protocol, Routing Protocol, Routing Policy, Algorithm.

I. INTRODUCTION

The Internet consists of thousands of distinct regions of administrative domains, each of which possesses one or several Autonomous Systems (ASs). These ASs exchange information about how to reach individual blocks of destination IP addresses via Border Gateway Protocol (BGP) [1], [2], [3]. BGP constructs AS paths by successively propagating updates between pairs of BGP routers that establish BGP *peering sessions* [3]. It allows each AS to choose its own policy on selecting the best routes, announcing and accepting routes.

One of the most important factors in determining routing policy is the commercial contractual relationship among administrative domains. As mentioned in [4], [5], the Internet topology alone does not imply reachability among ASs. The commercial contractual relationship among ASs is very important for us to understand the reachability and traffic flow in the Internet. Although such information is important for Internet research, not all ISPs are willing to publish it. In fact, an ISP might keep it as a commercial secret.

Several algorithms [4], [5], [6] have been proposed for inferring AS relationships in the recent years. The accurate information on AS relationships will help us to understand the Internet hierarchy and routing policies in the AS-level topology [7], [5]. Although these algorithms have been proposed for several years, there is no evaluation of these algorithms in terms of the accuracy of AS relationships. In this paper,

we present several techniques for obtaining partial AS relationships, and use this partial information to evaluate the AS relationships inferred from the existing algorithms in [4], [5]. From the AS relationship inferences on 07/10/2003, we find that although both algorithms in [4], [5] have good overall accuracy, the accuracy on the peer-to-peer relationships is quite poor. The algorithm in [4] achieves 49.08% accuracy on peer-to-peer relationships, while the algorithm in [5] achieves only 24.63% accuracy. Therefore, we propose a new algorithm for inferring AS relationships by taking advantage of the partial information. Our result shows that from only a small set of partial information, our algorithm achieves 96.37% overall accuracy, and 91.45% on peer-to-peer relationships, which is 1.8 ~ 3.8 times as much as the accuracy of existing algorithms.

The remainder of this paper is structured as follows. Section II introduces the background materials and reviews some related work. Section III describes the methodology for obtaining partial AS relationships. Section IV evaluates and discusses the accuracy of AS relationship inferences of existing algorithms. We propose a new algorithm for inferring AS relationships in Section V. Finally, we summarize the paper and discuss some challenges in Section VI.

II. BACKGROUND AND RELATED WORK

A. BGP Attributes

BGP is an interdomain routing protocol [8] that is used for exchanging routing information between ISPs. The information exchanged among BGP routers, called update message, carries one or more BGP attributes. Among these attributes, *BGP community* helps to achieve scalability and reduce management complexity in BGP configuration [9], [10]. An ISP can define the semantics on the usages of *BGP community* attribute for different routing policies. The usages of *BGP community* attribute has become more and more popular in the current Internet [11].

B. Routing Policy and Internet Routing Registry

BGP allows for policy-based routing, which means that BGP router will propagate update messages to its BGP peering routers according to its own policy. For simplicity, we put BGP routing policy into two categories, import policy and export policy. Upon receiving an update message, a router must decide whether or not to accept this path according to

its import policy. Once a router accepts paths from different neighbors, it must select one path among these candidates as the best one. Then, based on its export policy, the router decides whether or not to propagate this best path to its neighbors. In order to help maintaining the infrastructure of Internet routing and addressing, the Internet Routing Registries (IRR) [12] database is used to store the routing policies for each ISP. It uses Routing Policy Specification Language (RPSL) [13], [14] to define the common format and express a wide range of routing policies.

C. Related Work

The first study on inferring AS relationships is proposed by Gao in [4]. The paper summarizes the *valley-free* property of AS paths in BGP routing tables. Formally, an AS path is *valley-free* if and only if a provider-to-customer edge is followed by exclusively provider-to-customer or sibling-to-sibling edges, and a peer-to-peer edge is followed by exclusively provider-to-customer or sibling-to-sibling edges. The author proposes a heuristic algorithm for inferring AS relationships by extracting AS path information from public routing tables on Route Views routers [15]. Subramanian et al. [5] formulate AS relationship assignment as an optimization problem, *Type of Relationship (ToR)* problem. The authors present a heuristic algorithm for inferring the AS relationships by combining AS paths from multiple vantage points in the Internet. Battista et al. [6] analyze the complexity of *ToR* problem and present an approximately optimal solution to the *ToR* problem. However, it is not clear whether the optimal solution of *ToR* problem could certainly lead to an accurate inference of AS relationships in practice.

Although these algorithms have been proposed for several years, no study has been done on evaluating the performance of these algorithms. Given the facts on the incompleteness of AS-level topology from the limited sources [16], [17] and BGP misconfigurations [18] in the Internet, it becomes more and more necessary and interesting to evaluate the accuracy on the AS relationship inferences. Without any commercial agreements among ISPs, it is hard for us to evaluate the accuracy of these algorithms. Fortunately, the studies in [11], [19] suggest that some usages of *BGP community* indeed indicate the relationships among ASs. In this paper, we first introduce techniques for obtaining partial AS relationships and then use them to evaluate the inferences of existing algorithms. Furthermore, we propose a new algorithm for inferring AS relationships which outperforms existing algorithms.

III. METHODOLOGY FOR OBTAINING PARTIAL AS RELATIONSHIPS

One straight forward approach to evaluating the AS relationship inferences is to compare them with the commercial agreements among ISPs. It is infeasible to obtain all these commercial agreements in the Internet. However, we can at least obtain partial AS relationships of some ISPs from the usages of *BGP community* attribute and IRR databases. In this section, we introduce three different approaches to obtaining

```
TIME: 07/02/03 11:04:07
TYPE: BGP4MP/MESSAGE/Update
FROM: 64.200.199.3 AS7911
TO: 128.223.60.102 AS6447
ORIGIN: IGP
ASPATH: 7911 3561 1239
NEXT_HOP: 64.200.199.3
COMMUNITY: 3561:21000 7911:999 7911:7302
ANNOUNCE
141.238.0.0/16
```

Fig. 1. An Example of Update Message

this information. One is from the usages of *BGP community* attribute, one from the instances of AS-SET object in the IRR databases, and another is from the explicit description of routing policies in the IRR databases.

A. From the Usages of BGP Community Attribute

In this subsection, we first review the common usages of *BGP community* attribute in practice, and then analyze the persistency on its usages in the Internet. Finally, we describe how to obtain partial AS relationships from the usages of *BGP community* attribute.

1) *Usages of BGP Community Attribute*: *BGP community* attribute consists of a set of four-octet values. Each community value is denoted as two parts *xxxx:yyyy*. From previous studies [11], [19], the community value often implies the type of peer, geographic location or interconnection point in practice. So the usages of *BGP community* attribute make it possible to indicate AS relationships. For example, in the RIPE's IRR database archived on 07/14/2003, AS6320 uses the community "6320:21xxx" to denote its peers, use "6320:22xxx" to denote its providers and "6320:23xxx" to denote its customers.

Before we use *BGP community* to obtain AS relationships, we look at how persistent the usages of *BGP community* attribute are in the Internet and investigate whether this attribute is a good indication of the AS relationship.

2) *Persistency on the Usages of BGP Community Attribute*: We collect the BGP update messages from Route Views routers [15] and RIPE rrc00 [20] from 06/01/2003 to 07/10/2003. Fig. 1 illustrates the format of an update message translated by MRTd tools [21]. We parse all update messages in our data set and count the frequency of the usages of each community value. We define a tuple, $\{A, B, Next.Hop\}$, as a unit to maintain its community values between AS *A* and AS *B* that are observed from *Next.Hop*. For example, we can derive two tuples from the update message in Fig. 1. One is $\{AS7911, AS3561, 64.200.199.3\}$ with community value "7911:999 7911:7302", and the other is $\{AS3561, AS1239, 64.200.199.3\}$ with community value "3561:21000".

We use N_t to denote the total number of tuples that use one or more community values over the observation period, and N_p to denote the total number of tuples that use only one community value over the observation period. Table I shows the result of persistency on the usages of *BGP community* attribute. We see that almost 90% of tuples kept using one community value during the observation time.

TABLE I
PERSISTENCY ON THE USAGES OF BGP COMMUNITY ATTRIBUTE

Data Source	N_t	N_p
RIPE rrc00	5139	4593 (89.38%)
Route Views routers	31536	28256 (89.60%)
Overall	36675	32849 (89.57%)

From these observations, we argue that the usages of *BGP community* attribute are very persistent over a long period of time. It should be reliable to map AS relationships from the usages of *BGP community* attribute.

3) *Mapping AS Relationships from BGP Community Values*: Mapping AS relationships from *BGP community* values is not trivial because there is no standard syntax on the description of *BGP community* in IRR databases. We collect the semantics of *BGP community* attribute for each AS from the IRR databases, and put them in a simple database. In addition, we use additional public information [22], [23], [24], [25] to expand the knowledge of our database. Then, we parse the BGP update messages to get the usages of *BGP community* attribute. Finally, we query our database and map the *BGP community* value of AS pair into its relationship.

B. From AS-SET Object in the IRR Databases

In the IRR databases, many ISPs define their customers, providers and peers as different instances of AS-SET object. We parse the IRR databases, and find out the instances which contain the strings of “customer”, “provider” or “peer” in their AS-SET object. Then we manually check these usages and obtain partial AS relationships.

C. From Routing Policies in the IRR Databases

The description on routing policies in the IRR databases sometimes also imply the AS relationships. Similar to the method in AS-SET object, we extract AS relationships only when their relationships are explicitly described in their routing policies such as “Import from downstream transit customers”, etc.

D. Results of Partial AS Relationships

Based on the techniques we have presented in this section, we process the data on the RADB and RIPE’s IRR databases published on 07/14/2003 and BGP updates on Route Views routers [15] and RIPE rrc00 [20] from 06/01/2003 to 07/10/2003. To avoid using the obsolete information, we only analyze the records that were updated in 2003. We obtain the relationships for 4914 distinct AS pairs (2419 AS pairs from the usages of BGP Community, and 2916 AS pairs from the usages of AS-SET object and routing policies). 28 of them do not have a persistent AS relationship when we combine the results from these three approaches. We ignore these few special cases from our study. Among the remaining 4886 AS pairs, 3717 of them are provider-customer relationships and 1169 of them are peer-to-peer relationships. Since our partial information covers 3616 different ASs, we believe that it is a good sample to evaluate the AS relationship inferences from the existing algorithms.

IV. EVALUATION ON THE AS RELATIONSHIP INFERENCES

In this section, we use the partial AS relationships to evaluate the inference of the existing algorithms. We refer to the algorithm in [4] proposed by Gao as LG and the algorithm in [5] proposed by Subramanian et al. as SARK. We run the LG algorithm to obtain the inferences on the routing table from Route Views routers. The inferences of SARK algorithm are available online at [26]. We compare our partial information with the inferences of LG and SARK algorithms from the data on 07/10/2003.

A. Accuracy on the Type of Relationships

We calculate the accuracy based on the overlapping AS pairs appearing in both our partial AS relationships and the inferences of each algorithm. There are 2802 overlapping AS pairs in the inferences of the LG algorithm and 2821 overlapping AS pairs in the inferences of the SARK algorithm. We use these overlapping AS pairs as the samples to evaluate the accuracy of these two algorithms.

Table II shows the accuracy of inferences of the LG and SARK algorithms. We see that the accuracy of the LG algorithm is higher than that of the SARK algorithm on both provider-customer and peer-to-peer relationship inferences. Although the accuracy of both algorithms on provider-customer inferences is high, the accuracy on peer-to-peer inferences is very poor. On peer-to-peer relationship inferences, the LG algorithm achieves 49.08% accuracy and the SARK algorithm achieves only 24.63% accuracy.

B. Accuracy on the Degree of AS Pairs

In order to further investigate what AS pairs have incorrect inferences from the existing algorithms, we look at the distribution of the accuracy on the degree of AS pairs.

We define $\alpha(d)$ to be the accuracy of the inferences on the AS pairs in which both degrees are greater than or equal to d . From the definition, we know that $\alpha(1)$ denotes the overall accuracy of the inferences on all AS pairs.

Fig. 2(a) shows the relation between the accuracy of the inferences and the degree of AS pairs for both of the LG and SARK algorithms. Although the overall accuracy of the LG and SARK algorithms is high, the accuracy for the AS pairs with medium or large degrees is not good. For example, $\alpha(100)$ is 67% for LG algorithm and only 58% for SARK algorithm. If we only look at the accuracy on peer-to-peer relationships in Fig. 2(b), we find that neither of them can achieve a good accuracy for the AS pairs with small degrees.

C. Discussion on the Evaluation

From our evaluation results, we observe two limitations of the LG and SARK algorithms on the accuracy of the AS relationship inferences.

The first limitation of the LG and SARK algorithms is that both of them have a low accuracy on peer-to-peer relationship inferences. One interesting observation is that both algorithms achieve better accuracy on inferring provider-customer relationships than on inferring peer-to-peer relationships. The

TABLE II
ACCURACY OF AS RELATIONSHIP INFERENCES ON 07/10/2003

Inferring Algorithm	LG algorithm	SARK algorithm	PTE algorithm	
			Evaluated by 100% partial information	Evaluated by 80% partial information
Samples of AS pairs	2802	2821	2818	2254
Accuracy of Provider-Customer Relationships	99.13%	98.35%	96.92%	96.14%
Accuracy of Peer-to-Peer Relationships	49.08%	24.63%	91.45%	89.33%
Overall accuracy	94.25%	91.24%	96.37%	95.46%




Fig. 2. Accuracy on the Degree of AS Pairs

“PTE (100%)” denotes the accuracy evaluated by 100% partial information.
“PTE (80%)” denotes the accuracy evaluated by 80% partial information.

reason partly comes from the power-law distribution of AS degrees in the Internet topology [27]. More than two-thirds of ASs has only one or two degrees located at the edge of the Internet, most of which are the customers of other ASs. Therefore, in the graph of AS level topology, we can easily infer the edge ASs (those ASs with one or two degrees) to be the customers of their neighbor ASs. On the other hand, it is not obvious to determine the relationships among the ASs which are other than the edge ASs.

The second limitation is that neither the LG algorithm nor the SARK algorithm can achieve a consistent performance on the accuracy over the different degrees of AS pairs. Both algorithms have similar accuracy distribution on the degree of AS pairs shown in Fig. 2. For the large ISPs in the core of the Internet, both algorithms can correctly infer them as peer-to-peer relationships. So the accuracy for large degree AS pairs is high. For the small ISPs in the edge of the Internet, although both algorithms achieve good performance on overall accuracy, they are not consistent on the degree of AS pairs. Furthermore, the accuracy on peer-to-peer relationships is poor.

The LG and SARK algorithms do not perform well on inferring relationships for all AS pairs. One of possible reasons is that some AS paths in the routing table do not conform to the *valley-free* property because of BGP misconfigurations [18] or special routing policies. Without any prior information, existing algorithms cannot differentiate the *valley-free* and *non-valley-free* paths before they start inferring. Although the chances of *non-valley-free* AS paths are small, they do affect the accuracy of inferences in the existing algorithms. Motivated by these reasons, we propose a new algorithm in next section so that we can make use of our partial information to filter out these *non-valley-free* AS paths.

V. INFERRING AS RELATIONSHIPS FROM PARTIAL INFORMATION

Based on the discussion in section IV-C, we now propose a new algorithm for inferring AS relationships. The basic idea of our algorithm is to infer the entire AS relationships from partial information. We refer to our algorithm as PTE (Partialness To Entireness).

A. Algorithm for Inferring AS Relationships

Our algorithm consists of two major components. One is to filter *non-valley-free* paths, and the other is to infer AS relationships from partial information. We also have several techniques to improving our algorithm. Due to the limited space, we only describe major ideas in the algorithm. The detailed algorithm and analysis can be found in [28].

1) *Filtering Non-Valley-Free Paths*: Using our partial AS relationships, we examine the *valley-free* property for each AS path. We remove the AS paths which cannot be *valley-free* from our data because they are most likely caused by BGP misconfigurations or special routing policies.

2) *Inferring AS Relationships*: Based on *valley-free* property of AS paths, we define three inference rules on an AS path (a, b, c, d, e, f) and one refreshing rule as follows:

- *Inference Rule 1*: Provider-to-customer edge is followed by provider-to-customer edge. For example, if the edge (c, d) is a provider-to-customer edge, we mark (d, e) and (e, f) as provider-to-customer edges.
- *Inference Rule 2*: Customer-to-provider edge only follows customer-to-provider edges. For example, if the edge (c, d) is a customer-to-provider edge, we mark (b, a) and (c, b) into provider-to-customer edges.
- *Inference Rule 3*: The left part of peer-to-peer edge is customer-to-provider edge, and the right part of peer-to-peer edge is provider-to-customer edge. For example, if the edge (c, d) is a peer-to-peer edge, we mark (b, a) , (c, b) , (d, e) and (e, f) as provider-to-customer edges.
- *Refreshing Rule*: If both (u, v) and (v, u) are inferred to provider-to-customer edge, we mark (u, v) and (v, u) as sibling-to-sibling edge.

Note that if we start with our partial information and apply the above rules repeatedly, we can get more AS relationships identified. Once we cannot infer any AS relationships, we use the algorithm in [4] to infer the remaining relationships.

3) *Several Techniques for Improvement*: We have several techniques in our algorithm for further improvement. For example, we collect multiple routing tables to augment the AS-level connectivity and to get more AS paths. We filter the AS paths that are not frequently present in the collected routing tables so that we reduce the impact of BGP misconfigurations [18] on the accuracy of the inferences. In addition, we have a heuristic algorithm to eliminate all relationship cycles [4], [5], [6] in our inferences.

B. Data Sources and Inference Results

In our experiment, we download 12 routing tables from Route Views routers and 3 routing tables from RIPE rrc00 archived on 07/10/2003. Around 0.03% distinct AS paths have been filtered because of *non-valley-free* property.

We use partial AS relationships obtained in Section III-D and our PTE algorithm to infer the AS relationships on the above data. Among 4886 partial AS relationships, only 2819 AS pairs exist in our data set. Thus, in our experiment, we use these 2819 partial AS relationships as our initial information to infer the relationships on the entire collection of 35583 AS pairs. Our inference result is that 85.31% of them are provider-customer relationships, 14.38% of them are peer-to-peer relationships, and 0.31% of them are sibling-to-sibling relationships.

C. Evaluation on PTE Algorithm

In order to evaluate the PTE algorithm, we randomly split the partial AS relationships into two subsets of 20% and 80%. We take the subset of 20% partial AS relationships as initial information for the PTE algorithm and use the subset of 80% partial information or 100% partial information to evaluate the PTE algorithm.

1) *Accuracy on the Type of Relationships*: We run PTE algorithm for 10 times with different random seeds. We present the average accuracy of inferences in Table II. It shows that our algorithm achieves not only high accuracy on inferring provider-to-customer relationships, but also around 90% accuracy on inferring peer-to-peer relationships. Compared to the LG and SARK algorithms, our PTE algorithm improve the accuracy of peer-to-peer relationships significantly, which is 1.8 ~ 3.8 times as much as that of existing algorithms. On the overall accuracy, our PTE algorithm still outperforms the LG and SARK algorithms.

2) *Accuracy on the Degree of AS Pairs*: We also evaluate our PTE algorithm on the relation between the accuracy and the degree of AS pairs. The result in Fig. 2 shows that, no matter using 100% or 80% partial information to evaluate the PTE algorithm, the accuracy of PTE algorithm is consistently high on the different degree ranges of AS pairs.

From these experiments, we show that from only a small set of partial information, our PTE algorithm achieves high accuracy on both provider-customer and peer-to-peer relationship inferences. It also has a consistently high accuracy of the inferences on the degree of AS pairs.

VI. SUMMARY AND DISCUSSION

In this paper, we first introduce techniques on obtaining partial AS relationships, and then use this partial information to evaluate the accuracy of the existing algorithms. Then we propose a new algorithm on inferring AS relationships from partial information. The result shows that from only a small set of partial information, our algorithm is able to achieve better performance than existing algorithms.

With the rapid growth of the Internet in the recent years, inferring AS relationships becomes hard. First, the relationship

of two ISPs might not be consistent at different interconnect points [29]. For example, they might have peer-to-peer relationship in North American, but have provider-customer relationship in Europe. Second, without a standard syntax on describing the usage of BGP community and routing policies in IRR databases, it is hard to systematically abstract partial AS relationships from them.

ACKNOWLEDGMENT

This work is supported in part by NSF under Grant No. ANI-0208116 and Alfred Sloan Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation and Alfred Sloan Foundation. The authors also thank the anonymous reviewers for their helpful comments.

REFERENCES

- [1] Bassam Halabi. *Internet Routing Architectures*. Cisco Press, 1997.
- [2] John W. Stewart. *BGP4: Inter-Domain Routing in the Internet*. Addison-Wesley, 1999.
- [3] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). Request for Comments 1771, March 1995.
- [4] L. Gao. On Inferring Autonomous System Relationships in the Internet. In *IEEE/ACM Trans. Networking*, December 2001.
- [5] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet Hierarchy from Multiple Vantage Points. In *Proc. IEEE INFOCOM*, June 2002.
- [6] G. Di Battista, M. Patrignani, and M. Pizzonia. Computing the Types of the Relationships between Autonomous Systems. In *Proc. IEEE INFOCOM*, April 2003.
- [7] Z. Ge, D.R. Figueiredo, S. Jaiwal, and L.Gao. On the Hierarchical Structure of the Logical Internet Graph. ITCOM' 2001.
- [8] http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm.
- [9] R. Chandra, P. Traina, and T. Li. BGP Communities Attribute. Request for Comments 1997, August 1996.
- [10] E. Chen and T. Bates. An Application of the BGP Community Attribute in Multi-home Routing. Request for Comments 1998, August 1996.
- [11] B. Quoitin and O. Bonaventure. A Survey of the Utilization of the BGP Community Attribute, 2002.
- [12] Internet Routing Registry. <http://www.irr.net/docs/list.html>.
- [13] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg, and M. Terpstra. Routing Policy Specification Language (RPSL). Request for Comments 2622, June 1999.
- [14] D. Meyer, J. Schmitz, C. Orange, M. Prior, and C. Alaettinoglu. Using RPSL in Practice. Request for Comments 2650, August 1999.
- [15] Route Views Project. <http://www.antc.uoregon.edu/route-views/>.
- [16] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The Origin of Power Laws in Internet Topologies Revisited. In *Proc. IEEE INFOCOM*, June 2002.
- [17] J. Gast and P. Barford. Resource Deployment Based on Autonomous System Clustering. In *Proc. IEEE GLOBE INTERNET*, November 2002.
- [18] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP Misconfiguration. In *Proc. ACM SIGCOMM*, August 2002.
- [19] F. Wang and L. Gao. On Inferring and Characterizing Internet Routing Policies. In *Proc. ACM Sigcomm Internet Measurement Conference*, October 2003.
- [20] RIPE Routing Information Services. <http://www.ripe.net/ris/>.
- [21] Multi-Threaded Routing Toolkit. <http://www.merit.edu/~mrt/>.
- [22] Communities Used by EUnet. <http://www.jippii.net/communities.shtml>.
- [23] http://infopage.cary.cw.net/Routing_Registry/community_announce.htm.
- [24] http://cw-rr.cw.net/community_prepend.htm.
- [25] <https://robin.gblx.net/api/docs/customer-communities.html>.
- [26] <http://www.cs.berkeley.edu/~sagarwal/research/BGP-hierarchy/>.
- [27] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-Law Relationships of the Internet Topology. In *Proc. ACM SIGCOMM*, August 1999.
- [28] J. Xia and L. Gao. On the Evaluation of AS Relationship Inferences. *Technical Report*, March 2004.
- [29] Exchange Point Information. <http://www.ep.net/ep-main.html>.