A Methodology for Magnitude-based Inter-AS Distance Estimation

Hirochika ASAI† Hiroshi ESAKI†

† Graduate School of Information Science and Technology, The University of Tokyo

Abstract — The relationships among autonomous systems (ASes) can be categorized into three types: 1) *transit*, 2) *peering* and 3) *sibling*. Since customer ASes purchase the Internet access from provider ASes over transit links and any two ASes can exchange their traffic without any payment over peering links, transit links cost more than peering links from the economical point of view. We propose an inter-AS distance estimation method with eigenvalue analysis based on a traffic-flow/-volume assumption so that transit links are more distant than peering links. We evaluate this method with two topology datasets: 1) the topology generated by a topology generator *Inet-3.0* and the inferred Internet topology from the real Internet routing table. We show that the proposed method can estimate the inter-AS distance on the basis of types of AS relationships (i.e., transit is more distant and peering is closer).

Keywords: inter-AS distance, AS relationships, economics, eigenvalue analysis

1 Introduction

The Internet consists of thousands of autonomous systems (ASes) operated by different administrative domains such as Internet service providers (ISPs), companies and universities. The relationships between neighboring ASes can be categorized into three types [1]: 1) *transit*, 2) *peering* and 3) *sibling*. From the economical point of view, transit links cost more than peering links because customer ASes purchase the Internet access from provider ASes over transit links. Therefore, customer ASes are willing to reduce their traffic exchanged with their provider ASes. For the reduction of the traffic over transit links, a novel routing architecture on the basis of AS relationships is required.

In recent studies [2–5], topology-aware overlay routing architectures have been proposed to reduce the inter-AS traffic. We have used the types of AS relationships as routing metric by transforming them into cost in peer-to-peer content delivery networks to achieve efficient resource utilization; we have proposed a path control architecture which takes into account the types of AS relationships, and we have shown that the architecture can reduce inter-AS traffic, especially traffic through high-cost transit links [5]. However, most commercial ISPs cannot disclose the relationships among ASes. Hence, it is difficult to use inter-AS resources efficiently in the real Internet and the inter-AS distance estimation is required.

Gao [1] has proposed an algorithm to infer the relationships among ASes. The author has shown that the relationships can be inferred by comparing the number of neighbors between two neighboring ASes, analyzing the paths in Border Gateway Protocol (BGP) routing tables according to the valley-free path model. Battista et al. [6] have mapped the type-of-relationship problem into the weighted MAX2SAT problem to compute the orientation on transit. They have also referred to the WHOIS database [7] as the organization information database to infer the sibling relationship. However, these approaches do not estimate the relationships numerically and they require the AS paths retrieved from the routing table, and consequently, it is difficult to apply the relationships inferred by the algorithm as routing metric.

In this paper, we propose and evaluate an inter-AS distance estimation method based on a traffic-flow/-volume assumption by using the AS adjacency matrix, not the AS paths, so that transit links are more distant than peering links. The traffic over transit links is highly asymmetric and that over peering links is nearly symmetric, and consequently, we can estimate the inter-AS distance from the amount of traffic exchanged between two neighboring ASes. We assume larger ASes exchange a larger amount of traffic. According to this assumption, we can estimate the magnitude of each AS recursively, and we can also estimate the inter-AS distance. In the proposed method, we do not use BGP routing tables but the adjacency matrix of AS graph and the eigenvalue analysis to calculate the exchanged traffic.

We show that the proposed method can estimate the inter-AS distance with distinction between peering and transit, and moreover, several times recursion in the estimation delivers more accurate inter-AS distance estimation. The contribution of this paper is to quantify AS relationships as the inter-AS distance for the economical overlay routing without BGP routing tables.

2 Autonomous Systems

The routing among ASes is determined by the inter-domain routing protocol such as BGP [8]. We describe the routing policy, economics and the overview of the

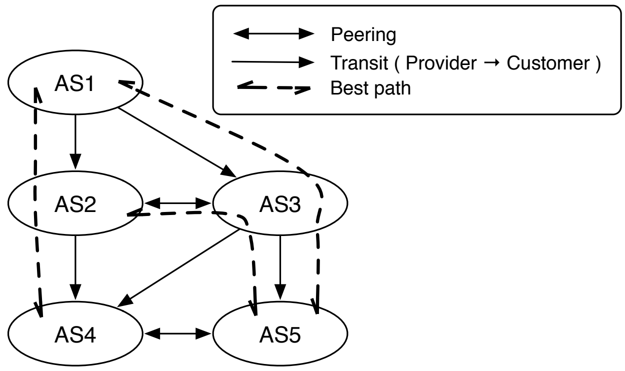

Figure 1: Example of AS relationships and best paths in BGP. The inter-domain routing policy results in the valley-free path topology.

AS relationships inference in the following subsections.

2.1 Routing Policy and Economics

The relationships between neighboring ASes can be categorized into three types, and there is a certain routing policy for each relationship [1, 9]. We describe the relationships, economics and the routing policy below.

1. *Transit*: An AS purchases the Internet access from another AS by paying some amount of money according to the bandwidth usage. This is also called customer-to-provider (c2p) relationship or provider-to-customer (p2c) relationship. *A provider AS* exports its routes and its customer routes, and as well as its provider or peer routes. On the other hand, *a customer AS* exports its routes and its customer routes, but usually does not export its provider or peer routes [1].

2. *Peering*: A pair of neighboring ASes can exchange traffic directly and the traffic exchanged between two peering ASes is free of charge. However, when the traffic becomes highly asymmetric, one party *will* start charging the other according to the bandwidth usage [10]. This is also called peer-to-peer (p2p) relationship. *A peer AS* exports its routes and its customer routes, but usually does not export its provider or peer routes [1].

3. *Sibling*: Multiple ASes can belong to the same organization. Even though each AS might be managed separately from the perspective of network administration, traffic can be exchanged among them *without* any payment. This is also called sibling-to-sibling (s2s) relationship. *A sibling AS* exports its routes and routes of its customers, and as well as its provider or peer routes [1].

This policy results in *valley-free* paths [8] in the inter-domain routing. We show an example of AS relationships and the best paths according to this policy in Figure 1. A *valley-free* path means that a path between two any ASes traverses first uphill (c2p) links and goes across one peering link at most and then traverses downhill (p2c) links when we ignore the sibling links.

2.2 AS Relationships Inference

The relationships among ASes can be inferred by comparing the number of neighbors (degree) between two neighboring ASes, analyzing the paths in BGP routing tables, and referring to the organization information database. We describe the overview of AS relationships inference methods as follows;

1. Transit inference: All the routes in the Internet are constructed according to the valley-free path model. Valley-free paths can be mapped into the 2SAT problem [6]. Additionally, the degree (i.e., the number of neighbor ASes) of a provider AS is generally larger than that of the customer AS on transit [1, 11]. The provider-customer relationship can be inferred by solving the weighted MAX2SAT problem mapped from the AS graph [12], so that we eliminate the contradictions against the valley-free path and provider ASes have larger degrees.

2. Peering inference: According to the valley-free path model, only one at most peering link can exist in a path, so we can replace one link belonging to the top provider into peering without a new contradiction. Moreover, the degrees of two peering ASes are nearly equal in most cases [1, 12]. From these heuristics, peering links can be inferred.

3. Sibling inference: Sibling ASes can be inferred from their organization information registered in public databases such as WHOIS database [7] and IRR (Internet Routing Registry) [13].

The relationships among ASes are more complex and not flat in the real Internet; for instance, some ASes pay some amount of money to peering ASes when the exchanged traffic is highly asymmetric. Hence, it is essential to estimate the relationships numerically when we use the type of AS relationships as a routing metric. However, these inference methods cannot estimate the relationships numerically.

3 Inter-AS Distance Estimation

We propose an inter-AS distance estimation method adopting eigenvalue analysis based on a traffic-flow/-volume assumption. The inter-AS distance denotes the distance between two adjacent ASes based on the type of the AS relationship, and we define it so that transit links are more distant than peering links. We assume the egress traffic from AS i to AS j is in proportion to the magnitude of AS j, and the magnitude of AS i is in proportion to the total ingress traffic to the AS. We calculate the steady-state of traffic according to this assumption, and define the total ingress/egress traffic as the magnitude of each AS recursively. We describe the AS magnitude estimation method and inter-AS distance estimation method based on the magnitude below.

An AS relation matrix $A^{(n)}$ is defined in Equation (1). Each diagonal element of $A^{(n)}$ is 0, and the other elements are defined by Equation (2) in the base case, and by Equation (3) in the other cases. $\psi^{(n)}$ denotes the magnitude of each AS calculated from the matrix $A^{(n)}$ recursively. The base case is $n = 0$, and the matrix $A^{(n)}$ is defined recursively for $n \geq 1 (n \in \mathbb{Z})$. That is, $A^{(0)}$ is the adjacency matrix of AS graph.

$$A^{(n)} := \begin{pmatrix} a_{11}^{(n)} & \dots & a_{1j}^{(n)} & \dots & a_{1m}^{(n)} \\ \vdots & \ddots & \vdots & & \vdots \\ a_{i1}^{(n)} & \dots & a_{ij}^{(n)} & \dots & a_{im}^{(n)} \\ \vdots & & \vdots & \ddots & \vdots \\ a_{m1}^{(n)} & \dots & a_{mj}^{(n)} & \dots & a_{mm}^{(n)} \end{pmatrix} \tag{1}$$

(i) $n = 0$

$$a_{ij}^{(n)} = \begin{cases} 1 : \text{if AS } i \text{ and AS } j \text{ are adjacent} \\ 0 : \text{otherwise} \end{cases} \tag{2}$$

(ii) $n \geq 1 (n \in \mathbb{Z})$

$$a_{ij}^{(n)} = \begin{cases} \psi_j^{(n-1)} : \text{if AS } i \text{ and AS } j \text{ are adjacent} \\ 0 : \text{otherwise} \end{cases} \tag{3}$$

Here, $a_{ij}^{(n)}$ denotes the egress traffic from AS i to AS j. In order to equalize the ingress and egress traffic on each AS, we define normalized AS relation matrix (i.e., traffic transition matrix) $B^{(n)}$ in Equation (4).

$$B^{(n)} = \left(\frac{a_{ij}^{(n)}}{\sum_k a_{ik}^{(n)}} \right) \tag{4}$$

Finally, we calculate the steady-state of the traffic by eigenvalue analysis. The steady-state of the traffic processed by AS i is the ith element of the left eigenvector $\boldsymbol{\psi}^{(n)}$ corresponding to the maximum eigenvalue, and we define $\boldsymbol{\Psi}^{(n)}$ as the magnitude in log-arithmetic scale by Equation (5).

$$\boldsymbol{\Psi}^{(n)} := \log_{10} \left(\boldsymbol{\psi}^{(n)} \right) \tag{5}$$

Table 1: Evaluation 1: The number of links and the proportion for each relationship (*Inet-3.0* dataset)

relationship	#links	proportion
peering (p2p, $\Delta rank = 0$)	1,915	22.0%
transit (p2c/c2p, $\lvert\Delta rank\rvert = 1$)	6,775	78.0%

Since peering is the connection among equal magnitude ASes and a smaller-magnitude (customer) AS purchases the Internet access from a larger-magnitude (provider) AS by paying some amount of money according to the actual bandwidth usage on the transit, the difference of magnitude between two neighboring ASes roughly represents the traffic charge.

We define the difference of magnitude between AS i and AS j as the inter-AS distance in Equation (6). Here, $\Psi_i^{(n)}$ is the magnitude of AS i.

$$\Delta\Psi_{ij}^{(n)} = \Psi_i^{(n)} - \Psi_j^{(n)} \tag{6}$$

4 Evaluation

We evaluate the proposed inter-AS distance estimation method with two topology datasets: 1) the topology generated by *Inet-3.0* [14] and 2) the inferred Internet topology. We show the relation between AS relationships and the inter-AS distance $\Delta\Psi^{(n)}$. Moreover, to make a quantitative evaluation, we adopt ROC[1] analysis on inferring the relationships from $\Delta\Psi^{(n)}$ by Equation (7) with sliding the threshold.

$$\begin{cases} \Delta\Psi^{(n)} > th^{(n)} & : \text{p2c} \\ \Delta\Psi^{(n)} < -th^{(n)} & : \text{c2p} \\ -th^{(n)} \leq \Delta\Psi^{(n)} \leq th^{(n)} & : \text{p2p} \end{cases} \tag{7}$$
$$(th^{(n)} : \text{threshold})$$

4.1 Evaluation 1: Topology Generator

First, we evaluate the proposed method with a topology generated by a topology generator. We employ *Inet-3.0* [14] as a generator, and generate a topology dataset with the following parameters; the total number of nodes in the topology is 5,000, the fraction of degree-one nodes is 0.3, the size of the plane used for node placement is 10,000, and the seed to initialize the random number generator is 0. This dataset includes 5,000 ASes and 8,690 inter-AS links.

We first calculate the spanning tree which root has the maximum degree in the topology, and define the hop count ($hops_{(root,i)}$) from the root to AS i as $rank_i$ in Equation (8). Then, we define the difference of $rank$ between two associated ASes (AS i and AS j) as $\Delta rank_{ij}$

[1]Receiver Operating Characteristic

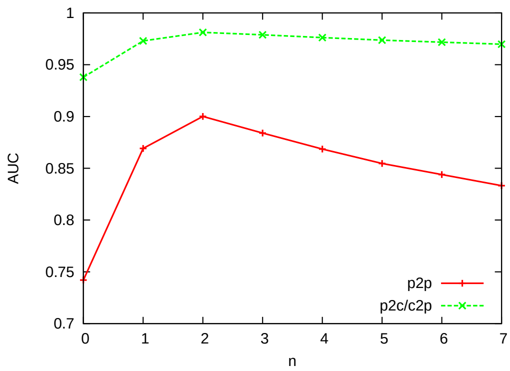

Figure 4: Evaluation 1: Recursion times and area under the ROC curve on inferring each relationship. This figure shows that the proposed estimation method for $n = 2$ improves most in accuracy both of p2p and p2c/c2p relationships inference.

in Equation (9).

$$rank_i := hops_{(root,i)} \tag{8}$$

$$\Delta rank_{ij} = rank_i - rank_j \tag{9}$$

In the generated topology, $\Delta rank_{ij}$ is -1, 0 or 1, so we assume that AS relationships are determined as follows:

$$\Delta rank_{ij} = \begin{cases} -1 & : \text{p2c} \\ 0 & : \text{p2p} \\ 1 & : \text{c2p} \end{cases} \tag{10}$$

Accordingly, we show the proportion for each relationship in Table 1.

Figure 2 represents the distribution of $\Delta\Psi^{(n)}$ normalized by the area for each relationship at $n \in \{0, 1, 2, 3\}$. From Figure 2, we can see p2p links are distributed around the point of $\Delta\Psi^{(n)} = 0$ compared to p2c and c2p links. This figure also shows that several times recursion in the estimation in magnitude improves the accuracy on inferring peering links (i.e., $\Delta\Psi^{(n)}$ of peering links is strongly distributed around 0).

We evaluate the proposed method quantitatively with ROC analysis. Figure 3 represents the ROC curve on inferring peering (p2p) and transit (p2c or c2p), and Figure 4 represents the area under the ROC curve on inferring each relationship. Figure 4 shows the proposed estimation method for $n = 2$ improves most in accuracy of the p2p relationship inference. The proposed method for $n = 2$ increases the area under the ROC curve by 21.3% on inferring peering, and by 4.58% on inferring transit compared to that for $n = 0$. Besides, we evaluate false positive rate (fpr) and true positive rate (tpr) at the optimal threshold. We choose the farthest point

Table 2: Evaluation 1: False positive rate (fpr) and true positive rate (tpr) on inferring each relationship for the optimal threshold

threshold	p2p		c2p/p2c	
	fpr	tpr	fpr	tpr
$th^{(0)} = 1.42$	0.452	0.822	0.0326	0.548
$th^{(1)} = 2.05$	0.241	0.836	0.0297	0.759
$th^{(2)} = 2.89$	0.165	0.860	0.0254	0.835
$th^{(3)} = 3.01$	0.0921	0.754	0.0445	0.908

Table 3: Evaluation 2: The number of links and the proportion for each relationship (CAIDA dataset)

relationship	#links	proportion
sibling (s2s)	209	0.367%
peering (p2p)	3,684	6.47%
transit (p2c/c2p)	53,003	93.2%

on the ROC curve from the line $y = x$ as the optimal threshold, then we show false positive rate (fpr) and true positive rate (tpr) on inferring each relationship in Table 2 in adopting the optimal threshold. Table 2 shows that the proposed estimation method can infer the relationship in higher accuracy with the recursive estimation.

4.2 Evaluation 2: Inferred Internet Topology

Second, we evaluate the proposed method with "The CAIDA AS relationships dataset (09/06/2008)" [15] as a correct AS relationships and topology. The relationships in this dataset are inferred by analyzing collected AS paths according to Gao's AS relationships inference algorithm [1]. This dataset includes 28,253 ASes and 56,896 inter-AS links. We write up the number of inter-AS links and the proportion for each relationship included in the dataset in Table 3. Table 3 shows that 93.2% of inter-AS links are transit, and only 0.367% are sibling.

In the same way as the evaluation with the topology generated by *Inet-3.0* (Evaluation 1), we evaluate the proposed method with ROC analysis. Since only 0.367% of inter-AS links are s2s links, we do not evaluate the estimation on s2s links. Figure 5 represents the ROC curve on inferring peering (p2p) and transit (p2c or c2p), and Figure 6 represents the area under the ROC curve on inferring each relationship. We could not evaluate for n larger than 5 due to the restriction of the accuracy in the floating-point number calculation.

Figure 6 shows the proposed estimation method for larger n improves in accuracy of the p2p relationship

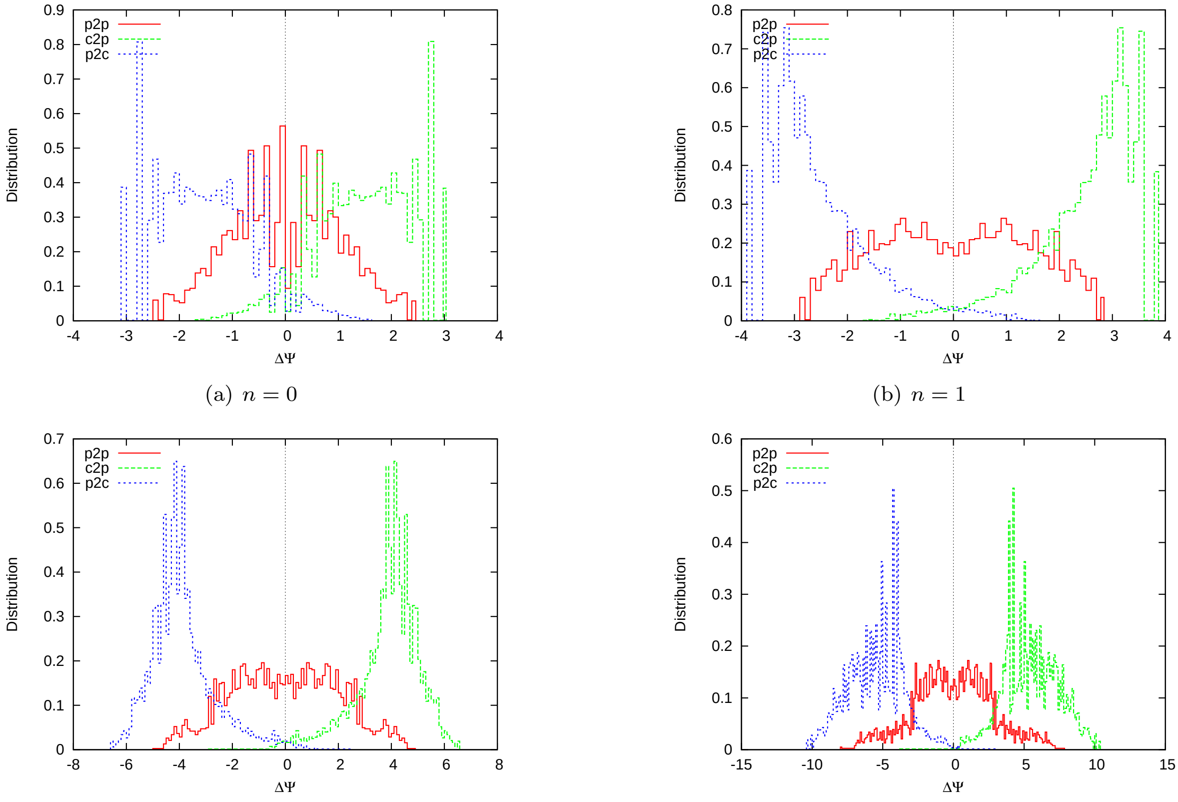

(a) $n = 0$ (b) $n = 1$

(c) $n = 2$ (d) $n = 3$

Figure 2: Evaluation 1: Distribution of $\Delta\Psi^{(n)}$ normalized by the area for each relationship ($n \in \{0, 1, 2, 3\}$). This figure shows that the accuracy on inferring peering links is improved by several times recursion in the estimation in magnitude; $\Delta\Psi^{(n)}$ of peering links is distributed around 0 in comparison with that of transit links.

Table 4: Evaluation 2: False positive rate (fpr) and true positive rate (tpr) on inferring each relationship for the optimal threshold

threshold	p2p		c2p/p2c	
	fpr	tpr	fpr	tpr
$th^{(0)} = 1.97$	0.614	0.953	0.00350	0.386
$th^{(1)} = 2.70$	0.402	0.879	0.00897	0.599
$th^{(2)} = 3.62$	0.304	0.883	0.0114	0.694
$th^{(3)} = 4.35$	0.260	0.865	0.0182	0.731

inference. The proposed method for $n = 2$ increases the area under the ROC curve by 21.0% on inferring peering, and by 0.102% on inferring transit compared to that for $n = 0$. Besides, we show false positive rate (fpr) and true positive rate (tpr) on inferring each relationship in Table 4 in adopting the optimal threshold.

5 Discussion

The difference between the topology generated by *Inet-3.0* and the inferred Internet topology

We evaluated the proposed inter-AS distance estimation method with two topology datasets: 1) the topology generated by a topology generator *Inet-3.0* and 2) the inferred Internet topology. The evaluation delivers the proposed method suits for both topologies. What is the difference between these two topologies? Since the inferred Internet topology is measured at just a part of the Internet, it does not include unobservable links due to the valley-free path topology. Therefore, the result on the evaluation with the generated topology shows that the proposed method is applicable for both an AS topology graph and a spanning subgraph of it.

Complementing the inter-AS distance of unobservable links from the spanning subgraph

What is the advantage of the proposed method in comparison with the other AS relationships inference methods based on path analysis? When we use the type of AS relationships as an overlay routing metric, nodes

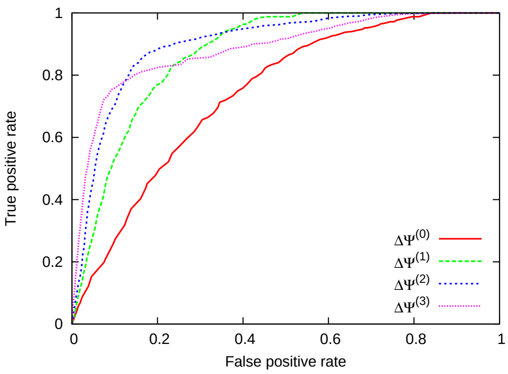

(a) On inferring peering (p2p relationship)

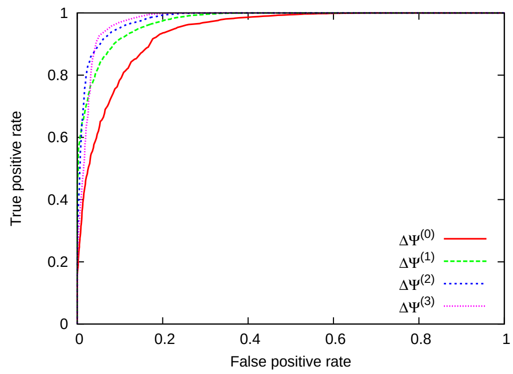

(b) On inferring transit (p2c/c2p relationship)

Figure 3: Evaluation 1: ROC curve on inferring peering and transit for $n \in \{0, 1, 2, 3\}$. Both peering and transit links can be distinguished by the recursive estimation in magnitude. However, the recursive estimation for larger n ($n \geq 3$) makes the accuracy of p2p relationship inference worse.

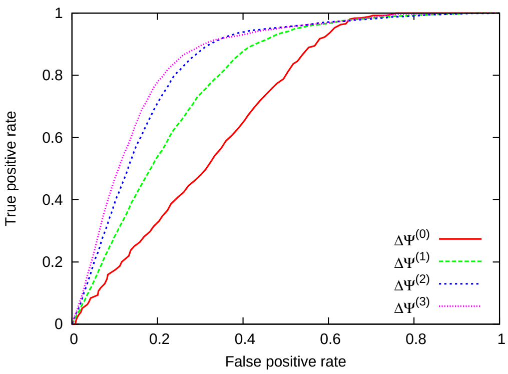

(a) On inferring peering (p2p relationship)

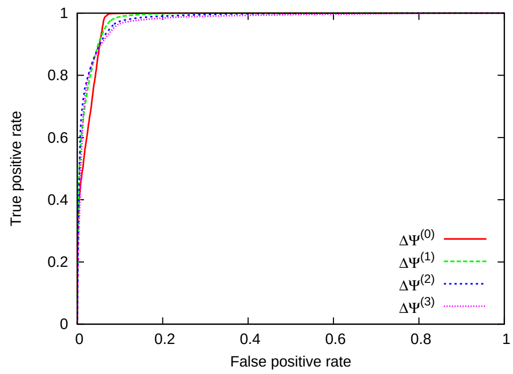

(b) On inferring transit (p2c/c2p relationship)

Figure 5: Evaluation 2: ROC curve on inferring peering and transit for $n \in \{0, 1, 2, 3\}$. Peering links can be distinguished by the recursive estimation in magnitude. However, the recursive estimation for larger n makes the accuracy of the p2c/c2p relationship inference worse.

in the overlay network can find the unobservable links. Since these nodes cannot refer to the BGP routing table and it is hard to gather enough paths, we cannot infer the type of AS relationships of the unobservable links. The proposed method does not require the routing table (AS paths), and the magnitude can be estimated from the spanning subgraph. Hence, we can complement the inter-AS distance of unobservable links from the estimated magnitude. We show an example of the complement in Figure 7. We will evaluate the difference between the distance estimated from the full-graph and the complemented distance estimated from the spanning subgraph in future.

The mismatch between the routing policy and the traffic charge

We adopt "The CAIDA AS relationships dataset (09 /06/2008)" [15] as a correct AS relationships and topology dataset for the evaluation. The relationships in this dataset are inferred by Gao's algorithm [1], and consequently, this dataset can include erroneous inferences and the mismatch between the routing policy and the traffic charge such as so-called *paid peering relationship*. The proposed method is not based on the routing policy but based on the traffic volume, so the proposed method can annotate the mismatch. We will evaluate this point in future.

The validity of the traffic-flow/-volume assumption and modeling the inter-AS traffic

We assume that larger ASes exchange a larger amount of traffic, and the egress traffic is in proportion to the magnitude of the neighbor AS. However, we have not evaluated the validity of this assumption itself. If this assumption is not valid, the results of AS relationships inference from the inter-AS distance would be more in-

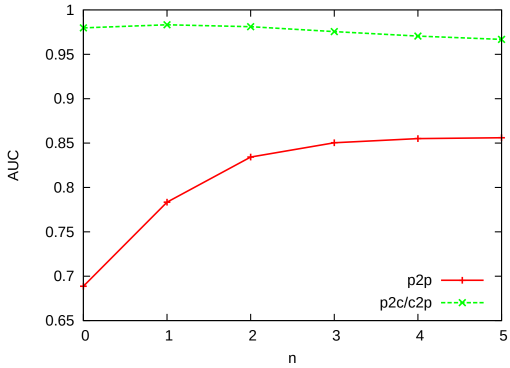

Figure 6: Evaluation 2: Recursion times and area under the ROC curve on inferring each relationship. This figure shows that the proposed estimation method for larger n improves in accuracy of the p2p relationship inference, but it makes the accuracy of the p2c/c2p relationship inference worse.

accurate. We will evaluate the assumption, then we will reconsider it to improve the inter-AS distance estimation.

6 Conclusion

In this paper, we proposed and evaluated a mathematical inter-AS distance estimation method on the basis of AS relationships. The proposed method achieves to estimate the inter-AS distance without BGP routing tables (AS paths) but with the adjacency matrix of AS graph. The contribution of this paper is to quantify AS relationships as the inter-AS distance for the economical overlay routing without BGP routing tables. The proposed inter-AS distance estimation method is applicable for both the topology generated by *Inet-3.0* and the real Internet topology, and several times recursion in the estimation delivers more accurate inter-AS distance estimation.

As future work, we will evaluate so-called *paid peering relationship*; one peer AS pays some amount of money to the other when the exchanged traffic is highly asymmetric on peering. We project that the proposed inter-AS distance estimation method can infer paid peering relationship by evaluating the difference of magnitude around the threshold.

Acknowledgment

We would like to thank Kensuke Fukuda, Yosuke Himura for their support and valuable advice.

References

[1] L. Gao, "On inferring autonomous system relationships in the internet," *Networking, IEEE/ACM Transactions on*, vol. 9, no. 6, pp. 733–745, Dec 2001.

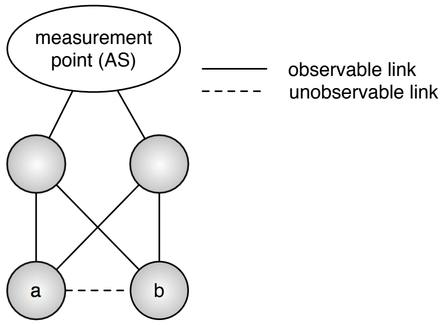

Figure 7: An example of the complement of inter-AS distance of unobservable links from the spanning subgraph. It is possible that the link between AS a and AS b cannot be observed from the measurement point. Even if this link is not observable from the measurement point, we can estimate the magnitude of AS a and AS b. From the estimated magnitude, we can obtain the inter-AS distance between AS a and AS b.

[2] W. Li, S. Chen, and T. Yu, "Utaps: An underlying topology-aware peer selection algorithm in bittorrent," *Advanced Information Networking and Applications, 2008. AINA 2008. 22nd International Conference on*, pp. 539–545, March 2008.

[3] V. Aggarwal, O. Akonjang, and A. Feldmann, "Improving user and isp experience through isp-aided p2p locality," April 2008, pp. 1–6.

[4] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4p: provider portal for applications," in *SIGCOMM '08: Proceedings of the ACM SIGCOMM 2008 conference on Data communication*. New York, NY, USA: ACM, 2008, pp. 351–362.

[5] H. Asai and H. Esaki, "A content delivery path control architecture for the multi-domain P2P CDN system (in japanese)," in *Internet Conference 2008*, Oct 2008, pp. 78–84.

[6] G. D. Battista, T. Erlebach, A. Hall, M. Patrignani, M. Pizzonia, and T. Schank, "Computing the types of the relationships between autonomous systems," *IEEE/ACM Trans. Netw.*, vol. 15, no. 2, pp. 267–280, 2007.

[7] "ARIN WHOIS Database Search," https://www.arin.net/whois/arinwhois.html.

[8] T. Erlebach, A. Hall, A. Panconesi, and D. Vukadinovi, "Cuts and disjoint paths in the valley-free path model of internet BGP routing," in *Combinatorial and Algorithmic Aspects of Networking*, ser. Lecture Notes in Computer Science, vol. 3405/2005. Springer Berlin / Heidelberg, 2005, pp. 49–62.

[9] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple isps," *IEEE/ACM Trans. Netw.*, vol. 14, no. 6, pp. 1233–1245, 2006.

[10] C. Huang, J. Li, and K. W. Ross, "Can internet video-on-demand be profitable?" in *SIGCOMM '07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*. New York, NY, USA: ACM, 2007, pp. 133–144.

[11] B. Huber, S. Leinen, R. O'Dell, and R. Wattenhofer, "Inferring AS relationships beyond counting edges," D-INFK, Tech. Rep. Nr.446, 2004.

[12] X. Dimitropoulos, D. Krioukov, B. Huffaker, kc claffy, and G. Riley, "Inferring AS relationships: Dead end or lively beginning?" in *Experimental and Efficient Algorithms*, ser. Lecture Notes in Computer Science, vol. 3503/2005. Springer Berlin / Heidelberg, 2005, pp. 113–125.

[13] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley, "As relationships: inference and validation," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 1, pp. 29–40, 2007.

[14] J. Winick and S. Jamin, "Inet-3.0: Internet topology generator," EECS, University of Michigan, Tech. Rep. UM-CSE-TR-456-02, 2002.

[15] "The CAIDA AS Relationships Dataset (09/06/2008)," http://www.caida.org/data/active/as-relationships/.