Estimating AS Relationships for Application-Layer Traffic Optimization

Hirochika Asai and Hiroshi Esaki

The University of Tokyo, Japan
panda@hongo.wide.ad.jp, hiroshi@wide.ad.jp

Abstract. The relationships among autonomous systems (ASes) on the Internet are categorized into two major types: *transit* and *peering*. We propose a method for quantifying AS' network size called *magnitude* by recursively analyzing the AS adjacency matrix converted from a spanning subgraph of the AS-level Internet topology. We estimate the relationships of inter-AS links by comparing differences in magnitude of two neighboring ASes, while showing differences in the magnitude, representing AS relationships appropriately through three evaluations. We also discuss the applicability of this method to AS relationships-aware application-layer traffic optimization.

1 Introduction

The Internet consists of thousands of autonomous systems (ASes) operated by distinct administrative domains such as Internet service providers (ISPs), companies and universities. There are commercial relationships between interconnected ASes, and the relationships are categorized into two major types [1]: *transit* and *peering*. *Transit* relationships are also called provider-customer relationships, and customer ASes purchase Internet access from their transit providers by paying some amount of money. On the contrary, *peering* relationships are equal relationships between interconnected ASes, and traffic exchanged between peering ASes is free of charge. Therefore, transit traffic exchanged with provider ASes costs more for customer ASes compared to that exchanged with customer ASes or traffic exchanged over peering links from the economical viewpoint. Note that we refer to a transit link from a customer AS to a provider AS and a link with opposite orientation as customer-to-provider (c2p) link and provider-to-customer (p2c) link, respectively. We also refer to a peering link as peer-to-peer (p2p) link.

Researches and discussions regarding application-layer traffic optimization have been conducted [2, 3]. We propose a path selection method that takes into account the types of AS relationships in content delivery networks utilizing peer-to-peer technologies [4]. We show that the proposed method has reduced high-cost transit traffic for residential ISPs, which provide their network to consumers hosting content delivery network peers, by assigning link cost onto inter-AS links and avoiding selecting high-cost paths. In the proposed method, we have used the types of the relationships to assign cost into inter-AS links. However, there

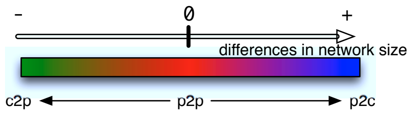

Fig. 1. AS relationships representation by differences in the network size

still exists a problem that most commercial ISPs do not want to disclose their relationships because the interconnections are established by their commercial contracts.

Several AS relationships inference algorithms [5–8] have been proposed. These algorithms infer the relationships by analyzing AS paths in Border Gateway Protocol (BGP) routing tables according to the valley-free path model [9]. However, these algorithms cannot infer the relationships of these invisible links though there are lots of invisible inter-AS links in the set of AS paths extracted from the BGP routing tables because the set of AS paths produce spanning subgraphs (i.e., parts) of the Internet topology and the number of ASes which provide their BGP routing tables to public are limited. Here, we refer to links which are not contained in the set of AS paths in publicly available BGP routing tables as *invisible* inter-AS links. On the other hand, applications on the Internet possibly utilize paths containing invisible inter-AS links as well for their communications because the routing tables of ASes which provide their network to these applications are usually different from the publicly available BGP routing tables. Consequently, it is essential for AS relationships-aware application-layer traffic optimization to estimate the relationships of these invisible links as well as visible links. We note that invisible inter-AS links in an AS path which an application utilizes for its communication can be found by the application with a network management tool (e.g., "traceroute" tool), and the existing algorithms cannot infer the relationships from the found AS path due to lack of AS paths.

In this paper, we propose a method for quantifying the AS' network size, which we call *magnitude*, by recursively analyzing the AS adjacency matrix approximated from a measured spanning subgraph of the AS-level Internet topology according to inter-AS connectivities and a traffic flow model. We show the differences in magnitude appropriately represent AS relationships through three evaluations.

2 AS Relationships Estimation

Relationships between interconnected ASes are characterized by the exchanged traffic volume and the network size [10–12]. The traffic volume exchanged over transit links is highly asymmetric, and the traffic volume from a transit provider to the customer is generally larger than that from the customer to the provider. On the other hand, the traffic volume exchanged among peering ASes is nearly symmetric. From the viewpoint of the network size, transit providers are larger than their customers and peering ASes are nearly equal in size as well as exchanged traffic volume. Fig. 1 shows the AS relationships representation by dif-

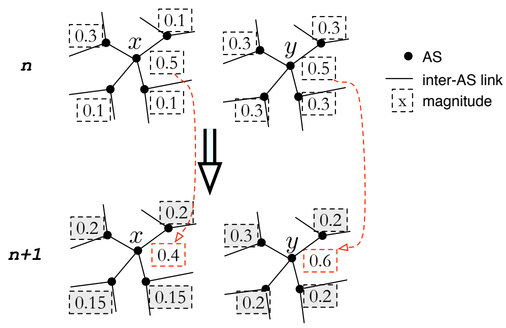

Fig. 2. The concept of recursive definition of magnitude

ferences in the network size. Links with positive and those with negative values are considered p2c ones and c2p ones, respectively, and links with values around 0 are considered p2p ones. For example, degree, the number of interconnected ASes, can be used as one of the indicators which represent AS' network size [13]. It is said that differences in degree numerically represent the relationships, and consequently, it has been used for AS relationships inference algorithms based on the path analysis [5–8].

We propose a method for quantifying AS' network size called *magnitude* which represents AS relationships better than degree. The magnitude is computed recursively by taking into account the magnitude of neighboring ASes to improve the representation of AS' magnitude. Fig. 2 shows the concept of recursive definition of magnitude. In this figure, AS x and AS y have equal magnitude (0.5) where the recursion level is n, but the magnitude of their neighbors is different. This difference makes the magnitude of AS x and AS y different where the recursion level is $n + 1$. Since the neighbors of AS y are larger than those of AS x, AS y become larger than AS x by taking account the magnitude of their neighbors. We then propose a method for estimating the relationships from the quantified magnitude.

2.1 AS Magnitude and Inter-AS Traffic Flow Model

We use a measured spanning subgraph of the AS-level Internet topology for the AS magnitude quantification. Let a graph $G_I = (V_I, E_I)$ be the whole AS-level Internet topology, i.e., a set V_I of ASes contains all ASes on the Internet and a set E_I of inter-AS links contains all inter-AS links though some of them may be invisible from a measurement. We can measure subgraphs of the whole Internet AS graph G_I from AS paths in BGP routing tables. A measured subgraph is generally a quasi spanning subgraph[1] of the AS-level Internet topology. Here,

[1] A *quasi* spanning subgraph denotes the subgraph containing *almost all* vertices (ASes) even though it does not contain some edges (inter-AS links). The measured AS paths constitute a quasi spanning subgraph due to the existence of default route configuration and so on; i.e., there exist a few invisible ASes as well as some invisible inter-AS links.

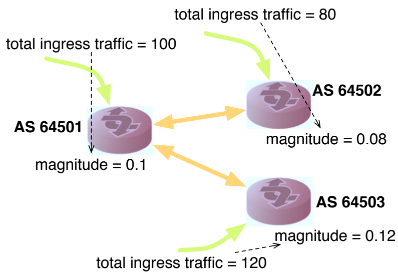

Fig. 3. Definition of AS magnitude

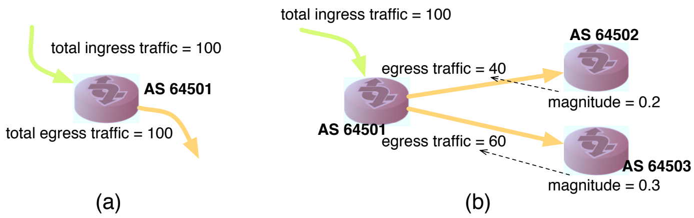

Fig. 4. Traffic flow assumptions

we define a graph $G_S = (V_S, E_S)$ as the measured spanning subgraph. Since the graph G_S is a quasi spanning subgraph, the set V_S of ASes is a subset of the set V_I (i.e., $V_S \subseteq V_I$) and the set E_S of inter-AS links is a subset of the set E_I (i.e., $E_S \subseteq E_I$). A complementary set $\bar{E}_S$ which is the set of invisible inter-AS links are represented by an equation of the form: $\bar{E}_S = E_S \backslash E_I$.

From the measured spanning subgraph G_S, we compute AS' magnitude according to the traffic flow model. In our AS magnitude quantification method, we define that the magnitude is proportional to the total ingress traffic to the AS at the steady state in the traffic flow model as shown in Fig. 3. In another word, the magnitude represents the traffic density at the steady state in the traffic flow model. Let $t_{v_i v_j}$ be traffic from AS v_i to AS v_j, the magnitude of AS v_i is defined by Equation (1). A symbol ρ_{v_i} denotes the magnitude of an AS v_i, subject to equations: $\sum_{v_k \in V_I} \rho_{v_k}^2 = 1$ and $\rho_{v_k \in \bar{V}_S} = 0$.

$$\rho_{v_i} := C \sum_{v_k \in \mathrm{nbr}(v_i)} t_{v_k v_i} \tag{1}$$

$$\text{s.t. } C = \mathit{const.}, \sum_{v_k \in V_S} \rho_{v_k}^2 = 1$$

Here, the function $\mathrm{nbr}(v_i)$ returns a set of neighbor ASes of AS v_i, provided that the links between AS v_i and the neighbors are in the set E_S.

We then introduce the simple traffic flow model to compute the magnitude as shown in Fig. 4. We describe the assumptions in the model, as follows:

(a) The total amount of ingress traffic to AS v_i is equal to the egress traffic of AS v_i: $\sum_{v_k \in \mathrm{nbr}(v_i)} t_{v_k v_i} = \sum_{v_k \in \mathrm{nbr}(v_i)} t_{v_i v_k}$.

(b) The amount of egress traffic from AS v_i to AS v_j is proportional to the magnitude of AS v_j (ρ_{v_j}): $t_{v_i v_j} = \frac{\rho_{v_j}}{\sum_{v_k \in \mathrm{nbr}(v_i)} \rho_{v_k}} \sum_{v_k \in \mathrm{nbr}(v_i)} t_{v_i v_k}$.

Since the magnitude is used in these assumptions and it is computed from these assumptions, the magnitude is computed recursively. We compute the steady state of ingress/egress traffic in these assumptions with fixed values of magnitude (e.g., $\boldsymbol{\rho} = [1, \cdots, 1]^t$ for initial case), and then we recursively redetermine the magnitude from the traffic distribution at the steady state.

2.2 AS Magnitude Computation

The steady-state of traffic according to the traffic flow model described in the previous subsection is solved by eigenvalue analysis of a traffic transition matrix. We first define a weighted AS adjacency matrix ${}^{n}A$ by the equation: ${}^{n}A := \left({}^{n}a_{v_i v_j}\right)$, where $v_i, v_j \in V_S$. Here, the left superscript ${}^{n}\bullet$ ($n \geq 0, n \in \mathbb{Z}$) denotes the recursion level. This matrix is extracted from the measured spanning subgraph G_S. Each diagonal element of the matrix ${}^{n}A$ is 0, and other elements are defined by Equation (2a) for initial case ($n = 0$), and by Equation (2b) for other cases.

(i) $n = 0$

$$ {}^{n}a_{v_i v_j} = \begin{cases} 1 & \text{: if AS } v_i \text{ and AS } v_j \text{ are adjacent} \\ 0 & \text{: otherwise} \end{cases} \tag{2a}$$

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

$$ {}^{n}a_{v_i v_j} = \begin{cases} {}^{(n-1)}\rho_{v_j} & \text{: if AS } v_i \text{ and AS } v_j \text{ are adjacent} \\ 0 & \text{: otherwise} \end{cases} \tag{2b}$$

The matrix ${}^{n}A$ where $n \geq 1$ is defined recursively from the vector of magnitude ${}^{(n-1)}\boldsymbol{\rho}$. The matrix ${}^{n}A$ where $n \geq 1$ is also represented by the equation: ${}^{n}A = I\ {}^{(n-1)}\boldsymbol{\rho}\ {}^{0}A$, where I denotes the identity matrix.

Next, we equalize the ingress and egress traffic on each AS by considering that ${}^{n}a_{v_i v_j}$ represents the egress traffic from AS v_i to AS v_j. We define a traffic transition matrix ${}^{n}T$ by Equation (3).

$$ {}^{n}T := \left(\frac{{}^{n}a_{v_i v_j}}{\sum_{v_k} {}^{n}a_{v_i v_k}} \right) \tag{3}$$

We note that the traffic transition matrix is represented by a form of the stochastic matrix. Finally, we compute the steady state of traffic by eigenvalue analysis of the traffic transition matrix. The steady state is determined by calculating the left eigenvector of ${}^{n}T$ corresponding to the maximum eigenvalue. We define this left eigenvector as the vector of magnitude: ${}^{n}\boldsymbol{\rho} = [{}^{n}\rho_{v_1}, \cdots, {}^{n}\rho_{v_m}]^t$ (s.t. $\|{}^{n}\boldsymbol{\rho}\| = 1$, $m = \|V_S\|$). Here, we note that the magnitude of an AS where $n = 0$ results in a value of the AS' degree multiplied by a constant, though we omit the proof.

2.3 AS Relationships Estimation

From the quantified magnitude, we estimate the relationships of inter-AS links. We define the difference in logarithmic magnitude for an inter-AS link e_x from an AS v_i to an AS v_j as the magnitude distance ${}^n\delta_{e_x}$. The magnitude distance ${}^n\delta_{e_x}$ is defined by Equation (4).

$$\begin{aligned} {}^n\delta_{e_x} &:= \log_{10} {}^n\rho_{v_i} - \log_{10} {}^n\rho_{v_j} \\ &\text{s.t. } e_x = (v_i, v_j),\ e_x \in E_I,\ v_i, v_j \in V_I \end{aligned} \tag{4}$$

The distribution (e.g., the minimum and maximum values) of magnitude distances is different for each recursion level n. We can compare the magnitude distances for the same recursion level, but we cannot do it for different recursion levels. To normalize the distribution of magnitude distances to uniform distribution, we define the ranked magnitude distance ${}^n\delta'_{e_x}$ for a inter-AS link e_x from the magnitude distances. Let a set ${}^n\boldsymbol{\delta}_S$ be a vector of the magnitude distances for $\forall e \in E_S$, the ranked magnitude distances are defined by Equation (5).

$$\begin{aligned} {}^n\delta'_{e_x} &:= 2\frac{\text{rank-of}({}^n\delta_{e_x}) - 1}{\|E_S\| - 1} - 1 \\ &\text{s.t. } e_x \in E_S \end{aligned} \tag{5}$$

Here, the function rank-of returns the rank of the magnitude distance ${}^n\delta_{e_x}$, sorting magnitude distances in the elements of the vector ${}^n\boldsymbol{\delta}_S$ in ascending order; i.e., the returned value should be distributed uniformly in the range $[1, \|E_S\|]$. We note that the ranked magnitude distances are hardly applied to application-layer traffic optimization because the ranked magnitude distances are defined only for visible inter-AS links (i.e., $\forall e_x \in E_S$).

The magnitude distance ${}^n\delta_{e_x}$ numerically represents AS relationships as shown in Fig. 1. For example, if the absolute value $|{}^n\delta_{e_x}|$ is around 0, two neighboring ASes are nearly symmetric and the inter-AS link e_x is estimated as p2p. Since the magnitude distances numerically represent the relationships, we can infer the type of the relationships from these magnitude distances by setting a threshold in Equation (6).

$$\begin{cases} {}^n\delta > {}^n\tau & \rightarrow \text{p2c} \\ {}^n\delta < -{}^n\tau & \rightarrow \text{c2p} \\ -{}^n\tau \le {}^n\delta \le {}^n\tau & \rightarrow \text{p2p} \end{cases} \quad \text{s.t. } {}^n\tau \ge 0\ ({}^n\tau\text{: threshold}) \tag{6}$$

We note that the magnitude distances can be directly used for AS relationships-aware application-layer traffic optimization without inferring the types of the relationships by Equation (6) as well because they numerically represent the relationships, despite the fact that the inferred types of the relationships can be helpful to give some guidelines to applications.

Table 1. The number of inter-AS links and the proportion by type of relationships

type of relationships	#links	proportion
sibling (s2s)	219	0.302%
peering (p2p)	6142	8.47%
transit (p2c/c2p)	66181	91.2%

3 Evaluation

We make three evaluations on the proposed AS relationships estimation method. In the first evaluation, we evaluate the accuracy of the inference of types of the relationships, which are inferred by Equation (6). We show that the types of the relationships are inferred appropriately by the magnitude distances without analyzing AS paths. We also show that the recursive computation of magnitude improves the accuracy of peering inference. In the second evaluation, we show the characteristics of the magnitude distances among well-known tier-1 ISPs. Since the relationships between any two tier-1 ISPs are considered peering, we show these peering links are characterized better by recursive computation of magnitude. In the third evaluation, we show the proposed method can estimate the invisible links by counting the number of the paths which follow the valley-free path model.

3.1 Datasets

We use two types of datasets for the evaluation; 1) *CAIDA's dataset* and 2) *RIB datasets*. CAIDA's dataset defines inter-AS links and the relationships, and RIB datasets define AS paths. CAIDA's dataset is used for both the magnitude computation (i.e., as a quasi spanning subgraph) and the verification (i.e., as a correct AS relationships dataset). RIB datasets are used only for the verification. We describe these datasets below.

We employ "The CAIDA AS relationships dataset (10/08/2009) [14]" as a quasi spanning subgraph for the magnitude computation and a correct AS relationships dataset for the verification. The relationships in this dataset are inferred by the algorithm [7, 8]. In this paper, we call this *CAIDA's dataset*. This dataset contains 32281 ASes and 72542 inter-AS links. We write up the number of inter-AS links and the proportion by type of relationships in Table 1.

We also use Routing Information Base (RIB) datasets (archives: 01/08/2009–05/08/2009) from "Route Views Project [15]" and "RIPE NCC Projects Routing Information Service [16]". We call these *RIB datasets*. We extract AS paths from these datasets, excluding the paths which include private AS numbers, four-octet AS numbers and *AS23456*[2]. We summarize the measurement points, the data

[2] Four-octet AS numbers can be translated into *AS23456* when BGP routers do not support four-octet AS numbers. Since we do not analyze BGP options in this paper, we also exclude four-octet AS numbers to identify ASes.

Table 2. Measurement points, the data sources, the number of unique AS paths and the number of unique inter-AS links

	measurement point	abbr.	source	#unique paths	#unique links
a)	Oregon IX	oregon-ix	RV	1,641,700	69,246
b)	Equinix Ashburn	eqix	RV	257,630	57,726
c)	ISC (PAIX)	isc	RV	433,861	60,641
d)	LINX	linx	RV	784,053	65,774
e)	DIXIE (WIDE)	wide	RV	208,542	51,328
f)	RIPE NCC, Amsterdam	rrc00	RIS	641,324	64,151
g)	Otemachi, Japan (JPIX)	rrc06	RIS	96,951	45,040
h)	Stockholm, Sweden (NETNOD)	rrc07	RIS	242,386	56,563
i)	Milan, Italy (MIX)	rrc10	RIS	291,297	56,241

"RV" and "RIS" stand for "Route Views Archive Project [15]" and "RIPE NCC Projects Routing Information Service [16]", respectively.

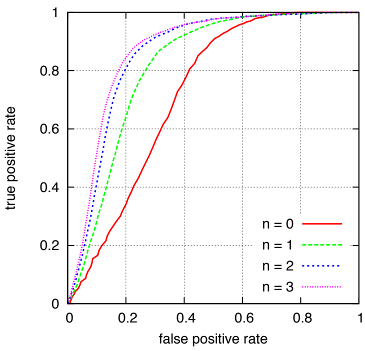

(a) peering inference

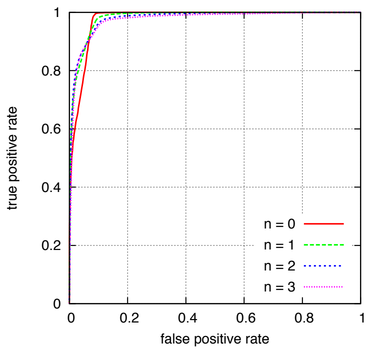

(b) transit inference

Fig. 5. ROC curve on inferring peering and transit relationships where $n \in \{0, 1, 2, 3\}$

sources, the number of unique AS paths and the number of unique inter-AS links in Table 2. We use AS paths in these datasets for the verification.

3.2 Evaluation 1: Accuracy of AS Relationships Inference

In this evaluation, we use CAIDA's dataset for both the magnitude computation and the verification. We compute the magnitude for all ASes in the dataset and magnitude distances for all inter-AS links in the dataset. We infer the relationships from the magnitude distances by Equation (6) with sliding the threshold ${}^n\tau$, and verify inferred relationships by those relationships defined in CAIDA's dataset.

We draw a Receiver Operating Characteristic (ROC) curve on inferring peering and transit relationships in Fig. 5; we plot the false positive rate and the true positive rate at x-axis and y-axis, respectively, with sliding the threshold. It is commonly said that the area under the curve (AUC) represents the accuracy of the inference because points of lower false positive rate and higher true

Table 3. Well-known tier-1 ISPs

AS no.	Name	AS no.	Name
7018	AT&T	3549	Global Crossing (GBLX)
3356	Level 3 Communications (L3)	2914	NTT Communications (Verio)
209	Qwest	1239	Sprint
6453	Tata Communications	701	Verizon Business
3561	Savvis	1299	TeliaSonera
6461	AboveNet	2828	XO Communications

positive rate increases the AUC. This figure shows that the magnitude distances represent both transit and peering relationships appropriately, and the recursive computation improves the accuracy of peering inference. By comparing the AUC on inferring peering, the values of AUC where $n = \{0, 1, 2, 3\}$ are 0.720, 0.814, 0.856 and 0.871, respectively, i.e., the recursive computation improves the accuracy of peering inference. By comparing the AUC on inferring transit, the values of AUC where $n = \{0, 1, 2, 3\}$ are 0.977, 0.982, 0.980 and 0.974, respectively, i.e., the recursive computation does not change the accuracy of transit inference much. From these results, the proposed method with the recursive computation represents the network size and the relationships better than the method without recursive computation (i.e., degree-based one). We note again that the magnitude where $n = 0$ is degree multiplied by a constant.

3.3 Evaluation 2: Characteristics of Ranked Magnitude Distances among Tier-1 ISPs

We showed the recursive computation improves the accuracy of peering inference in the previous subsection. In that evaluation, we assumed that the types of AS relationships defined in CAIDA's dataset are correct, but the relationships in CAIDA's dataset may include inaccurate inferences. In this evaluation, we do not use the relationships defined in CAIDA's dataset to eliminate the influence of inaccurate inferences in CAIDA's dataset while we use CAIDA's dataset for both the magnitude computation. Instead, we evaluate the links among well-known tier-1 ISPs. The links between any two tier-1 ISPs are considered peering. We list the well-known tier-1 ISPs in Table 3.

We show the characteristics of AS relationships of the inter-AS links among well-known tier-1 ISPs in Fig. 6. Each line represents a link between two neighboring tier-1 ISPs. The relationships between any two tier-1 ISPs are considered peering. Hence, Fig. 6 shows the characteristics of magnitude differences of peering links. This figure shows that the recursive computation of magnitude decreases the absolute value of ranked magnitude distance $|{}^{n}\delta'|$. This means the recursive computation of magnitude improves the accuracy of peering estimation. the maximum values among the ranked magnitude distances of the links between tier-1 ISPs are 0.178 where $n = 0$ and 0.0732 where $n = 5$.

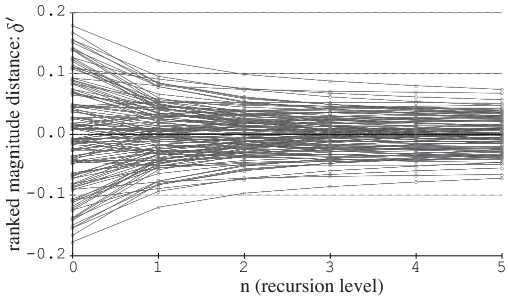

Fig. 6. Characteristics of AS relationships of inter-AS links among well-known tier-1 ISPs

3.4 Evaluation 3: AS Relationships Estimation for Invisible Links

In this evaluation, we use RIB datasets for the verification based on the valley-free path model [9]. We annotate AS relationships to all the inter-AS links in RIB datasets by Equation (6) with the threshold: ${}^{n}\tau = 0$. To show the advantage of the proposed method, we annotate AS relationships to them by the relationships defined in CAIDA's dataset as well. We then count valid paths (i.e., paths following valley-free path model). When we consider that paths between any two ASes follow the valley-free path model, the number of valid paths represents the accuracy of AS relationships inference.

We show the percentage of valid paths for each RIB dataset in Fig. 7. A legend *caida* denotes the paths are annotated by the relationships defined in CAIDA's dataset, and the other legends $n = \{0, \cdots, 5\}$ denote the paths are annotated by the inferred relationships from the magnitude distances. We note that larger values, in this figure, represent higher accuracy on inferring AS relationships. Excluding the paths annotated by the relationships defined in CAIDA's dataset, the percentage of valid paths annotated by the magnitude distances where $n = 1$ is highest for every RIB dataset. This means the orientations of transit links are represented the best by the magnitude distances where $n = 1$. Additionally, including the paths annotated by the relationships defined in CAIDA's dataset, the percentage of valid paths annotated by the magnitude

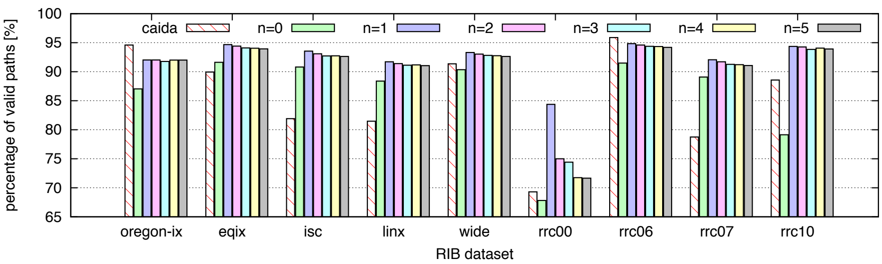

Fig. 7. Percentage of valid paths, which follow the valley-free path model

distances where $n = 1$ is highest for all of the datasets except oregon-ix and rrc06. Though the relationships in CAIDA's dataset are inferred so as to maximize the number of valid paths, the percentage is not highest. This is because there are some links which are not contained in CAIDA's dataset, but the magnitude distances can be computed from the quantified magnitude if the edge ASes are contained in the magnitude computation procedure.

4 Discussion

Inter-AS traffic flow model: We introduce a simple inter-AS traffic flow model in Section 2.1 to compute the magnitude, but the actual inter-AS traffic flow is not so simple. For example, we assume the total ingress traffic volume is equal to the total egress traffic volume but the total ingress traffic volume is generally larger than the total egress traffic volume at residential ISPs. To justify the traffic flow model, we discuss the meaning of the model for the recursion level $n = 0$. For the recursion level $n = 0$, the traffic transition matrix ${}^{0}T$ becomes a stochastic matrix; i.e., this model implies random walk-like transition of traffic. We have described that the magnitude of an AS where $n = 0$ results in a value of the AS' degree multiplied by a constant. For the recursion levels $n \geq 1$, the traffic transition matrix ${}^{n}T$ is weighted by the magnitude of neighbors, and the weighting procedure matches the hierarchical routing on the Internet; i.e., traffic tends to go to transit providers because the transit providers (larger ASes) relay traffic to other ASes. Therefore, the traffic flow model is justified by the definition of degree and the hierarchical routing. There is no doubt that the traffic flow model and the weighting procedure can be modified to improve the proposed method. We will work on this modification in future.

The recursion level: The recursion level means the hop count to which the method takes into account the network size for the magnitude computation. For example, $n = 1$ means the method takes into account the network size of neighbors, and $n = 2$ means the method takes into account the network size of neighbors and that of two-hop neighbors. From Evaluation 1 and 2, we show that the accuracy of peering inference is improved by increasing the recursion level n. On the other hand, from Evaluation 3, the orientations of transit links are represented the best where the recursion level is $n = 1$. These results show that the orientations of transit links are represented by the network size from the local viewpoint (i.e., at most one-hop neighbors' size), and the peering links are represented by the network size from the global viewpoint.

AS relationships-aware application-layer traffic optimization: We have described that the proposed method is applicable to AS relationships-aware application-layer traffic optimization. Suppose, for instance, there are two content mirror servers s_1 and s_2, and a client c, and the path from s_1 to c and that from s_2 to c are $\{s_1 \rightarrow \text{p2p} \rightarrow c\}$ and $\{s_2 \rightarrow \text{c2p} \rightarrow \text{p2c} \rightarrow c\}$, respectively. To reduce high-cost transit traffic, the client should select the server s_1. The relationships of each inter-AS link between s_1 and c, and s_2 and c are required to be inferred to enable AS relationships-aware server selection. The magnitude distances or the inferred relationships by Equation (6) can be used for it. As described in Section 2.3, the ranked magnitude distances are hardly applied to application-layer traffic optimization. The magnitude distances are easily applied to applications by using values of these distances as metric directly. When we use the inferred relationships, the threshold should be tuned for each application; e.g., some applications permit false positive and the others do not.

Paid peer consideration: On the Internet, there are so-called *paid peering* relationships, which are intermediate relationships between peering and transit. The proposed method can quantify the network size well, and the relationships are characterized by the magnitude distances. Hence, we do consider the possibility of estimating the paid peer relationships as well as the applicability applicability of the proposed method for estimating these complex relationships in future.

5 Related Work

Gao [5] has proposed an algorithm to infer AS relationships. The author has shown that the relationships can be inferred by comparing the number of neighbors (i.e., degree) between two neighboring ASes, analyzing the AS paths in BGP routing tables based on the valley-free path model [9]. Battista et al. [6] improved Gao's algorithm. They mapped this problem into weighted MAX2SAT (maximum-2-satisfiability) problem to compute the orientation of transit links. However, on the real Internet, there are lots of invisible inter-AS links. Therefore, it is difficult to apply these AS relationships inference algorithms based on path analysis to AS relationships-aware application-layer traffic optimization because applications often utilize links which relationships are not annotated by these algorithms from the spanning subgraphs.

6 Summary

We proposed a method for quantifying the AS' network size called *magnitude* by recursively analyzing the AS adjacency matrix which is approximated from a measured spanning subgraph of the AS-level Internet topology. We showed that the differences in magnitude appropriately represent AS relationships by three evaluations. We also showed that the recursive computation of magnitude improved the accuracy of peering inference. The contributions of this paper are followings: 1) The proposed method can estimate AS relationships of any inter-AS links, and the estimated relationships are applicable to AS relationships-aware application-layer traffic optimization. 2) The proposed method uses AS adjacency information which is more highly available information than AS paths which have been commonly used in the previous works.

We will apply the estimated magnitude distances to AS relationships-aware application-layer traffic optimization, and design an architecture to utilize these distances as a traffic control metric.

References

1. Shakkottai, S., Srikant, R.: Economics of network pricing with multiple ISPs. IEEE/ACM Trans. Netw. **14**(6) (2006) 1233–1245
2. Xie, H., Yang, Y.R., Krishnamurthy, A., Liu, Y.G., Silberschatz, A.: P4P: provider portal for applications. In: SIGCOMM '08: ACM SIGCOMM 2008 conference on Data communication, New York, NY, USA, ACM (2008) 351–362
3. The Internet Engineering Task Force (IETF): Application-layer traffic optimization (alto). http://datatracker.ietf.org/wg/alto/charter/
4. Asai, H., Esaki, H.: A content delivery path control architecture for the multi-domain P2P CDN system (in japanese). Internet Conference 2008 (Oct 2008)
5. Gao, L.: On inferring autonomous system relationships in the Internet. Networking, IEEE/ACM Transactions on **9**(6) (Dec 2001) 733–745
6. Battista, G.D., Erlebach, T., Hall, A., Patrignani, M., Pizzonia, M., Schank, T.: Computing the types of the relationships between autonomous systems. IEEE/ACM Trans. Netw. **15**(2) (2007) 267–280
7. Dimitropoulos, X., Krioukov, D., Huffaker, B., kc claffy, Riley, G.: Inferring AS relationships: Dead end or lively beginning? In: Experimental and Efficient Algorithms. Volume 3503/2005 of Lecture Notes in Computer Science., Springer Berlin / Heidelberg (2005) 113–125
8. Dimitropoulos, X., Krioukov, D., Fomenkov, M., Huffaker, B., Hyun, Y., kc claffy, Riley, G.: As relationships: inference and validation. SIGCOMM Comput. Commun. Rev. **37**(1) (2007) 29–40
9. Erlebach, T., Hall, A., Panconesi, A., Vukadinovi, D.: Cuts and disjoint paths in the valley-free path model of Internet BGP routing. In: Combinatorial and Algorithmic Aspects of Networking. Volume 3405/2005 of Lecture Notes in Computer Science., Springer Berlin / Heidelberg (2005) 49–62
10. Norton, W.B.: Internet service providers and peering (2000)
11. Weiss, M.B., Shin, S.J.: Internet interconnection economic model and its analysis: Peering and settlement. Netnomics **6**(1) (2004) 43–57
12. Huang, C., Li, J., Ross, K.W.: 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) 133–144
13. Tangmunarunkit, H., Doyle, J., Govindan, R., Willinger, W., Jamin, S., Shenker, S.: Does AS size determine degree in AS topology? SIGCOMM Comput. Commun. Rev. **31**(5) (2001) 7–8
14. Cooperative Association for Internet Data Analysis: The CAIDA AS Relationships Dataset (10/08/2009). http://www.caida.org/data/active/as-relationships/
15. University of Oregon: Route Views Project. http://www.routeviews.org/
16. Réseaux IP Européens Network Coordination Centre: Routing Information Service (RIS). http://www.ripe.net/projects/ris/

Acknowledgment

We especially thank Kensuke Fukuda and Yosuke Himura for their valuable advice on the AS graph analysis. We also thank Burkhard Stiller for shepherding us to improve this paper.