Autonomous Systems From The View Of Multiple Vantage Points

Sr.Asst.Prof.G.Srinivasa Rao*, Asso.Prof.G.Venkateswara Rao*, Asso.Prof.M.Venkateswara Rao*, Asso.Prof.G.Appa Rao**, Asst.Prof.D.kishore Kumar*, Asst.Prof.G.Anuradha**, Asst.Prof.D.Veerabadra Rao*, Asst.Prof.P.Surya Chandra*, K.Krishna Chaitanya*,T.Anand*

* Dept of Information Technology, GITAM University, Visakhapatnam,
*Dept of computer Science And Engineering, GITAM University, Visakhapatnam,

Abstract

The delivery of IP traffic through the internet depends on the complex interactions between thousands of autonomous systems (ASs) that exchange routing information using the Border Gateway Protocol (BGP). This paper investigates topological structure of the internet in terms of customer _provider and peer-peer relationship between AS's, as manifested in BGP routing policies. We describe a technique for inferring as relationships by exploring partial views of the AS graph available from different vantage points. Next we are apply the technique to a collection or ten BGP routing tables to infer the relationship between neighboring ASs .Based on these results , we analyze the hierarchical structure of the Internet and purpose a five-level classification of ASs. Our characterization differs from previous studies by focusing on commercial relationships between ASs rather than simply the connectivity between the nodes.

Keywords:

Computer Networks: Autonomous Systems from the view of multiple vantagepoints

Introduction:

TODAY's internet is divided into more than 10,000 Autonomous systems (ASs) that interact to coordinate the delivery of IP traffic. An AS typically falls under the administrative control of single institution, such as university, company or internet service provider (ISP).Neighboring ASs use the Border gate way protocol (BGP) to exchange information about how to reach the individual blocks of destination IP address (prefixes).

AS relationships and the associated routing polices have a profound influence on how traffic leaves through the internet. An understanding the structure of the internet in terms of these relationships facilitates a wide range of important applications. For example, consider a content distribution network (CDN) that can place replicas of the web site in data centers hosted by different Ass. The company can identify the best locations for its replicas.

As another example, consider a new regional ISP that want to connect to a small number of upstream providers. An accurate view of the AS topology and the relationships can help the ISP determine which ASs would provide the best connectivity to and from the rest of the internet.

In this work, we propose a technique for combining data from multiple vantage points in the internet to construct a more complete view of the topology and the AS relationships. Each vantage points offer a partial view of the internet topology as viewed from one source node. Due to the presence of complex routing polices, these partial views are not necessarily shortest path trees and may, in fact, include cycles. We generate a directed AS level graph from each vantage point and assign a rank to each AS based on its position. Then, each AS is represented by the vector that contains its rank form each of routing table dumps.

We infer the relationship between two ASs by comparing their vectors. Based on these relationships, we construct a new directed AS graph and examine the AS level hierarchy of the internet. We present a five-level classification of ASs with the top most level that consists of a rich set of peer–peer relationships between 20 so called tire-1 providers.

The work we describe is novel in three ways. First, we analyze AS path seen from multiple locations to from a more complete views of the graph. Second, rather than simply combining the data from the various vantage points, we propose a methodology for exploiting the uniqueness of view to infer the relationships between AS pairs. Third we characterize the hierarchy of Ass based on the

Manuscript received October 5, 2009
Manuscript revised October 20, 2009

commercial relationships, rather than simply the connectivity of the graph. We evaluate our technique on a collection of ten BGP routing tables and summarize characteristics of the AS relationships.

Problem formulation:

In this section, we present a brief overview of how AS relationships affect BGP export polices and formally define the type of relationship (ToR) problem. Then we discuss the practical challenges that arise in solving this problem and validating potential solution.

A. Type of Relationship problem:

The relationships between ASs arise from contracts that define pricing model and the exchange of traffic between domains. ASs typically have a provider-customer or peer-peer relationship.

In a provider-customer relationship, the customer is typically a smaller AS that pays a larger AS for access to the rest of the interne. The provider may in turn be a customer of an even larger AS. In a peer-to-peer relationships of the two peers are typically of comparable size and find it mutually advantageous to exchange traffic between their respective customers. These relationships translate directly into policies for exporting route advertisements via BGP sessions with neighboring Ass.

a. Exporting to a provider: In exchanging routing information with a provider, an AS can export its routes learned from other providers or peers.

b. Exporting to a peer: In exchanging routing information with a peer, an AS can export its routes and routes of the customers, but cannot export routs learned from other providers or peers.

c. Exporting to customers: An AS can export its routes, routes of its customers, but routes learned from other-providers and peer to its customers.

B. Practical challenges:

Indentifying the commercial relationships between Ass in challenging in practice. First, peer-peer relationships are difficult to classify. As pair may have different relationships for certain blocks of IP address. For example, an AS in Europe may be a customer of an AS in the US for some destinations and a peer for others. Router misconfiguration also causes violations in the export rules. For example, a customer may mistakenly export advertisements learned from one provider to another.

AS Ranking:

Our algorithm assigns a rank to each AS for each vantage point. Let X denote the source AS of a particular view of the AS graph and let P(X) denote the set of AS paths seen from X .Since each p ε P(X) consists of a sequence of notes starting with X, we construct a directed graph G routed at X from P(X). Let c(Gx)

denote the set of all vertices in Gx and let levels (Gx) ⊂

v(Gx) denote the leaves of the graph. We assign a ranking rank(u) to each vertex u ε v(Gx) by applying the reverse pruning algorithm in figure1 . At each edge the algorithm identifies the leaf node, assigns them a rank and remove these nodes (and their incident edges) from the graph. In the end of the remaining nodes (if any) form the connected component of the original graph G\: these nodes are assigned the same (highest) rank.

```
G = Gx;
r = 1;
while (leaves(G) = φ)
{
        for all u ∈ leaves(G)
               rank(u) = r;
        v  = v(G) − leaves(G);
        r = r + 1;
        G = Gv  ;
}
for all u ∈ v(G)
      set rank(u) = r;
```

Fig1: Reverse pruning algorithm on Graph Gx

Multiple vantage points:

If we continuing the pruning in figure2(a) the eventual leaf C will be inferred as a customer of D, even though the two ASs have a pear-pear relationship. Identifying the boundary points the uphill and downhill portions of the path is tricky. The structure of partial view of the AS graph depends on the position of the AS in the internet hierarchy. In the figure3, the boundary is between C and D (peer-peer relationship) not at E as suggested by this partial view in the figure2(a). Now consider the view from AS B in figure2(b) .This view confirms that A is a customer of D and F is a customer of E, however, the graph contradicts the previous view in that d is a customer of C. clearly D and C cannot be customers of each other. This contradiction suggests that the two ASs may have a peer-peer relationship.

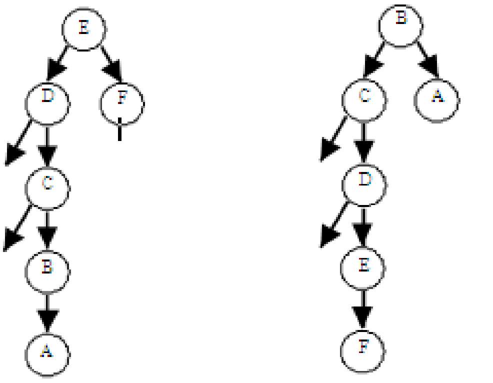

Fig2:(a)Partial view from AS E;(b) Partial view from AS B

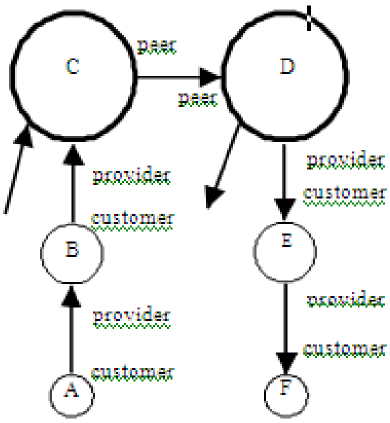

Fig 3: Path following Export rules

Inference rules of the ToR problem:

1. Inferring peer-peer relationships:

We use equivalence rule below to identify the peer-peer edges that are visible from many views. An AS relationship may not be visible from same partial views because some ASs may assign to low preference to path that traverse in this edge. We use the probalistic equivalence rule to find peering edges where the relationship between two ASs is not visible from many partial views.

- Equivalence two ASs i and j are said to be equivalent if e(i,j)>N/2. This rule considers two ASs that have the same rank in more than half the vantage points. In these ASs share an edge, they are likely to be peers.

- Probabilistic equivalence two ASs are probably 1/δ1 δ1 close to 1. We use this rule to infer peering relationships between ASs when visibility is poor across the partial views.

2. Inferring provider-customer relationships:

We use the dominance rule to determine if an edge between two ASs is a provider customer relationship because one AS tends to have a higher rank than the other in many of the partial views. Typically, in the graphs from the vantage point of J or its customers, it is the probable that rank (j) >rank (i) even if is a provider of j. to avoid an incorrect inference in such cases, we use the probabilistic dominance rule.

- Dominance an AS i is said to be dominate AS j if 1(j,j) >= N/2 and 1(j,i)=0. If dominates j, then we can infer that i is the provider of j, if the two ASs share an edge.

- Probabilistic dominance If for the high value of then i probably dominates j, and thus i is a provider of j. δ_0 should be greater thanδ_1. We use the value of 3 for δ_0 in our experiments

The orthogonal equivalence and the dominance rules infer peer-peer and provider customer relationships with a high degree of confidence. We apply those rules first in our inference algorithm, followed by the two probabilistic rules. Those AS relationships which are not infer using these have the values of max($\frac{l(i,j)}{l(j,i)}$,$\frac{l(j,i)}{l(i,j)}$) δ_1 between δ_0.

Experimental Results:

This section evaluates our inference techniques on a collection to ten publicly available BGP Routing tables. We classify our relationships between ASs and identify a small number of AS paths that of inconsistent with a relationship agreement. The most common anomalies seam to steam from resent acquisitions and mergers, suggest that some AS pairs have a sibling relationship.

Internet hierarchy:

In addition to relationship between AS pairs it is useful to identify the position of each AS in the internet hierarchy. Previous work has classified ASs based on the node degree: ASs with a large number of neighbors are placed above ASs with small node degree. However, a simple degree based approach may not capture the essence of the tiers in the hierarchy. Instead, we classify ASs based on the commercial relationships. Typically, a customer should be at a low level in the hierarchy than provider(s) we represent the AS topology as directed graph, where the direction of an edge indicates the type of relationship between two ASs. In our graph a provider-customer relationship between A and B is represented by a directed edge from A to B and a peering relationship between A and B is represented by two directed edges, one from A to B and other from B to A. in such a graph representation has also been independently proposed. An important difference between our approaches is the procedure used for determining the internet hierarchy. The work in maps the internet topology in to a strict hierarchy based on provider-customer edges while our classification also uses the distribution of peering links as identify the top levels of the hierarchy.

1. Customer and small regional ISPs

Customers are the easiest class of ASs that can be classified from these directed graph structure of AS topology. Customers are those sub networks which are origin and sinks of traffic and which do not carry any transit traffic. From the very definition of the direction of the edges in our graph we can infer the customer ASs to be the leaves of the directed graph. In a directed graph, a leaf is a node with out- degree 0. Since and undirected graph makes no distinction between out-degree and in-degree, customers with multiple providers would have a degree more than 1 and hence would not appear as leaves of the graph. Modeling the topology as a directed graph provides a more precise characterization of the bottom-most level in the AS hierarchy, in the directed graph constructed from the BGP dumps, 8898 of the 10915 ASs are leaf nodes. The rest of the graph contains just 18.5% of the ASs.

2. Dense core:

i). Identifying the dense core:

First we order vertices based on "greedy" notation on connectivity, following the heuristic in figure4. Let G represents the directed graph representation of the core let p(G) and E(G) represents the vertices and the edges of the

graph G. let d(x,y) for x ∈ v(G) denote the number of the

edges of the form(x,y) where z ∈Y. connectivity from a

node to a given set of nodes refer to the number of directed edges from the node to any of the nodes in the set. Assume that k of the N nodes are already ordered, for each of the remaining N-k nodes, we determine the connectivity to the k node and the pick the node with maximum connectivity as the(k+1). When multiple nodes have the same connectivity, with higher outer degree. In figure4, pos(x) denotes the position of the node X in the final ordering.

```
compute z ∈ v(G) with maximum out-
degree;
X = {z};
pos(z) = 1; r = 1;
while (X = v(G)) {
        compute y ∈ v(G) − X with max
        d(y, X)
            (selecting the y with the max
            out-degree)
        X = X ∪ {y};
        maxindegree(r) = d(y, X);
        r = r + 1;
        pos(y) = r;
}
```

Fig:4

ii). Transit core:

After removing the dense core, we notice the presence of other large national providers and hosting companies that have peering relationships with many of the ASs in the dense core. To identify these ASs, we define the notation of a transit core. Nodes in the transit core peer with each other and with ASs in the dense core, but they do not tend to peer with many out ASs. In our directed graph representation, these peering links are essentially the incoming directed edges from vertices outside this set to vertices within the set. We define such a set of edges to be in-way cutoff graph indicated by the given set of vertices. Using the property, we define the transit core as a smallest set of ASs containing the dense core which indicates a weak in-way cut. That is one having small number of edges compared to the total number of ASs in the transit core.

1. Identifying the transit core: given $X \subset v(G)$, let

$cut_{in}(X)$ denote the set of all edges of the form (y,z). We define a cut X of the vertex set v(G) to be a weak cut in $|cut_{in}(X)| > X/2$. The problem of finding weak cuts in the graph is NP-complete and no good approximation algorithms are known for that problem. Given the transit core is a superset of the dense core and that the dense core is derived by the greedy ordering, we apply some ordering to find the transit core as was used to find the dense core. A natural way of using this ordering to find the transit core is to find the smallest value of k such that $|cut_{in}(X_k)|<k/2$. Surprisingly we find that the value of k at which |cut in $(X_k)|< k/2$ also satisfy the property that conn(k+1)=1. This means that no two edges in the $cut_{in}(X_k)$ have the same source. A weak cut also means that more than 50% of the ASs in Xk do not have any peering relationship with any of the ASs in $v(G)-X_k$. Hence by this definition, X_k should intended contain all the transit providers.

2. Properties of the transit core: Applying the in-way cut algorithm to our graph, we discover a transit core consisting of 129 ASs. These 129 ASs have 183 peering links with the ASs in the dense core. We found many of the top providers in Europe and Asia to be present in our transit core.

iii). Outer core:

We classify all the remaining ASs in the core as the outer core. The members of the outer core typically represent regional ISPs which have a few customer ASs and a few peering relationships with other such regional ISPs. The outer core consists of 897 ASs that have 29 peering sessions with ASs in the dense core and 145 peering sessions with ASs in the transit core. We observed that many members of our outer core are regional ISPs.

iv). Summary:

Table1, shows summarized number of ASs at each level in the hierarchy dense core(level 0), transit core(level 1), outer core(level 2), small regional ISPs(level 3) and customer (level 4). Table2, summarizes the connectivity between various levels in the AS hierarchy . each number in the table is the total number of edges from level 0 to level 1. The table shows several key properties of the internet topology.

- The ASs in dense coer are veru well connected.
- As we move from the dense core towards customer, the inter-level and intra-level connectivity graphs significantly.

The large number of customers ASs have their providers distributed across levels is not strictly hierarchal.

- The number of edges within the outer core is less than the total number of vertices in the outer core. This indicates the presence of multiple disconnected graphs of ASs in the outer core;

Level	# of Ass
Dense core(0)	20
Transit core(1)	129
Outer core(2)	897
Small regional ISPs	971
Customers	8898

ASs in the different groups communicates via ASs in the dense core and transit core.

Distribution of ASs in the Hierarchy(Table1)

Interconnectivity Across levels(Table2)

The graph in figure5, explores the relationship between node degree and the levels in the hierarchy. We define node degree as number of neighboring ASs without regard to the relationships. The graph plats the cumulative distribution of node degree as a logarithmic scale. In general level 0 and 1 ASs have high degree, and level3 and level4 ASs tends to have low degree. However, this is not universally true. Some customers at level 4 have a large number of upstream providers, and some ASs in the dense core at level 0 have a relatively small number of neighbours. A hierarchy based solely on degree distribution would not be able to make this distribution.

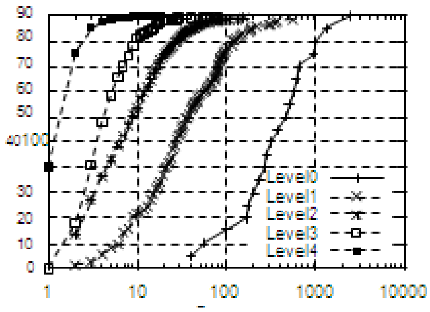

Fig5: Cumulative distribution of AS degree by level

Summary:

The relationships between ASs has a significant impact on the flow of traffic through the internet. Our work makes two important contribution toward understanding the structure of the internet in terms of these relationships:

- An algorithm for inferring AS relationships from partial views of the AS graph from different vantage points.
- A mechanism for dividing the Internet hierarchy into levels based on AS relationships and node connectivity.

The complete structure of the internet is unknown and difficult if not impossible to obtain. Our approach is comprised of many heuristics with certain limitations.

- We draw our inferences based on only ten vantage points available. Ideally we would have a larger collection of routing tables from more diverse vantage points, including small customers.

Level	0	1	2	3	4
0	312	626	1091	958	6732
1	183	850	1413	665	3373
2	29	145	1600	543	3752
3	0	0	0	212	2400

- We treat the route views routing tables as a view from a single AS.

Future work:

We plan to extract a separate view for each AS participating in the route views project. Multiple ASs may fall under the administrative control of a single institution, due to historical artifacts and market forces. We plan to extend our methodology to incorporate more complex routing policies that are captured by the traditional customer-provider and peer-peer relationship.

Despite of these limitations, we have shown that our approach provides a detailed view of the Internet topology in terms of the relationships between ASs.

References:

[1] John W. Stewart, BGP4: Inter-Domain Routing in the Internet, Addison-Wesley, 1998.

[2] Sam Halabi and Danny McPherson, Internet Routing Architectures, Cisco Press, second edition, 2001.

[3] Geoff Huston,“Interconnection, peering, and settlements,” in Proc.INET, June 1999.

[4] Alaettinoglu, “Scalable router configuration for the Internet,” in Proc. IEEE IC3N, October 1996.

[5] B. Krishnamurthy and J. Wang, “On network-aware clustering of Web clients,” in Proc. ACM SIGCOMM, August/September 2000.

[6] R. Govindan and H. Tangmunarunkit, “Heuristics for Internet map discovery,” in Proc. IEEE INFOCOM, March 2000.

[7] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger,"Network topologies, power laws, and hierarchy," Tech. Rep. 01-746, Computer Science, University of Southern California, June 2001.

[8] "CAIDA Web Site," http://www.caida.org.

[9] R. Govindan and A. Reddy, "An analysis of interdomain topology and route stability," Proc. IEEE INFOCOM, April 1997.

[10] L. Gao, "On inferring automonous system relationships in the Internet,"IEEE/ACM Trans. Networking, vol. 9, no. 6, December 2001.

[11] "University of Oregon RouteViews project,"http://www.routeviews.org/.

[12] "BGP tables from the University of Oregon RouteViews Project,"http://moat.nlanr.net/AS/Data/.

[13] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos, "On power- law relationships of the Internet topology," in Proc. ACM SIGCOMM, August/September 1999, pp. 251–262.

[14] Jin, Q. Chen, and S. Jamin, "Inet: Internet topology generator," Tech. Rep. CSE-TR-433-00, U. Michigan, September 2000.

[15] Z. Ge, D. Figueiredo, S. Jaiwal, and L. Gao, "On the hierarchical structure of the logical Internet graph," in Proc. SPIE ITCOM, August 2001.

[16] Lixin Gao, Timothy G. Griffin, and Jennifer Rexford, "Inherently safe backup routing with BGP," in Proc. IEEE INFOCOM, April 2001.

[17] L. Gao and J. Rexford, "Stable Internet routing without global coordination," IEEE/ACM Trans. Networking, vol. 9, no. 6, pp. 681–692, December 2001.

[18] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "On inferring AS-level connectivity from BGP routing tables," 2001.

[19] "RADB Whois Server," whois.radb.net.

[20] "IXC Buys Three ISP's,"

Mr. G.Srinivasa Rao, M.Tech,(Ph.D)., Sr.Asst.Professor. He has Submitted Ph.D thesis in M.U., Over 7 Years of teaching experience with GITAM University, handled courses for B.Tech, M.Tech. Research areas include Computer Networks And Data Communications. published 4 papers in various National and International Conferences and Journals.

Mr. G.Venkateswara Rao Received M.Sc, M.Phil, and M.Tech in Computer Science Engineering from Andhra University during 1993, 1996 and 1999 respectively. He is working as faculty member in GITAM University in the Dept. of Information Technology from last 10 years. His research interest includes Mobile computing and Networks. He was awarded UGC Major Research Project in 2007 and published 8 papers in various National and International Conferences and Journals.

Mr. M.Venkateswara Rao., M.Sc, M.Phil (Ph.D)l, and M.Tech in Computer Science Engineering from Andhra University. Published 4 papers in various National and International Conferences and Journals.

MR. G.Appa Rao., M.Tech., M.B.A., (Ph.D)., in computer science and Engineering form Andhra Universiy. He has Submitted Ph.D thesis in M.U., Over 12 Years of teaching experience with GITAM University, handled courses for B.Tech, M.Tech. Research areas include Data Mining and AI. Published 8papers in various National and International Conferences and Journals.

Mr. D.Kishore Kumar M.Tech in computer science and engineering. Over 4 Years of teaching experience with GITAM University, handled courses for B.Tech.

Mr. K.Krishna Chaitanya M.Sc., M.Tech(IT) from GITAM University

Smt. S.Anuradha., M.Tech., (Ph.D). ,in computer science and engineering. Over 4 Years of teaching experience with GITAM University, handled courses for B.Tech and M.Tech .

Mr. D.Veerabadra Rao., M.Tech., in computer science and engineering. Over 4 Years of teaching experience with GITAM University, handled courses for B.Tech and M.Tech .

Mr. P.Surya Chandra., M.Tech in computer science and engineering. Over 3 Years of teaching experience with GITAM University, handled courses for B.Tech and M.Tech.

Mr. T.Anand ., M.Tech(IT) from GITAM University