A Social Network Approach for the Ranking of the Autonomous Systems of the Internet

Fabrice Clérot [(1)], Quang Nguyen [(2)]

(1) France Télécom, Recherche & Développement, TECH-SUSI-TSI, 2 avenue Pierre Marzin, 22307 Lannion Cedex, France

(2) France Télécom, Recherche & Développement, TECH-LEI-MOT, 38 rue du Général Leclerc, 92794 Issy-les-Moulineaux Cedex, France

Abstract

The worldwide scale transit of information flows in the Internet is governed by trade agreements between autonomous systems; these agreements are translated into routing policies by the Border Gateway Protocol (BGP). The negotiation of these trade agreements implicitly relies on a hierarchy of the autonomous systems and the relative position of two systems leads to an agreement of the customer-provider type (one of the systems, the provider, is ranked higher than the other, the client, and the client pays the provider for the transit of information flows) or to a no cost agreement of the "peering" type (two service providers that agree to exchange traffic between their respective customers) when both systems consider their rankings to be equivalent.
In spite of its importance, there is no official hierarchy of the Internet (the commercial clauses of the agreements between autonomous systems are not necessarily public, it is usually a bilateral arrangement) nor a consensus on the way of establishing such a hierarchy. We propose a simple heuristics inspired of the concept of "spectral centrality" borrowed from the social networks analysis to analyze the relative positions of the autonomous systems of the Internet starting from their connectivity information only.

Introduction

The worldwide scale transit of information flows in the Internet is governed by trade agreements between autonomous systems. The negotiation of these trade agreements implicitly relies on a hierarchy of the autonomous systems (AS) and the relative position of two systems leads to an agreement of the "customer-provider" type (one of the systems, the provider, is ranked higher than the other, the client, and the client bpays the provider for the transit of information flows) or to an agreement of the "peering" type (two service providers that agree to exchange traffic between their respective customers) when both systems consider their rankings to be equivalent.
These agreements are translated into routing policies by the Border Gateway Protocol (BGP). Thus, the establishment of the routing paths on a worldwide scale obeys the rules of economic effectiveness deduced from a relative ranking of autonomous systems (a route cannot, for example, "go down" from a provider to his client and then "go up" towards another provider: which client would agree to pay and carry the traffic of its providers?) [GSW02], [GW02]. Such rules are quite different from the engineering-based rules which govern the routing inside the autonomous systems.
In spite of the importance of this relative ranking of the autonomous systems, for the understanding of the large scale routing behaviour in the Internet and for the autonomous systems themselves for negotiation purposes, there is not publicly available ranking reference (the commercial agreements between autonomous systems are not necessarily public) nor even of consensus on the means of establishing such a ranking (each Internet Service Providers have their proper decision-making rule).

Importance of an autonomous system and spectral centrality

Importance of an autonomous system and centrality in the social networks

The concept of importance of an autonomous system *AS1* within the global routing architecture of the Internet relies on its capacity to mediate communication between other autonomous

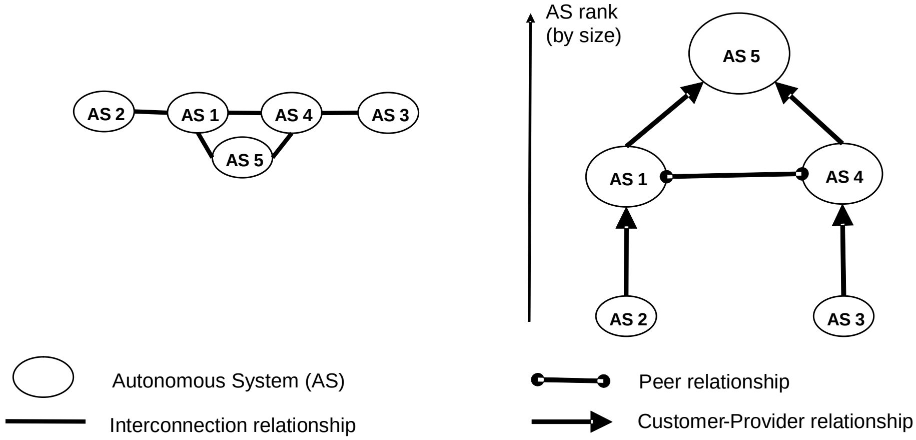

(a) Physical relationship (b) Logical relationship

Fig 1: *Interconnection graph at the autonomous system level*

systems which do not have direct connectivity (*AS2* and *AS3*); this capacity does not necessarily imply a "direct" mediation AS2-AS1-AS3 between the autonomous systems but can also rely on an "indirect" mediation by the means of the new connectivities offered through *AS1*: AS2-AS1-AS4-AS3 where it is understood that *AS2* does not have connectivity with *AS4*, nor *AS3* with *AS1* (see Figure 1).
In this later example, the importance of *AS1* depends on the importance of *AS4*: being connected to *AS4*, *AS1* "inherits " a part of the importance of *AS4* as it can provide *AS2* with a wider connectivity (the systems like *AS3* which *AS1* does not reach directly but to which *AS4* has access); the same is true for *AS4* which also inherits a part of the importance of *AS1*.

This concept of building the importance of a node in a graph by "inheriting" part of the importance of its neighbours has been thoroughly studied in the realm of social networks analysis and we shall borrow the concept of "spectral centrality" (introduced below) from this realm in order to establish a ranking between autonomous systems of the Internet, starting from their connectivity graph only.
This connectivity graph can be easily inferred from BGP routing tables [CCGSW2002]. We remark that, by construction, this connectivity graph only includes links which are part of at least one BGP path.

Spectral centrality in the social networks

We follow a generalization of the concept of spectral centrality for the asymmetrical graphs; we give below a very short presentation of this concept and refer to [BL2001] for a comprehensive description and motivation.

The vector X of the node centralities in a graph (given by its weighted asymmetrical adjacency matrix A) has two origins of different nature, an intrinsic term E which depends only on the node taken in isolation and a term coming from the effect of network (inheritance of the importance of the neighbours). This results in a fixed-point equation $X=\alpha AX+E$, where α must approach (by lower values) the reverse of the principal eigenvalue of A so that the result obtained by this method is consistent with the result obtained by the usual spectral method in the case of undirected graphs [BL2001].

Technically, the solution is obtained by iteration until convergence of $X_{i+1}=\alpha AX_i+E$, the direct inversion $X=(I-\alpha A)^{-1}E$ leading to a very large non-sparse matrix.

This method reveals explicitly the contributions to the centrality and makes it possible to choose the intrinsic importance attached to each node (the vector E).

The default choice is $E_i=1$ for all nodes but other choices are possible; the only condition imposed on E is that the values should independent of the effects of network (for example, it would not be consistent to use the node degree as a measure of its "intrinsic" importance).

Calculation of the spectral centrality from the connectivity graph

Data sources and inference results

The real Internet topology is unknown; however, since BGP, the inter-autonomous system routing protocol in the Internet, is a path-vector protocol (advertises sequence of autonomous system numbers to the destination network), the graph of autonomous system connectivity can be easily inferred from BGP routing tables. In this study, we used data from the Oregon Route Views server [RV] and from about twenty Looking Glass sites [TR] to obtain BGP tables, and thus, to built the adjacency matrix containing autonomous system connectivity [DMN2004]. The autonomous system graph is measured on September 1st, 2004. It has 17886 vertices and 42123 edges. Table 1 shows some connectivity properties of the graph.

Note that, the degree is a good parameter to know the *physical connectivity* of a node, but it is not enough to provide information about AS-level *reachabitity*. This later also depends of *logical relationship* between autonomous systems. As introduced in [G2000], three main logical types can be considered: customer-provider, peering and sibling. The first correspond to transit service offered by a supplier AS to another AS (the customer pays its provider for that); the second type corresponds to a peering agreement between two ASs (whereby traffic is exchanged between their respective customers free of charge); the third is a special case of mutual transit between two ASs. Those types of relationships have a huge impact on the AS-level hierarchy (because inter-AS routing policies depending on the economic relationships determine which routes an AS could be used). Typically, a customer should be at a lower rank in the hierarchy than its providers; and the bigger a provider, the higher it negotiates settlement deals (particularly peering agreements) with other providers.

To assert peering type of relationships from BGP routing tables, we follow the method proposed in [G2000]. Table 2 summarizes the relationships inferred for the AS-level graph.

Table 1: *Inferred topological properties of the graph*

Number of nodes (ASs)	17886
Number of edges (undirected AS links)	42123
Average degree	4.9
Maximum degree	2413
Percentage of nodes with degree ≤ 2.	72.7 %
Percentage of nodes with degree ≥ 50.	0.9 %

Table 2: *Inferred relationships of the graph*

	No. of AS pairs	Percentage
Customer-Provider	36472	86,5 %
Peer	4996	11,8 %
Sibling	655	1,5 %

Centrality based on the degree

The concept of centrality based on the degree is the simplest which includes an effect of network; the experts estimate that this classification over-estimates the importance of the autonomous systems having many "final" neighbours (nodes with one neighbour only). This is consistent with the distinction between connectivity and reachability. The 20 top AS according to this ranking can be found in the first two columns of Table 3.

It should also be noticed that this degree-based centrality index does not allow to investigate the consequence of adding or removing a link between two AS beyond the trivial consequence on the connectivity of both AS. This is clearly a shortcoming of this approach since we aim at investigating the consequences of such addition or removal on the reachability of a given AS.

Spectral centrality from the graph of connectivity

The graph of connectivity between autonomous systems such as deduced from publicly available routing information is obviously a symmetrical graph.

Table 3: *Top 20 AS according to the centrality index. Centrality indexes are normalised to their maximum for easier comparison.*

DEGREE		PROPOSED HEURISTICS p=1,0		SYM. CONNECTIVITY GRAPH	
AS	Centrality	AS	Centrality	AS	Centrality
701	1,00	701	1,00	4513	1,00
1239	0,75	1239	0,79	6461	0,93
7018	0,73	7018	0,74	3303	0,92
3356	0,47	3356	0,59	3356	0,92
209	0,46	209	0,48	701	0,83
174	0,29	174	0,41	4589	0,79
8220	0,27	6461	0,38	1239	0,78
3549	0,26	4513	0,36	174	0,75
2914	0,26	8220	0,36	13237	0,75
6461	0,24	2914	0,35	8220	0,69
702	0,23	3549	0,35	8210	0,67
4513	0,22	3303	0,34	12956	0,62
7132	0,21	702	0,28	13129	0,62
3303	0,21	4589	0,28	6939	0,61
4323	0,19	13237	0,26	7018	0,61
4589	0,18	7132	0,24	13030	0,61
3561	0,17	4323	0,23	6320	0,57
13237	0,16	3561	0,23	3491	0,56
2828	0,13	6939	0,21	286	0,56
3786	0,13	2516	0,20	12859	0,54

When the traditional notion of centrality for the symmetrical graphs is applied to this graph, the following classification is obtained for the first ten systems: Globix (4513), Abovenet (6461), Swisscom (3303), Level 3 (3356), Uunet (701), Easynet (4589), Sprint (1239), Cogent (174), Lambdanet (13237), Colt (8220); see the last two columns in Table 3 for more details.

This ranking is clearly unsatisfactory for the experts of the field; it is common knowledge that Uunet (701) is the most important autonomous system of the Internet and one does not even find ATT Worldnet (7018) in these ten first systems. A similar approach is proposed in [GMZ2003].

A simple heuristics to direct the connectivity graph according to the ranking of the autonomous systems

The two approaches above have an obvious weakness: they do not take into account the difference in centrality between nodes so as to direct the graph according to client/provider relationships.

Since the graph is undirected and the ranking unknown at the beginning, the analysis should rely on an iterative heuristics allowing to progressively introduce an asymmetry in the graph in agreement with the asymmetry of the ranking. Such a simple heuristics is proposed below:

Initialization:

- the graph of interconnection is regarded as a balanced directed graph, each undirected edge giving rise to two weighted directed edges in opposite directions and with weight ½;
- all nodes are given the centrality score;

Calculation of the centrality:

- starting from a centrality score calculated at the preceding stage, one modifies the weights of the edges by reinforcing the asymmetry of the relationship between

two nodes according to their difference in centrality;
- computation of a centrality score starting from this new asymmetrical adjacency matrix.

The two steps above are iterated until convergence of the centrality score.

In the first step, several modifications of the weighting of the edge w_{ab}^n between the nodes *a* and *b* according to their centrality scores c_a^n and c_b^n at the preceding iteration *n* are possible; we choose a modification in the form:

$$w_{ab}^{n+1} = (1-\beta) w_{ab}^n + \beta \Delta_{ab}^n$$

with

$$\Delta_{ab}^n = \frac{1}{2}\left[1 + sign\left(c_b^n - c_a^n\right)\left|\frac{c_b^n - c_a^n}{c_b^n + c_a^n}\right|^p\right]$$

A low value of the parameter β allows a progressive adaptation of the orientation of the edges according to the centrality structure. In the experiments reported below, β is set to 10^{-3}.

The parameter p allows to investigate the importance attached to a weak variation of centrality between nodes: a p>1 value gives a small importance to weak variations of centrality, allowing to explore the concept of "peering" between nodes.

Experimental results

Ranking results

With p=1, the classification obtained with convergence for the first ten systems is as follows: Uunet (701), Sprint (1239), ATT WorldNet (7018), Level 3 (3356), Qwest (209), Cogent (174) , Abovenet (6461), Globix (4513), Colt (8220), Verio (2914). This classification is in agreement with the knowledge of the experts of the field.

Figure 2(a) shows the variation of this ranking according to the ranking deduced from the degree. The correlation is strong, at least for the best classified autonomous systems, but significant differences are nevertheless obvious. It must be noticed that the great variations of ranking observed for the high rank-systems are hardly significant, all these systems being practically ranked at the same level with centrality scores very close to 1.

Figure 2(b) shows that changing the parameter p from 1 to 3 does not significantly modify the ranking, at least for the best classified systems.

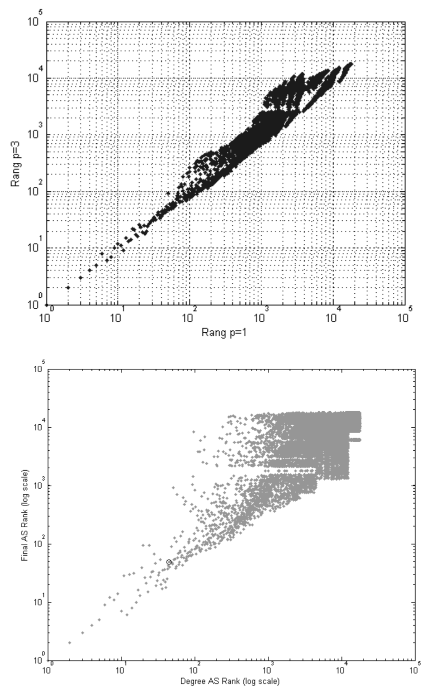

Fig 2:
(top) ranking obtained by the heuristics vs ranking deduced from the degree;
(bottom) ranking obtained with p=3 vs ranking obtained with p=1

Figures 3(a) and 3(b) show the weights of the outgoing edges of the OPENTRANSIT autonomous system (5511, France Telecom Worlwide IP Backbone), according to the rank of the other systems. A value significantly above (resp. below) 0.5 indicates that OPENTRANSIT is in the position of a client (resp. provider). A zero value indicates that there is no connectivity between OPENTRANSIT and the system considered.

With p=3 the weight values exhibit a "flat" area around 0.5 corresponding to a range of autonomous systems which are in a quasi-symmetrical relation with OPENTRANSIT and could therefore tentatively be identified as "peers" of OPENTRANSIT, at least in the sense of the proposed ranking. This is detailed below.

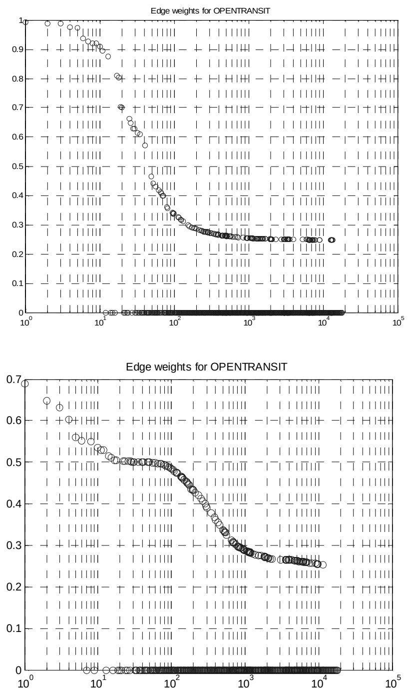

Fig 3:
Weights of the outgoing edges of OPENTRANSIT (top) for p=1; (bottom) for p=3

Peering versus non-peering classification of AS relationships

A crucial aspect of ranking lies in the ability to infer the peering relationships between the important autonomous systems. We show below that a proper setting of the parameter p allows a very good performance for the peer versus non-peer classification of a connection. We restrict our analysis to connections with at least an autonomous system with a centrality larger than 2. The peer versus non-peer labelling is given by an analysis of the real BGP routes as in [G2000]. Peering and sibling relationships form the "peering" class; client to provider and provider to client form the "non-peering" class. The set of connections is divided in two equal parts for training and testing purposes.

For each value of p, we report above the performance respectively obtained on the training and on the test set for the best threshold t from the simple rule:

- if $|w_{ij}-1/2|<t$, nodes i and j are peers
- if $|w_{ij}-1/2|>t$, nodes i and j are non-peers

The best threshold is determined from the training data only.

The performance index is defined as the half sum of the correct classification rate in the peer and non-peer classes; this index of performance is chosen so as to give the same importance to the rather small class of peering relationships and to the very large class of non-peering relationships.

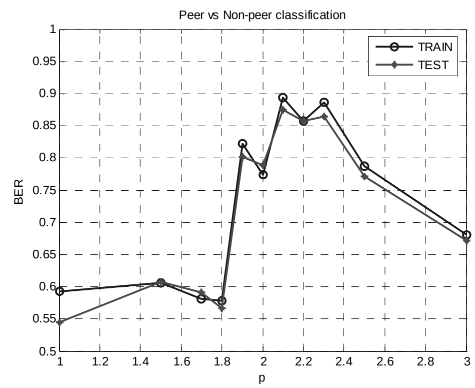

Fig 4: *Peer versus Non-peer classification performance as a function of p*

As can be seen on Figure 4, the results are quite close on the test set and on the training set; this shows that the classifier designed above has a good generalization capacity.

The best result is obtained by setting the parameter p to p = 2.1, with an optimal threshold t set at t = 0.34. The performance index is about 0.87 with a good classification rate of 0.96 for the peering class and of 0.77 for the non-peering class on the test set.

Table 4 gives the top 20 AS according to the proposed heuristics for p=1 and p=2.2; both rankings are quite close.

Table 4: *Top 20 AS for p=1 et p=2.2 with the proposed heuristics*
Centrality indexes are normalised to their maximum for easier comparison.

PROPOSED HEURISTICS p=1,0		PROPOSED HEURISTICS p=2,2	
AS	Centrality	AS	Centrality
701	1,00	701	1,00
1239	0,79	1239	0,80
7018	0,74	7018	0,74
3356	0,59	3356	0,61
209	0,48	209	0,48
174	0,41	174	0,43
6461	0,38	6461	0,43
4513	0,36	4513	0,42
8220	0,36	3303	0,39
2914	0,35	8220	0,38
3549	0,35	3549	0,36
3303	0,34	2914	0,36
702	0,28	4589	0,31
4589	0,28	13237	0,29
13237	0,26	702	0,29
7132	0,24	3561	0,24
4323	0,23	6939	0,23
3561	0,23	7132	0,23
6939	0,21	4323	0,23
2516	0,20	12956	0,23

Conclusion

The approach suggested in this communication relies on the spectral centrality concept as introduced in the social network analysis area and identifies the centrality deduced from the connectivity graph to the importance of the autonomous system.

Starting from an undirected (connectivity) graph, the proposed heuristics progressively induces an orientation of the graph in agreement with the asymmetry of the current centrality scores. Centrality scores are calculated according to the spectral centrality for asymmetric graphs.

The ranking results are in good agreement with expert knowledge. The proposed method does not rely on actually observed BGP paths and therefore allows significant practical extensions such as

1. simulating the consequences of the addition or the withdrawal of a connection;
2. studying the different contributions to the importance of any autonomous system by the simple analysis of $X_\infty=\alpha AX_\infty+E$ at convergence;
3. setting the intrinsic importance E_s attached to each autonomous system according to the problem, for instance according to their strategic importance or to a more objective criteria as the number of addresses which they can reach directly.

This approach also provides a weighting of the edges and the peering versus non-peering classification of the connections has been addressed using this weighting. A good agreement with a classification method based on observed BGP paths has been obtained with a proper choice of the parameters.

The ranking results of the proposed heuristics are also quite close from the simple degree-based centrality. However, it should be noted that point (3.) above and the peering versus non-peering classification cannot be addressed from the degree-based centrality approach.

More generally, this study also illustrates how concepts introduced in the area of social network analysis can be successfully applied to man-made complex systems.

Acknowledgements

We thank Mickael Meulle and Michel Decima for sharing their code *SpyBGP* (a tool for extracting AS paths from BGP table dump) with us, and for their help and assistance with our experiments and data collection.

References

[BL2001] P. Bonacich, P. Lloyd, "Eigenvector-like Measures of Centrality for Asymmetric Relations", Social Networks, 23(3), July 2001

[CCGJSW2002] Q. Chen, H, Chang, R. Govindan, S.Jamin, S. Shenker, W. Willinger, "The Origin of Power Laws in Internet Topologies Revisited", in Proc. INFOCOM, 2002

[DMN2004] M. Decima, M. Meulle, Q. Nguyen, "SpyBGP: Mapping the Internet from BGP routing table", Technical Report, France Telecom, 2004

[GMZ2003] C Gkantsidis, M. Mihail, E Zegura, "Spectral Analysis of Internet Topologies", in Proc. Infocom IEEE, 2003

[G2000] L.Gao, "On inferring Autonomous System relationships in the Internet", in Proc. Total IEEE Internet, 2000

[GSW2002] T. G. Griffin, F. B. Shepherd, G. Wilfong, "The stable paths problem and interdomain routing". IEEE/ACM Transactions on Networking, 10.2:232–243, April 2002

[GW2002] L. Gao, F. Wang, "The extent of AS path inflation by routing policies", in Proc. IEEE Global Internet Symposium, 2002

[RV] The Oregon Route Views server connects to several operational routers within the Internet for the purpose of collecting BGP routing tables. http://www.routeviews.org/

[SARK2001] L. Subramanian, S. Agarwal, J. Rexford, RH. Katz, "Characterizing the Internet Hierarchy from Multiple Vantage Points", in Proc. IEEE INFOCOM, 2002

[TR] The Looking Glass sites are maintained by individual Internet Service Providers to help troubleshoot Internet-wide routing problems. http://www.traceroute/org