Inter-Domain Access Volume Model: Ranking Autonomous Systems

Yixuan Wang, Ye Wang, Maoke Chen, and Xing Li

Network Research Center, Tsinghua University.
`{wangyx,sando}@ns.6test.edu.cn,{mk,xing}@cernet.edu.cn`

Abstract. Exploring topological structure at Autonomous System (AS) level is indispensable for understanding most issues in Internet services. Previous models of AS graph involve address or connectivity information separately, and neither of these information can serve as a comprehensive metric for evaluating an AS's contribution to the global routing. In this paper, we propose a new model for AS ranking named IDAV (Inter-Domain Access Volume). IDAV introduces the quantity of routed addresses into AS graph, and enriches the methodology of Internet marco structure inference. In IDAV, the magnitude of AS is measured with the primary eigenvector of the access volume matrix (Carriage Matrix) for AS graph. We construct the AS graph by parsing the forwarding information bases in border gateways. The computation, compared with previous approaches, demonstrates that IDAV model results in more accurate AS rankings. We believe the IDAV model is promisingly useful for studying inter-domain routing and Internet service behavior.

Key words: Internet topology, Inter-domain routing, AS ranking

1 Introduction

Internet connects thousands of Autonomous Systems[1] (ASs) operated by many different administrative domains. Each domain maintains one or more ASs and hides intra-domain information from others. Macro observation of almost all the fundamental issues of Internet services, such as inter-ISP settlements, troubleshooting, and traffic engineering, strongly relies on the knowledge of Internet structure at AS granularity. Since AS level network topology is determined by inter-domain routing protocol, namely BGP[2] in today's Internet, study on BGP tables is the general way of exploring AS graph. Forerunners have made great efforts to gather BGP routing information. The RouteView project[3] periodically dumps BGP FIBs from multiple vantage ASs. Internet Routing Registries (IRR)[4] provide a database facility for sharing inter-domain routing policies.

With help of these resources, people are trying to get a comprehensive and accurate AS graph. Ge[5], Vazquez[6] and Siganos[7] examined the hierarchical Internet topology by observing the power-law[8] distribution of AS degrees. Policy-based AS graph was first studied by Gao[9]. She classified the AS relationship into 4 categories, and discovered the valley-free principle to infer AS relationships from BGP routing table. Subramanian[10], Battista[11] and Dimitropoulos[12][13] brought in more elaborate methods and extensive resources to solve the AS relationship problem. However, Internet structure is much more complicated than we have already captured. Besides degrees and edge-types, people continue adding other parameters to strengthen AS graph. There are many choices, such as address space issued by AS, customer cones affiliated to AS, etc. This initiated a new question: which parameter is better?

We believe AS ranking is a fundamental benchmark of modeling Internet topology with augmenting AS graph. By AS ranking, people characterize each AS's importance in Internet as a quantitive weight. There have been quite a few approaches to rank the major backbone ASs, such as CAIDA[14], FixedOrbit[15], Renesys[16], etc. Different metrics can result in different ranking criteria, any of which might be reasonable. However, the common sense of AS ranking is that carrying more traffic always equals ranking higher, because the main role of an AS is a packet-carrier. Unfortunately, the impossibility of measuring all the traffic between each pair of ASs has prevented such kind of precise ranking. Overwhelming majority of known AS rankings employ the degree-based or relationship-based AS graphs due to the lack of new theories.

In this paper, we originate a new model to evaluate the contribution of AS to Internet routing: Inter-Domain Access Volume (IDAV). Engineering experience tells us inter-domain routing policy is designed and implemented per prefix. We define *access volume* as the quantity of routed addresses propagated between neighbor ASs. This is a counter indicating how much service one AS is providing for the other. Theoretically, we rank the ASs by the primary eigenvector of the access volume matrix for AS graph. Then in practice, we construct a real AS graph, compute AS rankings, and verify the results from several different sources.

To the best knowledge of the authors, this is the first approach to model and measure Internet structure by investigating both connectivity and routed prefixes information. The rest part of this paper is organized as follows: section 2 presents the theories of IDAV model; section 3 tells the computation methods; results inferred by our model and verification of them is shown in section 4; finally we summarize the paper and briefly specify future works in progress.

2 IDAV model

2.1 Definitions

Definition 1 (AS Graph) *An AS graph is a directed graph $G(V, E)$, where V is the set of ASs and E is the set of directed edges between two endpoint ASs. It may contain loops but must not contain multiple edges.*

The reason an AS graph may contain loops but no multiple edges is that we will add augmenting information to edges denoting the number of addresses routed. An AS can route its own addresses for itself, and if there are more than one pairs of routers connecting 2 ASs, their contribution to routing is counted together.

Definition 2 (Carry & Transit) *In inter-domain routing, if* ASy *reaches prefix* $N1$ *through* ASx*, we say that* ASx *carries* $N1$ *for* ASy*, written as* $N1 \in ASx \xrightarrow{c} ASy$*. If multiple prefixes are carried, they could be bracketed into comma-separated lists, like* $\{N1, N2\}$*. Since there is no need for an* AS *to access itself through other* ASs*, we think every* AS *carries its own prefixes for itself.*
ASx *transits prefix* $N1$ *for* ASy*, if and only if: (1) x, y and z are 3 different ASs; (2)* ASy *carries prefix* $N1$ *for* ASx*; (3)* ASx *carries prefix* $N1$ *or* $N1$*'s less-specific for* ASz*. Written as* $N1 \in ASx \xrightarrow{t} ASy$*.*

The prefix $N1$ is presented in IPv4 (or IPv6) CIDR[17] format. In Internet routing, the definition of transit implies ASx redistributes prefixes received from ASy to ASz, so it complies with the common notion "transit AS".

Definition 3 (Access Volume) *The access volume of an edge* $ASx \to ASy$ *in the AS graph is a numeric parameter that equals the total number of unique addresses in the prefixes carried by* ASx *for* ASy*. Written as* $C(ASx, ASy)$*.*

As previously defined, ASx "carry" prefix $N1$ for ASy not only means ASx announces route $N1$ to ASy, but also requires ASy does reach network $N1$ through ASx, e.g. ASy must install route $N1$ into its own forwarding table. We shall use each AS's FIB to calculate access volumes, which implies that $C(ASx, ASy)$ is much less than the total number of addresses that ASx announces to ASy.
Duplicated addresses should be counted only once, e.g., if $AS1 \xrightarrow{c} AS2 = \{1.0.0.0/8, 1.1.0.0/16, 2.0.0.0/8\}$, then $C(AS1, AS2) = 2 \times 2^{24}$, not $2 \times 2^{24} + 2^{16}$.

Definition 4 (IDAV model) *The Inter-domain Access Volume (IDAV) model is a weighted AS graph* $G(V, E, W)$*, where* $G(V, E)$ *is an AS graph as previously defined, and* W *is the set of numerical weight on each edge, suffices that* $W(ASx \to ASy) = C(ASx, ASy)$*.*

Definition 5 (CM) *The Carriage Matrix (CM) of a given IDAV model is an n-square matrix* M*, where n is the number of nodes in the model, and* $m(i, j)$ *equals* $C(ASi, ASj)$*, the access volume on each corresponding edge. If edge* $ASi \to ASj$ *does not exist, or that* $C(ASi, ASj) = 0$*, then* $m(i, j) = 0$*.*

In the case that $i = j$, the diagonal element $m(i, i)$ equals the total addresses owned by ASi itself, complying with the definitions of carry and access volume. This mathematical description creates a bijective mapping between IDAV model and CM: given a CM, we can draw the IDAV model, and vice versa.

2.2 AS ranking

Unlike previous models, we use virtual traffic throughput as the ranking metric. The problem of AS ranking is: *given an IDAV model denoted by carriage matrix* M*, find an n-by-1 vector* R*, whose elements are virtual traffic flows* $r_k, 1 \leq k \leq n$*, sufficing to that* $\forall i, j, 1 \leq i, j \leq n$*, if* $r_i \leq (\geq) r_j$*, the real traffic throughput of ASi is less (more) than that of ASj.* Moreover, besides relatively reflecting real traffic amount, could the virtual traffic be proportional to it?

Theoretical analysis Ideally, the IDAV model and carriage matrix should suffice to the following conditions:

1. If $N1 \in ASi \xrightarrow{c} ASj$, then $N1 \notin ASk \xrightarrow{c} ASj,\ \forall k,\ 1 \leq k \leq n\, and\, k \neq i$.
2. $\sum_{i=1}^{n} m(i,j) = Const,\ \forall j,\ 1 \leq j \leq n$.

The first condition requires unique path selection, and the second condition means every AS should know routes to the whole Internet. While *Const* in condition 2 equals the total address spaces that have been allocated by IANA[18] and utilized by ASs. In ideal situation, we have the following theorem:

Rank Theorem Given an IDAV model denoted by its carriage matrix M, the ranking vector R is determined by the equation: $M \cdot R = \rho(M) \cdot R$, where $\rho(M)$ is the spectral radius (maximum modulus of complex eigenvalue) of M.
Proof:

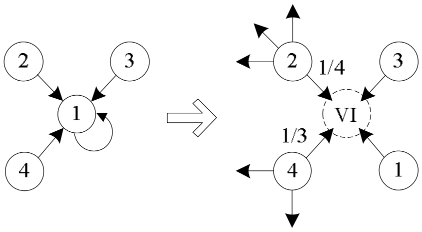

Fig. 1. VI: Proof of Rank Theorem

1. *Existence.*
According to Perron-Frobenius theorem[19], the carriage matrix M derived from IDAV model is a non-negative matrix, so that $\rho(M)$ is an eigenvalue of M and there is a non-negative vector R, $R \neq 0$, holds the statement $M \cdot R = \rho(M) \cdot R$. Thus the ranking vector R shall always exist. Moreover, if $\sum_{i=1}^{n} m(i,j) = Const,\ \forall j,\ 1 \leq j \leq n$, the spectral radius $\rho(M) = Const$.

2. *Rationality.*
First we normalize the carriage matrix M with $Const$: let $P(i,j) = m(i,j)/\sum_{i=1}^{n} m(i,j),\ \forall i,j,\ 1 \leq i,j \leq n$. $P(i,j)$ means the portion of access volume carried by ASi for ASj in the total address spaces. Let us consider the example in figure 1. AS1 has 3 neighbors, each of which has some access volume carried by AS1. AS1 also carries its own addresses for itself. Replace AS1 by adding a virtual intersection VI, the traffic incoming to VI equals the sum of traffic outgoing from VI's neighbors, including AS1 itself. For each neighbor, such as AS2, its egress traffic is most likely to be evenly distributed among all the destination addresses, so that traffic from AS2 to VI equals AS2's total egress traffic times $P(1,2)$. Let E_j and I_j be the egress and ingress traffic of ASj respectively, we have this formula: $I_j = \sum_{i=1}^{n} E_i \cdot P(i,j)$, e.g. $I = M/\rho(M) \cdot E$. Finally, because

VI associated with AS1 is a pure exchange point, leaving no remaining traffic within itself, its egress and ingress traffic vector must be equal, thus we get the ranking vector $R = E = I$. #

Let us illustrate the process of AS ranking using IDAV model and show its advantage over traditional single-metric driven approaches. In figure 2, each AS owns only one address and provides a complete route forwarding for others. If we count the number of addresses owned, every AS would be viewed equally important. If we use the degree-based criterion, AS3 and AS5 are both bigger than AS4. Neither of the judgement is convincing. In reality, AS4 may be a global exchange point, while AS3 and AS5 are only national ISPs yet with more customers. Using IDAV, the CM of this AS graph is M on the right of the structure. Solving equation $(M - 7 \cdot I) \cdot R = 0$, we get: $R = [0.08, 0.08, 0.51, 0.67, 0.51, 0.08, 0.08]$. As expected, although AS3 and AS5 have more customers, as long as AS4 is doing transit for both of them, AS4 should rank higher. This example shows obvious advantage of our method because IDAV studies the functional properties of AS and use routing contribution as the rank criterion.

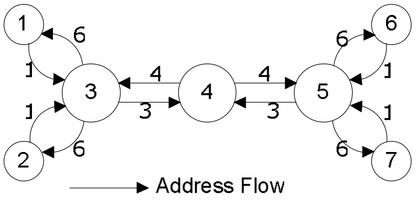

$$M = \begin{pmatrix} 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 6 & 6 & 1 & 3 & 0 & 0 & 0 \\ 0 & 0 & 4 & 1 & 4 & 0 & 0 \\ 0 & 0 & 0 & 3 & 1 & 6 & 6 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix}$$

Fig. 2. IDAV-based AS Ranking

Practical considerations In actual situation, the two conditions in previous section will no longer hold. Anyway, from rank theorem, we can still get R as the traffic estimation (albeit quite approximately): on the premise of the two following rules, elements of R is at least positively related to potential traffic.

1. Capability Rule: The more access volume does an AS carry for other ASs, the higher it ranks.
2. Chain Rule: Suppose ASx carries some access volume for ASy, the higher does ASy rank, the higher ASx ranks.

Although we can not prove them, the two rules are compliant with people's common sense. The first rule says that bigger ASs are always more busy, while the second rule says that the contribution of an AS to the Internet depends on the importance of the node it serves. These rules ensured that even in non-ideal situation, we can employ rank theorem to calculate AS ranking. For each ASj in the model, $r_j = \sum_{i=1}^{n} r_i \cdot m(i, j)$, if either the access volume $m(i, j)$ or another AS's ranking r_i gets bigger, the ranking r_j of ASj will go higher consequently.

It will nicely meet the two rules' requirements.
Admittedly, there are some limitations in applying our ranking theory, and the model would have to be refined with respect to the following factors: 1) It is almost impossible to gather all the route forwarding paths from every AS. 2) BGP signaling path is sometimes inconsistent with the actual packet forwarding path[20]. 3) Egress traffic of an AS is not evenly distributed among addresses[21][22]. The refinement to our model is another problem, which will be investigated in future works. To demonstrate some useful results, we use the simple form of IDAV and consider the influences of above drawbacks negligible.

3 Computations

3.1 Construction of IDAV

To build up a comprehensive Internet AS graph in IDAV, we use the BGP dump data collected at 02:43 AM PDT on 2005-12-24 by university of Oregon's RouteViews project, many thanks to their generosity. The routes are presented in either binary zebra MRT[23] or human readable Cisco CLI[24] format, both of which can be parsed into BGP attribute tuples as shown in table 1. Of all these attributes of a BGP table entry, NETWORK and AS_PATH are the most important to our study. Our aim is to parse each route into some $N \in ASx \xrightarrow{c} ASy$ atoms. We use perl to analyze text, and do the construction and computation of IDAV model with C programming language.

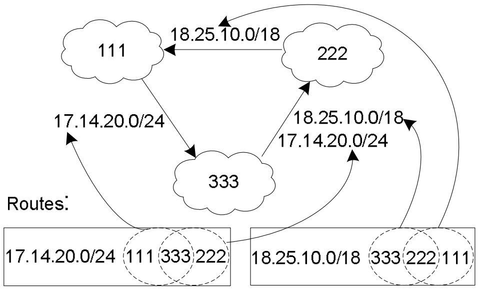

Fig. 3. Build IDAV model from BGP

In the AS_PATH attributes, each pair of neighbor AS numbers shows that the route is sent from the rightside AS to the leftside AS. As in tuple 1, we can get $8.8.9.0/24 \in AS27646 \xrightarrow{c} AS1239$ and $8.8.9.0/24 \in AS1239 \xrightarrow{c} AS1668$, etc. Sometimes the AS_PATH contains *repeated-prepending* or *aggregated AS_SETs* as in tuple 2 and 3, we just erase the duplicated AS numbers and decompose the aggregated ASs to count the routes separately. Being able to parse BGP routing tables into (network, AS-link) atoms, we create an IDAV model by scanning the Oregon data linearly, then count each edge's access volume, and finally perform the following iterative algorithm to calculate ranking vector.

Table 1. Sample BGP Table

NETWORK	NEXT_HOP	METRIC	LOCPRF	AS_PATH	ORIGIN
8.8.9.0/24	66.185.128.48	514	0	1668 1239 27646	?
166.111.0.0/16	66.185.128.48	514	0	1668 1239 4538 4538 4538	i
24.223.128.0/17	66.185.128.48	575	0	1668 10796 {11060,12262}	i

3.2 Ranking algorithm

In ideal situation, the IDAV model is a strong-connected graph, thus the CM is primitive, we have the following limit theorem[19]:

$$\lim_{m\to\infty}[\rho(M)^{-1}\cdot M]^m = L_{n\times n} > 0$$

thus the ranking vector R can be found through simple iteration starting from any initial non-zero vector $R(0) = l$, and $R(t) = M/\rho(M)\cdot R(t-1)$.
But in non-ideal situation, the limit theorem does no longer hold, for the CM is only a general non-negative matrix. Thanks to previous works like Pagerank[25][26], under these conditions, we can still employ an iteration algorithm to solve the equation $M\cdot R = \rho(M)\cdot R$:

$$R(t) = (1-\varepsilon)\cdot W\cdot R(t-1) + e(t-1)\cdot \ell_n$$

where $R(t)$ is the ranking vector after each iteration; W is the CM which is column normalized to 1; a small real number ε close to 0 and $\ell_n = [1,\ldots,1]^T$ are used to guarantee convergence; finally, for each t,

$$e(t-1) = \frac{1-\|(1-\varepsilon)\cdot W\cdot R(t-1)\|_2^2}{2\|(1-\varepsilon)\cdot W\cdot R(t-1)\|_1}$$

forces the condition $\|R(t)\| = 1$.

4 Results

4.1 Top 10 Ranking

Applying our ranking method to the final IDAV model, we calculated a numeric weight value for each of the 21,454 ASs appeared in the routing table. The weight vector is normalized to 1. As people are usually more interested in the bigger ones, in table 2 we show the top 10 ASs we found.

The table contains most of the well-known global backbone ASs, such as AS1239 and AS3356, as well as those very large but not so famous ISPs, such as Cogent, TeliaNet. From the degree values listed in the 4th column, we can see our ranking is quite different from degree-based rankings. A typical example is AS2914: although its degree is much less than many lower-ranked ASs, based in America, peering with the world's top ISPs like Sprint, Level 3 and MCI, it

Table 2. Top 10 ASs Found

Rank	ASN	Weight	Degree	Organization
1	1239	0.500845	1755	Sprint
2	3356	0.443078	1281	Level 3
3	209	0.397985	1155	Qwest
4	7018	0.342186	2019	AT&T
5	701	0.314771	2420	MCI/UUNET
6	2914	0.207631	496	NTT
7	3549	0.140935	700	Global Crossing
8	3561	0.137122	583	SAVVIS
9	174	0.125386	1364	Cogent
10	1299	0.125087	370	TeliaNet

Table 3. Other Top 10 AS Rankings

	Caida	FO.	RS.	IDAV
1	701	3356	1239	1239
2	7018	6461	3356	3356
3	1239	1239	701	209
4	174	3303	7018	7018
5	3356	2914	2914	701
6	209	8075	3549	2914
7	3549	4637	209	3549
8	7132	209	3561	3561
9	4323	3549	1299	174
10	3303	2497	6453	1299

is doing global transit for most APNIC ASs, thus its ranking is still fairly high. These results emphasizes the significance of our IDAV-based ranking theory. Also note that in ideal situation the ranking weight value is the same as normalized traffic throughput, but our computation is in the non-ideal situation, so AS209's weight is twice of that of AS2914's doesn't necessarily mean the actual traffic also is. The weight is only a relative measure here.

To verify our results, we gathered some academic or commercial AS rankings from other sources in the Internet. These different versions of top 10 ASs including ours are listed in table 3 for comparison. Caida offers hybrid criteria, we choose to sort by degree first. AS701 is the biggest because it has the most (2,420) connections, refer to table 2. FixedOrbit uses average hops, the weight calculated to indicate the average number of AS hops that must be traversed from inside the network to any other IP addresses in Internet. In this sense, AS3356 is the largest because it reaches other networks through the shortest AS_PATHs. ReneSys, a commercial statistics that puts the customers' quantity and quality first, has the most in common with IDAV, which indicates the success of our model. The only one disagreement is that we replaced AS6453 with AS174 in the top-10. In our result, AS6453 ranks the 16th because, unlike AS174, it is not carrying much access volumes for the other top 10 ASs. As the two rules suggested in section 2.2, an AS can promote its ranking by carrying access volumes for other big ASs. It is the chain rule that works here.

Although all above ranking approaches are meaningful, our method is a little superior because, according to IDAV theories, the calculated weight reflects each AS's contribution to the Internet routing. Others may be more connected, nearer to reach, have greater number of customers or even earn larger portion of money, but without AS1239, the Internet will suffer the severest and most disastrous route failure, or redirection.

4.2 Validation

Due to the lack of authoritative source of AS ranking information, to validate our results, we look for some indirect approach. We employ the "Valley-Free" principle of inter-domain routing behavior. Introduced by Gao in [9], the general rule for inter-domain path selection is: one route should only take the paths first going from smaller ASs to bigger ASs, followed by some steps between equal-sized ASs, then from bigger ASs to smaller ASs, e.g. no part of the path could be a "valley".

Although originally, the relative magnitude of AS is defined by means of customer, provider, and peer relationships, we can slightly modify it using our ranking result without changing the thinking behind valley-free principle. A threshold number "2" is used to classify the AS relationship: if an AS's ranking weight is twice or more than the next-step's, identify the step is a downhill one, the opposite step is an uphill one accordingly. If neither of the two ASs is sufficiently larger than the other, identify the step as flat. We use this law to examine each AS_PATH to see whether it is compliant with the valley-free pattern. According to our examination, in the RouteViews data we used, 99.3% AS_PATHs obey the valley-free rule, only 0.7% $(164,032/23,147,531)$ AS_PATHs have violations, which are mainly attributed to the following reasons:

1. Sibling ASs may transit arbitrary route for each other.
2. Some special address blocks have prefix-level traffic engineering.
3. Underestimation of certain ASs because of insufficient data.

Despite of these slight technical imperfections, this is an ultimate proof of our IDAV model and AS ranking theories. The wonderful consistency is a clue that both carriage matrix AS ranking and valley-free path selection are basic underlying characteristics of the Internet structure.

5 Conclusion

In this paper, we originate the Inter-Domain Access Volume (IDAV) model. This model investigates each individual AS's routing contributions to the Internet by adding access volume to the topological AS graph. Besides the formalization and construction of IDAV model, we also put forward a systematical AS ranking methodology using IDAV. The inferred AS magnitudes show more truthfulness than traditional approaches. And in return, the positive AS ranking results, as a benchmark for AS graph modeling, have validated the basic principles behind IDAV model. By leveraging information about both connectivity and the prefixes exchanged between neighbor ASs, the authors believe that this model will enable much more compelling applications, especially related to Internet routing architecture.

Future research topics, including visualization of the augmenting AS graph and applying the model to router-level intra-domain routing are ongoing. In addition, we are attempting to help traffic-engineering and smart-routing by employing this model, for example, we could make predictions of node congestions and link loads according to the network topology and routing solicitation, in order to adjust router configurations for better performance. The idea of traffic estimation upon IDAV model largely depends on establishing the relationship between real traffic flow and access volume. Some refinement to the model will be made and we believe such a mapping should be do-able.

References

1. B. Halabi, D. McPherson: *"Internet Routing Architectures".* Cisco Press, 2000.
2. RFC 1771, A Border Gateway Protocol 4 (BGP-4).
3. University of Oregon RouteViews project, http://www.routeviews.org/.
4. Merit Internet Routing Registry Services, http://www.irr.net/.
5. Z Ge, DR Figueiredo, S Jaiswal, L Gao: *"On the Hierarchical Structure of the Logical Internet Graph".* Proceeding of SPIE ITCOM, 2001.
6. A Vazquez, R Pastor-Satorras, A Vespignani: *"Internet Topology at the Router and Autonomous System Level".* Arxiv preprint cond-mat/0206084, 2002.
7. G Siganos, M Faloutsos, P Faloutsos, C Faloutsos: *"Power Laws and the AS-Level Internet Topology".* Transactions on Networking, vol. 11, no. 4, Aug. 2003.
8. Barabasi AL, Albert R: *"Emergence of Scaling in Random Networks".* Science 286 (5439): 509-512 Oct. 15 1999.
9. Lixin Gao: *"On inferring Autonomous System Relationships in the Internet".* IEEE/ACM Transactions on Networking, vol. 9, no. 6, Dec. 2001.
10. Subramanian L, Agarwal S, Rexford J, Katz RH: *"Characterizing the Internet Hierarchy from Multiple Vantage Points".* IEEE INFOCOM 2002.
11. Di Battista G, Patrignani M, Pizzonia M: *"Computing the Types of the Relationships between Autonomous Systems".* IEEE INFOCOM 2003.
12. X. Dimitropoulos, D. Krioukov, et. al: *"Inferring AS Relationships: Dead End or Lively Beginning?"* Lecture Notes in Computer Science, 2005.
13. X. Dimitropoulos, D. Krioukov, et. al.: *"AS Relationships: Inference and Validation."* In submission to ACM CCR (Computer Communication Review), 2006.
14. Caida, *"AS ranking",* http://as-rank.caida.org/.
15. FixedOrbit, *"Knodes Index",* http://www.fixedorbit.com/metrics.htm.
16. Renesys, *"Internet Index,"* http://www.renesys.com/products_services/market_intel/rankings/.
17. RFC 1519, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy.
18. Internet Assigned Numbers Authority (IANA), http://www.iana.org/.
19. R.A. Horn, C.R. Johnson: *"Matrix Analysis".* Cambridge University Press, 1990.
20. Mao ZM, Johnson D, Rexford J, Jia Wang, Katz R: *"Scalable and Accurate Identification Of AS-level Forwarding Paths".* IEEE INFOCOM 2004.
21. A. Soule, K. Salamatian, N. Taft, R. Emilion, and K. Papagiannaki, *"Flow Classification by Histograms or How to Go on Safari in the Internet".* Sigmetrics, 2004
22. J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, *"BGP routing stability of popular destinations".* Internet Measurement Conference, 2002.
23. MRT format: http://www.zebra.org/zebra/Packet-Binary-Dump-Format.html.
24. CLI format: Cisco IOS IP Command Reference, Volume 2 of 4: Routing Protocols, Release 12.3. http://www.cisco.com/.
25. Lawrence Page, Sergey Brin, Rajeev Motwani, Terry Winograd: *"The PageRank Citation Ranking: Bringing Order to the Web".* Proceedings of ASIS'98, 1998.
26. TH Haveliwala: *"Efficient Computation of PageRank".* Stanford Technical Report, 1999.