Exposing a Nation-Centric View on the German Internet – A Change in Perspective on AS-level

Matthias Wählisch[1], Thomas C. Schmidt[2],
Markus de Brün[3], and Thomas Häberlen[3]

[1] Freie Universität Berlin, Institut für Informatik, Berlin, Germany
[2] HAW Hamburg, Department Informatik, Hamburg, Germany
[3] Federal Office for Information Security (BSI), Bonn, Germany
{waehlisch, t.schmidt}@ieee.org,
{markus.debruen, thomas.haeberlen}@bsi.bund.de

Abstract. The Internet has matured to a critical infrastructure in many countries. The national importance of routing motivates us to study the AS-level subgraph that is relevant for a country. In this paper, we report on a methodology and tool chain for identifying and classifying a 'national Internet', and evaluate detailed results for the example of Germany. Our contribution (a) identifies the ASes that are important for the country, (b) classifies these ASes into functional sectors, (c) constructs the AS routing graph of a country as well as subgraphs of specific sectors, and (d) analyzes structural dependencies between key players. Our methods indicate the importance of examining individual IP-blocks held by individual organizations, as this reveals 25% more stakeholders compared to only looking at prefixes. We quantify the centrality of ASes with respect to specific sectors and the robustness of communication communities. Our results show that members of sectoral groups tend to avoid direct peering, but inter-connect via a small set of common ISPs. Even though applied for Germany here, all methods are designed general enough to work for most countries, as well.

1 Introduction

The Internet was originally shaped to offer open transmission services on a global scale, but has now turned into a mission-critical infrastructure of local relevance for most countries and dedicated players. The coherence of the Internet is defined by peering relations at the AS-level. Analyzing mutual impact, vulnerability and efficiency of the backbone requires the identification of ASes and corresponding transits between them.

Today, the global Internet is composed of more than 30,000 ASes with significantly more links, which challenge a clear picture on dependencies. Similar to traditional infrastructures, a country, its population as well as organizations share an obvious interest that the internal data exchange does not rely on weak third parties (cf., [1, 243 ff.]). In Internet terms, AS transits connecting key players of a country should be part of an apprehensible Internet ecosystem. However, the Internet is a globally distributed network without boundaries, which makes the identification of nationally relevant subparts hard and leads to the following questions: Which Autonomous Systems are important for a reliable interconnection of the Internet infrastructure of a country? How do sectors of a country communicate? Rigorous insight into the country-wise nature of the Internet thus carries fundamental importance and it is somewhat surprising that only recently the inter-network structures of nations attracted attention [2, 3].

This paper contributes with the following first steps to answer these questions: (1) A promising methodology to derive a country-centric view on the Internet structure. This is exemplarily verified for our home country Germany. The approach starts with IP-blocks. Compared to pure prefix-based techniques, we are able to identify approximately 25 % more members. (2) A novel, non-hierarchical AS taxonomy, as well as a heuristic sectoral classification technique. Both allows us to identify ASes with national relevance. By adding routing information, we are able to generate, visualize, and analyze the structure of communication flows between relevant public and business sectors. This has not been evaluated before. (3) The evaluation of the German inter-AS structure based on common and new graph metrics. This reveals for example that most eyeball providers peer dependent on the target AS, whereas the financial sector operates on static paths. (4) Finally, extracted and visualized data will be provided to the community for subsequent research and analysis.

The remainder of this paper is structured as follows: § 2 discusses the current state of Internet backbone analysis in the context of nation-state routing. Our methodology and corresponding toolchain, which allows for a nation-centric view on the Internet, is described and evaluated for Germany in § 3. § 4 analyses AS (sub-)graphs from Germany. Finally, § 5 concludes this paper.

2 Related Work

Research on the Internet AS structure continues to attract significant interest since more than one decade. This includes the analysis of structural properties of the inter-AS graph [4], the (mainly hierarchical) classification of ASes [5], and the inference of the relationship between them [6]. Active measurements within dedicated countries [7] reveal geographic reachability of ASes and thus follow a direction distinct from our work. Until now, there is only little work on a nation-state understanding of the Internet backbone routing, as well as on a horizontal classification of ASes that aims to identify key players of relevance for the Internet services of a country.

Dimitropoulus *et al.* [8] introduce a broader AS taxonomy (large/small ISPs, customer ASes, universities, IXPs, and network centers), but this does not include a detailed decomposition of customer ASes into dedicated sectors. The proposed inference algorithm analyzes the description value of the Internet Routing Registry [9] and follows a text classification technique. The authors focus on a complete mapping of ASes to classes. This differs from our perspective, as we do not intend to classify *all* ASes, but concentrate on *important* players of a country, viewed in further differentiated sectors. Cai *et al.* [10] introduce the interesting idea of an Internet AS ecosystem, which is based on a novel AS to organization map. The authors normalize contact records of the RIR WHOIS data to cluster AS numbers that belong to the same organization. Although this work could be applied to peering analysis and planning as well as threat analysis, our work is orthogonal as we classify ASes according to roles in a country.

The first paper that proposes a nation-state view on the Internet routing measures the impact of countries on the global data forwarding [2]. Karlin *et al.* start from IP prefixes, which they map to ASes. Routing paths are derived from an approximation of active traceroute measurements. We will show that IP prefixes are too coarse-grained to obtain an in-depth picture of a country and miss 25 % of German ASes. It is important to note that the base set of the national classification should be as complete as possible to judge on relevance.

Roberts and Larochelle [3] present a mapping project that visualizes and quantifies the relevance of ASes for countries. The AS to country mapping is based on the external service Team Cymru [11], which starts from IP prefixes. The authors introduce network maps of countries. Each map abstracts connections to the outside by a single Autonomous System that subsumes all foreign ASes. The relevance of an AS increases with the number of prefixes reachable via this AS. ASes with multiple upstream peers share routes equally among parents. This model oversimplifies common practice in Internet backbone routing. Normally, countries do not have a single entry point apart from China with the exclusive entry China Telecom. Furthermore, multilateral peering allows for path selection depending on the target AS. This diversity is not reflected and causes a weight distortion of AS importance.

To the best of our knowledge, current approaches do not provide sufficient mechanisms for identifying country-specific ASes, categorize ASes in business sectors, nor analyze importance of inter-AS communication between network domains for a country including international inter-connects. We will address these topics below.

3 Deriving Nation-Centric Subsets of the Internet

We want to identify all Autonomous Systems of the global Internet that host organisations from a specific country. Many organizations are normal ISP customers and do not own a prefix or AS. Thus we argue that the appropriate granularity must be IP blocks. An IP block is a subset of a prefix and will be assigned internally by the prefix owner to departments or customers. An organization and thus its hosting AS is coined to a country if the organization or its administrative contact person of an IP block is located therein. Consequently, we include also ASes in our view with primary base outside of the investigated country, as well as national organizations with IT infrastructure outsourced to foreign countries.

To demonstrate and validate our approach, we choose our home country Germany (*DE*). The introduced methodology, however, can be applied to other countries, as well. Its implementation is easily extendable and thus provides a good base for the community and subsequent work. In this section, we present the data sources, our inference and classification algorithm to derive a nation-state view, and the the fully automated construction of existing interconnects. These results enable us to analyse the composition of the nationally relevant part of the Internet in detail.

1. From Internet Members to ASes Regional Internet Registries (RIR) maintain network and contact details of their region. IP addresses and AS information related to Germany are registered at the RIPE database (DB). We start by extracting all `inetnum` records, which represent IP-blocks, from the RIPE DB that carry the mandatory `country` attribute of either *DE* or *EU*. Additionally, we collect address data for the associated `admin-c` and `org` objects. Based on the latter, we created keyword lists of synonymic country codes (e.g., Germany, DE), local city names, and international dialing codes (e.g., 0049, +49) representative for Germany. Applying the keywords on the contact record allows us to further resolve *EU* IP-blocks to *DE* and to verify the *DE* classification of IP-blocks. The result is a list of all IP-blocks allocated by organizations in Germany.

Next, we determine the longest covering IP-prefix for each IP-block. Prefix lengths are subject to aggregation and thus depend on the point of observation. Using passively measured BGP data from distant route collectors would be too coarse grained and yield less specific prefixes. Assuming that RIRs provide the most detailed prefix mapping, it is reasonable to query the RIPE DB. The inter-AS route is specified in the `route` record, which is referred by the `inetnum` object.

Finally, we map the prefixes to origin Autonomous System Numbers (ASNs) by the `route` object. However, in this step using the RIPE DB alone would lead to several unresolved mappings. Therefore, we also consider data of Team Cymru [11] and the route collector RRC12 of the RIPE RIS [12]. The latter peers at the largest German Internet Exchange Point (DE-CIX) and thus provides localized data. We apply the different data sources in the following order to maximize the number of resolvable ASNs: (1) RIPE DB, (2) Team Cymru, and (3) RIS RRC12. In cases of Multiple Origin ASes, we keep all discovered ASNs. The resulting list contains the ASes that compose the nation-centric part of the Internet for the example of Germany.

Our fully automated tool chain was applied in Oct. 2010 and yielded 246,861 German IP-blocks. Thereof 240,237 are embedded in 6,278 IP-prefixes that belong to 1,471 ASes, ≈ 2 % could not be resolved to a prefix. To estimate errors of our IP-block-to-country mapping, we checked back with the well-known MaxMind GeoLite Country service [13]. Deviations were found below 0.2% for both false positives and false negatives.

Our method of starting from IP-blocks rather than IP-prefixes identifies significantly more prefixes that carry relevance for the country Germany (cf., Table 1). When considering prefixes alone, only ≈ 84% can be identified as 'German' using RIPE-DB, while Team Cymru yields ≈ 70% *DE* prefixes. Thus a significant fraction

Approach	**DE**	**EU**	**other**
IP-Block	6,278	–	–
Prefix (RIPE DB)	5,243	–	1,035
Prefix (Team Cymru)	4,395	947	936

Table 1. Number of identified prefixes

of prefixes that route traffic relevant for Germany is not directly associated to country or address values from the this country.

Providers from outside Germany are also selected by our scheme. The corresponding 301 ASes ($\approx$ 20 %) are classified relevant for nation-state routing and internationally distributed as follows. More than a third (110) ASes originate from direct geographical neighbors, another third (107) from the remaining Europe, thereof 57 British, and 18 % (54) are North American ASes. These classifications are again based on RIR databases and Team Cymru with an estimated error of about 15 %.

2. Tier and Sector Classification of Autonomous Systems Having categorized the nationality of the stakeholders, we add two further classifications to the selected ASes. First, we harvest the topological hierarchy (tier1, large/small ISP, and stub) from [14]. Additionally, we investigate the role of ASes within those public or business sectors that are operationally relevant for the country. As there is no AS classification available that describes the professional role of an organization in relation to the global BGP routing, we introduce a sectoral categorization. This extends the taxonomy of critical infrastructure published by the Federal Office for Information Security (BSI), Germany. We determine sectoral classification by applying an optimized and manually verified keyword spotting to names, descriptions, and address fields of the previously derived AS data. Our approach uses general terms such as "bank", but also specific company names (e.g., Siemens, Daimler) as keywords associated with classes to identify important ASes. Keywords are correlated to enhance the identification. Thereof we obtain an additional list of the 'relevant national ASes' including branch tags such as *financial services* (cf., Table 3).

99 % of the classified ASes belong to exactly one sector, five ASes are assigned to two sectors. Companies may attain multiple roles. For example, AS 31438 is a municipal utility responsible for waste water and DSL access in the City of Marburg. Overall 279 ASes have been selected as 'systemically relevant' with sectoral attribute attached.

We admit that this step includes manual pre-definition of keywords for sectors. However, mapping ASes to sectors is a fine-grained process, which requires specific information that cannot be derived completely automatically. Our taxonomy, methodology, and tools can be applied to other countries based on an updated keyword list. The creation of the list needs local knowledge.

3. Constructing Spanning AS Routing Graphs Following the identification of all ASes relevant for the German Internet infrastructure and a classification of key players, we derive their interconnects. We limit the building of AS graphs to the construction of inter-AS paths without considering individual prefixes. This modeling step is meaningful for our purposes: Even though BGP policies and regional optimizations may lead to varying paths for different IP prefixes announced by the same origin AS, recent studies [15] show that multiple prefixes are reachable via the same AS path for 75 % of origin ASes. Additionally, we

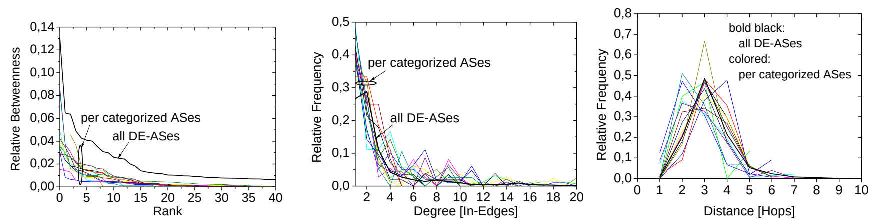

(a) Relative betweenness (b) In-degree (c) Distance distribution

Fig. 1. Properties of the Internet relevant for Germany and its categorized subgraphs

focus on a regionally bound network, which is densely meshed by peering points. International redirections within service provider networks are mainly outside the scope of our perspective. From this point of view, restricting the routing on the AS level is a valid approximation.

We identify an AS routing graph for each sector, the bilateral exchange between two sectors, as well as the AS graph of all DE ASes based on the weighted next hop matrix provided by the NECLab topology project [16]. This data is calculated using the continuously updated measurements by the UCLA [14] and reflects BGP policy decisions [15]. To exclude incomplete paths, we omit row column values of -1 for distinct indices during matrix processing. Note, this occurs very rarely ($\ll 0.3\,\%$). Naturally, the set of ASes in the routing graphs has been extended by intermediate ASes that we have not assigned before to the nationally relevant part of the Internet. These transit nodes are required to link nation-state subsets that would remain isolated otherwise.

4 Analysis of the AS-Structure

In this section, we investigate structural properties of the derived AS routing graphs and measure the relevance of members in sectors and in the overall DE AS graph. It is worth noting that we keep BGP policy modeling by a per path analysis, each path derived from the *weighted* next hop matrix of the NEC topology project. All measurements are relative to allow for comparing sectors of different sizes. Unfortunately, the underlying next hop matrix does not provide edges to connect the five ASes of the medical sector.

4.1 Node Centrality

Intermediate nodes between source and receiver attain a relevant role from serving as transits. The number of shortest paths passing through a node m is quantified by the betweenness $B(m)$. If the total number of shortest paths between two nodes i and j is $B(i,j)$, and the number of these paths going through node m is $B(i,m,j)$, then the betweenness of m is defined as the ratio: $B(m) = \sum_{i \neq m \neq j, i \neq j} \frac{B(i,m,j)}{B(i,j)}$. This measurement quantifies also the load at intermediate ASes. The betweenness is normalized by $(|V|-1)(|V|-2)$.

All DE		Large ISPs		Research	
Ranked AS	Betweenness	Ranked AS	Betweenness	Ranked AS	Betweenness
1. DTAG	0.131	1. DTAG	0.031	1. DFN	0.087
2. Level 3	0.065	2. Lambdanet	0.008	2. Verizon	0.037
3. Lambdanet	0.064	3. Telekom–AT	0.007	3. Manda	0.030
4. Colt	0.049	4. France Telecom	0.006	4. BELWUE	0.021

Table 2. Relative betweeness of the top ranked ASes for selected sectors

The *term shortest path* refers to the routing path that is actually taken. Our underlying BGP routing model reflects policies [15]. Using the NEC matrix, there exists exactly one effective path between two ASes. However, as discussed in Section 3, in our context of locally bound routing this is not a restriction. Independent of the nation-state view, BGP policies may lead to a violation of the triangle inequality. As the routing paths are based on a weighted graph, this property is preserved.

We calculate the betweenness of a node for the routing graphs under discussion. Figure 1(a) shows the relative betweenness, where ASes are ranked in decreasing order. Details for selected sectors are listed in Table 2. In 80% of the cases, this measurement exhibits sharp peaks at the transition from the top most important AS to the second one. This means that in the selected category a dedicated AS is part of a significant number of shortest paths and thus attains a major role in data forwarding. However, the decay from the top most ranked ASes is less steep in the overall German AS graph, showing a more evenly distributed relevance due to increased peering links. Looking at the actual rank orders reveals a relatively stable number of ASes among the top five in each category. For example, AS 3320 (Deutsche Telekom) has in 80% of the cases at least rank 5 and in 48% the highest betweenness.

4.2 Degree Distribution

The degree of a node denotes the number of its one-hop neighbors. Figure 1(b) shows the in-degree distribution. For visibility, we cut the tail at 20 edges. Overall, the relative frequency decays polynomially for all networks. Thus, there is a higher probability to maintain only a quite limited number of peering relations, but a non-vanishing likelihood for high peering numbers. The distribution of the full DE AS graph decays smoothly, while sectoral groups exhibit systematic peaks for selected node degrees between four and 13. Consequently, specific networks within the sectoral subgraphs are more densely connected than the full graph. These additional weights indicate regional star topologies in sectoral networks.

When comparing the topology within sectors to the complete DE network, we find more pronounced betweenness' and irregular peaks at increased node degrees. Jointly, these two structural metrics indicate that individual ASes provide enhanced connectivity within the specific communities as opposed to direct interconnects. A closer look on the corresponding AS graphs supports this observation. The majority of financial services, for example, tend to peer via Deutsche Telekom (AS 3320) and Colt (AS 8220), while no mutual peering is visible at

Sector (# ASes)	$\langle X \rangle$	σ_X	Sector (# ASes)	$\langle X \rangle$	σ_X
Transit providers (55)	2.41	0.92	Industry (28)	3.19	0.87
Trading (10)	2.69	0.87	Financial services (32)	3.21	0.89
Science & Culture (22)	2.77	1.12	Shipping & transportation (15)	3.22	0.8
Eyeball ISPs (23)	2.83	1.16	Public administr. & justice (14)	3.22	1.14
Peering points (8)	2.87	0.97	*DE All (1,471)*	3.23	1.04
Public services (4)	3.00	0.6	Energy (11)	3.34	0.79
Media & publishers (19)	3.08	0.94	Other public services (7)	3.40	0.73
Software and systems (31)	3.18	0.89	Medical services (5)	–	–

Table 3. Absolute number of ASes per sector and DE graph as well as mean ($\langle X \rangle$) and standard deviation (σ_X) of the distance distributions for corresponding routing graphs

all. Surprisingly, the governmental federation follows the same pattern. Governmental organizations are mainly interconnected by Deutsche Telekom and Versatel (AS 8881), but a small group uses Plusline (AS 12306) as upstream provider. The latter organizations require the external tier1 ISP AT&T to serve as inter-connect to the remainder of this sector.

4.3 Distances

The distance distribution of shortest paths measures the probability that two randomly selected nodes of a network are connected at distance k. This metric describes routing performance and usually follows a Gaussian law. This observation is also reflected in the analysis of the sectors and DE routing (cf., Fig. 1(c)) with average values between 2.4 and 3.4 (cf., Table 3). Routing distances, thus, largely depend on the sector under investigation. Naturally, connections between ISPs are shorter as compared to other branches. Surprisingly, ASes of the trading sector are significantly short, as well. In this group, the majority of members are connected via the same ISPs. Deutsche Telekom (AS 3320) and Vodanet (AS 3209) play a dominant role for transit. Even though there is no bilateral peering within the sector, many traders (e.g., Ebay AS 6907) maintain extensive peering relations. Paths that consist of only one transit hop are easily established. In general, our results show a similar behaviour compared to the global AS topology [17]. Even in the relatively small sectors, interconnects are not significantly denser and path lengths are not generally reduced. Most of the members from the same sector seem eager to stay at distance to each other.

4.4 Context Dependent Peer Selection

To analyze the peering behaviour in more detail, we answer the following question: How likely does a member of a sectoral group chooses its upstream peer dependent on the sector it communicates with? For each member of a sector, we count the number of different upstream peers relatively to the overall number of paths towards members of distinct sectors. We quantify the relative frequency of corresponding diversity classes over all members of a sector. For ASes that peer with many ISPs, it indicates high probability for high first hop neighbor diversity.

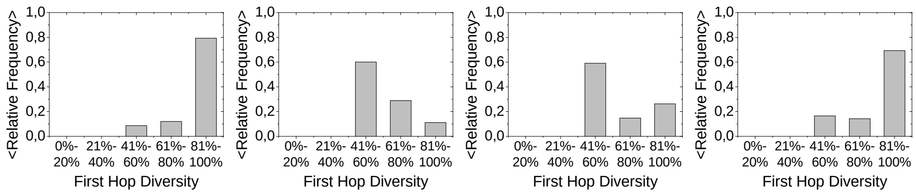

(a) Transit ISPs (b) Financial Srv. (c) Industry (d) All DE

Fig. 2. Relative upstream diversity for selected sectors

The calculated upstream peer diversity is very heterogeneous and may appear as a characteristic feature of the sectors. Figure 2 presents the measurements for selected sectors. Members of the financial services, for example, exhibit constant paths in about 50 % of the cases with enhanced probability (cf., Fig. 2(b)). In contrast to this, 80 % of the transit providers select above 80 % of the time neighbors dependent on the target. In general, target specific peering is dominant, as 8 of 10 ASes choose their one-hop neighbor with respect to the destination.

Combining the results with our previous findings indicates that multilateral peering has dominant routing effects on the Internet subpart relevant for Germany. For sectors, however, this higher amount of interconnections does neither result in more densely meshed inter-AS links nor in shorter paths.

5 Conclusions & Outlook

A clear understanding of the inter-AS structure at the country-level is needed to cope with the interdependencies and intrinsic vulnerability of the Internet. In this paper, we presented a methodology to identify and classify the relevant ASes of a country. This led to a fine-granular view onto meaningful subsets of Internet stakeholders and a detailed analysis. To the best of our knowledge, this is the first inspection of a country and its key players at the Internet routing substructure. The evaluation was exemplified for our home country Germany (DE), and created a list of relevant DE ASes including administrative data and sectoral classification, which will be publicly available.

We associated Autonomous Systems with a country whenever they host IP address space for an organization from there. Our approach outperformed prefix-based techniques by identifying 25 % more ASes. In particular, we were able to spot parts of the public sector hosted by international providers. Our analysis further revealed that members of the same public or business sector tend to not peer with each other, but interconnect via some few national and also international ASes. Deutsche Telekom, Level 3, Lambdanet, Colt, and Versatel were found to be the most important transit ASes for intra-DE communication. Multilateral peering was seen to have dominant routing effects on the German Internet, but the degree of variable upstream selection strongly depends on the sector.

Our future work will extend the current results in both directions, structural analysis and further countries, including their interdependencies. We expect structural properties on a fine-grained basis. In addition, we will extend our analysis towards IPv6. Regarding integration, we will employ current aggregation techniques for ASes belonging to the same organization to derive a condensed national Internet AS ecosystem. Finally, we will also concentrate on the application of our work to existing monitoring systems (e.g., [18]), which may help to reduce complexity due to selected observation points.

Acknowledgements

We would like to thank Lixia Zhang, Jennifer Rexford, Rolf Winter, and the DENOG community for enlightening discussions on the intrinsic topic of Internet backbone routing, as well as Steve Uhlig for advices on a previous version of this paper. Sebastian Meiling and Andreas Krohn have importantly contributed to the creation and verification of the toolchain. This is gratefully acknowledged.

References

1. D. M. Slane, C. Bartholomew *et al.*, "2010 Report to Congress," U.S.–China Economic and Security Review Commission, Annual Report, November 2010. [Online]. Available: http://www.uscc.gov/annual_report/2010/annual_report_full_10.pdf
2. J. Karlin, S. Forrest, and J. Rexford, "Nation-State Routing: Censorship, Wiretapping, and BGP," arXiv.org/CoRR, Tech. Rep. abs/0903.3218, March 2009.
3. H. Roberts and D. Larochelle, "Mapping Local Internet Control," Berkman Center, Harvard University, Tech. Rep., 2010. [Online]. Available: http://cyber.law.harvard.edu/netmaps/mlic.pdf
4. M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-Law Relationships of the Internet Topology," in *Proc. of ACM SIGCOMM'99*. New York, NY, USA: ACM Press, 1999, pp. 251–262.
5. R. Govindan and A. Reddy, "An Analysis of Internet Inter-Domain Topology and Route Stability," in *Proc. of the IEEE INFOCOM'97*. Washington, DC, USA: IEEE Computer Society, 1997, pp. 850–857.
6. L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, 2001.
7. S. Zhou, G.-Q. Zhang, and G.-Q. Zhang, "Chinese Internet AS-Level Topology," *IET Communications*, vol. 1, no. 2, pp. 209–214, 2007.
8. X. Dimitropoulos, D. Krioukov, G. Riley, and K. Claffy, "Revealing the Autonomous System Taxonomy: The Machine Learning Approach," in *Proc. of the PAM Conf. 2006*, M. Allman and M. Roughan, Eds. Web, 2006, pp. 91–100. [Online]. Available: http://www.pamconf.net/2006/papers/pam06-proceedings.pdf
9. "Internet Routing Registry," http://www.irr.net, 2010.
10. X. Cai, J. Heidemann, B. Krishnamurthy, and W. Willinger, "Towards an AS-to-Organization Map," in *Proc. of the 10th ACM IMC*. New York, NY, USA: ACM, 2010, pp. 199–205.
11. "Team Cymru." [Online]. Available: http://www.cymru.com/
12. "RIPE Routing Information Service (RIS)." [Online]. Available: http://www.ripe.net/projects/ris/rawdata.html

13. "MaxMind – GeoLite Country." [Online]. Available: http://www.maxmind.com/app/geoip_country
14. B. Zhang, R. Liu, D. Massey, and L. Zhang, "Collecting the Internet AS-level Topology," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 1, pp. 53–61, 2005.
15. R. Winter, "Modeling the Internet Routing Topology – In Less than 24h," in *Proc. of the 2009 ACM/IEEE/SCS 23rd Workshop on Principles of Advanced and Distributed Simulation (PADS '09)*. Washington, DC, USA: IEEE Computer Society, 2009, pp. 72–79.
16. "Internet AS-level topology construction & analysis." [Online]. Available: http://topology.neclab.eu/
17. P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, K. C. Claffy, and A. Vahdat, "The Internet AS-Level Topology: Three Data Sources and One Definitive Metric," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 17–26, January 2006.
18. Y.-J. Chi, R. Oliveira, and L. Zhang, "Cyclops: The AS-level Connectivity Observatory," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 5, pp. 5–16, 2008.