

Towards a Nation-Centric Understanding of the Internet

Matthias Wählisch
Freie Universität Berlin, Informatik

Thomas C. Schmidt
HAW Hamburg, Informatik

Sebastian Meiling
HAW Hamburg, Informatik

Markus de Brün & Thomas Häberlen
Bundesamt für Sicherheit in der Informationstechnik (BSI)

ABSTRACT

Data communication in the Internet is based on the borderless interplay of Autonomous Systems (ASes). An AS abstracts one or multiple IP networks, which may be globally distributed, but is owned by an organization located in a country. Current research on the Internet structure mainly focuses on a global perspective or considers local, intra-domain properties. In this paper, we analyze nation-centric subsets of the Internet taken from the AS-level graph of Germany. Each subset reflects a public or industrial sector. Based on a classified set of relevant German ASes, we are able to perform detailed investigations of structural dependencies for the critical Internet infrastructure. We identify and visualize the importance of dedicated ASes within specific sectors, and quantify robustness of the communication communities. Our preliminary results indicate that members of sectoral groups tend to avoid direct peering, but connect via a small group of common ISPs. This results in an enhanced dependence (betweenness) on selected hubs as compared to the characteristics of larger networks.

Categories and Subject Descriptors

C.2.1 [Computer-Comm. Networks]: Network Architecture and Design—*Network topology, Internet*

Keywords

AS-level subsets, Internet measurement and modeling, German Internet Graph

1. INTRODUCTION

Large companies, universities, governmental organizations etc. maintain their own Autonomous Systems (ASes) to participate in Internet peering and to enhance flexibility and robustness in connectivity. Members of the Internet who do not operate an AS depend on a service provider for announcing their IP-block within a prefix on the AS-level. The analysis of the corresponding Internet structure continues to attract significant interest since more than one decade (e.g., Faloutsos et al. 1999). Understanding the Internet from the perspective of its routing paths reveals scalability issues and the influence of different stakeholders, i.e., ASes. It may help to identify important 'hubs' and to estimate the reliability of the connected parties, for example.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM CoNEXT Student Workshop, November 30, Philadelphia, USA.
Copyright 2010 ACM 978-1-4503-0468-9/10/11 ...\$10.00.

Conceptually, the Internet is a global conglomerate of interconnected players and only physical infrastructure causes locality. However, this traditional Internet paradigm changes and there is an increasing tendency to deploy policies in Internet routing based on political rules of individual countries [1]. Attempts to foresee the impacts of such civil actions require a detailed understanding of nation-centric network structures. Analyzing the complete [2] global Internet on the inter-domain level, though, does not reveal local characteristics. Even intra-domain properties are usually too coarse-grained to observe relevant interdependencies of dedicated Internet members in complete communication contexts.

This paper aims to facilitate an understanding of the Internet from a nation-centric perspective. We discuss the Internet structure on two levels of reduction: (a) We look on the German Internet, i.e., the links of all ASes that announce IP prefixes which include IP blocks owned or registered by bodies located in Germany. (b) We filter ASes out of this subset with respect to a sectoral classification. We construct and analyse AS-level graphs of shortest paths that are based on public data (RIPE, UCLA, NEC, TeamCymru) and incorporate BGP policies [3]. First steps of our toolchain and measurement methodology are briefly described in [4].

In the remainder, we present and discuss our preliminary results for the German AS-level graph, and exemplary industry sectors (§ 2). We conclude with an outlook (§ 3).

2. ANALYSIS & DISCUSSION

We analyze the nation-centric AS structures w.r.t. two metrics that *quantify the transit properties* within the (sub-) graph and the degree of inter-connections between the ASes. Note, the derived graphs reflect realistic BGP-policies [3].

Betweenness Intermediate nodes between source and receiver attain a relevant role. 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 the importance of a node in data exchange, and the load on such intermediate vertex. For comparison of different sized, directed networks of $|V|$ nodes, the betweenness is normalized by $(|V| - 1)(|V| - 2)$.

Figure 1(a) presents the relative betweenness of each categorized AS compared to the betweenness of the overall German AS graph. ASes are ranked in decreasing

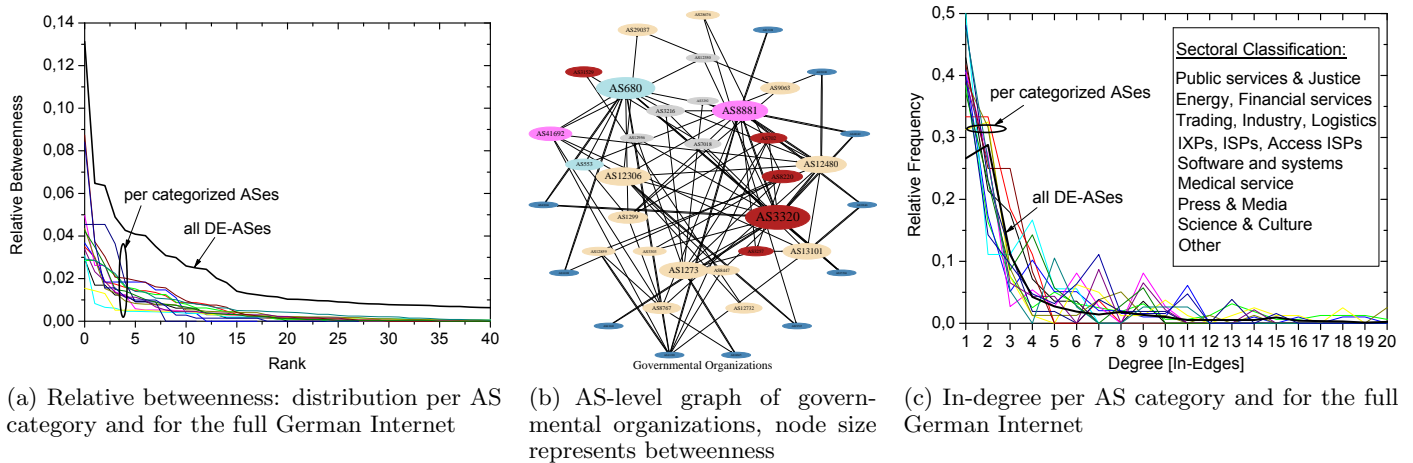


Figure 1: [Best viewed in color] Properties of the full German Internet and its categorized subgraphs

order. In 80% of the cases, this measurement exhibits sharp peaks from the transition of the top most important AS to the second important AS. This means that in the selected category a dedicated AS is part of a significant number of shortest paths and thus involved in the data forwarding. However, the general importance of the top most ranked AS decreases in the overall German AS graph due to increased peering links.

Looking on the actual ASes associated with the rank shows that there is a stable number of ASes, which belong to the top five in each category. For example, AS3320 (Deutsche Telekom) has in 88% of the cases at least rank 5 and in 63% the highest betweenness.

An exemplary AS-level graph for the identified German governmental organizations is presented in Figure 1(b). Each node color represents a category (e.g., dark blue for the government agencies) and vertex sizes are scaled according to the value of the AS betweenness.

Degree distribution The degree of a node denotes the number of its one-hop neighbors. Figure 1(c) shows the in-degree distribution of the sectoral ASes compared to the in-degree distribution of the full German Internet. For visibility, we cut the degree at 20 edges. In general, the relative frequency decays polynomially for all networks. The overall German AS graph decays smoother, where sectoral ASes exhibit peaks along smaller degree values. Thus, there is a higher probability to maintain only a quite limited number of peering relations. Still, for the subgraphs induced by the sectoral members, selected networks are more densely connected than the full graph. It is worth noting that each link represents an active routing path in the cluster, as we analyze the AS graph of shortest paths under BGP-policies.

Discussion These measurements are taken from a fine-granular view onto meaningful subsets of Internet stakeholders and allow for a detailed interpretation. When comparing the topology within business sectors to the national network, we find decreased 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 all. Surprisingly, the governmental federation follows the same pattern (see Figure 1(b)). Public and governmental networks 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 tier1 network of AT&T to serve as inter-connect to the remainder of the governmental organizations.

3. FUTURE STEPS

Our future work will detail out the current results. In the structural analysis, we will incorporate additional metrics, in particular to reflect clustering properties. Finding more comprehensive measurements to represent the different dimensions of the subgraphs will be a challenge, as well. In addition, we will focus on the connection between different sectors and countries answering the question, which transit ASes are important to allow the inter-domain communication. It will yield structural properties on a fine-grained basis. Finally, we will extend our analysis towards IPv6.

4. REFERENCES

- [1] J. Karlin, S. Forrest, and J. Rexford, "Nation-State Routing: Censorship, Wiretapping, and BGP," arXiv.org/CoRR, Tech. Rep. abs/0903.3218, 2009.
- [2] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (in)Completeness of the Observed Internet AS-level Structure," *IEEE/ACM Trans. Netw.*, vol. 18, no. 1, pp. 109–122, 2010.
- [3] R. Winter, "Modeling the Internet Routing Topology – In Less than 24h," in *Proc. of the 23rd PADS*. IEEE ComSoc, 2009, pp. 72–79.
- [4] M. Wählisch, S. Meiling, and T. C. Schmidt, "A Framework for Nation-Centric Classification and Observation of the Internet," in *Proc. of the ACM SIGCOMM CoNEXT. Student Workshop*, 2010.