Networking Science

RESEARCH ARTICLE

Discovering the geographic properties of the Internet AS-level topology

Enrico Gregori[1], Alessandro Improta[1,2*], Luciano Lenzini[2], Lorenzo Rossi[1] and Luca Sani[3]

1. Institute of Informatics and Telematics, National Research Council, via G. Moruzzi 1, Pisa, Italy
2. Information Engineering Department, University of Pisa, via Diotisalvi 2, Pisa, Italy
3. IMT, Institute for Advanced Studies, piazza S. Ponziano 6, Lucca, Italy

Abstract Research to date has analyzed the Internet AS-level topology at a worldwide level of detail. Every inference found for an AS is extrapolated from the global set of AS paths gathered from public monitors, independently of the geographic location of the ASes. This approach is useful when the Internet is analyzed at a very coarse level. However, it may be misleading if the analysis is more focused on a specific geographical region. The risk is that the particular characteristics that the Internet has in that region may be lost. An AS connection that has been identified in a global analysis may hide multiple connections located in different geographic regions, each with its own characteristics. Moreover, a couple of ASes may establish different economic relationships in each geographic region where they are connected. In this paper we propose an innovative technique to geolocate the AS connections retrieved from BGP raw data, in order to highlight the Internet characteristics both at a continental and national level. The analyses that we performed revealed some regional characteristics, in terms of graph properties and inter-AS economic relationships, that should be taken into account in a future analysis of the Internet.

Keywords Internet topology; Geography; Autonomous system; BGP

1 Introduction

The Internet is a complex system that evolved over the last few decades from a small network confined to the U.S. (i.e. ARPANET, 1969) to the current worldwide network of networks. It now consists of thousands and thousands of networks, under the administrative control of about 40,000 Autonomous Systems (ASes). This pervasive evolution did not occur homogeneously around the world for obvious historical, economic and political reasons. The result is that the Internet today is the composition of loosely connected groups of networks identifiable by some geographic boundaries, each with its particular pricing models, business contexts and regulatory environments [1]. Typically each AS has a particular role and specific economic behavior in each region of the world where it is present, which strictly depend on the connectivity and performance that it can provide for its customers in that region. For example, an intercontinental AS may be widespread in its home region − in terms of the number of connections and services offered − while it may be not competitive outside that region. This different level of pervasiveness may lead the same AS to establishing economic relationships in those regions with different criteria. Most research in the Internet topology analysis have considered ASes as homogeneous entities, each with a global set of metrics and characteristics, regardless of their heterogeneity. Those works (e.g. [2,3,4]) that tried to

---

* E-mail: alessandro.improta@iet.unipi.it

get a better insight in the geographic distribution of ASes relied on *traceroute* and packet delays to infer geographic information but, to the best of our knowledge, none of them explicitly focused on inferring regional AS topologies. Moreover from *traceroute* data is not trivial to obtain an AS-level topology, due to dealiasing and router-to-AS mapping issues [5], and it is almost impossible to have a precise view of the dynamic evolution of each AS path, that is fundamental to correctly infer the economic nature of each connection [6]. These problems can be easily solved exploiting the BGP data collected by Route Collectors (RCs) deployed by projects such as Route Views[1], RIPE RIS[2] and PCH[3]. The main source of information in BGP data about the AS-level topology is represented by the well-known mandatory `AS_PATH` attribute. This attribute contains the sequence of ASes that the traffic crosses to reach the announced subnet, but no information on the geographical regions where the traffic actually flows. Our aim is to search for a methodology able to infer geographic information from AS paths and to analyze local properties of the Internet, with a special focus on AS topology properties and economic relationships. Initially we infer the geolocation of each AS by exploiting the geolocation of its prefixes, which can be collected from BGP data. These data are exploited to produce a set of geographically tagged AS paths in which each connection is geolocated. This set is used to derive regional AS-level topologies that are analyzed both from a statistical and economic perspective. The economic analysis requires an economic tagging algorithm (e.g. [6,7] and [8]), but none of the available algorithms can deal with geographic information without any modification. In this paper we show how to adapt the algorithm proposed in [6] to deal with geographic AS paths. We found that the Internet actually consists of regional and independent ecosystems, which differ greatly in terms of both topological and economic properties and introduce particular characteristics that are hidden in a global-level analysis. These differences should be taken into serious consideration by all research based on the AS-level topology of the Internet − e.g. Internet modeling evolutions, protocol analyses, worm spread analyses − in order to rely on a more realistic structure of the Internet, instead of on a coarse and potentially misleading representation.

This paper is organized as follows. Section 2 details the process of geolocating each AS from raw BGP data. Section 3 introduces the methodology aimed at producing geographically tagged AS paths. Section 4 shows the graph properties found by the undirected analysis of the regional topologies extracted from geographically tagged AS paths. Section 5 shows how to modify the economic tagging algorithm described in [6]. Finally, section 6 presents the economic analysis of the economic regional topologies extracted using the economic tagging algorithm and Section 7 concludes the paper.

2 AS Geolocation

Knowledge regarding the geographic range of an AS is one of the fundamental parameters for decisions concerning the establishment of a settlement-free peering or a transit type of relationship between ASes. Several Tier-1 (T1) ASes include in their peering requirements at least one geographic constraint for candidate peers that need to be fulfilled. Just to name a few, AT&T[4] requires a list of the countries served by the candidate peer in the peering request submission; Verizon[5] requires a minimum number of served countries in the region where the peering is requested and that candidates own a “geographically-dispersed network”; and TeliaSonera[6] requires that the candidate peer is present and able to exchange traffic and to be interconnected in a minimum number of cities in two out of three regions (Europe,

---

[1] http://www.routeviews.org

[2] http://www.ripe.net/data-tools/stats/ris/routing-information-service

[3] http://www.pch.net

[4] http://www.corp.att.com/peering/

[5] http://www.verizonbusiness.com/terms/peering/

[6] http://www.teliasoneraic.com/Ourservices/IP/IPTransit/index.htm

North America and/or Asia Pacific/Oceania).

To geolocate each AS we start from its formal definition, given in RFC 1930:

> *"an AS is a connected group of one or more IP prefixes run by one or more network operators which has a single and clearly defined routing policy."*

Given this definition, it is straightforward that an AS is geolocated if its own prefixes are geolocated. The list of all the active prefixes of a given AS can be collected by parsing the BGP raw data provided by RouteViews, RIS and PCH. Each prefix can be geolocated in turn by geolocating each IP address inside it, using one of the IP geolocation databases available [9]. Consider a generic route `x.y.z.0/24-A B C D`. It is possible to claim that the last element of the AS path[7] owns at least a network − and thus it may or may not be present − in the region(s) where the prefix is geolocated. This approach is correct for any given geographic region (e.g. countries, continents) iff the granularity of the geolocation tool is fine enough and iff the route does not carry the `AGGREGATOR` and `ATOMIC_AGGREGATE` attribute. The `AGGREGATOR` is an optional transitive attribute and the `ATOMIC_AGGREGATE` is a well-known discretionary attribute of the BGP protocol and may be included in UPDATE messages by a BGP speaker which performs route aggregation. If one of these attributes is present, it is possible that part of the real AS path is missing, hidden by the aggregating router. In this case, it is not possible to state that the considered prefix belongs to the last element of the AS path, but additional confirmation is needed from the `WHOIS` service provided by the Internet Routing Registries[8]: the prefix is considered to belong to the last AS of the AS path if that AS is the owner of the announced prefix also according to the `WHOIS` response. For example consider the route shown in Figure 1 − which is the textual representation in MRT data of a route collected by the RC `rrc12` of RIPE RIS − where the prefix is entirely geolocated in Europe. Given the presence of the `AGGREGATOR` attribute, we need to query the `WHOIS` service. Since the response state that the prefix belongs to AS 2597, we can conclude that AS 2597 is present in Europe.

```
TIME: 10/01/11 08:00:06
TYPE: TABLE_DUMP_V2/IPV4_UNICAST
PREFIX: 192.12.193.0/24
SEQUENCE: 241676
FROM: 80.81.192.98 AS9189
ORIGINATED: 08/17/11 00:23:52
ORIGIN: IGP
ASPATH: 9189 8422 3356 2597
NEXT_HOP: 80.81.192.98
MULTI_EXIT_DISC: 100
AGGREGATOR: AS2597 217.29.66.79
COMMUNITY: 9189:1003 9189:1102
```

Fig. 1 Textual representation of a route in MRT format

In this paper we will show details about the Internet infrastructure both at continental and national level by exploiting the Maxmind GeoIPLite database[9], which has been proved to be reliable at both level [14]. In the continental level analysis, we divide the Internet in five macroregions: Africa, Asia-Pacific (Asia and Oceania), Europe, Latin America (the Caribbean, Central America, Mexico and South America) and North America (Bermuda, Canada, Greenland, Saint Pierre and Miquelon, U.S.A.). In the national level analysis, for sake of readiness, we consider only those nations that are currently participating in the G8 forum, i.e. U.S.A., Japan, Germany, France, United Kingdom, Italy, Canada and Russia. Note that the same information related to the missing countries can be found on our website[10].

3 Introduction of geography in BGP data

Geolocation of ASes by itself is not enough to extract geographic information from AS paths. An AS can have a

---

[7] The rightmost AS of an AS path is the AS that has originated the BGP UPDATE.

[8] A list of available `WHOIS` locations can be found at http://www.irr.net/

[9] http://www.maxmind.com/app/geoip country

[10] http://www.isolario.it

geographic range that spread across multiple regions, thus it is not possible to infer where each AS connection forming an AS path is located. To overcome this problem we propose a three-step algorithm which, based on the geolocation of each AS, is able to geolocate each AS connection of the AS paths.

Enhanced routes from BGP raw data In this step we obtain an *enhanced route* − defined as the triplet `{SOURCE, ASPATH, DESTINATION}` − for each route available in the BGP data. `SOURCE` is the region where the BGP AS Border Router that announced that route to the RC is located and can be obtained by geolocating its IP address (`FROM` field in Figure 1). `ASPATH` is the content of the homonym BGP attribute cleaned of private/reserved/unallocated ASNs and the ASN 23456, i.e. `AS_TRANS` used by 4-octet capable BGP speakers to communicate with 2-octet capable BGP speakers (RFC 4893). `DESTINATION` is the region where the prefix announced is located. Since a prefix could be geolocated in more than one geographic region, more than one enhanced route could be created from a single route, one for each region where the destination is found to be located. Consider the route reported in Figure 1. Both the IP address of the BGP speaker (`80.81.192.98`) and the prefix announced (`192.12.193.0/24`) are located in Europe, thus we obtain the enhanced route `{EU, 9189 8422 3356 2597, EU}`.

Detection of Single Region Located Transit Points (SRLTPs) in each enhanced route In this step we extract from each enhanced route the set of SRLTPs. This set contains regional intermediate points where the traffic needs to flow. The `SOURCE` and the `DESTINATION` of each enhanced route are by definition part of this set, since they are both geolocated in a single region. This set also includes two classes of ASes that can be found in the `ASPATH`. The first class of ASes that fits in this definition is represented by ASes that own prefixes only located in a single region , i.e. ASes that do not own an inter-regional infrastructure. Another class of ASes that fits this definition are those ASes that have a single region in common with a neighboring AS.

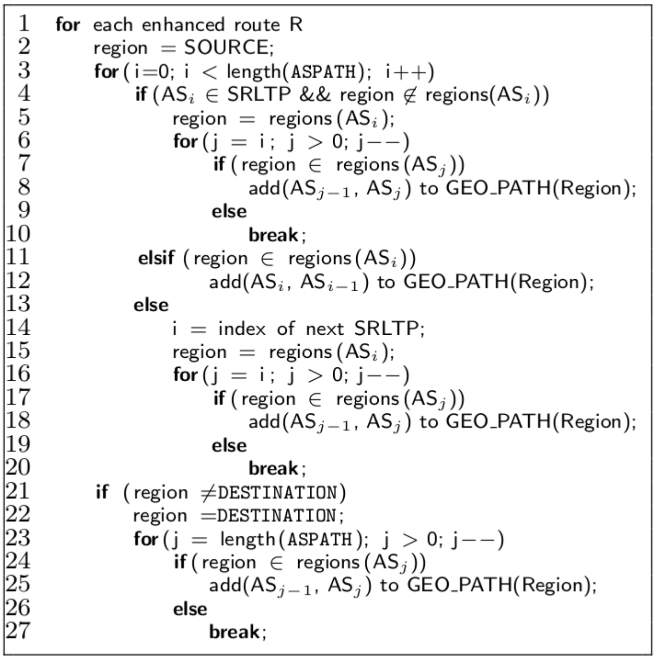

```
for each enhanced route R
    region = SOURCE;
    for (i=0; i < length(ASPATH); i++)
        if ($AS_i \in$ SRLTP && region $\notin$ regions($AS_i$))
            region = regions($AS_i$);
            for (j = i; j > 0; j−−)
                if (region $\in$ regions($AS_j$))
                    add($AS_{j-1}$, $AS_j$) to GEO_PATH(Region);
                else
                    break;
        elsif (region $\in$ regions($AS_i$))
                add($AS_i$, $AS_{i-1}$) to GEO_PATH(Region);
        else
            i = index of next SRLTP;
            region = regions($AS_i$);
            for (j = i; j > 0; j−−)
                if (region $\in$ regions($AS_j$))
                    add($AS_{j-1}$, $AS_j$) to GEO_PATH(Region);
                else
                    break;
    if (region $\neq$DESTINATION)
        region =DESTINATION;
        for (j = length(ASPATH); j > 0; j−−)
            if (region $\in$ regions($AS_j$))
                add($AS_{j-1}$, $AS_j$) to GEO_PATH(Region);
            else
                break;
```

Fig. 2 Geographic Tagging Algorithm

The basic idea is that typically ASes follow a *regional principle* to route their traffic. Inter-regional ASes tend to subdivide their ASes into different areas by exploiting the features of IGP protocols such as OSPF and IS-IS in order to maintain traffic as regional as possible to maximize the performances. Thus, an inter-regional AS does not exploit its inter-regional infrastructure when it is not needed, by representing a SRLTP under certain circumstances. For example consider the enhanced route extracted from the route shown in Figure 1. Geolocating each AS using the methodology described in Section 2, we obtain that `9189`, `8422` and `2597` are located only in Europe, while `3356` is located in every continent. Since also `DESTINATION` is located in Europe, each AS represent a SRLTP.

Geographic AS paths In this step we exploit the information just extracted to geolocate the connections of the AS path of each enhanced route. Following the same regional principle introduced above, inter-regional ASes are typically interconnected on every location where they are co-located, trying to maintain the inter-AS traffic as regional as possible. This means that, for example, if the traffic flows from a source to a destination both located in

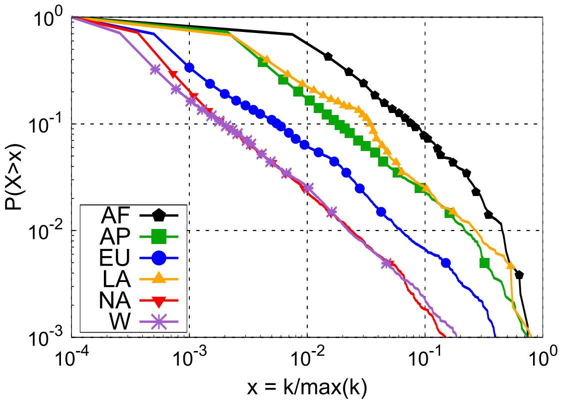

Fig. 3 CCDF of the node degree per continent

region *R* through ASes located in *R*, the traffic is very likely to remain confined in *R* even if these ASes are co-located in other regions. By exploiting these considerations and the SRLTPs identified in the previous step we are able to complete the geographic tagging algorithm, that is presented in Figure 2. The algorithm aims to create a set of geolocated AS connections from each enhanced route. We refer to each enhanced route together with its set of geolocated connections as a *Geographic AS path*. To achieve this, we need to analyze each AS in the `ASPATH` and check whether the connection with its neighboring ASes can be established in the considered region. We start by considering `SOURCE` as the initial region (line 2), and we add all those AS connections that may belong to the considered region (line 18) to the set of AS connections located in that region to the set of AS connections located in that region. The considered region will be changed if an SRLTP (line 4) or if a multi-regional AS (line 13) not located in that region is found. In this last case the change is preceded by a jump to the next SRLTP (line 14). The output of our geographic tagging algorithm is composed by the set of Geographic AS paths, each inferred from the related enhanced routes. Considering the route in Figure 1 and its characteristics shown so far, we infer that the full AS path is located in Europe.

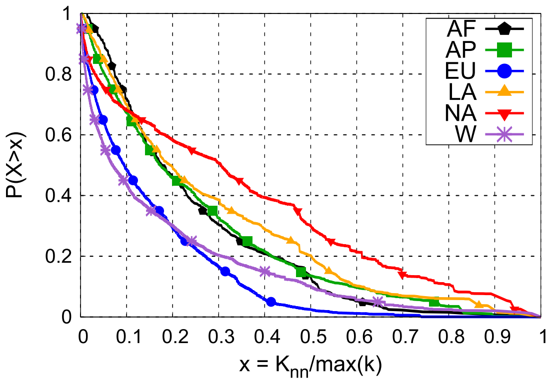

Fig. 4 CCDF of the normalized average neighbor degree per continent

4 Undirected Graph Analysis

We applied the algorithm to obtain geographic AS paths from BGP routes gathered by Route Views, RIS and PCH RCs during April 2012. Then we extracted from each geographic AS path the geolocated connections to create regional AS-level topologies. We also created a global AS-level topology that is composed by all the connections found in each AS path, independently from any geographic information. We found that 41,242 out of 41,778 ASes appear in at least one regional topology. The missing ASes are in 420 out of 536 cases due to missing IP prefixes in the Maxmind database. In the remaining cases the ASes, although being geolocated, do not appear in any regional topology because they do not share any region with the neighboring ASes in any AS path in which they appear. This may be due to the use of BGP multi-hop sessions − where the ASes are actually located in different regions − or due to the mistaken/partial geolocation of prefix(es) by the Maxmind database. In both cases the geographic tagging algorithm is not able to infer where the connection is geolocated. For the same reason, the geographic tagging algorithm is not able to assign a region to 9,480 connections out of 154,322.

The first important result is that the continental topologies

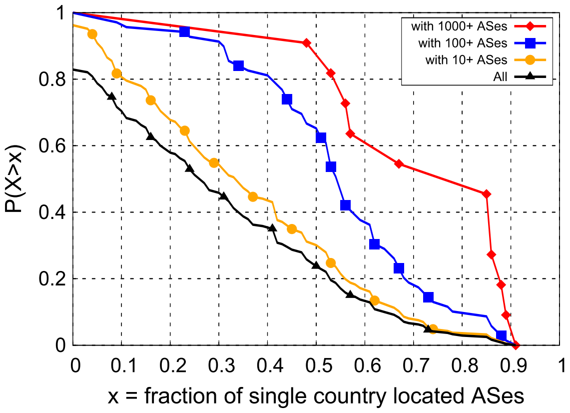

Fig. 5 CCDF of the number of single country located ASes

extracted overlap only slightly, as highlighted by the Jaccard similarity indices[11] computed between pairs of continents for nodes (J_{nodes}) and connections (J_{edges}) and reported in Table 1. This poor overlap is confirmed by the fact that only 5.04% of ASes are located in more than one continent and only 1.57% in more than two, and indicates that the Internet can be subdivided in five independent macro-regions with their own peculiar properties (Figure 3 and 4). This poor overlap also means that the regional principle has been applied by the algorithm only on a small set of connections and thus the largest part of connections is correctly geolocated at the continental level.

The situation changes slightly if the Internet is analyzed at a national level. Despite only 7.99% of ASes are located in more than one country and only 3.2% in more than two, the national topologies extracted by our methodology do not show the same degree of independence found in the continental scenario. In particular, the smaller is a national topology – in terms of nodes – the lower is the probability to find in that topology ASes whose networks do not cross the national borders (Figure 5). This is quite intuitive, since the more a national topology is populated by ASes, the higher is the probability that the related country is economically and technologically advanced and thus home of a large number of local-scoped organizations with a network infrastructure sufficiently developed to meet the AS number assignment criteria [19]. On the opposite, less developed countries are likely to demand their Internet reachability requirements to the few inter-regional ISPs, increasing thus the degree of overlapping with other country topologies.

Table 1 Jaccard similarities indices $J = (J_{nodes}, J_{edges})$

	AF	**AP**	**EU**	**LA**	**NA**
AF	-	0.027,0.026	0.014,0.010	0.046,0.034	0.014,0.015
AP	0.027,0.026	-	0.031,0.052	0.031,0.039	0.042,0.097
EU	0.014,0.010	0.031,0.052	-	0.017,0.017	0.037,0.079
LA	0.046,0.034	0.031,0.039	0.017,0.017	-	0.027,0.033
NA	0.014,0.015	0.042,0.097	0.037,0.079	0.027,0.033	-

A particular scenario is though depicted by countries that are representing critical communication hubs for the Internet and strategic points for IP transit market, such as Great Britain and Germany, that are currently hosting LINX and DE-CIX, two of the most populated Internet eXchange Points (IXPs) of the entire world. The topologies related to these countries are indeed showing a large number of ASes located in more than a single country as well as a large average node degree value (Table 3), reflecting the dense interconnectivity tendency of those ecosystems.

Table 2 Continental topology statistics

Continent	**Edges**	**Nodes**	**Average Degree**	**Max degree**	**Single continent located Ases**
AF	2,182	841	5.19	143	529 (62.90%)
AP	22,087	6,739	6.55	479	5,500 (81.61%)
EU	82,428	17,906	9.21	2,004	16,299 (91.03%)
LA	10,001	2,671	7.49	452	2,117 (89.26%)
NA	48,263	16,250	5.94	2,760	14,391 (88.56%)
W	154,322	41,778	7.39	3,899	-

Further evidence regarding the differences between regional topologies is provided by the properties of the regional graphs summarized in Table 2 and Table 3. The sizes of the topologies differ greatly in terms of nodes, reflecting the different degrees of economic and technological development of the regions.

---

[11] The Jaccard similarity index is a measure that allows to quantify the similarity between pair of sets and is defined as $J(A,B) = |A \cap B|/|A \cup B|$

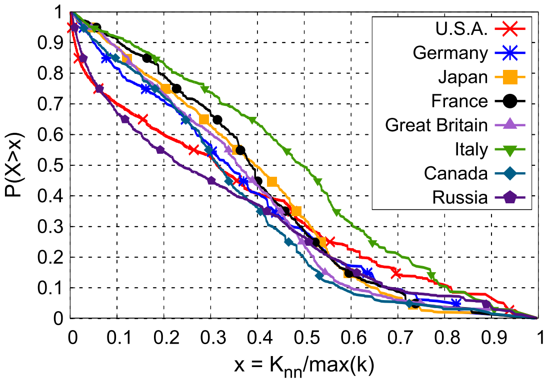

Fig. 6 CCDF of the normalized average neighbor degree per country

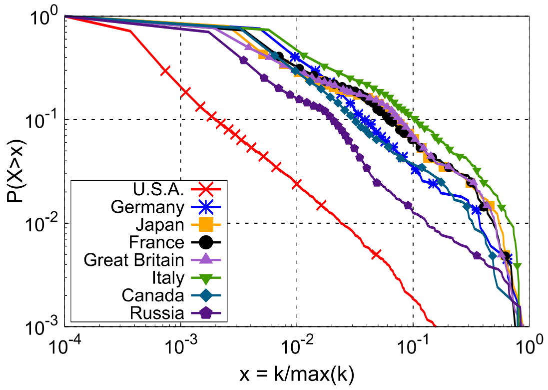

Fig. 7 CCDF of the node degree per country

It is particularly interesting to compare North American and European topologies, which have a quite similar number of nodes but differ significantly in terms of edges. The CCDF of the normalized node degree (Figure 3) shows that the European region is more densely connected than the North American region where, on the other hand, there are ASes with a larger degree and where the number of ASes with a small degree is higher. This suggests quite a hierarchical structure in North America versus a flatter structure in Europe. This is confirmed by the CCDF of the normalized Average Neighbor Degree (Figure 4), which shows that in Europe, ASes tend to connect to ASes with a similar degree, while in North America they tend to connect to ASes with a very large degree.

Table 3 National topology statistics

Country	Edges	Nodes	Average Degree	Max degree	Single country located Ases
U.S.A.	46,158	15,463	5.9543	2718	13,467 (88.86%)
Great Britain	16,909	2,034	16.433	524	1,006 (49.46%)
Russia	13,429	3,911	6.848	574	3,626 (92.71%)
Germany	10,407	1,633	12.516	388	965 (59.09%)
France	4,895	1,043	9.2796	287	620 (59.44%)
Canada	4,301	1,246	6.8162	311	677 (54.33%)
Italy	3,409	768	8.7298	179	502 (65.36%)
Japan	1,864	665	5.5476	208	487 (73.23%)

The differences between these ecosystems reflect the Internet's historical evolution in the respective regions. In North America, especially in the U.S., a small set of large ISPs (e.g. AT&T, Centurylink and Verizon Communications) provide connectivity to all the states. In Europe on the other hand, each country is typically characterized by the presence of a national telco (e.g. Deutsche Telekom and Telecom Italia) which usually own a large part of the national Internet infrastructure, and by the presence of at least one IXP that encouraged the establishment of settlement-free peering connections among local ISPs[12]. In particular, the impact of national telcos can be ascertained by comparing the European normalized knn curve with the normalized knn curves related to the European countries in the G8 forum (Figure 6). In the former, ASes seem to prefer to connect with ASes with a small-medium degree value rather than with the ASes with the largest degree. In the latter, it is possible to see that ASes are likely to connect to ASes with a degree value close to the maximum value in each country. Since the number of large degree ASes in each country is rather small (Figure 7), it is possible to conclude that there are ASes that are extremely important in the national connectivity, while they have not the same importance in

---

[12] More details on the role of IXPs in the development of the Internet can be found in [15], [16] and [17].

the continental scenario.

5 Geography and inter-AS business relationships

The analysis of the undirected graph of the Internet highlighted just part of the extreme complexity of the Internet ecosystem, since it completely lacks the real business model of each element. The Internet consists of a network managed by thousand of different organizations. Some of these organizations, e.g. small or large ISPs, live on the sale of Internet transit to other organizations, others, such as CDNs and search engines, just aim to offer content to end users. Most organizations, though, just care about Internet connectivity. In order to highlight that heterogeneity and to get a better insight into the Internet, some research studies ([6] and [10]) have proposed to annotate the Internet undirected graph with some economic tags that reflect the real business relationships established between different ASes, typically identified as provider-to-customer (p2c), customer-to-provider (c2p), peer-to-peer (p2p) and sibling-to-sibling (s2s). Several existing economic tagging algorithms can be used to obtain an economic tagged AS-level graph. The common approach of these algorithms is to infer economic relationships by exploiting the *valley-free* rule introduced in [6]. For example, [6,7] and [8] infer economic relationships by searching for AS(es) with no provider in the AS path and then applying this rule. On the other hand, [11] and [12] try to find a tag assignment that maximizes the number of valley-free AS paths. Although these approaches were devised for the global Internet, they are still valid in a regional analysis, since the Internet is nothing more than the sum of the regional economic ecosystems that it is made up of. We now present a methodology to infer the regional economic relationships between ASes, by adapting the algorithm proposed in [6]. The original algorithm requires a list of AS paths as input together with their maximum lifespan, defined as the time interval during which there is at least one active route in a RC that includes that AS path in its attributes. It consists of three steps. In the first step all the possible tags for each connection `(A,B)` are computed by applying the knowledge of the list of Tier-1 ASes to infer economic relationships in each of the AS connection of the path. In the second step, all the tags inferred for each connection `(A,B)` are merged to obtain a single economic relationship for each connection. Finally, in the last step, the economic relationship obtained in the second step for the specular connections `(A,B)` and `(B,A)` are merged to infer a single economic relationship between AS `A` and AS `B`. It should be noted that both the merging phases are based on the parameter N_{MAG}: economic tags with lifespans that differ by more than N_{MAG} orders of magnitude are not merged together in order to avoid transient – and potentially wrong [13] − AS paths from distorting the results. The lower the N_{MAG} value, the lower the probability that transient information will affect the results. On the other hand, a low N_{MAG} value also reduces the number of two-way validated[13] economic relationships. The input of the algorithm can no longer be a AS path together with its lifespan. An AS path may be gathered from multiple BGP routers located in different parts of the world and may be used to reach different locations. In this case, some of the AS connections that make up the AS path are likely to refer to different physical links and the lifespan of each AS path may differ depending on the destination region. For example, the AS Path `A B C` used to reach subnets in Asia may be affected by the failure in the link that interconnects `B` and `C` in Tokyo, while the same AS Path used to reach subnets in Europe may not, since `B` and `C` are also interconnected in Paris. To overcome this problem, we exploit the concept of geographic AS path introduced in Section 3 which, together with its lifespan, represents the enhanced algorithm input. The lifespan of a geographic AS path is computed similarly to the lifespan of an AS path, i.e. it is the maximum period of time in which there is at least one active route that includes the related `ASPATH` that is announced from a router in `SOURCE` and reaches at least one subnet in `DESTINATION`. To deal with the new input, we enhance part of the first step of the original algorithm

---

[13] An economic relationship between AS `A` and `B` is considered to be *two-way validated* if an economic tag for the connection `(A,B)` and for the reverse connection `(B,A)` has been found.

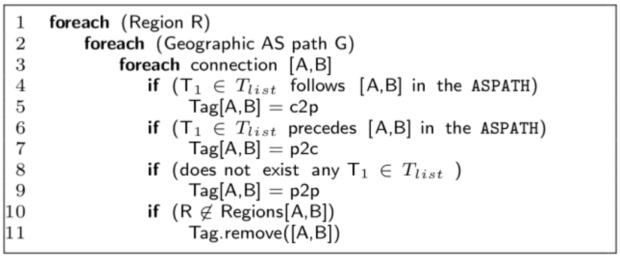

```
1  foreach (Region R)
2      foreach (Geographic AS path G)
3          foreach connection [A,B]
4              if ($T_1 \in T_{list}$ follows [A,B] in the ASPATH)
5                  Tag[A,B] = c2p
6              if ($T_1 \in T_{list}$ precedes [A,B] in the ASPATH)
7                  Tag[A,B] = p2c
8              if (does not exist any $T_1 \in T_{list}$ )
9                  Tag[A,B] = p2p
10             if ($R \notin$ Regions[A,B])
11                 Tag.remove([A,B])
```

Fig. 8 First step of the geographic economic tagging algorithm

in order to obtain only information concerning a specific region *R*, as is shown in Figure 8. The economic tagging of each AS connection that make up the considered Geographic AS Path is initially performed using exactly the same methodology proposed in [6] (line 4-9) exploiting the presence of T1 ASes in the AS path and their provider-free property. Note that the initial tagging phase needs to be performed on the full AS paths, since the presence of a T1 AS need to affect every connection of the AS path, irrespectively of any geographic information. Once the classic tagging phase is completed, we introduce an additional discarding phase (see line 10), in which all the economic tags related to AS connections not geolocated in *R* are removed from the partial results. As a consequence of the new version of the first step, the remaining steps of the original algorithm are fed only with economic tags related to AS connections established in *R* and, thus, the resulting economic tagged topology is related to *R*.

6 Economic analysis

The application of the enhanced economic tagging algorithm to the sets of geographic AS paths allows deeper insights of each regional ecosystem which reveal the real nature of the regional differences that were only deduced from the undirected analysis of the Internet (Section 4). Table 4 and 6 shows the results obtained by applying the enhanced economic algorithm with NMAG = 1, listing the number of economic tags inferred for each region. The choice of NMAG = 1 is justified by the higher reliability of the economic tags inferred, as shown in [6]. We also performed the analysis for other NMAG values, finding that the following inferences still hold. The most relevant characteristic highlighted by the economic analysis is the large proliferation of potential p2p connections in the European ecosystem, representing the 54.76% of the total. This feature is in contrast with the peering behaviors of other regions, where the amount of p2c connections is larger (around 70% of the total) than the amount of p2p connections.

Table 4 Economic Statistics (continental level)

Continent	P2C	P2P	S2S
W	84,696 (54.94%)	68,361 (44.39%)	1,019 (0.66%)
AF	1,383 (66.52%)	671 (32.27%)	25 (1.20)
AP	14,169 (65.31%)	7,342 (33.84%)	183 (0.84%)
EU	34,583 (42.09%)	47,284 (57.55%)	287 (0.34%)
LA	4,962 (51.12%)	4,682 (48.24%)	61 (0.62%)
NA	33,820 (70.47%)	13,838 (28.83%)	331 (0.68%)

Table 5 Economic relationships changes (continental level)

Continent	Tag Changes	Transit to Peering	Among multireg. ASes	Among non-stub ASes
AF	195 (9.5%)	147 (7.16%)	150 (7.31%)	132 (6.43%)
AP	831 (3.84%)	636 (2.94%)	558 (2.57%)	625 (2.88%)
EU	2,308 (2.81%)	1,662 (2.02%)	676 (0.82%)	1,614 (1.96%)
LA	441 (4.56%)	299 (3.09%)	280 (2.9%)	322 (3.33%)
NA	1,177 (2.45%)	904 (1.88%)	721 (1.5%)	904 (1.88%)

Together with the conclusions drawn in Section 4, this allows to understand the real nature of the flat structure of the European Internet ecosystem. This joint analysis shows that Europe is rich in small/medium transit providers that, in addition to offering transit to end-users and stub ASes, tend to establish settlement-free p2p connections among them. The proliferation of these small/medium providers is also the reason for the development in Europe of largely-populated IXPs which in turn facilitated the establishment of settlement-free relationships among ASes, helping to create the large amount of p2p connections just described. This is confirmed by the large percentage of p2p connections found in Germany and Great Britain, that are respectively hosting the DE-CIX and the LINX.

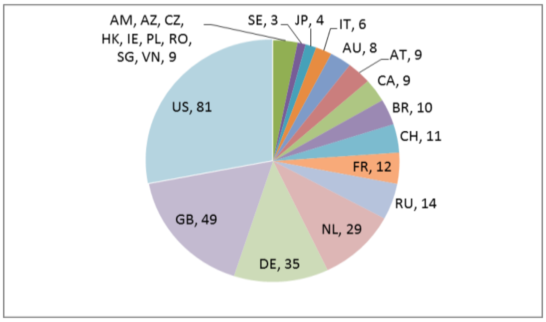

Fig. 9 BGP Full Feeder country distribution

Table 6 Economic Statistics (country level)

Country	P2C	P2P	S2S
U.S.A.	32,435 (70.46%)	13,289 (28.87%)	306 (0.66%)
Germany	3,480 (33.77%)	6,808 (66.06%)	17 (0.16%)
Japan	1,240 (67.06%)	605 (32.72%)	4 (0.21%)
France	1,888 (38.90%)	2,961 (61.01%)	4 (0.08%)
Great Britain	5,073 (30.18%)	11,698 (69.61%)	33 (0.19%)
Italy	1,435 (42.78%)	1,914 (57.06%)	5 (0.14%)
Canada	2,575 (60.47%)	1,649 (38.72%)	34 (0.79%)
Russia	7,106 (53.44%)	6,100 (45.87%)	90 (0.67%)

Table 7 Economic relationships changes (country level)

Country	Tag Changes	Transit to Peering	Among multireg. ASes	Among non-stub ASes
U.S.A.	1171 (2.54%)	894 (1.94%)	771 (1.67%)	899 (1.95%)
Germany	676 (6.6%)	610 (5.95%)	582 (5.68%)	590 (5.76%)
Japan	115 (6.25%)	105 (5.7%)	95 (5.16%)	92 (5%)
France	453 (9.44%)	439 (9.14%)	404 (8.42%)	386 (8.04%)
Great Britain	782 (4.67%)	697 (4.16%)	674 (4.02%)	681 (4.06%)
Italy	237 (7.2%)	213 (6.47%)	210 (6.38%)	201 (6.11%)
Canada	344 (8.18%)	306 (7.28%)	299 (7.11%)	303 (7.2%)
Russia	479 (3.6%)	277 (2.08%)	124 (0.93%)	319 (2.4%)

Another interesting analysis is to investigate how the economic relationships change when the focus is moved from the worldwide scenario to narrower geographic scopes. Inter-regional providers may indeed decide to establish regional p2p connections − exchanging only routes of regional customers − in those region where their pervasiveness is similar, while they may decide to establish a p2c agreement elsewhere. This is especially true for large transit ASes . In Table 4 are summarized the tag changes with respect to the global scenario (W), highlighting in particular the change from a global transit (p2c, c2p, s2s) to a continental peering (p2p). Although the number of these connections may look not relevant at a first glance, it should be considered that these tags are referred to AS connections that compose the core of the region. This is highlighted by the large number of tag changes that involve only non-stub ASes and that involve only multi-regional ASes. . It is interesting also to note the percentage of tag changes of the North American and European region is very similar and is lower with respect to the other continents, highlighting the importance of these two ecosystems in the definition of the BGP economic relationships. A similar analysis is reported in Table 5, where are summarized the number of economic tags changes comparing the national and the global inferences. The number of economic tags changes in every country scenario are slightly larger with respect to the changes experienced by the related continent, with the exception of the U.S.A., confirming its centrality in the Internet ecosystem.

It must be noted though that a part of the p2p connections of every economic topology is caused by the incompleteness of BGP data [18]. Due to the small number of BGP full feeders (FF)[14] currently connected to RCs and to the BGP decision processes on AS border routers, only a small subset of the possible AS paths are actually collected. In these conditions, the algorithm described in Section 5 may not record any AS path including a T1 AS when deciding the nature of a connection (A B). This means that the algorithm is unable to prove the existence of a transit connection between A and B, and thus may conclude that the economic relationship among them is p2p, i.e. A and B announce to each other only their networks and the networks

---

[14] A BGP full feeder is an AS that announce to the RCs an IPv4 space that can be assumed to be the whole Internet IPv4 space [18]

announced to them by their customers. Obviously, the narrower the geographic focus is, the larger is the impact of the incompleteness on the economic topologies inferred, leading the topologies of the countries that host at least a FF to be more complete than others, and the inferences made to be more reliable. Note that, the more the Internet is pervasive in a country, the higher is the probability that it hosts at least one FF, as can be seen in Figure 9, enforcing the quality of topologies and economic inferences made. Thus the topologies and economic inferences made for the G8 members can be considered more reliable than those made for other countries, since they host more than the 70% of the FFs of the RCs. Likewise, the European (136 FFs) and North American (89) topologies and inferences are more reliable than those related to Asia Pacific (17), Latin America (10) and Africa (0). Note that, despite the Europe is hosting a larger number of FFs than North America, the number of p2p connections is still much larger in Europe, thus the rich set of European p2p relationships cannot be imputed only to lack of information, highlighting that the current set of FFs is able to capture the different peering behaviors in the European and North American Internet regions.

7 Conclusions and future work

In this paper we proposed an innovative tagging algorithm in order to geolocate AS connections starting from BGP data. This algorithm allowed us to infer regional AS-level topologies, that we have analyzed both from an undirected graph and economic perspective. From this analysis, we found that the study of the Internet at a global level fails to take into account several characteristics that, on the other hand, play a fundamental role in the regional Internet connectivity. In particular we found evidence of structural differences between the European and the North American Internet topologies, that reflect different historical developments of the Internet in those regions. There is plenty of room for improvements in this study. For example, the geographic tagging algorithm proposed assumes that two ASes establish a BGP session in each region where they are co-located and, consequently, the connection between these ASes may be not correctly geolocated. Moreover, this coarse-grained assumption does not allow to distinguish between regional and worldwide peer-to-peer connections. To overcome these problems, we are currently working to enhance our methodology with a traceroute-based tool, which may discover the real location of connection of currently unknown or uncertain location, allowing to refine our algorithm.

References

[1] W.B. Norton, *The Internet Peering Playbook: Connecting to the Core of the Internet.* Palo Alto, CA: DrPeering Press 2011

[2] P. Matray, P. Haga, S. Laki, I. Csabai, G. Vattay, "On the Network Geography of the Internet," *IEEE INFOCOM*, pp. 126–130, Apr. 2011

[3] S. Laki, P. Matray, P. Haga, T. Sebok, I. Csabai, G. Vattay, "Spotter: A Model Based Active Geolocation Service," *IEEE INFOCOM*, pp. 3173–3181, Apr. 2011

[4] Y. Shavitt, N. Zilberman, "A Structural Approach for PoP Geo-Location," *IEEE INFOCOM Workshop on Network Science For Communication Networks*, pp. 1–6, Mar. 2010

[5] Y. Zhang, R. Oliveira, Y. Wang, S. Su, B. Zhang, J. Bi, H. Zhang, L. Zhang, "A Framework to Quantify the Pitfalls of Using Traceroute in AS-level Topology Measurement," *IEEE JSAC, Special Issue on Measurement of Internet Topologies*, vol. 29, this issue, Oct. 2011

[6] E. Gregori, A. Improta, L. Lenzini, L. Rossi, L. Sani, "BGP and Inter-AS Economic Relationships," *IFIP TC-6 Networking*, vol. 2, pp. 54–67, May 2011

[7] L. Gao, "On Inferring Autonomous System Relationships in the Internet,". *IEEE/ACM Transactions on Networking* vol. 9, no. 6, pp. 733–745, Dec. 2001

[8] R. Oliveira, D. Pei, W. Willinger, B. Zhang, L. Zhang, "The (in)Completeness of the Observed Internet AS-level Structure," *IEEE/ACM Transactions on Networking* vol. 18, no. 1, pp. 109–122, Feb. 2010

[9] Y. Shavitt, Y., N. Zilberman, "A Geolocation Databases Study," *IEEE JSAC* vol. 29, no. 10, pp. 2044–2056, Dec. 2011

[10] G. Huston, "Interconnection, Peering, and Settlements," Proc. INET, Jun. 1999

[11] L. Subramanian, S. Agarwal, J. Rexford, R.H. Katz, "Characterizing the Internet hierarchy from multiple vantage points," *IEEE INFOCOM* vol. 2, pp. 618–627, Jun. 2002

[12] G. Di Battista, M. Patrignani, M. Pizzonia, "Computing the Types of the Relationships between Autonomous Systems," *IEEE INFOCOM*, Mar. 2003

[13] R. Mahajan, D. Wetherall, T. Anderson, "Understanding BGP Misconfiguration," *Proceedings of the 2002 SIGCOMM conference* vol. 32, no. 4, pp. 3–16, Oct. 2002

[14] I. Poese, S. Uhlig, M.A. Kaafar, B. Donnet, B. Gueye, "IP Geolocation Databases: Unreliable?," *ACM SIGCOMM Computer Communication Review* vol. 41, no. 2, pp. 53–56, Apr.

2011

[15] E. Gregori, A. Improta, L. Lenzini, C. Orsini, “The impact of IXPs on the AS-level topology structure of the Internet,” *Computer Communications* vol. 34, pp. 68–82, Jan. 2011

[16] B. Augustin, B. Krishnamurthy, W. Willinger, “IXPs: mapped?,” *IMC 2009*, pp. 336–349, Nov. 2009

[17] K. Xu, Z., Duan, Z.L. Zhang, J. Chandrashekar, “On properties of Internet exchange points and their impact on AS topology and relationship”. In: *IFIP-TC6 Networking*, pp. 284–295, May 2004

[18] E. Gregori, A. Improta, L. Lenzini, L. Rossi, L. Sani, “On the Incompleteness of the AS-level graph: a Novel Methodology for BGP Route Collector Placement”. *IMC* 2012, to appear.

[19] J. Hawkinson, T. Bates, “RFC 1930 - Guidelines for creation, selection, and registration of an Autonomous System (AS)”, Mar 1996