Selecting new BGP Feeders to Address the Incompleteness of the Internet AS-level Graph

Enrico Gregori
IIT-CNR
enrico.gregori@iit.cnr.it

Alessandro Improta
IIT-CNR
Information Engineering Dept.
University of Pisa
alessandro.improta@iet.unipi.it

Luciano Lenzini
Information Engineering Dept.
University of Pisa
l.lenzini@iet.unipi.it

Lorenzo Rossi
IIT-CNR
lorenzo.rossi@iit.cnr.it

Luca Sani
IMT, Lucca
luca.sani@imtlucca.it

Abstract

In the last decade many studies have used the Internet AS-level topology to perform several analyses, from the discovery of its graph properties to its impact on the effectiveness of worm-containment strategies. Yet, the completeness of the BGP data used to reveal the Internet AS-level topology is questionable. In this work we analyze BGP data currently gathered by RouteViews and RIS route collectors, investigating the reasons of its incompleteness. This analysis lead us to define an ad-hoc metric, named p2c-distance, able to identify which ASes hide part of their connectivity from the route collectors due to BGP export policies and how much the information is filtered due to the best route selection applied by each BGP border router crossed before reaching a route collector. This metric allow us to create the basis of an innovative methodology able to select the minimum set of Autonomous Systems from which the route collectors should receive information in order to maximize the coverage of the core of the Internet. The large number of elements found to be required confirms the incompleteness of the current dataset.

1 Introduction

The AS-level topology of the Internet in the last years has been an hot research topic. The most common approach applied to infer the AS-level topology and to perform its analysis is to exploit the availability of data gathered by BGP route collectors (e.g. RouteViews[1], RIPE RIS[2]) and/or Traceroute monitors (e.g. Caida, DIMES). Despite several works have highlighted problems in using traceroute data to infer AS-level information – like aliasing [Key10], biasing [Ach05] and router-to-AS mapping [Huf10, Zha10] – only few works have investigated the limitations of the BGP approach, generally focusing on the incompleteness of data.

Without any doubt, BGP data is much easier to analyze, since the AS-level information is directly contained in the *AS Path* BGP attribute and no further heuristics have to be applied. To obtain such data, BGP route collectors create BGP connections with cooperating ASes – from now on named BGP feeders – to receive their routing information. In this way, it is possible to re-create the dynamics of the Internet as seen from a particular point of view (i.e. the BGP feeder) at any time of the BGP data collection, including all the possible AS-level information that could be

---

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

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

found via a traceroute campaign from the considered BGP feeder.

There are two main projects collecting BGP data freely available on the Internet: RouteViews and RIS. The former is a project developed at University of Oregon and, as the website itself states, was originally conceived as a tool for Internet operators to obtain real-time information about the global routing system from the perspectives of several different backbones and locations around the Internet. Since its birth in 1997, this project has provided an unvaluable amount of BGP data through several route collectors. Nowadays, there are 10 route collectors owned by RouteViews that collect data in MRT (RFC 6397) format that can be easily used in Internet path analysis studies. The latter is a project developed by RIPE that collects and stores Internet routing data from several locations around the globe. In addition to BGP data, RIS offers also several tools that allow the Internet community to easily read and use them. Currently, RIS provides data in MRT format from 13 different route collectors deployed on the largest IXPs. Despite the well-known aim of these projects, several of their BGP feeders do not provide any relevant contribution, as it will be highlighted in this work.

In this paper we firstly analyze BGP data currently gathered by RouteViews and RIS route collectors, highlighting and explaining the causes of its incompleteness. With this analysis, we show that the current view of the Internet is extremely *narrow* – due to the low number of ASes that are actively feeding the route collectors – and *biased* – due to the large size of the feeding ASes. In particular, this top-down view does not allow the route collector infrastructure to discover a large set of p2p connections that may be established among ASes that are part of the lower part of the Internet hierarchy, as already highlighted in [Oli10, HeS09, Coh06]. This classic analysis lead us to develop an innovative metric, named *p2c-distance*, that takes into account the presence of BGP decision processes and BGP export policies crossed by BGP UPDATE messages before reaching a route collector and that allow to better understand the amount of completeness of data gathered. Differently from other works, we are thus able to analyze and quantify the amount of incompleteness of BGP data and the quality of the route collector coverage without exploit any private data, but relying only on the route collector infrastructure. This metric allow us to create the basis of a Minimum Set Cover (MSC) problem aimed to select the minimum number of ASes that should provide full routing information to the route collector infrastructure. Even if the MSC has been proved to be NP-complete [Gar90], we are able to solve it leveraging on the extreme low densities of the covering matrices applying on them classic mathematical reduction techniques followed by a brute force approach on the remaining uncovered components of the original covering matrix.

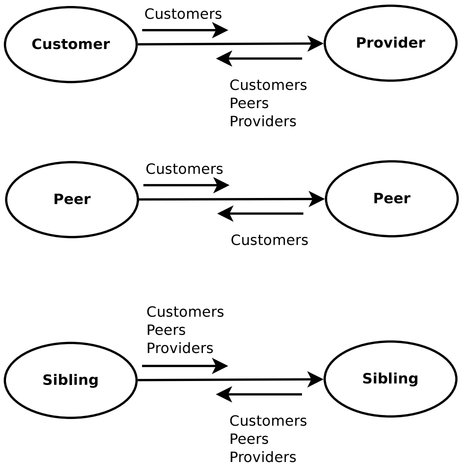

Figure 1: Inter-AS economic relationship

The paper is organized as follow. In Section 2 we briefly discuss the related work. In Section 3 we analyze in detail the incompleteness of data retrieved via BGP RouteViews and RIS route collectors and we introduce the p2c-distance metric. In Section 4 we describe the MSC problem and the algorithm that we developed to solve it. Finally, Section 5 concludes the paper.

2 Related work

Some hints about the incompleteness of BGP data were initially found in [Gov97], but only several years later was perfomed the first attempt [Cha04] to quantify such incompleteness. In this last work authors compared BGP and Internet Registries data highlighting that a large amount of connections (about 40%) was missing in BGP-derived topologies. Since then, this topic has been slightly shelved, and only recently has been brushed up with [Che09] and particularly [Oli10]. In [Che09] it is provided an analysis of the BGP data obtained from each AS participating to RouteViews and RIS finding that these contributions are heavily redundant and that it is possible to reduce the number of ASes participating to the projects obtaining a similar results.

A completely different approach was applied in [Oli10]. The AS-level topology inferred via BGP data was compared with a ground truth composed by proprietary router configurations of two major ISPs (a Tier-1 and a Tier-2 ISP), of two research networks (Abilene and GEANT) and several content providers.

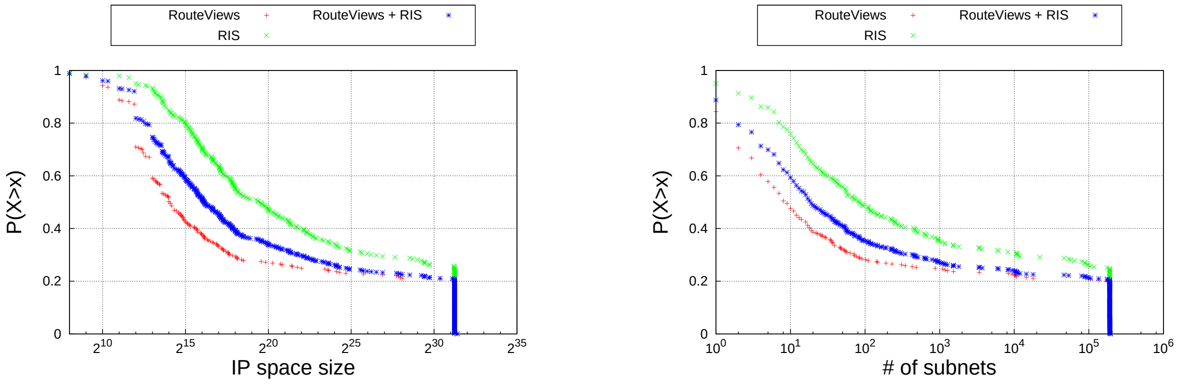

(a) (b)

Figure 2: Amount of information obtained by each project

Despite the results of this approach are not reproducible, this work can still be considered a milestone in the analysis of the incompleteness of BGP data. The most relevant result is that economic relationships established between ASes strongly affect the information that can be revealed by each monitor. The main rationale behind this relies in the valley-free principle introduced in [Gao01] (see Fig. 1), where are identified the most common used type of inter-AS economic relationships – provider-to-customer (p2c), customer-to-provider (c2p), peer-to-peer (p2p) and sibling-to-sibling (s2s) – and are described the related BGP export policy. In particular, an AS announces to its customers the routes obtained by its peers, customers and providers, while it announces to its providers and peers only the routes related to its customers. This means that p2p connections are hidden to any of the providers or peers of a given AS. Basing on that, authors claimed that the largest part of connections that are missing are p2p connections, and that monitors should be placed in the periphery of the Internet to discover them. A similar conclusion about the invisibility of p2p connections was also drawn by [HeS09] that provided an active measurement tool approach based on traceroute to discover missing connections established on Internet Exchange Points (IXPs).

Despite these strong evidences of incompleteness of BGP data, only a few works addressed the problem of placement of new BGP monitors to minimize the lack of information. A remarkable approach was proposed in [Rou08], where the authors extended a model based on techniques developed in biological research for estimating the size of populations to work on the Internet AS graph. Their results showed that thousands of connections are missing, and authors estimated that 700 route monitors would be able to see almost the totality of the connections. Anyhow, the heuristic used in this approach took into account only marginally the inter-AS economic relationships introduced by [Gao01], causing the optimal number of route monitors found to be an heavy underestimation of the real solution.

3 The incompleteness of BGP datasets

From previous works it is well-known that the AS-level information that can be extracted from RouteViews and RIS projects is far from being complete. In order to deeply understand the causes of this incompleteness we now analyze the BGP data provided by each BGP feeder, focusing particularly on the effects of BGP export policies and BGP decision processes. To collect data, each route collector behaves like a BGP border router maintaining its own BGP routing table according to the BGP update messages received from its BGP feeders, but not announcing anything. The BGP data provided by each route collector is composed by the UPDATE messages received during time from its BGP feeders and periodic snapshots of its BGP routing table. In order to perform the analysis of the data provided by each BGP feeder, we downloaded the first RIB snapshot and all the subsequent UPDATEs provided by each route collector during the month of February 2012[3], that will be used in this paper as month of reference.

Fig. 2 shows in detail the amount of information

---

[3] We do not downloaded *all* the BGP routing table snapshots, since the content of the routing table at any given instant of time is determined by the UPDATEs messages that the route collector has received

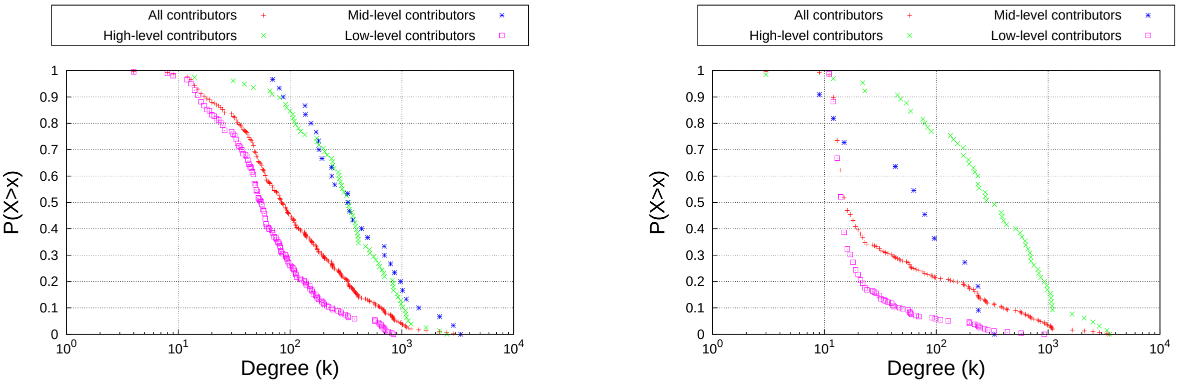

(a) RIS

(b) RouteViews

Figure 3: Degree of ASes connected to route collectors

obtained from BGP feeders of each project in terms of a) IPv4 space and b) number of subnets announced during the month of February 2012, for the analysis. It is interesting to note that, in both projects, only about one over three feeders announces the full routing table to the route collectors. This percentage slightly increases if an IPv4 space equivalent to the size of a /8 subnet is considered as relevant contribution. To better quantify the total contribution of BGP feeders, we subdivide them on the basis of the amount of information that each of them announce to the route collectors: *low-level contributors* announce an IPv4 space smaller than that of a single /8 subnet, *high-level contributors* announce an IPv4 space closer to the full Internet IPv4 space currently advertised[4] (more than 2 billions of IP addresses), while *mid-level contributors* include those ASes in between. Following this subdivision, the number of high-level contributors is 120, i.e. 0.29% of the total. More details about the number of BGP feeders and high-level contributors per each route collector can be found in Table 1. To better understand the nature of these ASes we can analyze Fig. 3, where it is shown the CCDF (Complementary Cumulative Distribution Function) of the node degree of BGP feeders. About 80% of the high-level contributors shows a node degree larger than 100 in both of the projects, and the largest part of them is represented by large ISPs. The Internet extracted from these project, thus, represents more the Internet viewed by routers owned by some of the most important ISPs in the world than the real Internet.

A view from the top of the AS hierarchy, as already pointed out in [Oli10], is not able to discover a large set of connections. Due to the valley-free property, a route collector can gather from an AS at the top of the Internet hierarchy mostly its customer cone (all the p2c connections that allow the BGP feeder to reach the networks announced by its customers), the set of its p2p connections and a mixed set of p2c and p2p connections reachable through its small number of providers. The lower in the hierarchy the contributor is located, the higher is the probability to gather information about an AS path involving a p2p connection via its providers.

Consider for example Fig. 5. In this case, if the route collector is connected to the AS E at the top of the hierarchy, it cannot reveal neither the p2p connection between A and B, nor the p2p connection between C and D. On the opposite, if the route collector is connected to C, then it can reveal every single connection in the scenario.

Location is not the only important factor though. Another fundamental requirement is to *obtain* full routing tables from BGP feeders. This may be obvious, but results shown in Table 1 highlight that neither RouteViews nor RIS achieve to gather full routing information from a large set of ASes.

More details about the amount of information provided by each BGP feeder can be found by analyzing in detail the difference between the direct node degree and the inner node degree. We define as *direct* node degree of a BGP feeder X the cardinality of the set of neighbors of X discovered using only BGP data directly announced by X to a route collector, and as *inner* node degree of a BGP feeder X the cardinality of the set of neighbors of X discovered using BGP data announced by every BGP feeder but X. It is interest-

---

[4] More information about the current IPv4 space advertised can be found at http://www.potaroo.net/tools/ipv4

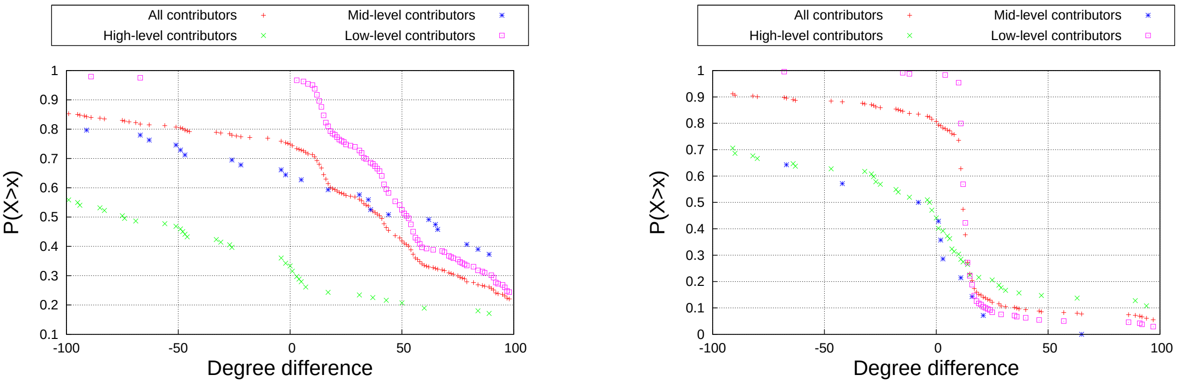

(a) RIS

(b) RouteViews

Figure 4: Degree difference of ASes connected to route collectors

ing to note that the degree difference rarely is equal to zero, as can be seen by the absence of steps in the CCDF plot of Fig. 4. In detail, this difference is equal to zero only for 1 AS out of 394 in RIS and 3 out of 363 in RouteViews. Analyzing the degree difference is also possible to delineate two different classes of behavior of ASes connected to the route collectors: a) ASes that announce just a partial view of the Internet, passing either only routes obtained from its customers (i.e. ASes that consider the route collectors as peers and not as customers) or filtered AS Paths (see negative values of degree difference), and b) ASes that contribute with connections not announced by any other BGP feeders (see positive values), like p2p and p2c connections that are hidden to the other feeders due to the presence of BGP decision processes applied by each AS of the AS path. In detail, the presence of BGP decision processes along the AS path limits the completeness of data collected, since ASes select and announce only the best route per-destination (RFC 4271). In other words, the information that a BGP feeder announces to the route collector is the result of its BGP decision process which, in turn, is fed with routes that are the result of the BGP decision processes of its neighboring ASes, and so on. From the AS-level measurement point of view BGP decision processes are route filters, that can potentially reduce the amount of connectivity information received by each route collector. As consequence, the higher is the distance of an AS from the BGP route collectors, the higher is the probability that some BGP decision processes along the path filter some of its connections. This situation is worsened if we consider also the inter-AS economic relationships and the related BGP export policies applied, since they dramatically lower the set of best routes that an AS propagates towards a certain class of neighbors. Depending on the type of economic relationship established with its neighbor [Gao01], an AS can decide to announce the best routes towards its own networks and the networks owned by its customers or the best routes towards every network of the Internet. The former approach is typically used in announcements made by an AS towards its providers (c2p) or its peers (p2p), while the latter is typically used in announcements made by an AS towards its customers (p2c). Summarizing, an AS announces the full set of routing information only to its customers. It is thus possible to claim that information arriving to a route collector traversing only p2c connections may be richer than information that arrives traversing (also) other types of links.

We exploited these concepts to define a new metric able to indicate the actual completeness of data gathered by the current set of route collectors. We define as *p2c-distance* of an AS X towards another AS Y as the minimum number of consequent p2c connections that connect X to Y in every AS paths involving X or, likewise, the minimum number of consequent c2p connections that connect Y to X. This metric can be easily applied to quantify the distance and the amount of transit connections crossed by each AS to reach a route collector, allowing to reveal which part of the In-

p2c-distance	*# ASes*
1	122
2	366
3	275
3+	40,353

Table 2: Distribution of p2c-distances

Project	*Route collector*	*# BGP feeders*	*# high-level contributors*
RouteViews	route-views2 (Eugene, US)	33	30
	route-views4 (Eugene, US)	15	13
	route-views6 (Eugene, US)	12	0
	route-views.eqix (Ashburn, US)	14	10
	route-views.isc (Palo Alto, US)	13	10
	route-views.kixp (Nairobi, KE)	1	0
	route-views.linx (London, UK)	23	18
	route-views.saopaulo (Sao Paulo, BR)	231	4
	route-views.sydney (Sydney, AU)	8	4
	route-views.wide (Tokyo, JP)	4	2
RIS	rrc00 (Amsterdam, NL)	19	17
	rrc01 (London, UK)	70	10
	rrc03 (Amsterdam, NL)	71	5
	rrc04 (Geneva, CH)	12	8
	rrc05 (Wien, AT)	40	5
	rrc07 (Stockholm, SE)	14	2
	rrc10 (Milan, IT)	17	3
	rrc11 (New York, US)	26	8
	rrc12 (Frankfurt, DE)	45	11
	rrc13 (Moscow, RU)	19	9
	rrc14 (Palo Alto, US)	16	5
	rrc15 (Sao Paulo, BR)	12	6
	rrc16 (Miami, US)	6	1

Table 1: February 2012 monitor contribution

ternet is well-monitored and which part is still a dark zone.

Consider for example the connectivity scenario depicted in Fig. 5. In this case, the route collector R has a p2c-distance of 1 from AS A and B, and of 2 from AS E, while the p2c-distance of C, D and F is not defined. This means that R has a non-negligible probability to discover every p2p connection established by A, B and E. This probability is lower the higher the p2c-distance value of the considered AS is, due to the presence of an higher number of BGP decision processes. On the other hand, it also means that R is not able to reveal the p2p connectivity of C, D and F in any way. Anyhow, R can discover the p2c (c2p) connectivity of each AS of the scenario.

In Table 2 is briefly shown the current distribution of p2c-distance values of the ASes from the route collector infratsructure, computed using the economic tagging algorithm introduced in [Gre11]. Note that we consider to be ∞ the p2c-distance of ASes that cannot reach any route collector using only p2c connections. The largest part of ASes is either too far or not reachable via only c2p connections from any route collector. This means that each of those ASes is representing a possible source of hidden information that have to be better investigated.

4 Monitor placement

Given the definition of p2c-distance, it is straightforward that the complete view of the Internet is obtained only connecting a route collector to each stub AS[5], as already concluded in [Oli10]. Stub ASes are typically owned by Content Delivery Networks (CDNs), local access providers (that provide connectivity to end users but not to other ASes) and organizations that do not have the IP transit as a core business (e.g. banks and car manufacturers). Given their nature, these ASes tend to be customers in the economic relationships established with other ASes, representing a perfect starting point to minimize the p2c-distance of every AS composing the Internet. However, only a small amount of these ASes is really interested in developing p2p connections, as can be inferred by the small percentage (2,941 out of 33,848 stub ASes) of them that is participating in at least an IXP[6], where ASes typically interconnect with settlement-free p2p connections to reduce the amount of their traffic directed to their providers (see [Aug09, Gre10] and [HeS09] for more

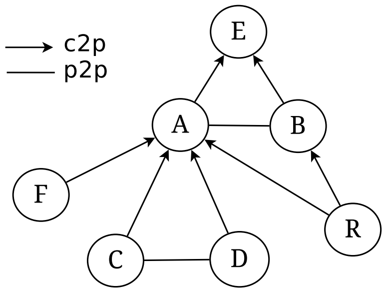

Figure 5: Connectivity scenario I

---

[5] A stub AS is an AS that never appear in the middle of any AS path, i.e. an AS that do not offer transit service to any other AS, thus, from a hierarchical point of view a stub AS is located at the bottom.

[6] We collected the set of ASes participant to at least an IXP by downloading and parsing the participant list webpage of 190 IXPs.

details on IXPs). Since p2c connections are already discovered from route collectors connected to the top of the hierarchy [Oli10], this means that BGP data that could obtained by the largest number of these ASes would represent only redundant data to the measurement system. Since it is not possible to infer a priori which stub AS is actually interested in establishing p2p connections, the measurement system has to introduce each Stub AS (~35,000 elements) as BGP feeder, establishing with each of them a new BGP connection. This is a practically unfeasible task. Given the typical small interest of stub ASes in establishing p2p connections, we believe that a good trade-off between the possibility to discover p2p connections and the feasibility of obtaining such data is represented by not stub ASes, i.e. transit ASes. In particular, we aim to connect each not stub AS to at least a route collector using only sequences of consecutive p2c connections, in order to have at least a non-zero probability to reveal p2p connectivity of the real core of the Internet.

4.1 Problem description

The choice to focus only on transit AS connectivity allow us to reduce the number of required BGP connections from about 35,000 (one per stub AS) to about 8,000 (one per not stub AS). This number can be lowered even more by exploiting the multihoming setup of several ASes, which typically establish c2p connections to multiple providers to improve the reliability of their Internet reachability. To formally obtain the optimal number of BGP feeders we reformulate this problem as a MSC problem that can be described with the following Integer Linear Programming (ILP) formulation:

$$\text{Minimize} \quad \left(\sum_{AS_i \in \mathcal{U}} x_{AS_i} \right) \tag{1}$$

subject to

$$\sum_{AS_i : n \in S^{(d)}_{AS_i}} x_{AS_i} \geq 1 \qquad \forall n \in \mathcal{N} \tag{2}$$

$$x_{AS_i} \in \{0, 1\}, \qquad \forall AS_i \in \mathcal{U} \tag{3}$$

where $\mathcal{U} = \{AS_1, AS_2, \ldots, AS_n\}$ is the set of ASes, $\mathcal{N} \subset \mathcal{U}$ is the set of not stub ASes, $\mathcal{S}^{(d)} = \{S^{(d)}_{AS_1}, S^{(d)}_{AS_2}, \ldots, S^{(d)}_{AS_n}\}$ is a collection of *covering sets* $S^{(d)}_{AS_i}$ that represents the set of ASes in $\mathcal{N}$ that AS_i can reach in at most d hops using only c2p connections, and x_{AS_i} is 1 if $S^{(d)}_{AS_i}$ is part of the final solution, 0 otherwise. Note that AS_i belongs to $S^{(d)}_{AS_i}$ for any $d \geq 0$.

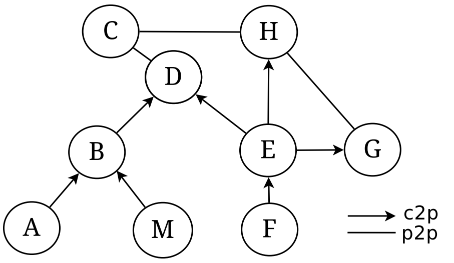

Figure 6: Connectivity scenario II

Consider for example the scenario depicted in Fig. 6. In this case, $\mathcal{U} = \{A, B, C, D, E, F, G, H, I\}$ and $\mathcal{N} = \{B, D, E, G, H\}$. Thus we compute $S^{(1)}_A = \{B\}$, $S^{(1)}_B = \{B, D\}$, $S^{(1)}_C = S^{(1)}_I = \{\emptyset\}$, $S^{(1)}_D = \{D\}$, $S^{(1)}_E = \{E, D, G, H\}$, $S^{(1)}_F = \{E\}$, $S^{(1)}_G = \{G\}$, $S^{(1)}_H = \{H\}$.

The goal of the MSC is to select the minimum number of BGP feeders from $\mathcal{U}$ such that the p2c-distance of any not stub AS from at least one of them is at most d. For example, one of the optimal solutions to cover the not stub ASes in Fig. 6 is $\mathcal{C}^{(1)} = \{S^{(1)}_B, S^{(1)}_E\}$. The parameter d defines the maximum number of BGP decision processes[7] that the information announced by each not stub AS will traverse before reaching a BGP feeder and, thus, indicates the number of filters encountered that can cause loss of information. Notice that imposing $d = 0$ implies that the solution is composed by the entire set of not-stub ASes (~7,200), which in practice means that the measurement system will receive the full-routing table directly from each of them, minimizing the number of BGP decision processes. The larger the value of d gets, the heavier is the effect of BGP decision processes but the smaller is the number of required BGP feeders and, thus, the numbers of BGP connections that have to be established, making the solution more feasible.

4.2 Positioning algorithm

We can model this problem with a directed graph $G^{(d)}(\mathcal{V}, \mathcal{E})$ where the set of nodes $\mathcal{V}$ is the set of ASes, while an edge in $\mathcal{E}$ directed from node i to node j represents that the not stub AS_j is contained in $S^{(d)}_{AS_i}$. The adjacency matrix related to the graph $G^{(d)}$ is a $|\mathcal{V}| \times |\mathcal{E}|$ matrix $A^{(d)}$ such that $A^{(d)}_{ij} \in \{0, 1\}$, and $A^{(d)}_{ij} = 1$ if $(i, j) \in \mathcal{E}$. With our data set, each adjacency matrix has 41,116 rows – one per AS – and 7,268 columns – one per not stub AS. Since this is a MSC problem, proven to be NP-hard, the only effective method to ob-

---

[7] The number of BGP decision processes encountered before reaching a route collector is $d + 1$, since BGP feeders introduce an additional BGP decision process before providing BGP data to the route collectors.

```
1  
2  Input: distance $d$, $S^{(d)}_{AS_i}$ $\forall AS_i \in \mathcal{U}$
3  
4  $Pool = \mathcal{U}$
5  $\mathcal{P}$ =<set of high-level contributor ASes>
6  $Old\_Pool = \emptyset$
7  
8  while $Pool \neq Old\_Pool$
9      $Pool = Old\_Pool$
10     foreach $AS_i \in \mathcal{N}$
11         if $|\{AS_k \mid AS_k \in Pool \wedge AS_i \in S^{(d)}_{AS_k}\}| == 1$
12             foreach $AS_j \in Pool$
13                 $S^{(d)}_{AS_j} = S^{(d)}_{AS_j} - S^{(d)}_{AS_k}$
14                 if $|S^{(d)}_{AS_j}| == 0$
15                     remove $AS_j$ from $Pool$
16             remove $AS_k$ from $Pool$
17             insert $AS_k$ into $\mathcal{P}$
18 
19     foreach $AS_i \in Pool$
20         foreach $AS_j \neq AS_i \wedge AS_j \in \mathcal{P}$
21         if $S^{(d)}_{AS_i} \subset S^{(d)}_{AS_j}$
22             remove $AS_i$ from $Pool$
23         else if $S^{(d)}_{AS_j} \subset S^{(d)}_{AS_i}$
24             remove $AS_j$ from $Pool$
25         else if $S^{(d)}_{AS_i} = S^{(d)}_{AS_j}$
26             remove $AS_j$ from $Pool$
27 
28 $subgraphs$ = find_components($Pool$)
29 foreach $component$ in $subgraphs$
30     $\mathcal{P}_{component}$ = brute_force($component$)
31     $\mathcal{P} = \mathcal{P} \cup \mathcal{P}_{component}$
32 
33 Output: solution set $P$, interchangeable AS set $A$
34 
```

Figure 7: MSC resolver algorithm

tain the exact solution is to use a brute-force approach. However the size of the adjacency matrix makes this approach unfeasible, since it would require to test an huge number of combinations.

Nevertheless, we can still find the cardinality of the optimal solution by exploiting the extremely low density of the adjacence matrices $A^{(d)}$, some logical considerations and mathematical techniques to reduce the problem in small sub-problems solvable with the brute-force approach. Details about this methodology are depicted in Fig. 7. First of all, it is possible to consider the current high-level contributors as part of the initial solution, since it is not needed any additional cost to obtain their routing information with minimum distance (see line 5). Another useful consideration (see line 10) to reduce the problem is that all the AS_j in $\mathcal{U}$ found to be the sole covering solutions of at least an AS_i are certainly part of the solution. This means that it is possible to reduce the number of columns that have to be investigated and, as consequence, also the number of rows (see line 11-17), since all the elements found in $S^{(d)}_{AS_j}$ can be considered to be covered by them. The last consideration concerns the elements covered by each set in $\mathcal{S}^{(d)}$ (see line 19). If a set $S^{(d)}_{AS_j}$ is fully included in another set $S^{(d)}_{AS_i}$, it is possible to claim that every optimal solution with AS_i, imply an optimal solution with AS_j. Thus, it is possible to continue the algorithm dropping the dominated AS without affecting the cardinality of the final result [Mec04]. The combination of these two steps may lead to new possible scenarios on which the two considerations can be still applied, leading to new possible reductions. Once it is not possible to reduce the problem using these considerations anymore, the problem can be further reduced applying the usual mathematical techniques to reduce the adjacence matrix $A^{(d)}$ (see line 28), leading to a set of small and independent MSC sub-problems that can be solved via brute-force approach (see line 30). The output of the algorithm is the number of additional BGP feeders that are required to have at least a non-zero probability to discover each p2p connection established by not stub ASes.

4.3 Results

We applied the above methodology to find the cardinality of the optimal set of ASes candidate to become BGP feeders. Results are shown in Table 3. Obviously, the number of new BGP feeders required to have a distance equal to 0 from each not stub ASes is equal to the number of not stub ASes that are not yet connected to any route collector. Increasing the value of d, the number of new BGP feeders required decrease, but at the same time increase also the amount of decision processes that may filter AS-level connectivity information. The large number of BGP feeders required gives a perfect snapshot of the current status of the data collected by route collectors deployed by RouteViews and RIS. As shown in the previous sections, the measurement system is unable to retrieve complete information from the largest part of the transit ASes, as confirmed by the still large number of new BGP feeders required even using large values of d. In detail, RouteViews and RIS should establish about 30 times

d	*# BGP feeders*
0	7,268
1	4,352
2	3,985
3	3,920

Table 3: Results of the positioning algorithm

as many BGP connections with high-level contributors as those currently established to become a measurement system with a non-zero probability to discover currently invisible connections. As a future work we plain to find and characterize the whole set of ASes that could be part of at least an optimal solution.

5 Conclusions

BGP route collector projects developed so far are extremely valuable for researchers, representing the most reliable source of information to gather data about the real infrastructure of the Internet. Nevertheless, BGP data nowadays is collected only from a small amount of ASes, limiting the goodness of the inferences that can be drawn from its analysis. Studies in this topic must be fully aware of the high level of incompleteness that BGP data shows, since a topological analysis of the Internet as viewed from these monitors can be compared to an analysis of the roadmap of a given country in which are known only the highways, while almost all the roads that interconnect part of the highways are missing.

In this paper, we highlighted and quantified the amount of BGP data obtained, showing that the current BGP route collector projects are able to reveal the complete information about a very small number of ASes. The current BGP feeders are typically large ASes such as provider-free large ISPs, implying that the current vision of the BGP route collector projects cannot catch any of the p2p connections that small or medium sized ASes may establish. The only solution to face this incompleteness is to increase dramatically the number of BGP feeders of these projects. In this paper we propose a systematic methodology to infer the number of ASes that the route collector projects should introduce as BGP feeders in order to maximize the amount of information collectable and minimizing the costs to obtain such information. We are aware that this kind of data is extremely hard to obtain, but we also believe that the largest problem in data gathering is that ASes are not stimulated enough to join any of the running projects. In particular, it might be a good idea to create services that would be valuable for ASes in change for full routes, following the *do ut des* principle. Otherwise, it might be useful to exploit alternative tools to improve the amount of available data. In particular, some of traceroute-based projects (e.g. [Sha05] and [Che09]) are able to bypass the reluctance in disclosing the routing information of AS owners by placing agents directly on user applications and, thus, obtaining data that would not be collected otherwise.

References

[Ach05] D. Achlioptas, A. Clauset, D. Kempe, C. Moore. On the Bias of Traceroute Sampling: or, Power-law Degree Distributions in Regular Graphs. *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing* – Volume 2, pp 694–703, 2005.

[Aug09] B. Augustin, B. Krishnamurthy, W. Willinger. IXPs: mapped? *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference (IMC)* – pp. 336–349, 2009

[Gre10] E. Gregori, A. Improta, L. Lenzini, C. Orsini. The Impact of IXPs on the AS-level Topology Structure of the Internet. *Computer Communications* – Volume 34, issue 1, pp. 68–82, 2010

[Cha04] H. Chang, R. Govindan, S. Jamin, S. Shenker, W. Willinger Towards Capturing Representative AS-level Internet Topologies *Computer Networks* – Volume 44, issue 6, pp. 737–755, 2004

[Che09] K. Chen, C. Hu, W. Zhang, Y. Chen, B. Liu, On the Eyeshots of BGP Vantage Points *Proceedings of the 28th IEEE conference on Global telecommunications* – pp. 3558–3563, 2009

[Gao01] L. Gao On Inferring Autonomous System Relationships in the Internet *IEEE/ACM Transactions on Networking (TON)* – Volume 9, issue 6, pp. 733–745, 2001

[Gar90] M. R. Garey, D.S. Johnson *Computers and Intractability; A Guide to the Theory of NP-Completeness.* W. H. Freeman & Co., 1990.

[Gao01] L. Gao On Inferring Autonomous System Relationships in the Internet *IEEE/ACM Transactions on Networking (TON)* – Volume 9, issue 6, pp. 733–745, 2001

[Gre11] E. Gregori, A. Improta, L. Lenzini, L. Rossi, L. Sani. BGP and Inter-AS Economic Relationships. *Proceedings of the 10th international IFIP TC 6 conference on Networking* – Volume 2, pp 54–67, 2011.

[Huf10] B. Huffaker, A. Dhamdhere, M. Fomenkov, kc. Claffy, L. Sani. Toward Topology Dualism: Improving the Accuracy of AS Annotations for Routers. *Proceedings of the 11th international conference on Passive and active measurement (PAM)* – pp 101–110, 2010.

[Key10] K. Keys Internet-Scale IP Alias Resolution Techniques *ACM SIGCOMM Computer Communication Review* – Volume 40, issue 1, pp. 50–55, 2010

[Mec04] S. Mecke, D. Wagner Solving Geometric Covering Problems by Data Reduction *Proceedings of the 12th Annual European Symposium on Algorithms* – Volume 3221, pp. 760–771, 2004

[Oli10] 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* – Volume 18, pp. 109–122, 2010

[HeS09] Y. He, G. Siganos, M. Faloutsos, S. Krishnamurthy Lord of the Links: A Framework for Discovering Missing Links in the Internet Topology *IEEE/ACM Transactions on Networking* – Volume 17, pp. 391–404, 2009

[Coh06] R. Cohen, D. Raz The Internet Dark Matter - on the Missing Links in the AS Connectivity Map *Proceedings of IEEE INFOCOM 06. 25TH IEEE International Conference on Computer Communications* – pp. 1–12, 2006

[Gov97] R. Govindan, A. Reddy An Analysis of Internet Inter-Domain Topology and Route Stability *Proceedings IEEE INFOCOM 97. 16TH IEEE International Conference on Computer Communications* – pp. 850–857, 1997

[Rou08] M. Roughan, S.J. Tuke, O. Maennel. Bigfoot, Sasquatch, the Yeti and Other Missing Links: What We Dont Know About the AS Graph *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement conference* – pp. 325–330, 2008

[Zha10] Y. Zhang, R. Oliveira, H. Zhang, L. Zhang Quantifying the Pitfalls of Traceroute in AS Connectivity Inference *Proceedings of the 11th international conference on Passive and active measurement (PAM)* – pp. 91–100, 2010

[Sha05] Y. Shavitt, E. Shir DIMES: let the internet measure itself *ACM SIGCOMM Computer Communication Review* – pp. 71–74, 2005

[Che09] K. Chen, D. R. Choffnes, R. Potharaju, T. Chen, F. E. Bustamante, D. Pei, Y. Zhao Where the Sidewalk Ends: Extending the Internet AS Graph Using Traceroutes from P2P Users *Proceedings of the 5th international conference on Emerging networking experiments and technologies (CoNEXT)* – pp. 217–228, 2009