

Quantifying ASes Multiconnectivity Using Multicast Information

Pascal Mérindol, Virginie Van den
Schrieck, Benoît Donnet, Olivier
Bonaventure
Université catholique de Louvain
Louvain-la-Neuve – Belgium
firstname.name@uclouvain.be

Jean-Jacques Pansiot
Université de Strasbourg
Strasbourg – France
pansiot@unistra.fr

ABSTRACT

Redundant connectivity (or multiconnectivity) between adjacent autonomous systems (ASes) is important for inter-domain traffic engineering and fast recovery in case of failures. However, the redundancy of ASes business relationship links has not been quantitatively studied, mainly due to the difficulty of obtaining relevant data. In this paper, we show that the `mrinfo` multicast monitoring tool can provide useful data about the Internet topology and such redundant links in particular. Our analysis relies on more than four years of daily queries to about ten thousand routers mapped into more than two hundred ASes. We demonstrate that peering links between ASes are frequently redundant. In particular, our analysis shows that more than half of the studied ASes pairs are connected through multiple physical links. We then refine our analysis by considering the different types of ASes and their business relationships. A particular result of our analysis is that at least 75% of the peer-to-peer relationships between adjacent Tier-1 ASes are redundant, i.e., the connections between these ASes involve several physical links. Our analysis is conservative, providing so a lower bound, as some links might not be seen by `mrinfo` due to ISPs filtering policies.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Network topology

General Terms

Measurement

Keywords

network topology, `mrinfo`, ASes connectivity, multicast

1. INTRODUCTION

Today's Internet is composed of more than 30,000 administratively independent entities called *Autonomous Systems* (ASes). Many researchers have developed techniques to discover its topology [1, 2] at two different levels: the *router level*, describing how routers are interconnected, and the *AS level*, defining the logical interconnection between ASes. While the router level graph is mostly built using *traceroute* (see Donnet and Friedman for a survey [1]), the AS level graph is typically constructed through BGP dumps (see references cited by Haddadi et al. [2]).

In the AS level graph, two ASes are interconnected with a single logical link, hiding the potential multiple physical links between their routers. With this AS level graph, we thus lose information on how each pair of ASes is connected. Note that ISPs peering agreements often include requirements for establishing physical disjoint links [3].

Such a redundancy of interdomain links is of the highest importance. Indeed, interdomain links fail [4] and it is important to recover from these failures. Solutions have been proposed to recover from such failures locally [4] or by using an alternate AS-level path [5, 6]. Furthermore, several researchers have proposed different techniques to exploit such redundant links [7, 8, 9]. However, up to now and to the best of our knowledge, this redundancy has not been studied in details. This is mainly due to the difficulty in obtaining a data set relevant for that purpose. Indeed, standard ways of collecting topological information, such as *traceroute* and BGP dumps, are known to be inaccurate [10, 11, 12].

In this paper, based on a large and unexploited dataset obtained with the `mrinfo` [13, 14] IP multicast tool, we tackle the problem of quantifying the *multiconnectivity* of ASes. By multiconnectivity, we mean the redundancy between neighbor ASes. Note that ASes multiconnectivity is different from *multihoming* [15], where a stub AS is connected to several providers.

Our quantification is based on a four year daily dataset collected by `mrinfo` [14], a tool that silently discovers all interfaces belonging to an IPv4 multicast enabled router and the IP addresses of its neighbors. In contrast to *traceroute*-like probing, the major advantage of `mrinfo` is to capture links that are not used in the normal forwarding plane such as backup links. Indeed, even if a link has a high IGP weight or a low BGP local preference, `mrinfo` is able to capture it. Further, `mrinfo` is able to discover all interfaces on a router, avoiding so the risk of missing links or approximations during the alias resolution.

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.

IMC'09, November 4–6, 2009, Chicago, Illinois, USA.

Copyright 2009 ACM 978-1-60558-770-7/09/11 ...\$10.00.




Figure 1: `mrimfo` output

Our analysis demonstrates that more than half of the studied ASes pairs are connected through multiple physical links. In addition to this general analysis, we provide a classification of multiconnectivity that is based on the classical AS hierarchy (Tier-1, Transit, Stub) and AS business relationships (peer-to-peer, customer-to-provider, provider-to-customer). Among others, we show that roughly 75% of the peer-to-peer relationships between adjacent Tier-1 ASes are composed of multiple links. The results of our evaluation should allow for more accurate topology generation, by providing router-level information on the interconnections between ASes.

The remainder of this paper is organized as follows: Sec. 2 discusses the `mrimfo` tool and the dataset we collected; Sec. 3 analyzes the ASes multiconnectivity; Sec. 4 positions this paper regarding the state of the art; finally, Sec. 5 concludes this paper by summarizing its main contributions and discussing future directions.

2. MRINFO

In the late 1980s, the developers of IP Multicast designed the MBone, an overlay network composed of tunnels that interconnected workstations running an implementation of the Distance Vector Multicast Routing Protocol [16] (DVMRP). Several tools have been developed to monitor and debug the MBone [17]. Most of these tools have been deprecated with the replacement of DVMRP by the Protocol Independent Multicast (PIM) family of multicast routing protocols with one notable exception : `mrimfo` [13].

`mrimfo` uses the Internet Group Management Protocol (IGMP) [18]. IGMP was designed to allow hosts to report their active multicast groups to a multicast router on their LAN. Most IGMP messages are sent with a Time-to-Live of 1. However, DVMRP has defined two special types of IGMP messages that can be used to monitor routers [16]. Although current IPv4 multicast routers do not use DVMRP anymore, they still support these special IGMP messages. Upon reception of an IGMP ASK_NEIGHBORS message, an IPv4 multicast router will reply by sending an IGMP NEIGHBORS reply message that provides the list of all its local interfaces with some information about their state. Cisco and Juniper routers also report in the IGMP NEIGHBORS message the version of their operating system. Fig. 1 shows an example of the usage of `mrimfo` to query the router R₂ (1.1.0.2 is the responding interface of R₂). `mrimfo` reports that this router is directly connected to R₀ (through the in-

terface 1.1.0.1) and to two AS border routers (ASBRs) R₃ (through the interface 2.2.4.2) and R₄ (through the interface 2.2.1.2). This information is obtained by sending a single IGMP message. In practice, `mrimfo` provides similar information as a `show` command on the router’s command line interface.

Compared to traceroute and its variants, `mrimfo` has both drawbacks and advantages. The main drawback of `mrimfo` is that it can only be used on routers having IPv4 multicast activated. IPv4 multicast is not always enabled in IP networks, but thanks to the deployment of video or television services that rely on IP multicast, more and more ISP networks have enabled multicast. IGMP, like ICMP which is used by traceroute, can be disabled or rate limited on routers. More and more operators filter or rate limit ICMP, but fortunately for our measurements not all of them are aware of this special usage of IGMP.

The main advantage of `mrimfo` is that, in a single IGMP reply, a router lists all its interfaces, their IP addresses, and the IP addresses of its neighbor routers. Thus, `mrimfo` does not suffer from the alias resolution problems affecting traceroute. Second, all links of a responding router are captured, even if the IGP weight of a link is high and no data packet are forwarded over it. Furthermore, the IGMP monitoring load is very small compared to traceroute. Indeed, with `mrimfo` it is possible to collect the topology of a multicast enabled network by sending a single packet to each router.

Standard traceroute is only able to discover a single path from the source to the destination. To discover more topology information, it is required to increase both the number of destinations and vantage points. However, standard systems that extensively trace the Internet, such as the recent *Archipelago* [19], mostly rely on a small set of sources (roughly 20 machines spread around the world). Paris Traceroute [20] is able to discover load balancing routers, as well as the set of paths joining those load balancing routers. However, this works mostly for intra-domain routers, BGP load balancing being much more difficult to detect. `mrimfo` is able to discover all links between routers from a single source if domains authorize multicast. In particular, `mrimfo` is able to report backup links inside and between domains. Intra-domain backup links are links configured with a high IGP weight. They are not used to forward data packets and thus cannot be discovered by traceroute. Interdomain backup links can be configured by using different types of BGP policies. A frequent policy is to attach a low BGP local-preference attribute to the routes learned from a backup link. For instance, on Fig. 1, if the link between R₂ and R₃ has a low local preference, it will never be discovered by traceroute-like exploration but `mrimfo` will report it.

2.1 Dataset and Collection Methodology

Since May 1st, 2004, we collect the `mrimfo` data from a host located at the University of Strasbourg, France. In this paper, we consider the data collected until October 31st, 2008. The entire dataset is publicly available [14]. The collection script maintains a list of known multicast routers. Each day, it sends an IGMP query to each of these routers to collect their list of interfaces. If a new router is discovered in a received answer (in the list of outgoing interfaces of an already known router), this router is also queried in turn. For instance, in Fig. 1, when R₂ replies to `mrimfo`, in the next iteration of our collection script, an IGMP query will




Figure 2: Data captured by `mrinfo` over time

be sent to R_0 , R_4 , and R_3 . These recursive queries stop at unresponsive routers or when all known routers have been queried. Each router is queried at most once per day. It is worth to notice that a router not replying to an `mrinfo` probe during a given day will not be queried the days after except if it appears again in a list of captured interfaces. Additional information about the collection script and the raw `mrinfo` dataset may be found in [21].

Fig. 2 provides some information on data captured by `mrinfo` during the four years of collection. In particular, Fig. 2(a) provides the number of routers replying to `mrinfo` during the four years of the data gathering, while Fig. 2(b) provides the number of IP interfaces listed by `mrinfo` on these routers every day. We see that the best period for `mrinfo` was between 2006 and 2007. On average, `mrinfo` was able to daily discover roughly 10,000 different routers while scanning 100,000 interfaces, and 1,000 ASBRs belonging to 200 ASes (not shown in Fig. 2). In particular, we collect data about six Tier-1 ASes such as Sprint, UUnet, and Level3. For instance, during the whole period analyzed, we received responses from, on average, 684 routers belonging to the Sprint network. Note that for Fig. 2(b) and the remainder of this paper, we remove interfaces with non publicly routable IP addresses. The addresses of those interfaces are a subset of the special-use IPv4 addresses described in RFC 3330 [22]. Specifically, we eliminate the private IP address blocks 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. We also remove the loopback address block 127.0.0.0/8 and the 0.0.0.0 address. On average, 25% of the interfaces collected by `mrinfo` fall within this category. We also remove all multicast tunnel interfaces to focus on actual physical links.

3. AS MULTICONNECTIVITY ANALYSIS

In this section, we quantify the ASes multiconnectivity based on the `mrinfo` dataset. We first explain the methodology we follow throughout this quantification (Sec. 3.1). We next provide a broad and general view of the ASes multiconnectivity (Sec. 3.2). We then deepen our analysis by evaluating the impact of peering relationships on ASes multiconnectivity (Sec. 3.3).

3.1 Methodology

We first merge the daily traces in weekly traces to cope with the fact that some routers might have been missed on a given day but unlikely during an entire week. Each weekly trace contains the interfaces of routers that replied at least

once during the week. Since the number of routers responding with a different list of interfaces during a week is not significant, we also merge the list of captured interfaces of each router during a week.¹ For information, less than 1% of the interfaces collected by `mrinfo` during a given week correspond to new interfaces appearing in previously known routers. Among the four years of data, we select 100 weekly traces uniformly distributed on the whole `mrinfo` data set. With this subset of weekly traces, as mentioned later in Sec. 3.2, we achieve a sufficient accuracy for the purpose of our study.

Second, we map each IP address with the corresponding AS based on a longest prefix matching. This mapping is done by using a BGP table from the corresponding week obtained from Routeviews (the collector located in Oregon). We remove from weekly data traces the IP addresses that could not be mapped (less than 1% of the IP addresses collected) as well as those advertised by multiple ASes (less than 0.3% of the addresses fell in this category). We identify a responsive router as an ASBR if its outgoing interfaces are mapped to several ASes or if there exists another replying router (belonging to a neighbor AS) connected to it through one or several of its interfaces. It is worth to notice that the prefix assigned to a link connecting two ASBRs from different ASes is taken from the address space of one of the ASes. In order to affect an AS number to a given ASBR, we use a simple election process. We select the AS with the maximum number of mapped interfaces. However, in less than 0.5% of the cases (i.e., equality case), we arbitrarily select the first AS by lexicographical order. Developing smarter algorithms for assigning AS number to ASBR is subject to further work.

In addition, based on CAIDA data [23], we assign a level to each AS: *Tier-1*, *Stub* and *Transit*. An AS is identified as a Tier-1 if it does not maintain any customer-to-provider relationship. Similarly, an AS is labeled as Stub if it does not have any provider-to-customer relationship. With the CAIDA AS ranking data set [23], we then label business relationships: *customer-to-provider* (c2p), *provider-to-customer* (p2c), *peer-to-peer* (p2p), and *siblings* (s2s). In contrast to Routeviews data, the CAIDA data set is not refreshed daily. For instance, between 2004 and 2005, it was only monthly refreshed, whereas since 2006, it is weekly refreshed. Therefore, we decide to use the data which is chronologically the closest to the BGP table provided by Routeviews.²

Note that the two mapping phases (the IP to AS and the business relationships) act as a filter to suppress from our analysis the connections between ASes that do not correspond to actual BGP peering sessions. Indeed, when an Internet Exchange Point (IXP) is captured by `mrinfo` queries, it will find the IP addresses of all routers attached to the IXP. Although present on the IXP, all these routers do not necessarily peer which each other. To avoid over-estimating the multiconnectivity in this case, we decide to eliminate ASes connections that do not appear in BGP table dumps. This filtering may also eliminate actual BGP peering sessions as a false negative but ensures that results presented

¹We assume that there are fewer router configuration changes during a week than the number of routers and interfaces not captured by `mrinfo` during a day.

²We assume that business relationships do not change sufficiently frequently to introduce a significant bias in our analysis.

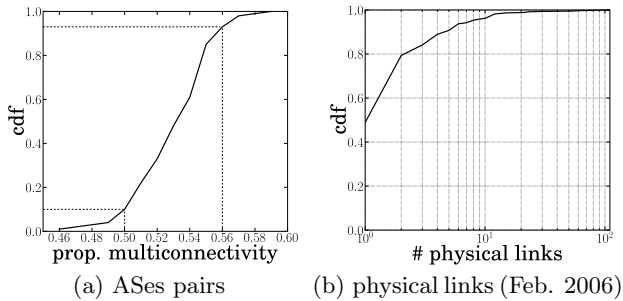


Figure 3: Multiconnectivity

here are lower bounds of the actual connectivity redundancy between two neighbor ASes.

In the following of the paper, we present multiconnectivity results as cumulative distribution function (CDF) plots. Instead of evaluating the multiconnectivity evolution over time, we plot the proportion of ASes pairs that are multiconnected (i.e., the ASes pairs are connected with, at least, two physical links) as CDF. On each plot, the horizontal axis gives the proportion of multiconnected ASes pairs while the vertical axis gives the CDF. Note that in the following we analyze ASes pairs multiconnectivity as directed connections, such that each directed connection between neighbor ASes is considered independently.

3.2 Overall Analysis

Fig. 3 provides a global overview of the ASes multiconnectivity. In particular, Fig. 3(a) quantifies the ASes pairs multiconnectivity. We notice that, for 90% of the processed weeks, more than 50% of the ASes pairs captured with `mrinfo` present a redundant connectivity (see the dashed lines). In more than 80% of the cases, between 50% and 56% of the ASes pairs are redundant. On average, we measure that 54% of ASes pairs present a redundant connectivity. This result is based on the 100 weeks we processed with a confidence interval of 0.04% using a confidence level of 95%.

Generally speaking, Fig. 3(a) emphasizes the high degree of multiconnectivity between ASes pairs. This result is a lower bound on the actual redundancy. `mrinfo` might miss some inter-AS connections but cannot introduce non-existing ones. This limitation is partially due to some ISPs packet filtering policies and the fact that some routers are not multicast native even if they belong to ASes providing multicast capabilities. Moreover, we observed that this result is not time dependent (not shown in this paper) and reflects the operators goal of being strongly connected.

Fig. 3(b) quantifies the number of physical links involved in the connectivity redundancy. This result is plotted on the week during which the largest number of routers responded to `mrinfo`. First, we notice that, obviously, the general result presented in Fig 3(a) is consistent with Fig 3(b). Indeed, in more than 50% of the cases, ASes pairs are connected through at least two physical links. We also notice that the number of involved physical links is not really high: 40% of ASes pairs being multiconnected have a degree of connectivity greater or equal to three physical links.

In the following, we refine this analysis to consider a taxonomy highlighting the different kinds of redundancy. In particular, we choose to differentiate multiple AS connec-

Level (# AS)	Relationship (# AS pairs)				Total	
	p2p	c2p	p2c	?		
Tier-1	6 \pm 0	16 \pm 1	0	15 \pm 1	7 \pm 1	38
Transit	129 \pm 4	35 \pm 2	126 \pm 8	97 \pm 4	135 \pm 7	393
Stub	93 \pm 3	5 \pm 0	80 \pm 3	0	39 \pm 2	124
Total	228	56	206	112	155	555

Table 1: Taxonomy general view

tions involving several routers from those implying only physical links between two routers, each one being in one of the two ASes. Indeed, for fast recovery mechanisms such as those proposed by Bonaventure et al. [4], the protection level (preventing from links or routers failures) depends on the multiconnectivity nature. In the next section we introduce a taxonomy allowing to consider this distinction. For instance, during the first week of February 2006, our `mrinfo` queries have captured that Sprint and Level3 are connected through five routers and eight physical links in the direction Sprint \rightarrow Level3.

3.3 Impact of Policies

In Sec. 3.2, we saw that, on average, 54% of the ASes pairs are composed of multiple connections. In this section, we look closer at the ASes multiconnectivity. We propose a taxonomy based on two categories. The first one, labeled *multiple routers* (*mr*), describes a situation in which two ASes are connected through multiple links and through multiple routers. This category has the advantage of protecting an AS against router and link failures. The second one, labeled *single router* (*sr*), illustrates situations in which the pair of ASes is connected through multiple links, but with a single router on each side. This comes with the drawback of protecting only against link failures. If `mrinfo` captures a connection in both directions (i.e., it detects ASBRs involved in this connection in the two neighbor ASes), we label this connection as *mr* if, at least, one of the two ASes has several ASBRs ensuring the multiconnectivity of the relationship. If there exists several links from an AS_x to several addresses belonging to an AS_y, but these addresses are not responding to `mrinfo`, then they are counted for only one router (*sr*), since they could belong to multiple interfaces of the same router. Note, however, that it might be possible that AS_y provides several routers. This, actually, might lead to an overestimation of the actual *sr* redundancy compared to the *mr* redundancy. *sr* redundancy (parallel links) may be motivated by throughput gain.

In this section, we study the distribution of the multiconnectivity based on this taxonomy and on a hierarchical basis, i.e., for Tier-1, Transit, and Stub ASes. For all ASes in a given hierarchy level, we consider the business relationships, as discussed in Sec. 3.1. Note that we assume the s2s relationship as being a particular case of p2p, then we aggregate both into p2p.³

Table 1 shows the number of ASes falling into each level, as well as the number of business relationships identified in each level. Note that each number in the table is the mean over the 100 weeks processed (in small font, we also provide

³s2s relationships are not significantly represented in our dataset, less than 0.01% of captured connections fall in this category.




Figure 4: Tier-1

the 95% confidence interval around the mean). Obviously, all possible business relationships are not valid at all levels. For instance, a c2p relationship is prohibited at the Tier-1 level. Such a prohibited relationship is denoted \emptyset .⁴ The column labeled “?” indicates the number of relationships that we were not able to map using CAIDA data. It is worth to notice that we discarded such unmapped relationships from our analysis. This implies that several relationships are not considered to ensure that the connectivity results presented here correspond to actual BGP peering sessions.

First, we should emphasize a bias in Table 1 due to two reasons. First, the *mrinfo* AS probing vision is not necessarily symmetric. For instance, let us assume that a link between AS_x and AS_y has been captured by *mrinfo* through an outgoing interface of an ASBR belonging to AS_x . If AS_y has not enabled IPv4 multicast, it will not reply to *mrinfo* queries, leading to a link discovered from AS_x to AS_y but not from AS_y to AS_x . Second, since our router-to-AS election mechanism (see Sec. 3.1) is somewhat naive, we only consider ASBRs whose outgoing interfaces are mapped to several ASes. Furthermore, c2p relationships are more represented than p2c because, generally, the AS provider gives IPs from its address space to its customers. Those reasons lead to the asymmetry observed in Table 3.1 between p2c and c2p relationships.

Thus, in our analysis, c2p relationship of Transit ASes is the most represented category. We notice that our *mrinfo* dataset captures, on average, more than 200 ASes and more than 500 connections between them.

Fig. 4 to Fig. 6 plot the CDF of the multiconnectivity proportion as explained in Sec. 3.1. On all those figures, according to our classification, we plot two curves: the first one, labeled *mr* shows the fraction of multiconnectivity involving multiple routers, while the second one, labeled “total”, provides the proportion of multiconnected relationships whatever its label (*mr* or *sr*). These results must be read in parallel with Table 1. To interpret this set of figures, it is necessary to consider the gap between the two curves in order to determine the proportion of multiconnected relationships not involving multiple routers (*sr*).

Fig. 4 shows the Tier-1 ASes multiconnectivity. The mean of the multiconnected Tier-1 pairs proportion is within the range [0.6; 0.9] (Fig. 4(a)). The vast majority (roughly 75%) of p2p relationship between Tier-1 ASes are redundant, particularly through the *mr* way (roughly 70%). We

⁴Note that our classification results are consistent with this observation.

however observe that p2c relationships (Fig. 4(b)) are less protected in general, approximately 65% of these relationships are multiconnected. These p2c relationships involve fewer routers than the p2p ones (roughly 55% belongs to the *mr* category). Compared to ASes mapped to lower level of hierarchy, the Tier-1 ASes are much more multiconnected between themselves and with their customers than the other categories, as discussed below in this section.

Transit ASes (Fig. 5) seem to be less multiconnected than Tier-1 ASes. On average, the multiconnected Transit ASes proportion is larger than 50%. Fig. 5 shows that p2p and c2p relationships are more redundant than p2c relationships. This might be explained by the *mrinfo* bias discussed earlier in this section. Note that the set of p2c relationships belonging to Transit ASes is the lowest multiconnected category in all the analysis, partially because, in this study, we do not consider ASBRs whose all outgoing interfaces are mapped to the same AS. Moreover, it is worth to notice that the lower multiconnectivity fraction of p2c relationships in Fig. 5(c) is probably due to our underestimation of p2c relationships. Consequently, it might appear that multiconnectivity is not significantly different for Transit and Stub ASes. Finally, the *sr* and the *mr* redundancy seem to be well balanced whatever the category of relationships analyzed.

For the Stub ASes, we present the c2p relationships (on average, only five p2p/s2s relationships of Stub ASes are captured). Stub ASes (Fig. 6) seem to be more multiconnected than Transit ASes. Globally, 60% of the c2p relationships for Stub ASes are redundant, and as with Transit ASes, it is well balanced between *mr* and *sr*.

The classification and taxonomy results presented in this section allow us to conclude that Tier-1 ASes often achieve a high degree of redundancy involving several ASBRs to establish relationships between them as well as with their customers. Transit and Stub ASes results indicate that the multiconnectivity seems to be a less important requirement for ASes which are not in the top of the AS level graph hierarchy. In particular, the measured redundancy for intermediate and leaf ASes involves fewer ASBRs than with Tier-1 indicating that the level of protection achieved seems to be lower.

4. RELATED WORK

During the past decade, a large amount of efforts have been done for mapping the Internet topology, mostly using traceroute-like exploration [1]. However, traceroute is unable to capture links that are not currently used to forward packets, such as backup links. On the contrary, *mrinfo* is able to discover all the links between routers, even those having a high IGP weight or a low BGP local preference. Note that Paris Traceroute [20] is able to detect load-balancing routers and, by extension, multiple paths between a source and a destination [24, 25]. But this is done by injecting a large quantity of packets in the network. Other extensions to traceroute use the IP Record Route option to improve topology discovery, mostly for the alias resolution [26, 27]. But this does not fix the issue of revealing non-active links.

Magnien et al. recently showed the emergence of new IP addresses during repeated traceroute explorations [28]. However, they do not quantify multiple paths between ASes, neither propose a classification.

Teixeira et al. evaluate the IP-level path diversity between *Points of Presence* (PoPs) [12]. Using the real Sprint ISP




Figure 5: Transit




Figure 6: Stub (c2p)

network, Teixeira et al. show that all pairs of PoPs have, at least, two disjoint paths between them. However, based on active measurement techniques, they notice that the resulting topologies over-estimate the path diversity between PoPs within an ISP. They suspect that the large number of false links is due to imperfect alias resolution. Nevertheless, this could not explain the false PoP level edges. Recent developments by Pai et al. [11] show that DNS misnaming appears to be a major source of false edges and offer ways to fix them.

5. CONCLUSION

To the best of our knowledge, prior to this paper, there did not exist accurate measurements analyzing the ASes multiconnectivity. Indeed, existing tools such as traceroute do not allow to correctly measure the degree of connectivity redundancy between ASes pairs. To avoid this drawback, we use the IPv4 multicast tool `mrinfo` whose main specificity is to capture all interfaces attached to replying routers. Therefore, we are also able to capture backup links that are not discovered by traceroute probing.

In this paper, using a consistent data set based on more than four years of `mrinfo` traces, we quantify the multiconnectivity between neighbor ASes. As a general result, we demonstrate that, on average, more than half of the ASes pairs are connected by, at least, two inter-domain links. In addition, we propose a taxonomy of the multiconnectivity and classify each relationship according to an accurate AS level hierarchy.

As further work, we will investigate other advantages provided by `mrinfo` exploration. In particular, we are interested in the alias resolution issue and the possibility to provide better topology model for both intra- and inter-domain levels.

Acknowledgements

This work has been partially supported by the European-founded 216372 Trilogy project. Mr. Donnet’s work is supported by the FNRS (Fonds National de la Recherche Scientifique, rue d’Egmont 5 – 1000 Bruxelles, Belgium.).

Authors would like to thank anonymous reviewers, Cristel Pelsser, Steve Uhlig, and Jennifer Rexford for their helpful comments on the paper.

6. REFERENCES

- [1] B. Donnet and T. Friedman, “Internet topology discovery: a survey,” *IEEE Communications Surveys and Tutorials*, vol. 9, no. 4, pp. 2–15, December 2007.
- [2] H. Haddadi, G. Iannaccone, A. Moore, R. Mortier, and M. Rio, “Network topologies: Inference, modeling and generation,” *IEEE Communications Surveys and Tutorials*, vol. 10, no. 2, pp. 48–69, April 2008.
- [3] N. Feamster, Z. M. Mao, and J. B. Rexford, “Borderguard: Detecting cold potatoes from peers,” in *Proc. ACM SIGCOMM Internet Measurement Conference (IMC)*, October 2004.
- [4] O. Bonaventure, C. FilsFils, and P. Francois, “Achieving sub-50 milliseconds recovery upon bgp peering link failures,” in *Proc. ACM SIGCOMM CoNEXT*, December 2005.
- [5] N. Kushman, S. Kandula, D. Katabi, and B. M. Maggs, “R-BGP: Staying connected in a connected world,” in *Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, April 2007.
- [6] Y. Liao, L. Gao, R. Guerin, and Z.-L. Zhang, “Reliable interdomain routing through multiple complementary routing processes,” in *Proc. Re-Architecting the Internet Workshop (ReARCH)*, December 2008.
- [7] J. He and J. Rexford, “Toward Internet-wide multipath routing,” *IEEE Network*, vol. 22, no. 2, pp. 16–21, March/April 2008.
- [8] D. Wischik, M. Handley, and M. Bagnulo Braun, “The resource pooling principle,” *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 5, pp. 47–52, October 2008.
- [9] W. Xu and J. Rexford, “MIRO: Multi-path interdomain routing,” in *Proc. ACM SIGCOMM*, September 2006.
- [10] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. L. Zhang, “In search of the elusive ground truth: The Internet’s AS-level connectivity structure,” in *Proc. ACM SIGMETRICS*, June 2008.
- [11] M. Zhang, Y. Ruan, V. Pai, and J. Rexford, “How DNS misnaming distorts internet topology mapping,” in *Proc. USENIX Annual Technical Conference*, May/June 2006.
- [12] R. Teixeira, K. Marzullo, S. Savage, and G. Voelker, “In search of path diversity in ISP networks,” in *Proc. ACM SIGCOMM Internet Measurement Conference (IMC)*, October 2003.

- [13] V. Jacobson, "mrinfo," 1995, see http://cvsweb.netbsd.org/bsdweb.cgi/src/usr.sbin/mrinfo/?only_with_tag=MAIN.
- [14] J.-J. Pansiot, "mrinfo dataset," see <http://svnet.u-strasbg.fr/mrinfo/>.
- [15] S. Agarwal, C.-N. Chuah, and R. H. Katz, "OPCA: Robust interdomain policy routing and traffic control," in *Proc. IEEE Conference on Open Architectures and Network Programming (OPENARCH)*, April 2003.
- [16] T. Pusateri, "Distance vector multicast routing protocol version 3 (DVMRP)," Internet Engineering Task Force, Internet Draft (Work in Progress) draft-ietf-idmr-dvmrp-v3-11, October 2003.
- [17] P. Sharma, E. Perry, and R. Malpani, "IP multicast operational network management: Design, challenges, and experiences," *IEEE Network*, vol. 17, no. 2, pp. 49–55, March 2003.
- [18] S. Deering, "Host extensions for IP multicasting," Internet Engineering Task Force, RFC 1112, August 1989.
- [19] k. claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, "Internet mapping: from art to science," in *Proc. IEEE Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH)*, March 2009.
- [20] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding traceroute anomalies with paris traceroute," in *Proc. ACM/USENIX Internet Measurement Conference (IMC)*, October 2006.
- [21] J.-J. Pansiot, "Local and dynamic analysis of Internet multicast router topology," *Annals of Telecommunications*, vol. 62, no. 3–4, pp. 408–425, March 2007.
- [22] IANA, "Special-use IPv4 addresses," Internet Engineering Task Force, RFC 3330, September 2002.
- [23] CAIDA, "AS relationships," 2009, see <http://www.caida.org/data/active/as-relationships/index.xml>.
- [24] B. Augustin, R. Teixeira, and T. Friedman, "Measuring load-balanced paths in the Internet," in *Proc. ACM/USENIX Internet Measurement Conference (IMC)*, October 2007.
- [25] D. Veitch, B. Augustin, T. Friedman, and R. Teixeira, "Failure control in multipath route tracing," in *Proc. IEEE INFOCOM*, April 2009.
- [26] R. Sherwood and N. Spring, "Touring the internet in a TCP sidecar," in *Proc. ACM/USENIX Internet Measurement Conference (IMC)*, October 2006.
- [27] R. Sherwood, A. Bender, and N. Spring, "DisCarte: a disjunctive Internet cartographer," in *Proc. ACM SIGCOMM*, August 2008.
- [28] C. Magnien, F. Ouedraogo, G. Valadon, and M. Latapy, "Fast dynamics in Internet topology: Preliminary observations and explanations," in *Proc. 4th International Conference on Internet Monitoring and Protection (ICIMP)*, May 2009.