Extracting Intra-Domain Topology from `mrinfo` Probing

Jean-Jacques Pansiot[1], Pascal Mérindol[2],
Benoit Donnet[2], Olivier Bonaventure[2⋆]

[1] Université de Strasbourg – Strasbourg – France
[2] Université catholique de Louvain – Louvain-la-Neuve – Belgium

Abstract. Active and passive measurements for topology discovery have known an impressive growth during the last decade. If a lot of work has been done regarding inter-domain topology discovery and modeling, only a few papers raise the question of how to extract intra-domain topologies from measurements results.
In this paper, based on a large dataset collected with `mrinfo`, a multicast tool that silently discovers all interfaces of a router, we provide a mechanism for retrieving intra-domain topologies. The main challenge is to assign an AS number to a border router whose IP addresses are not mapped to the same AS. Our algorithm is based on probabilistic and empirical IP allocation rules. The goal of our pool of rules is to converge to a consistent router to AS mapping. We show that our router-to-AS algorithm results in a mapping in more than 99% of the cases. Furthermore, with `mrinfo`, point-to-point links between routers can be distinguished from multiple links attached to a switch, providing an accurate view of the collected topologies. Finally, we provide a set of large intra-domain topologies in various formats.

1 Introduction

The *Internet topology* discovery has been an extensive subject of research during the past decade [1]. While topological information can be retrieved from passive monitoring (using, for instance, BGP dumps in the case of AS level topology), router level topology is usually obtained from active measurements based on *traceroute*.

Nevertheless, if traceroute has been largely deployed in the last few years, it comes with some important drawbacks. Traceroute provides a partial view of the network as it is routing dependent. For instance, backup links (high IGP weighted links for intra-domain and low BGP local preference links for inter-domain) are rarely captured by traceroute. Furthermore, the alias resolution problem is a complex issue to fix [2]. This leads thus to an incomplete and biased view of the

⋆ This work is partially funded by the European Commission funded Trilogy ICT-216372 project. B. Donnet's work is supported by the FNRS/FRS (Fonds National de la Recherche Scientifique, rue d'Egmont 5 – 1000 Bruxelles, Belgium.).

network. Obtaining complete intra-domain topologies is further a daunting task, requiring extensive probing campaigns [3].

Recently, we used `mrinfo` [4], a management multicast tool, in order to collect topology information [5]. `mrinfo` has the advantage of sweeping out many of traceroute's limitations as it is able to silently discover all interfaces of a router. However, it requires multicast being enable within ISPs' networks and no filtering policies, limiting so its applicability range. Indeed, only IPv4 multicast enabled routers reply to `mrinfo`. Also, some ISPs filter the IGMP messages used by `mrinfo` (i.e., they do not propagate them).

In this paper, we take advantage of the `mrinfo` dataset [6] for extracting intra-domain router level topologies. Obtaining real data concerning intra-domain topologies is of the highest importance. Indeed, it allows one to study actual network characteristics (e.g, degree distribution, network connectivity, ...) and to obtain insights on the way operators build their network. Furthermore, real topologies are crucial inputs for network simulations in order to consider complex and realistic scenarios. By modeling the collected topologies characteristics, it can also contribute to building better topology generators.

The contributions of this paper are twofold. We first describe how to extract intra-domain topologies from raw `mrinfo` data. While it is pretty easy to map IP addresses to an autonomous system number (ASN), the challenge is to mark the boundary of a given autonomous system (AS). Then, it is necessary to assign the right ASN to an AS border router (ASBR) whose IP addresses are not mapped to a single AS. In this paper, we provide an efficient algorithm, called *router-to-AS mapping*, for fixing this issue. We evaluate our algorithm and show that it provides a consistent mapping in more than 99.5% of the cases. In addition, an interesting feature of `mrinfo` is that point-to-point links between routers may be distinguished from multiple links attached to a switch. On average, we discover that roughly 11% of the nodes, in probed networks, are actually switches. As depicted in Sec. 3, this is a fundamental issue to correctly analyze network characteristics. Second, based on our router-to-AS mapping, we provide a set of intra-domain topologies under various formats. Our set of topologies is composed of three kind of networks: Tier-1 (such as Sprint), Transit networks (such as TDC), and Stub networks (such as UNINETT).[3] An extended version of this paper provides more results and discussions [7].

The remainder of this paper is organized as follows: Sec. 2 discusses how we collected topology data using `mrinfo`; Sec. 3 explains and evaluates our router-to-AS algorithm; Sec. 4 positions our work regarding the state of the art; Finally, Sec. 5 concludes this paper by summarizing its main achievements and discussing further works.

2 Collection Methodology and Dataset

`mrinfo` messages use the Internet Group Management Protocol (IGMP [8]). IGMP was initially designed to allow hosts to report their active multicast groups

[3] See `http://inl.info.ucl.ac.be/content/mrinfo`

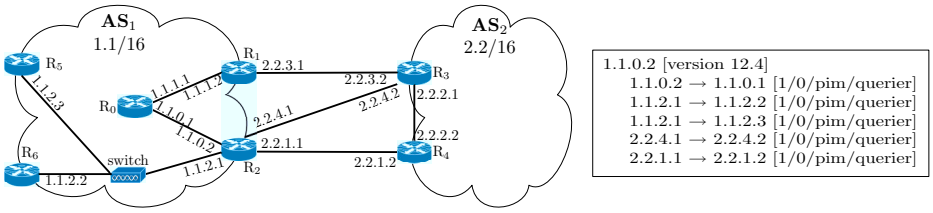

Fig. 1. `mrinfo` example with R_2 output

to a multicast router on their LAN. Upon reception of an IGMP `ASK_NEIGHBORS` message, an IPv4 multicast router replies with an IGMP `NEIGHBORS_REPLY` message providing the list of all its local interfaces with some information about their state. Fig. 1 shows an example of the usage of `mrinfo` to query the router R_2 (`1.1.0.2` is the responding interface of R_2). `mrinfo` reports that this router is directly connected to R_0 (through interface `1.1.0.1`) and two ASBRs, R_3 (through the interface `2.2.4.2`) and R_4 (through interface `2.2.1.2`). We can also notice that R_2 is connected to routers R_5 and R_6 through a switch because the interface `1.1.2.1` appears twice in R_2's reply. This information is obtained by sending a single IGMP message. In practice, `mrinfo` provides similar information to a `show` command on the router's command line interface.

Based on `mrinfo`, we build a recursive probing scheme, `mrinfo-rec`, to scan connected components of networks. Initially, `mrinfo-rec` is fed with a single IP address corresponding to the first router attached to the `mrinfo-rec` vantage point. `mrinfo-rec` probes this router and recursively applies its probing mechanism on all the collected IP addresses. These recursive queries stop at unresponsive routers or when all known routers have been queried. The same process is run every day. It is worth noticing that a router not replying to an `mrinfo` probe during a given day is not queried again afterwards except if it appears again in a list of captured addresses.

To illustrate this behavior, let us apply it on the topology depicted in Fig. 1. `mrinfo-rec` receives, as input, the IP address of router R_0. From R_0, `mrinfo-rec` collects a set of neighbor IP addresses, i.e., {1.1.1.2, 1.1.0.2}. For all IP addresses in this set that were not previously probed, `mrinfo-rec` sends an IGMP `ASK_NEIGHBORS` message and, if the probed router replies, it again runs through the set of neighbor IP addresses collected.

Since May 1st, 2004, we have been collecting the `mrinfo` data from a host located in the University of Strasbourg, France. In this paper, we consider the data collected until the end of December 2008. The entire dataset is publicly available [6]. During this period, on average, `mrinfo-rec` was able to daily discover roughly 10,000 different routers while scanning 100,000 interfaces. Note that we remove interfaces with non-publicly routable IP addresses, i.e., the special-use IPv4 addresses described in RFC 3330. We also remove all tunnel and disabled

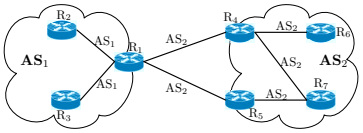

Fig. 2. Shared Addressing Space case on R_1

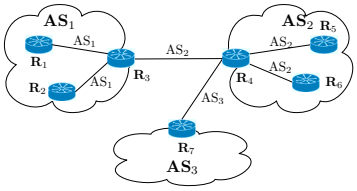

Fig. 3. Neighborhood empirical rule, N

interfaces. The IP-to-AS mapping is done using the last daily BGP table dump of the given day from the Routeviews dataset. We remove from our dataset IP addresses that cannot be mapped to an AS (0.5% on average) as well as those that are labeled to multiple origin ASes [9] (between 2 and 3% of IP addresses discovered each day by `mrinfo-rec`). We roughly identify between 400 and 650 different ASes every day of `mrinfo-rec` probing and we capture more than 850 ASes during the whole period of probing. Those ASes are distributed among Tier-1, Transit, and Stub networks, Transit being the most represented.

3 Router-to-AS Mapping

If it is easy to determine the ASN of a core router (each IP address of such a router is mapped to the same ASN[4]), the challenge is to accurately identify a router as an ASBR and assign it the right ASN. Fig. 2 and 3 illustrate the basics of our *router-to-AS* algorithm. The label attached on each link is the result of the IP-to-AS mapping (we assume that the two IP addresses on each directed link are necessarily mapped to the same AS). First, there are ASBRs whose IP addresses *do not all belong* to the same AS. In such a case, identifying them as ASBRs[4] is straightforward but assigning them an ASN is more difficult. This situation is illustrated in Fig. 2 where router R_1 has two interfaces mapped to AS_1 while the remaining two interfaces are mapped to AS_2. This *Shared Addressing Space* case must be solved to perform the router-to-AS mapping. As soon as all routers are mapped to their right ASN, it is possible to extract intra-domain topologies without falsely cutting between ASes. We denote $\mathbb{SAS}$ the subset of ASBRs falling into the Shared Addressing Space case.

Second, there are ASBRs whose all IP addresses are mapped to the same ASN. If the router-to-AS mapping is obvious, identifying them as ASBRs is a different ball game: their detection essentially relies on their relationships with ASBRs belonging to $\mathbb{SAS}$. This is illustrated in Fig. 2 with routers R_4 and R_5 as all their interfaces are mapped to AS_2. If R_1 is correctly assigned to AS_1, then R_4 and R_5 are ASBRs mapped to the AS corresponding to the address space of links $R_1 \leftrightarrow R_4$ and $R_1 \leftrightarrow R_5$. This issue is thus trivial if the $\mathbb{SAS}$ case has been previously correctly solved.

[4] Note that there exists specific cases more difficult to solve (see [7]).

At this point, it is already worth noticing (see Sec. 3.2 for further details) that the vast majority (almost 90%) of routers are directly mapped to the right ASN because they do not belong to the $\mathbb{SAS}$ set. Thus, our router-to-AS algorithm is applied to only 10% of routers ($\sim$ 1,000 routers on average per day).

3.1 Router-to-AS algorithm

Our router-to-AS algorithm is based on two families of rules: *probabilistic* and *empirical* rules. The main idea behind our algorithm is to quickly converge to a single and consistent mapping for each router. For that purpose, our algorithm verifies the consistency of the results returned by each rule.

We start by assigning a candidate ASN to any router. This is done using our first probabilistic rule (called *global election*, or *elec*). It works as follows: each router is mapped to the ASN assigned to the largest number of its IP addresses (with an IP apparition order for tie-breaking equality cases). Let S_r be the set of the occurrence of each IP-to-AS mapping computed on addresses belonging to the router r. If r is initially mapped to AS n, it means that n appears $max(S_r)$ times in the IP-to-AS mapping of r. We attribute a confidence level to the ASN mapped to r in such a way: $c(r) = 1 - \frac{max(S_r \setminus \{max(S_r)\})}{max(S_r)}$. Closer to one, higher the confidence in the mapping. Note that $r \in \mathbb{SAS}$ if $c(r) < 1$. For instance, regarding the link between R_3 and R_4 in Fig. 3, it means that IP addresses involved in both directions of the link $R_3 \leftrightarrow R_4$ belong to AS_2. Then, looking at Fig. 2, $c(R_4) = 1$ whereas $c(R_1) = 0$, we know that R_1 is an ASBR because it belongs to the set $\mathbb{SAS}$, but we have to figure out whether it belongs to AS_1 or to AS_2 and respectively whether R_2, R_3 or R_4, R_5 are ASBRs. In contrast, in Fig. 3, $c(R_3) = 0.5$ meaning that R_3 seems to belong to AS_1. The *elec* rule is the primary block of all our analysis: other rules aim at confirming or disproving it when $c < 1$.

The second probabilistic rule relies on the detection of LAN interfaces. The *lan* rule concerns a subset of $x \rightarrow 0.0.0.0$ interfaces. Those 0.0.0.0 interfaces usually describe leaf LAN without transit and multicast capabilities (more details are given in [7]). We assume that intra-domain LANs are more frequent than inter-domain LANs, and considering that a local LAN interface uses the address space of the internal domain, a router has a higher probability of belonging to the AS assigned to LAN interfaces. Note that we only consider cases where all LAN interfaces are mapped to the same AS. Furthermore, we do not take into account $x \rightarrow 0.0.0.0$ interfaces if the AS to which IP x belongs is not multicast enabled (e.g., an AS where `mrinfo-rec` does not obtain replies). We observe that, probabilistically talking and, on average, $x \rightarrow 0.0.0.0$ interfaces produce fewer *Shared Addressing Space* cases than *elec*: fewer than 4% compared to the 10% produced by the global *elec* rule. This observation reinforces our hypothesis on intra-domain LAN detection.

As already mentioned, we also use empirical rules consisting in a set of usual rules. The first empirical rule takes into account the loopback interface of a router (*lb* rule). When configuring a loopback interface, an ISP uses an address

belonging to its own IP address space. Thus, identifying loopback interfaces using their DNS name and performing a standard IP-to-AS mapping on this address resolves the router-to-AS mapping.

We also use an empirical rule consisting of a neighborhood analysis (N rule): we assume that inter-domain links are mapped to the address space of one of the ASes it interconnects. This rule can be applied when a $\mathbb{SAS}$ is mapped to a given AS thanks to another rule. For instance, in Fig. 3, let us assume that R_3 has been mapped to AS_1 with the *lb* rule, then R_4 necessarily belongs to AS_2 and iteratively R_7 is then mapped to AS_3. In practice, we apply this rule iteratively until no more new AS assignment is recorded and use it at each step of the router-to-AS algorithm. We also apply this rule iteratively during the last steps of the router-to-AS algorithm.

As already mentioned by claffy et al. [10], a provider generally allocates IP addresses from its own address space to its customers links (*c2p* rule). In a simple case, this means that if two routers denoted R_1 and R_2 are connected through an inter-domain link mapped to AS_x, and such that R_1 also uses the address space allocated by another AS_y ($y \neq x$) while AS_y is a customer of the AS_x, then R_1 is mapped to AS_y (and R_2 to AS_x). In Fig 2, if we know that AS_1 is a customer of AS_2, the c2p rule allows us to map R_1 to AS_1. To perform such a relationship mapping, we use the AS ranking data set provided by CAIDA [11]. Note that this rule seems relatively consistent with the *elec* rule: on average, in more than 70% of the cases, this rule is verified when focusing on routers with a confidence level superior to 0.5. This is our penultimate rule, so that it tie-breaks remaining equality cases when the rule can be used (c2p or p2c relationships between involved ASes).

Finally, to perform the router-to-AS mapping we need a global order between our pool of rules to characterize the confidence we attribute to each of them. We use the following order:
$elec > lb > N_0 > lan > N_1 > H_{0.9} > N_2 > H_{0.8} > N_3 > \ldots > N_{10} > c2p > N_{11}$
where H_β stands for a β-confident assignment rule. According to the confidence threshold $0 < \beta < 1$, if $c(r) > \beta$ (for a given router $r \in \mathbb{SAS}$ mapped to the ASN n with the *elec* rule), H_β maps definitively r to n by attributing to r a confidence level of 1. In order to take advantage of AS assignments produced by the decreasing level of confidence of our set of rules, we apply the neighboring rule between each other rules' application.

Moreover, we use a threshold $0 < \alpha < 1$ to decide whether other rules can overwrite the candidate assignment (the result of *elec*). For all routers $r \in \mathbb{SAS}$, if a given rule is not in concordance with the *elec* rule (i.e., the ASN returned by the given rule differs from the candidate one given by *elec*), we select the ASN returned by the tested rule only if $c(r) \leq \alpha$. Otherwise we ignore the result provided by the tested rule. In practice we choose $\alpha = 0.5$[5]. Sec. 3.2 describes the consistency of the mapping we obtain using this ordered set of rules.

[5] It means that there are at least two more IP addresses mapped to the candidate ASN compared to any other ASN

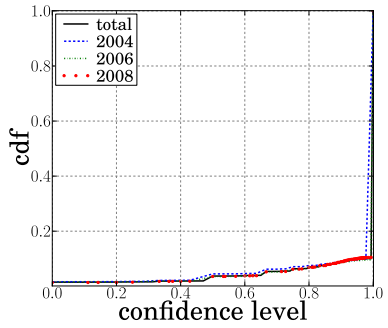

Fig. 4. *elec* rule efficiency

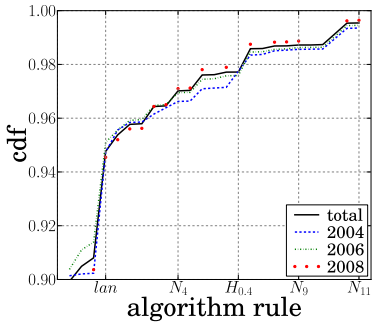

Fig. 5. Algorithm convergence time

3.2 Evaluation

From our four years daily dataset, we arbitrarily select the largest `mrinfo` raw data file of each month, leading to 56 files. We then evaluate our router-to-AS algorithm on those files.

Fig. 4 provides the cumulative distribution of the confidence level c (the horizontal axis) assigned during the first step (*elec*) of the router-to-AS mapping. The first observation is that, on average, 90% of the routers have addresses mapped to a single AS. We identify roughly that only 1.5% of the routers have a confidence level equal to 0, whereas more than 95% of routers have a confidence level superior to 0.5. According to our threshold $\alpha = 0.5$, only 5% of the router assignments are really problematic.

Fig. 5 shows the cumulative distribution of the required number of rules (the horizontal axis) to converge to a positive decision, i.e., $c = 1$, in the router-to-AS algorithm. In this figure, we observe that the decision process quickly converges: more than half of $\mathbb{SAS}$ cases are already treated after the *lan* rule application (i.e., after the 4^{th} rule). This means that a large subset of the critical cases are treated at the beginning of our set of rules which are ordered depending on the confidence we attribute to each of them. At the end of the process, using the N_{11} rule, we can notice that fewer than 0.46% of the routers remain unmapped.

Fig. 6 gives a more detailed overview of the actions taken during the convergence of our algorithm. Each point represents the mean over the 56 files we consider. We determine 95% confidence intervals for the mean but intervals are typically too tight to appear on Fig. 6.

We can see that in most cases, our pool of rules confirms the candidate assignment of the global election, *elec*: at the end of the process, 88.9% of the candidate assignment are confirmed with one or another rule. We also count the number of contradictions produced by our set of rules against *elec*, and divide them in two categories according to the threshold α that we use. We notice that on average, 12% of the AS assignments are concerned, and mainly when $c > 0.5$ (6.4% compared to 5.6% when $c \leq 0.5$). However, using our threshold $\alpha = 0.5$,

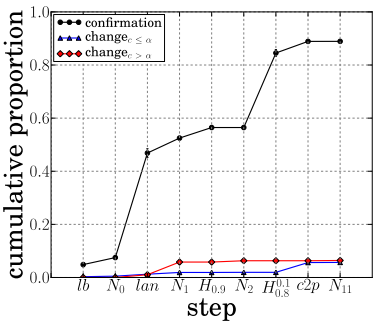

Fig. 6. A closer look at each step

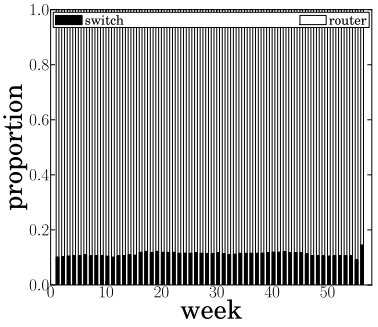

Fig. 7. Switches and routers proportion

we only effectively record the 5.6% of changes that are not strongly inconsistent ($c \leq 0.5$) to stay consistent with the *elec* rule. We have also noticed that less than 6% of routers among the $\mathbb{SAS}$ set are subject to real inconsistencies (see [7]).

To summarize, we have seen that 90% of the routers are assigned, directly at the first step of the algorithm, to an ASN with the highest level of confidence. The remaining 10% of routers belongs to $\mathbb{SAS}$ and represent thus critical cases. Our algorithm is able to quickly solve a large subset of those cases. At the end of the process, only 0.46% of the routers remain unmapped.

A more detailed discussion on the algorithm evaluation as well as on particular cases may be found in [7]. In particular, we have empirically verified that the AS where our probing host is located (AS2259) is correctly and fully discovered by our algorithm.

3.3 Point to point links and switches

In addition to our router-to-AS algorithm, we also provide a way to distinguish point-to-point links from switch inter-connections. As previously mentioned (see Fig. 1), replies collected with `mrinfo-rec` allow us to easily discover switch pseudo-nodes and extract them from our raw data.

This point is of the highest importance since it provides accurate information on the real network connectivity. Using traceroute-like probing, switches are not easily detectable and this bias leads to produce false interpretations: a set of nodes may appear to be fully meshed whereas they are actually connected through a simple switch. Identifying switches in `mrinfo` output is straightforward as it is enough to capture outgoing IP addresses appearing several times on the same router (see interface `1.1.2.1` on router R_2 in Fig. 1).

Note that a switch inter-connection discovered with `mrinfo-rec` can hide a switch cascade, i.e., several switches might be connected together. It can also hide some other types of level 2 inter-connections. Moreover, when possible, we verify that all routers connected through a switch share the same vision

of the inter-connection (e.g., one IP address pointing towards the same set of addresses).

Fig. 7 provides the distribution of switches and routers over the 56 weeks. On average, we identify that 11% of inter-connection points discovered in the networks are switches (or cascade of switches), while the remaining 89% are actual routers. Note that the same distribution occurs when we distinguish inter-domain from intra-domain connections. Only 1% of the whole set of discovered nodes are inter-domain switches (Internet exchange points, IXPs) whereas 9% of them are ASBR. Note that we do not apply the neighboring rule N for IXPs.

4 Related Work

A tool like *Rocketfuel* [3] has been used to infer ISP topologies. However, inferring topologies in a non-cooperative and heterogeneous environment has proven to be extremely difficult, and results obtained have to be carefully evaluated in terms of validity [12, 13]. The recently introduced *DisCarte* [14] pushes the accuracy of collected data a few steps further but it requires the "record route" option being enable and does not entirely sweep out standard traceroute limitations.

Mao et al. provide mechanisms for improving the IP-to-AS mapping [15, 16]. Their techniques are based on several information sources: traceroute, BGP update, BGP table dumps, and reverse DNS lookup. In addition, they propose heuristics for identifying IXPs, sibling ASes as well as ASes sharing address space. Their work differs from ours as they focus only on IP-to-AS mapping and not on router-to-AS mapping.

The recent work done by claffy et al. is probably the most relevant compared to this paper [10]. For assigning ASes to routers, claffy et al. assume that a provider always gives IP addresses belonging to its own address space for connections to their customers [10]. Given that assumption, the router-to-AS mapping becomes straightforward when focusing on customer-to-provider links (and reciprocally). Otherwise, the router is assigned to the AS with the smallest outdegree. Note that no evaluation of this technique has yet been made in [10].

5 Conclusion

We provide a mechanism for extracting intra-domain topologies from raw data collected by `mrinfo`, a multicast based tool that is able to silently discover all interfaces of a router. The main challenge is to mark the boundaries of each AS. The goal of our algorithm is to assign an AS number to a router, performing the so called *router-to-AS mapping*. We demonstrate that our router-to-AS mapping is able to efficiently assign an AS number to a router with a high confidence level. In addition, our AS extraction mechanism is able to discover connections through layer-2 switches, providing a more accurate view of the topology than with traceroute probing. Finally, we provide, in various format,

several intra-domain topologies for Tier-1, Transit, and Stub networks all along the four years of collected data.

We believe the technique described in this paper as well as the whole `mrinfo` dataset are valuable for the research community. Indeed, the next steps of this work would be to deeply study intra-domain topologies and improve `mrinfo` based probing using complementary topology discovery methods.

References

1. Donnet, B., Friedman, T.: Internet topology discovery: a survey. IEEE Communications Surveys and Tutorials **9**(4) (December 2007) 2–15
2. Gunes, M.H., Sarac, K.: Importance of IP alias resolution in sampling Internet topologies. In: Proc. IEEE Global Internet Symposium. (May 2007)
3. Spring, N., Mahajan, R., Wetherall, D.: Measuring ISP topologies with Rocketfuel. In: Proc. ACM SIGCOMM. (August 2002)
4. Jacobson, V.: `mrinfo` (1995) see `http://cvsweb.netbsd.org/bsdweb.cgi/src/usr.sbin/mrinfo/?only_with_tag=MAIN`.
5. Mérindol, P., Van den Schriek, V., Donnet, B., Bonaventure, O., Pansiot, J.J.: Quantifying ASes multiconnectivity using multicast information. In: Proc. ACM USENIX Internet Measurement Conference (IMC). (November 2009)
6. Pansiot, J.J.: `mrinfo` dataset see `http://svnet.u-strasbg.fr/mrinfo/`.
7. Pansiot, J., Mérindol, P., Donnet, B., Bonaventure, O.: Internet topology discovery through `mrinfo` probing. TR 2009-01, Université catholique de Louvain (UCL) (October 2009) See `http://inl.info.ucl.ac.be/content/mrinfo`.
8. Deering, S.: Host extensions for IP multicasting. RFC 1112, Internet Engineering Task Force (August 1989)
9. Zhao, X., Pei, D., Wang, L., Massey, D., Mankin, A., Wu, S.F., Zhang, L.: An analysis of BGP multiple origin AS (MOAS) conflicts. In: Proc. ACM SIGCOMM Internet Measurement Workshop (IMW). (October 2001)
10. claffy, k., Hyun, Y., Keys, K., Fomenkov, M., Krioukov, D.: Internet mapping: from art to science. In: Proc. IEEE Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH). (March 2009)
11. CAIDA: AS relationships (2009) see `http://www.caida.org/data/active/as-relationships/index.xml`.
12. Zhang, M., Ruan, Y., Pai, V., Rexford, J.: How DNS misnaming distorts internet topology mapping. In: Proc. USENIX Annual Technical Conference. (May/June 2006)
13. Teixeira, R., Marzullo, K., Savage, S., Voelker, G.: In search of path diversity in ISP networks. In: Proc. ACM SIGCOMM Internet Measurement Conference (IMC). (October 2003)
14. Sherwood, R., Bender, A., Spring, N.: DisCarte: A disjunctive Internet cartographer. In: Proc. ACM SIGCOMM. (August 2008)
15. Mao, Z.M., Rexford, J., Wang, J., Katz, R.H.: Towards an accurate AS-level traceroute tool. In: Proc. ACM SIGCOMM. (August 2003)
16. Mao, Z., Johnson, D., Rexford, J., Wang, J., Katz, R.: Scalable and accurate identification of AS-level forwarding paths. In: Proc. IEEE INFOCOM. (April 2004)