

Quantifying and Mitigating IGMP Filtering in Topology Discovery

Pietro Marchetta*, Pascal Mérindol[‡], Benoit Donnet[†], Antonio Pescapé*, Jean-Jacques Pansiot[‡]

* University of Napoli Federico II – Italy

[‡] Université de Strasbourg – France

[†] Université de Liège – Belgium

Abstract—Recent developments in router level topology discovery have suggested the introduction of IGMP probing in addition to standard techniques such as traceroute and alias resolution. With a single IGMP probe, one can obtain all multicast interfaces and links of a multicast router. If such a probing is a promising approach, we noticed that IGMP probes are subject to filtering, leading so to the fragmentation of the collected multicast graph into several disjoint connected components.

In this paper, we cope with the fragmentation issue. Our contributions are threefold: (i) we experimentally quantify the damages caused by IGMP filtering on collected topologies of large tier-1 ISPs; (ii) using traceroute data, we construct a hybrid graph and estimate how far each IGMP fragment is from each other; (iii) we provide and experimentally evaluate a recursive approach for reconnecting disjoint multicast components. The key idea of the third contribution is to recursively apply alias resolution to reassemble disjoint fragments and, thus, progressively extend the mapping of the targeted ISP. Data presented in the paper, as well as reconstructed topologies, are freely available at <http://svnet.u-strasbg.fr/merlin>.

I. INTRODUCTION

Internet topology discovery has been an extensive subject of research during the past decade [1]. While topology information can be retrieved from passive monitoring (using, for instance, BGP dumps in the case of the AS level topology), router level topology is usually obtained from active measurements using *traceroute* and *alias resolution* [2] for gathering all IP addresses of a router into a single identifier.

Inferring the router level topology of IP networks is an important aspect, in particular to study routing characteristics. More specifically, understanding the design of an AS is crucial for analyzing intra-domain routing protocol performance. Network protocols designers should evaluate the performance of their proposals on realistic topologies in order to highlight their advantages and limitations. For example, performance of fast-rerouting schemes or multipath transport protocols may strongly depend on the underlying topology. Inferring AS at the router level may help in developing solutions able to perform well on various topology designs and standard patterns.

Two IGMP-based probing approaches have been recently introduced for topology discovery: first with `mrinfo-rec` [3] and then with MERLIN [4], [5]. `mrinfo-rec` sends multicast management requests that are able to retrieve, within a single probe, all multicast interfaces and links of a targeted router. IGMP probing can natively discover multicast topologies at the

router level with a low probing cost [3], avoiding so the use of any alias resolution techniques. While the resulting vision may be incomplete (because limited to the multicast part of the network), it is also less subject to false positives than common topology discovery techniques. MERLIN is an extension of `mrinfo-rec` that use both IGMP and traceroute-like probing.

When probing a multicast enabled AS with IGMP probing, we expect obtaining its complete *backbone* as it should be entirely multicast to ensure the correct multicast tree establishment by the PIM multicast routing protocol. By multicast backbone, we mean the AS areas where links and routers providing connectivity to non-multicast customers or peers are pruned. Unfortunately, some routers do not reply to IGMP probes sent by MERLIN, leading to an anonymous behavior that is similar to the one observed with traceroute [6], [7], [8]. We call this phenomenon *IGMP filtering*. As a consequence, the topology obtained using solely IGMP probing is incomplete and disconnected: the collected IGMP graph exhibits a set of disjoint components.

In this paper, we first experimentally investigate in Sec. II how IGMP filtering damages collected topologies with MERLIN. Based on a dataset jointly collected with Paris traceroute [9] and IGMP probing, we propose a hybrid graph analysis to understand the distances between IGMP components of a given AS. Since most distances are limited between IGMP fragments, our analysis suggests that reconnecting them is possible using their ICMP neighborhood. Then, in Sec. III, we propose and experimentally evaluate a recursive hybrid reconnection mechanism based on traceroute and alias resolution (able to keep the native router level view of MERLIN) for merging isolated IGMP components into a larger one. This mechanism significantly enhances the multicast reconnection strategy proposed in our previous work [5]. Apart a few marginal isolated routers, all collected and reassembled graphs exhibit a large connected component.

II. IGMP FILTERING

IGMP probing campaigns may suffer from the multicast graph “disconnection” due to IGMP filtering: some multicast routers do not reply to IGMP probes (*local filtering*) while some others do not forward IGMP messages (*transit filtering*). While the second problem can be reduced with the use of multiple vantage points in a cooperative distributed platform [5],

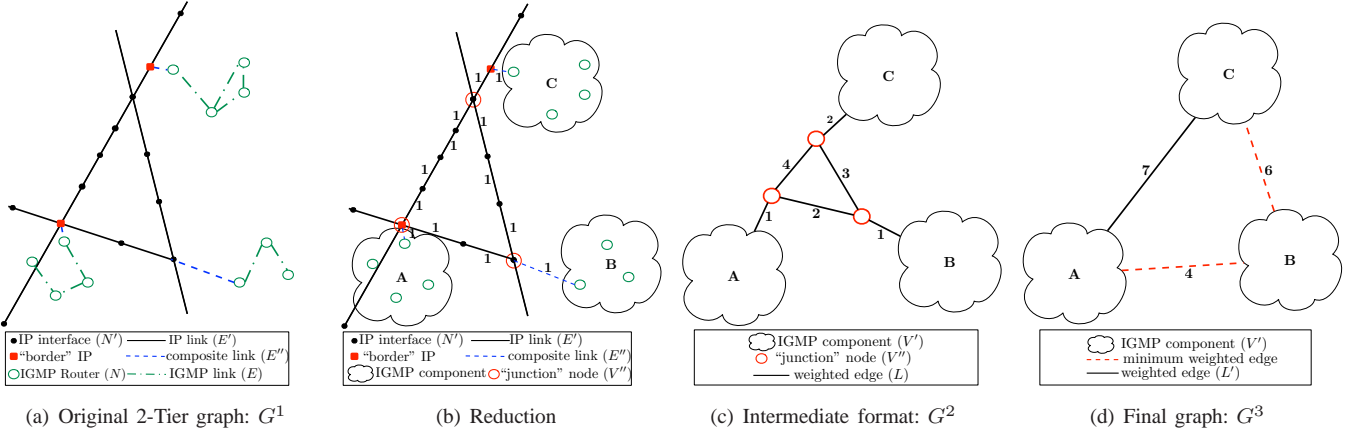


Fig. 1. Compute minimal distance between disjoint IGMP components

the first one is more challenging as it impacts the collected topologies. Indeed, multicast routers that do not respond to IGMP probes may divide the resulting collected multicast graph into disjoint components.

MERLIN [4] extends the recursive application of `mrinfo-rec` (that only probes connected routers) to improve its coverage by probing independent *seeds*, i.e. IP addresses belonging to various prefixes that are used as input to MERLIN for improving its coverage. However, due to IGMP filtering, it may result in non connected graphs. Given that the “globally accessible” multicast graph should be physically connected, we assume that scattered components and isolated routers result from non-responding multicast routers. Indeed, even a low proportion of non-responding routers may result in a huge disconnection of the multicast graph.

This “disjoint state” may be exacerbated by missing unicast adjacencies: the connectivity of the multicast graph can be lower than the unicast one. In practice, a multicast router can be configured at the interface granularity such that each interface can independently support multicast or not. Nevertheless, an ISP supporting IP multicast should enable multicast everywhere in its network to ensure the correct PIM tree establishment, although some exceptions may arise at inter-area border routers and AS border routers. An area border router does not need to support multicast adjacencies with routers belonging to non-multicast areas. Between ASes, the BGP routing protocol can use specific multicast forwarding entries to disseminate PIM messages. Thus, although it is likely that a multicast border router will not enable multicast on all its interfaces, it is also likely that the multicast graph should be connected. Even in presence of non-multicast adjacencies, there should exist at least one multicast path between each multicast component.

In this paper, the use of an intensive tracerouting campaign allows us to understand the IGMP disconnection state and so improve the MERLIN reconnection phase. Indeed, by constructing a 2-tier graph (both at the IP and the router level) and analyzing its connectivity/disconnection state¹, we will then be

¹Rather than simply considering traceroute paths crossing several IGMP components for distance computation as it is done in [5].

able to design an efficient reconnection scheme (see Sec. III).

A. A Hybrid Graph Transformation Procedure

In this section, we are interested in the connected components size distribution and in the connected components “distance distribution”. While evaluating the size of disjoint connected components is straightforward, obtaining the distance between the components requires a dedicated hybrid methodology. In addition to the IGMP probing phase of MERLIN, we performed a large scale Paris traceroute [9] campaign (one Paris traceroute per /24 prefix per router) targeting each IGMP router previously discovered with MERLIN. We also use preliminary traces used by MERLIN as static seeds [5]. The combination of IGMP and ICMP replies leads to a hybrid 2-tier graph where some nodes are routers (the IGMP view) and others are IP interfaces (the ICMP view), as illustrated in Fig. 1(a).

In the remainder of this section the notation (V, L) refers to an undirected graph composed of a set of vertices, V , and a set of links, L . Except when explicitly specified, the valuation of links is uniform such that the distance metric only relies on the number of hops.

We define a hybrid graph $G^1(\{N, N'\}, \{E, E', E''\})$ where:

- N is the set of IGMP routers;
- N' is the union of the ICMP IP interfaces set and the IGMP border IP interfaces set (IP addresses part of non responding multicast router which were captured with the collected IGMP replies as neighbor interfaces);
- E is the IGMP adjacencies set (router level links between nodes in N);
- E' is the IP level links set (links between nodes in N');
- E'' is the hybrid connections set: links connecting a router level node and an IP interface node. This corresponds to dashed lines in Fig. 1(a). The set E'' is the key point of the analysis since it describes the interaction between the two node levels, being therefore the starting point for reconnecting disjoint IGMP components.

Fig. 1 illustrates on a small example the graph transformation process we used. It basically works as follows:

- 1) Construct the initial hybrid graph G^1 gathering all ground data (both at the router and IP level).

- 2) Reduce it to a smaller weighted graph G^2 where solely ICMP interfaces of interest are kept. Those IP addresses act as junction nodes in our hybrid graph description.
- 3) Compute direct shortest paths between each pair of IGMP components to produce a distance oriented graph G^3 . This graph provides hop distances that characterize the connectivity between IGMP islands.
- 4) Compute the minimal weighted tree of G^3 : it gives the minimal distances required to reconnect the collected multicast data. The sum of resulting distances depicts the worst case for obtaining a minimal connected graph, i.e., a tree.

The remainder of this section introduces the detailed operation mode of our four steps transformation. As already mentioned, the set E'' is the basis for our hybrid reconnection scheme. An edge is added to $E'' = E''_b \cup E''_n$ according to two possible cases: (i) an IGMP router reports a neighbor IP address that is not locally attached to another IGMP router (this subset is denoted E''_b), (ii) a traceroute intersects a node belonging to N (this subset is denoted E''_n). Note that a node in N is a set of local IP interfaces, an IGMP alias, such that E''_n is almost equivalent to the intersection between IGMP and ICMP probing coverage. Moreover, it is worth to notice that we have no guarantee that G is connected (it mainly depends on the utility of traceroute traces), so that the distance distribution analysis may be incomplete.

For the purpose of our analysis, G^1 can be reduced to a weighted graph $G^2(V, L, w)$ where nodes in V are either connected components of IGMP routers in the graph (N, E) (such a connected component becomes a node in the set V') or IP interfaces in N' that are junction nodes (this set of nodes is denoted V'' , $a \in V'' \Rightarrow$ the degree of a in G^1 is greater than or equal to 3). Thus, we have $V = V' \cup V''$. The valuation w of an edge in L is the hop distance between nodes in the graph $(V, \{E', E''\})$. Since non junction nodes are removed from N' to form V'' , we keep track of this distance information: $\forall a, b \in V$, $w(a \leftrightarrow b) - 1$ is equal to the number of nodes $\in N'$ removed from the shortest path between a and $b \in (V, \{E', E''\})$ if any, $w(a \leftrightarrow b) = \infty$ otherwise. Note that this reduction operation preserves distances computed in the initial graph. Fig. 1(b) illustrates the reduction operation: after such an operation, nodes in N' whose degree is still greater than or equal to 3 become “junction nodes” of the new graph. Moreover, nodes belonging to the same IGMP connected component are merged so that they become an “IGMP cloud”. Fig. 1(c) provides the resulting graph G^2 : distances between nodes in V are updated to reflect the number of hops between them.

Then, the graph G^2 can be reduced to a third graph $G^3(V', L', w')$ where V' is the set of connected components of IGMP routers and L' are links between them. A link (a, b) is in L' if there exists a path between a and b in G^2 without intermediate nodes in V' . Thus, the weight $w'(a, b)$ of such a link is the minimal length among the set of existing paths in G^2 . Those paths only use intermediate nodes in $V'' = V \setminus V'$.

From the last reduced graph $G^3(V', L', w')$, we compute its resulting minimal weighted tree. This final computation permits distance estimation between disjoint IGMP compo-




Fig. 2. IGMP component size




Fig. 3. Distances distribution

nents. Fig. 1(d) illustrates the final result: $\{4, 6\}$ is the set of minimal distances for IGMP connected component A , B , and C . The weight of edges belonging to the resulting minimal weighted tree describes the *a priori* required minimal effort to reconnect the topology. This metric has several advantages but also suffers from the interface level view provided by traceroute. On the one hand, it offers insights on the required effort to reconnect the topology: the more important the distances, the more difficult the reconnection. On the other hand, although this metric is *a priori* stable to analyze the evolution of the topology reconnection (the reconnection of two disjoint components does not impact other distances than those between them), it may overestimate distances due to the lack of IP alias resolution. Indeed, nodes that describe different IP interfaces (and so different nodes in N') may belong to the same router, and thus falsely increase distances between disjoint components. Hence, this metric provides a worst case scenario to reconnect the topology when all IPs in N' belong to distinct routers.

In practice, the graph G^3 can be efficiently computed using a modified version of the *Dijkstra* algorithm where the minimal length extraction is limited to nodes in V'' . Finally, we use the *Kruskal* algorithm [10] to compute the minimal weighted tree of G^3 .

B. Evaluating the Impact of IGMP Filtering

To evaluate how IGMP filtering impacts the collected topologies, we considered several ASes. In this paper, we focus our efforts on three large ISPs: Sprint (AS1239), Level3 (AS3356), and Global Crossing (AS3549). In the remainder of the paper, all presented results are related to these three domains.² We select those ASes among our set of experiments because a large proportion of their routers replies to IGMP probes and, more importantly, they are representative of different difficulty levels to obtain a fully connected multicast map.

Fig. 2 provides the IGMP connected component size distribution for the three ASes of interest. The horizontal axis, in log-scale, is the component size (i.e., the number of routers included in a given IGMP component), while the vertical axis is the cumulative distribution. Although a very low proportion of IGMP components are quite large (larger than 200 for AS 3549), we see that the vast majority of IGMP components are made of a single router (70% for AS1239, 46% for AS3356, and 96% for AS3549). This means that, even if MERLIN is able to capture one or two reasonably

²Interested readers can find additional results at <http://svnet.u-strasbg.fr/merlin>.

		AS1239	AS3356	AS3549
IGMP cmp	#cmp : $ V' $	124	118	33
	largest cmp	153	58	276
G^1 graph	$ N $	328	386	308
	$ N' $	5,064	10,610	7,934
	$ E' $	6,859	15,856	12,667
	$ E'' $	2,342	3,158	1,342
graph reduction	$ V'' $	1,680	3,907	3,366
	$ V'' / N' $	0.33	0.37	0.42

TABLE I
GENERAL STATISTICS

large connected components within an AS, most of the time, MERLIN discovers information about isolated IGMP routers. Table I provides relevant information about graphs studied (for instance the total number of collected IGMP routers, N). It also gives information about the graph reduction process to describe the quantity of data of interest (e.g. junction nodes).

Analyzing the final graph G^3 , we observe two notable properties. First, on the three explored ASes, we notice that most of disjoint IGMP components are “reconnectable” thanks to our ICMP dataset, i.e., for the vast majority of nodes pairs in V' , there exists at least one path in G^3 connecting them. Only (respectively for AS3549, AS1239, and AS3356) 2, 6, and 8 IGMP components (each being made of a single router) are disconnected from the remainder of the graph (among 33, 118, and 124 nodes in V'): these completely isolated routers provide almost no useful routing information (non-publicly routable IP address or stale configurations) or we do not succeed to reach them using our ICMP campaign. Second, considering the minimal weighted tree obtained on G^3 , we discover that all edges involved in its construction have a weight of two. This is of the highest importance since it implies that we can reconnect multicast components using only ICMP neighbor and IGMP border IP addresses: those two hop distances correspond to two edges in the set E'' made of composite links.

In order to better understand distances and path diversity in the “meshed logical graph” G^3 before applying Kruskal, we also study the distance distribution between nodes in V' . Fig. 3 provides such a distribution. The horizontal axis gives the distance, while the vertical axis shows the cumulative mass. We observe different behaviors depending on the AS: for AS3549, all computed distances are lower than three hops but its density ($\frac{2 \times L'}{V' \times (V'-1)}$) is quite limited (0.14). In AS1239 and AS3356, the collected hybrid graphs are quite dense (0.95 and 0.88, respectively). Note that density values given here do not have the same meaning than in standard graph theory analysis. Indeed, this number rather means (when it tends to 1) that there exists an IP level path between almost each IGMP component pair but those paths may share a common subset of IP level links. On the one hand, it potentially implies that using such an additional ICMP information we are able to produce a qualitative inference of the backbone that is likely to be much more connected than a tree. On the other hand, considering the quite large distances (we observe paths up to ten hops long), it also potentially means that MERLIN possibly misses a quite significant part of the AS due to IGMP filtering.

III. RECURSIVE RECONNECTION

This section describes our strategy for dampening IGMP filtering. The objective is to merge a large number of disjoint IGMP components into a single one. For that purpose, MERLIN relies on an alias resolution phase: IP level links and *aliased* IP addresses - forming so routers - fill the gap between disjoint components discovered during the probing phase.

Considering the original graph G^1 described in Sec. II-A, our goal is to progressively “transfer” nodes from N' to N in order to qualitatively reconnect all original nodes in N between themselves. Thus, we use alias resolution mechanisms to gather IP level nodes in N' in order to provide a connected router level graph. Note that alias resolution allows for both checking and anti-checking a set of IP interfaces pairs so that we can also easily conclude when IP level nodes are independent in the router level graph. In order to reduce the alias resolution space, we decide to not consider all the IP addresses extracted from traceroute traces but only those that are located “close” to routers in the already discovered topology. Hence, our method starts by trying to alias ICMP neighbors and IGMP neighbors to generate new routers and, thus, expand each connected component. Then, considering the neighborhood of each new aliased routers, we recursively re-apply the same alias resolution mechanism on new formed borders, progressively expanding the new topology. Furthermore, each time a traceroute reveals a direct connection between two router level nodes (one hop distance), a new link is added to the topology since the neighborhood information obtained with IGMP queries could be incomplete (unicast lacks - see [4] - or even empty for ICMP aliases).

During the MERLIN probing phase, note that all the collected ICMP data is subject to IGMP probing so that we perform what we call “IGMP unicast alias resolution” on such IP addresses. Indeed, as mentioned in [4], even if IGMP probing does not reveal unicast interfaces, if one probes a unicast IP of a multicast router, MERLIN can deduce whether the unicast IP belongs or not to the router. Hence, final graphs resulting from the reconnection phase (using alias resolution) are validated in the sense that all their vertices has been proven independent: they do not belong to the same router.

Except the potential impact of their bias and overhead, any alias resolution techniques can be implemented in our modular reassembling strategy. Future work should reveal how a particular mechanism influences the resulting topology.

A. Experimental Evaluation

For this evaluation, we targeted the same set of ASes than in Sec. II: Sprint (AS1239), Level3 (AS3356), and Global Crossing (AS3549). Measurements were done between April, 4th 2011 and April, 9th 2011. Our measurement campaign was performed, for each, as follows: a MERLIN probing campaign is launched towards each AS from five vantage points: Strasbourg (France), Napoli (Italy), Louvain-la-Neuve (Belgium), Hamilton (New Zealand), and San Diego (USA). While Paris traceroute campaigns were launched from all the vantage points towards multiple interfaces of each router (a single IP address for each /24), the alias resolution phase itself,




Fig. 4. Recursive reconnection evaluation

that makes use of the information retrieved from traceroute, is performed for each AS by a single monitor. We prefer to avoid interferences between monitors when they try to infer alias in the same AS topology: it could result in the exceeding of the ICMP rate-limiting threshold and make the AS silent to our probes. In our implementation of the reassembling strategy, we make use of Ally [11] for performing the alias resolution. Recently, novel approaches have been proposed (see, for instance, Sherry et al. [12] and the survey by Keys [2]) to improve the alias resolution state of the art. Evaluating the impact of using a particular alias resolution technique is left for future work.

Fig. 4 shows the topologies evolution over the various iterations of the reconnection procedure. This evolution is given in terms of cumulative number of routers (Fig. 4(a)), of links (Fig. 4(b)) created at each step of our recursive process (the horizontal axis), and in terms of registered number of disjoint components (Fig. 4(c)). Note that “iteration 0” on Fig. 4 refers to the situation before applying the reassembling process: it provides the original multicast graph after adding some traceroute IP interfaces to IGMP routers (IGMP alias unicast resolution) and after correcting one hop distances as mentioned in Sec. III.

Fig. 4 shows that the number of new routers and links created at each iteration seems to speed down: in particular, for all evaluated ASes, at least as many links are introduced in the first iteration as in subsequent iterations. AS3356 shows a specific behavior: it reacts favorably at each alias generation iteration. For other ASes, the gain seems to become marginal after three or four iterations: the number of new aliases slows down and most of the links have been discovered earlier. Based on this observation and the cumulative bias introduced by Ally, we decide to definitively stop the analysis of the recursive process after five iterations. Note that a number of k iterations is able to ideally solve distances of $2 \times k$ hops. Intuitively, a distance of k corresponds to a potential reconnection path made of k hops (i.e., a path of k links allowing to merge several IGMP components). In order to further limit the number of false information potentially generated by Ally, and for practical and accuracy reasons, we decide to stop the graphs reconstruction at $k = 2$. Hence, the final topologies we consider hereafter are obtained after the second iteration.

New routers³ and links cause the reduction of disjoint components as depicted in Fig. 4(c). Considering the final

³In this analysis, note that a single IP proven independent from others is not considered as a new router.




Fig. 5. Final topologies properties

topologies, the number of IGMP components decreases from 33 to 24 for AS3549, 118 to 45 for AS3356, and 124 to 10 for AS1239. The reduction level highlights each network specificity regarding our measurements: AS1239 and, to a lower extent, AS3356 offer a good alias performance. Indeed, a significant number of alias is generated during the first step, allowing so to fix most of two hop distances. On the contrary, AS3549 does not provide such an efficient result: either a small amount of IP addresses we retrieved from aliases, either Ally does not work well within this AS. The impact of the alias resolution phase reveals the level of dependency among forwarding paths discovered through our traceroute campaigns. For AS1239, it seems that almost all two hop distances are subject to alias, favoring so the almost complete reconnection during the first iteration.

Fig. 5 provides a graph analysis of the final topologies. In particular, Fig. 5(a) shows the efficiency of the alias resolution process used for our reassembling technique. For instance, before applying it, the largest component in AS1239 was made of 153 routers. After the second iteration, the largest component is made of 393 nodes and only nine components (made of a single router) are still isolated. On other ASes such as AS3356, we can notice that some low distance reconnection paths do not seem to involve aliases so that we still have a significant number of isolated IGMP fragments after studying four hops paths. In practice, although there exists an ICMP IP level path connecting the vast majority of them, IP addresses involved are just still proven anti-checked with others tested.

Fig. 5(b) shows the impact of our recursive alias resolution approach on preliminary distances computed between native IGMP components. For this analysis, we consider the final resulting graph and apply the methodology described in Sec. II-A to obtain the G^3 graph. Although, most of IGMP components are now reconnected, we continue to distinguish IGMP native disconnected components from the rest of the

graph (newly introduced alias and IP level nodes). Compared to Fig. 3, we notice a great shift towards lower distances: even for the worst case (AS3356), we observe that almost 80% of distances are now lower than six hops instead of approximately 60% before alias computation. It is also worth to notice that the alias resolution phase allows one to compute new distances and can make the G^3 graph denser. When several IP addresses are merged into a given alias, the distance resulting from a combination of traceroute traces may decrease. On the contrary, when it results from a unique direct forwarding trace, the distance is unchanged. On AS3356, although most of distances decreases, maximal distances are incompressible: they result from direct and unique forwarding traces crossing distinct devices.

IV. RELATED WORK

To the best of our knowledge, our work is the first attempt to study and solve the issues coming from the IGMP filtering in the context of topology discovery. Being the first attempt for composing ICMP- and IGMP-based approaches, it shows superior performance than previous homogeneous techniques. Compared to the simple reconnection strategy we introduced in [5] that made solely use of IGMP border IPs, in this paper, we also take advantage of the ICMP neighborhood. Moreover, our novel approach is hybrid and recursive while, in our previous work, the reconnection was basically limited to 2-hop-multicast distant components.

IGMP filtering explained in Sec. II is somewhat equivalent to ICMP filtering encountered by traceroute. Indeed, a router along a traceroute path might not reply to probes because the ICMP protocol is not enabled, or the router employs *ICMP rate limiting*. In order to circumvent such an ICMP limitation, the traceroute vantage point activates a timer when it launches the probe. If the timer expires and no reply was received within the timeframe, then, for that TTL, the distant hop is considered as *non-responding*. Such a non-responding router is called an *anonymous router*. In the literature, techniques have been proposed to infer more accurate topologies in the presence of anonymous routers [6], [7], [8]. Those techniques are mostly passive since they do not require additional probing: Yao et al. proposes a graph minimization approach [6], Gunes and Sarac a graph based induction technique [7], while Jin et al. suggested an ISOMAP-based dimensionality reduction approach [8].

V. CONCLUSION

In this paper, we quantified how IGMP local filtering weakens the topology discovery approach based exclusively on IGMP probing. Similarly to anonymous routers with traceroute, IGMP filtering leads to the collection of topologies made by several disjoint components. Relying on both IGMP and ICMP probing, we deeply investigated the impact of IGMP local filtering on three large ISPs making use of a hybrid graph transformation. We accurately estimated the distances among the disjoint components. This analysis showed that it would be theoretically possible to reconnect almost all the fragments in a single large component. Based on a novel

recursive mechanism, our reconnection strategy is indeed able to strongly reduce the number of IGMP components of a given AS, making thus the resulting topology denser. While the knowledge acquired in this paper can be profitably used in MERLIN, the results we found are more general. More precisely, we are able to experimentally quantify the damages caused by IGMP filtering on collected topologies of large tier-1 ISPs: thanks to traceroute data, we construct a hybrid graph and estimate how far each IGMP fragment is from each other. Finally, based on this preliminary analysis, we design and experimentally evaluate a recursive approach for reconnecting disjoint multicast components. Our topologies are available at <http://svnet.u-strasbg.fr/merlin>.

ACKNOWLEDGEMENTS

A. Pescapé and P. Marchetta were partially funded by the PLATINO project financed by MIUR, and by the LINCE project of the FARO programme jointly funded by the Compagnia di San Paolo and by the Polo delle Scienze e delle Tecnologie of the University of Napoli, Federico II.

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] K. Keys, "Internet-scale IP alias resolution techniques," *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 1, pp. 50–55, January 2010.
- [3] P. Mérindol, V. Van den Schriek, B. Donnet, O. Bonaventure, and J.-J. Pansiot, "Quantifying ASes multiconnectivity using multicast information," in *Proc. ACM/USENIX IMC*, November 2009.
- [4] P. Mérindol, B. Donnet, J.-J. Pansiot, M. Luckie, and Y. Hyun, "MERLIN: MEasure the Router Level of the INternet," in *Proc. 7th Euro-NF NGI*, June 2011.
- [5] P. Marchetta, P. Mérindol, B. Donnet, A. Pescapé, and J.-J. Pansiot, "Topology discovery at the router level: a new hybrid tool targeting ISP networks," *IEEE Journal on Selected Areas in Communication, Special Issue on Measurement of Internet Topologies*, vol. 29, no. 6, pp. 1776–1787, October 2011.
- [6] B. Yao, V. R., F. Chang, and D. Waddington, "Topology inference in the presence of anonymous routers," in *Proc. IEEE INFOCOM*, April 2003.
- [7] M. H. Gunes and K. Sarac, "Resolving anonymous routers in the Internet topology measurement studies," in *Proc. IEEE INFOCOM*, April 2008.
- [8] X. Jin, W.-P. K. Yiu, S.-H. G. Chan, and Y. Wang, "Network topology inference based on end-to-end measurements," *IEEE JSAC, Sampling the Internet: Techniques and Applications*, vol. 24, no. 12, pp. 2182–2195, Dec. 2006.
- [9] 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 IMC*, October 2006.
- [10] J. B. Kruskal, "On the shortest spanning subtree of a graph and the traveling salesman problem," *American Mathematical Society*, vol. 7, no. 1, pp. 48–50, February 1956.
- [11] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with Rocketfuel," in *Proc. ACM SIGCOMM*, August 2002.
- [12] J. Sherry, E. Katz-Bassett, M. Pimenova, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy, "Resolving ip aliases with prespecified timestamps," in *Proceedings of the 10th annual conference on Internet measurement*, ser. IMC '10. New York, NY, USA: ACM, 2010, pp. 172–178. [Online]. Available: <http://doi.acm.org/10.1145/1879141.1879163>