Noname manuscript No.
(will be inserted by the editor)

Probing distribution in time and space for IP alias resolution

Santiago Garcia-Jimenez · Eduardo Magaña · Daniel Morató · Mikel Izal

Received: date / Accepted: date

Abstract The Internet is composed of thousands of networks, interconnected to provide end-to-end IP (Internet Protocol) connectivity. However, Very little public information is provided about these networks and their interconnections. The information needed to create an Internet map of the routers and the links between those routers must be derived from techniques for discovering IP addresses (traceroute) and for associating IP addresses that belong to the same router (IP aliases). Both processes IP address discovery and IP alias resolution require a large measurement infrastructure, and they introduce variable amounts of traffic into the network. Although systematic proposals have been made for creating a scalable IP address discovery system, the equivalent system for resolving IP aliases is far from being determined. In this paper, new proposals for obtaining a scalable IP aliasing system are evaluated and compared with existing solutions. Distributing measurements along multiple vantage points (spatial distribution) and extending probing tasks over time (temporal distribution) have been identified as the key methods for reducing the overhead of IP alias resolution.

Keywords Networks · Protocols · Performance evaluation, Quality of service

1 Introduction

Over the past decade, the interest in the Internet mapping has increased significantly. Since the creation of the Traceroute in 1987 [1], significant effort has

Santiago Garcia-Jimenez · Eduardo Magaña · Daniel Morató · Mikel Izal
Automatica y computacin (telemtica)
Universidad Pblica de Navarra
Campus de Arrosadia
31006 Pamplona
Spain
Tel.: +34 948 16 60 33
E-mail: santiago.garcia@unavarra.es

been put into creating a map of the Internet using only IP address discovery techniques. The resulting maps are graphs composed of nodes (IP addresses) and the links between these nodes (adjacencies detected in the Traceroute paths). And include as many router IP addresses as possible. Example of these efforts include Skitter [2], Rocketfuel [3], and more recently iPlane [4], ARK [5] and DIMES [6].

However, numerous applications require a router-level Internet map, a graph composed of nodes representing routers and links representing the connectivity between the interfaces of the various routers. Therefore, several IP addresses corresponding to different interfaces of the same router are aggregated into the same graph node. Those IP addresses are called *IP aliases* and the process is called *IP alias resolution* [7].

Router-level topology maps are rarely made available by Internet service providers (ISPs) due to security reasons or to a reluctance to share information with competing companies. However, this information is useful when attempting to improve Internet application scaling. For example, a router-level Internet map could help create a more realistic network simulation. Currently, most Internet simulations use techniques related to synthetic network generation based on analytical models [8][9] but it is unclear if they are representative, due to a lack of evidence [3].

Router-level Internet maps can also be used in P2P (Peer-to-Peer) balancing schemes [10] and routing protocols [11]. For P2P overlay networks, it is important to know the real network structure to send through the actual closest peers [12]. Another potential application for these maps is selecting the best ISP based on, for example, the path bandwidth [13]. For tasks related to network administration, this type of map can be used to reveal possible bottlenecks and can also be used to predict paths and latencies, as proposed in [14].

When facing security problems, a router-level Internet map can be used to locate the exact source point of a denial of service attack [15]. Finally, geolocation systems can also be improved using the geolocation of intermediate points in the network. The fact that two or more IP addresses belong to the same geographical point provides more information when attemping to determine the correct location for the rest of the IP addresses [16].

As stated before, the creation of router-level Internet maps involves two separate tasks: IP address discovery and IP alias resolution. IP address discovery has traditionally been performed with the traceroute tool, but the record-route option in the IP header can also be used. Recently, a variant called the paris-traceroute [17] was shown to obtain better results by avoiding the load-balancing effects in the Internet paths.

A more recent discovery tool called MERLIN [18] tries to solve some of the drawbacks related with the filtering of probe packets and the lack of responses in tools like traceroute. By mixing IGMP probes, traceroutes and the Ally resolution technique [3] (explained later), Merlin is able to discover the routers in different Autonomous Systems.

This paper focuses on the IP alias resolution task; The existing IP alias resolution techniques present serious scaling difficulties. Section 5 will show that the best identification results are obtained with those IP alias resolution techniques that have a quadratic cost in time and bandwidth with regard to the number of IP addresses. The application of these techniques to large-scale Internet maps is therefore not realistic. Other IP alias resolution proposals have linear costs with regard to the number of IP addresses, but suffer from greater rates of misidentification. This paper analyzes two alternatives for achieving scalability, even for IP alias resolution techniques with quadratic costs: the distribution of probing between several vantage points (spatial distribution) and extending probing campaigns over time (temporal distribution). These enhancements which help cover bigger network topologies, are analyzed in detail in this paper.

The main contribution of this paper is to prove that it is possible to perform an alias resolution in extensive networks by using those alias resolution techniques that provides the best identification results. Those techniques have just the drawback of quadratic cost like Ally-based techniques. However, this extra cost can be assumed thanks to the distribution of active probing in time and space. Extending the measurement campaigns in larger time intervals is possible thanks to the stability found in alias for Internet routers up to 30 days. At the same time, active probing measurements can be distributed between a larger set of vantage points. Therefore, the extra cost in the amount of probing traffic and the processing of these techniques can be assumed with the advantage of providing the best alias resolution results.

Another interesting finding of this work is related with the lack of guides to select the right probing parameters depending on the network scenario like for example the time between packet probes to avoid packet filtering in routers. In this paper, a design rule of RadarGun [19] parameters has been provided to optimize the identification results.

This paper is organized as follows: Section 2 presents the main IP alias resolution techniques and the concerns in obtaining a router-level Internet map. Section 3 presents the real network scenarios used in the evaluation. A study of router stability is performed in section 4. Section 5 presents the evaluation of IP alias resolution techniques with linear and quadratic costs. Section 6 considers the scalability concerns related to IP alias resolution in large-scale router-level Internet maps. Finally, the conclusions and future work are presented.

2 Previous studies on IP alias resolution

There are several proposed methods for IP alias resolution in the literature, but they may provide incomplete or even false identification results. Several types of results can therefore be distinguished. Two IP addresses can be tagged as *positive* (or true positive) when both IP addresses belong to the same router or *negative* (or true negative) when they do not belong to the same router.

When an insufficient number of response packets are received from the IP addresses to complete the identification, the technique will produce an *error.* Finally, if the technique has all the needed packets but the information required for identification cannot be retrieved (for example, the response packets are empty), the output is *unknown.* IP alias resolution techniques can also provide incorrect information. When a technique provides a positive result and both IP addresses actually belong to different routers, the result is called a *false positive.* The same occurs for the reverse situation; if a technique provides a negative result when both IP addresses actually belong to the same router, the result is called a *false negative.*

These possible results help evaluate IP alias resolution techniques using metrics such as accuracy, completeness, efficiency and distributability [20]. Accuracy measures the number of errors committed during the IP alias resolution. Completeness measures the ratio of aliases and non-aliases that have been identified, compared with the total number of IP addresses pairs. Efficiency is related to the intrusiveness of the method. Distributability indicates whether data collection and processing can be distributed among variouos vantage points (controllable nodes that generate probe packets), and is therefore scalable. These four metrics will be reviewed in the following sections.

One of the first proposals for IP alias resolution was the Mercator technique [21], which was used by the CAIDA research group in the well-known topology tool called Skitter [2]. Mercator is based on the default behavior of certain routers when some of its interfaces have to answer a UDP packet that was sent to an unused port. A common behavior in this situation is to return an “ICMP port unreachable” packet from one specific interface of the router with the shortest-path to the destination. UDP probing packets that are sent to different target IP addresses belonging to the same router are answered with “ICMP port unreachable” packets that share the same source IP address. One probing packet per IP address is enough to perform this technique, but the vantage point has to be the same for all probing. The identification rates using this technique are not good, because the response packets are mostly filtered out by the routers [22].

The next advance in the field came from the development of a technique called Ally [3], which is used in the Rocketfuel monitoring tool and is based on the IP identification field (IPID) of the IP header. Due to segmentation and reassembly issues in the IP protocol, the IPIDs from different IP packets should be different. A common implementation of the IPID is the use of an incremental counter. Each new packet increases the counter by one unit, and this value is used to fill up the IPID field in a new IP datagram generated by the router. The counter is shared by all router interfaces, and therefore the IPIDs from the different interfaces in the router will show an incremental, correlated pattern. Because this IPID counter is increased by any traffic generated (not forwarded) by the router, the probing has to be performed in short time intervals to avoid possible interference from traffic external to the probing.

Ally is based on sending three UDP packets to an unused port. The first two probe packets are sent back-to-back to the two IP addresses being checked for aliasing. One second later, a third probe packet is sent to the IP address whose "ICMP port unreachable" response was received first. The distance between the first and the last IPID must be less than 200 IPIDs to identify both IP addresses as aliases. Improvements to the Ally technique were presented in [23], where modifications were introduced, for example, in the type (ICMP instead of UDP), number and timing between probe packets. Significantly better identification results are obtained with these variants that will be referenced in this paper as *Ally-based techniques*.

Another proposal uses the prespecified timestamp option in the IP header, which selects up to four IP addresses and receive the timestamps from those IP addresses [24]. Typical implementations provide millisecond timestamps that check wether two IP addresses are aliases (aliases have the same timestamp).

Finally, there is a set of inference methods for IP alias resolution that is based on graph analysis without using any extra probe traffic. The methods use heuristics to join the expansion trees obtained by different traceroutes. The analytical Alias Resolver (AAR) [25] and the Analytical and Probe-based Alias Resolver (APAR) [26] are examples of these methods. However, their identification results are worse than those obtained with Ally-based techniques [23][27].

With Ally-based techniques [23], the probing must be performed by pairs of IP addresses, and therefore, as the number of IP addresses (N) in the topology increases, the cost of the Ally-based methods not only increases, but increases quadratically.

To solve this scalability problem, the RadarGun tool [19] was proposed. This tool uses the IPID values returned in response to UDP and TCP probe packets to derive the so-called "velocity modeling process" [19], which models the growth profile of IPIDs per target IP address. The model performs a linear regression, which describes the IPIDs received by each IP address. For each point, the error is the distance between the original line for the first IP address and the point of the second IP address. If the mean of the errors is more than a specific threshold, both IP addresses will be cataloged as non-aliases. Otherwise, the IP addresses will be considered aliases. RadarGun sends, back-to-back probing packets from a centralized vantage point to all target IP addresses simultaneously and uses 30 probe packets per target IP address to perform its velocity modeling. RadarGun probing is not performed by pairs of IP addresses, thus reducing the overall cost to a linear $O(N)$. However, RadarGun has problems with large network topologies because the time between probe packets to the same target IP address can be too long, thereby producing IPID wrapping [22].

MIDAR (Monotonic ID-Based Alias Resolution) [28] has been proposed by CAIDA as an evolution of RadarGun. The method is based on the same idea as RadarGun but changes the way the returned IPIDs are processed, keeping the overall cost to linear $O(N)$. MIDAR can collect IPID data from several vantage points simultaneously and aliases are identified using monotonicity tests rather than proximity tests. The main drawback of this tool is that it does not provide information about non-alias: pairs of IP addresses that do not belong to the same router. MIDAR results are good in identifying alias but it does not provide a good completeness ratio because of the lack of non-alias identification. This drawback also causes not knowing when the alias resolution process has finished because only when all possible pairs of IP addresses have been identified as alias or non-alias it can be said that the process has finalized.

Reduction methods are used to reduce the number of probe packets needed for IP alias resolution [29] and are based on a pre-selection of IP address pairs that are more likely to be aliases and the subsequent application of Ally-based techniques to just that subset of IP addresses pairs. These reduction methods significantly decrease the number of IP alias resolution tests that need to be performed, but they can also reduce the completeness of the analysis.

Three reduction methods are described in the literature, according to the attributes used to preselect the IP address pairs: TTL-based, IPID-based and IPoffset-based. A TTL-based reduction method [4] is based on the fact that two IP addresses that belong to the same router will face a close hop count when measured from a single vantage point. Although both IP addresses belong to the same router, probe packets can follow different paths to reach each interface and, therefore, some difference in the hop count is allowed. The IPID-based reduction method [30] determines that two IP addresses that belong to the same router will send packets using similar IPID numbers. This reduction method requires extra probing to obtain the IPID values, compared to a TTL-based method where the hop count information can be extracted from the traceroutes. For the IPoffset-based method [31], the IP offset metric is defined as the difference between two IP addresses considered as two unsigned integer numbers. The method is based on the fact that the probability of finding aliases increases significantly in IP address pairs belonging to some specific IP offsets. Without any extra probing traffic, the IPoffset-based reduction method allows researchers to reduce the number of IP address pairs that need to be checked for aliases to 10% [31], keeping the accuracy and completeness of the original IP alias resolution technique almost intact.

When creating extensive router-level Internet maps, IP alias resolution techniques are slow and need to introduce significant probe traffic into the network. Although Mercator is scalable (one probe packet per target IP address is enough), its identification results are very poor [23]. The Ally-based methods provide better identification results at the cost of making identification by pairs of IP addresses, which is not scalable. Reduction methods help to decrease the overall cost in processing time and the amount of probe traffic, but with uncertain effects in terms of accuracy and completeness. We will evaluate These effects and present improvements for Ally-based to obtain a level of scalabilitity closer to that of RadarGun.

3 Network scenarios

Four networks have been chosen in this paper to verify the performance of alias resolution techniques: Geant (pan-European data network [32]), Canet4

Table 1 Size and known data about selected network scenarios

Network	*Routers*	*IP addresses*	Number of aliases	Aliases percentage
Canet4	6	103	1225	23.78
GlobalNOC	16	569	13832	8.58
Geant	19	493	7441	6.11
PlanetLab subset 1	-	1971	-	-
PlanetLab subset 2	-	1282	-	-
PlanetLab subset 3	-	4844	-	-

(Canada's Research and Education Network [33]), GlobalNOC (Internet2 at Global Research Network Operations Center [34]) and PlanetLab [35].

The first three networks are NRENs (National Research and Educational Networks) with public information about their network topologies (routers and IP addresses per router interface). Details about the number of routers, the number of IP addresses and the number of pairs of IP addresses that are aliases for these network scenarios are presented in Table 1. The table also shows the 'Aliases percentage' defined as the percentage of pairs of IP addresses that are aliases, compared to the total number of pairs of IP addresses in each network scenario. In the table, not-available data are represented with a dash. These network scenarios were used to verify the accuracy and completeness of the IP alias resolution techniques.

Those three reference networks are not very large, as observed in Table 1, but they can be a reference to what would happen in bigger networks. In these networks, the discovery task is not needed because the actual IP addresses are already known.

The fourth network is provided by the PlanetLab measurement infrastructure. PlanetLab offers hundreds of nodes distributed around the world, but this time there is no information about the real network topology interconnecting those hosts. Therefore, this network cannot be used to evaluate accuracy or completeness, but it can be used to analyze efficiency and distributability. The paris-traceroute was used to discover the IP addresses of routers that interconnect three subsets of PlanetLab end-nodes: subset1 with 40 nodes, subset2 with 20 nodes and subset 3 with 56 nodes, distributed around the world. The paris-traceroutes were performed between each pair of nodes to discover a considerable number of possible aliases for each subset. This network provides a large topology with core and access routers.

An additional network scenario has been used to analyze router configuration stability on Internet, using the Etomic measurement infrastructure[36]. This infrastructure provides 18 nodes distributed across Europe, where various experiments can be launched using specialized hardware: clocks are GPS-synchronized and network interface cards provide high-precision functionalities. Etomic offers an open repository for which historical data of daily paris-traceroutes with 2,928 IP addresses is available since 2007. These data were used in the analysis of router configuration stability.

4 Evaluation of router configuration stability in the Internet

The creation of extensive router-level Internet maps requires significant time to perform the Ally-based techniques compared with the RadarGun approach. Thus, we need to know how long the IP addresses of each interface of a router (router configuration) remain unchanged. To quantify the maximum time available to perform IP alias resolution in a certain topology, a characterization of router configuration stability in the Internet is needed. The aliasing information changes over time as the router configuration change. The part of the router configuration that matters in alias resolution is the number of interfaces and the IP addresses assigned per interface. Therefore, these parameters were analyzed to look for changes over time.

To perform the analysis, Etomic routing data were considered. This database provides paris-traceroutes between 18 nodes distributed around Europe, running 3 times per day since 2007. This routing data has been processed to check for how long the IP addresses of routers appear in a specific path between source and destination: *IP address persistence*. IP address persistence is defined, for a pair of traceroute endpoints, as the time that an IP address is present in a certain path between a source and a destination. The same IP address in another path was considered separately. Therefore, if one IP address is present in several paths, an IP address persistence will be obtained per path.

Figure 1 presents the cumulative distribution function (CDF) of IP address persistence. It shows the probability of having an IP address that will remain in the system for at least a certain period of time. It can be observed that many IP addresses remain for quite a long time. Approximately 80% of the IP addresses remain for at least 200 days. Additionally, it can be observed that for 90% of the IP addresses the stability time is at least 44 days, and for 99% of the IP addresses this time is at least 11 days.

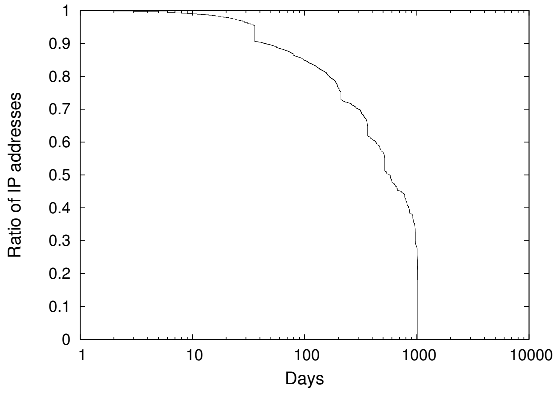

Fig. 1 CDF of IP address persistence in Etomic scenario

A similar analysis was performed with data obtained from the Internet Mapping Project [37]. These data were collected from only one host by doing traceroutes to 8 million IP addresses distributed around the world. It provides traces spanning 6 months, running once per day. The set of destination IP addresses analyzed varies each day, so only 129 of the 8 million destination IP addresses can be observed in all traces. The study was performed for this subset of destination IP addresses to avoid problems related to the absence of continuous data for the rest of the destination IP addresses. The results are shown in Figure 2, again using the CDF of IP address persistence. Persistence for 90% of the IP addresses was at least 34 days.

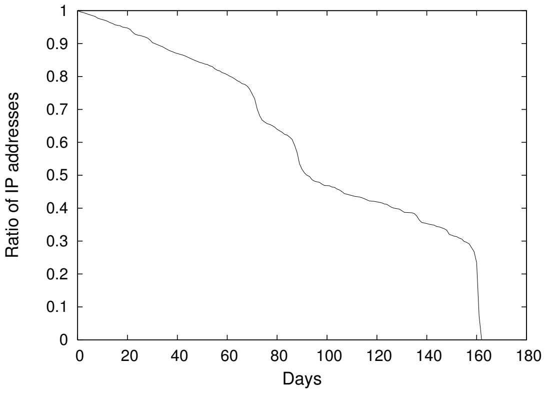

Fig. 2 CDF of IP address persistence in Internet Mapping Project scenario

As a result, the router configuration was considered stable for time intervals up to 30 days. This result is useful when distributing active probing over several weeks in IP alias resolution techniques and therefore is of great importance in IP alias resolution techniques with quadratic costs, such as Ally-based techniques.

5 Performance evaluation of RadarGun and Ally-based techniques

In the following subsections, RadarGun and Ally-based techniques have been chosen as representative IP alias resolution techniques with linear and quadratic costs respectively. They are evaluated for various scenarios, taking into account accuracy, completeness, efficiency and distributability metrics. RadarGun software was obtained from its authors' website, but today that webpage is not available and the original RadarGun software has been made available at [38]. The rest of the software and the data sets have also been made available at [38].

5.1 Accuracy and completeness

We evaluated the RadarGun and Ally-based techniques for accuracy and completeness for the three well-known networks Geant, Canet4 and GlobalNOC. For all three, details of the real topology are known, and therefore it is possible to obtain rates of positives (true positives), negatives (true negatives), false positives and false negatives as shown in Table 2 for RadarGun and Table 3 for Ally-based techniques. These tables show the IP alias resolution results per pair of IP addresses with respect to all pairs of IP addresses in each scenario. The column called 'Identified' is the sum of positives and negatives. It indicates the percentage of pairs of IP addresses that have been identified as alias or non-alias. The column called 'Aliases' represents the percentage of positives with respect to the real number of aliases present in each scenario (positives were considered with respect to the total number of pairs of IP addresses in the scenario). This column is another way to evaluate the identification performance. The meaning of columns titled 'Error' and 'Unknown' is the same as explained in section 2. Not available data are represented with a dash. The sum of the columns titled 'Positives', 'Negatives', 'Error' and 'Unknown' completes the 100% of occurrences.

The accuracy is shown in the low percentage of false positives and false negatives (errors in the identification). Ally-based techniques have neither one. RadarGun has some low percentages of false positives and false negatives, especially for the Canet4 network, in which false negatives reached 2.74%. The reasons for these differences between network scenarios are presented later.

The completeness can be observed in the total number of IP address pairs identified: the sum of positives and negatives (column labeled 'Identified'). This time, completeness in the Ally-based techniques clearly outperforms the results in RadarGun. Moreover, it can be observed that in the Geant and GlobalNOC networks, RadarGun was almost useless. Although false positive and false negative rates were low in RadarGun, they are significant compared to the low identification rates provided as positive and negative aliases.

In Tables 2 and 3, Planetlab scenarios are also shown. For them, there is no public information about the underlying topology. Therefore, for those scenarios the results for accuracy (false positives and false negatives) and percentage of real aliases are not available for Ally-based techniques, and in RadarGun are referenced with those obtained in Ally-based techniques [23]. For Canet4 and PlanetLab, RadarGun results in accuracy and completeness are better than in the GlobalNOC and Geant scenarios. However, as noted earlier, those results are clearly worse than those obtained with Ally-based techniques.

These results indicate that, from an accuracy and completeness point of view, Ally-based techniques provide better identification results. The analyzed scenarios are real and diverse, and therefore they can be considered representative of common network topologies. The reasons for this difference in performance are discussed in the following subsection.

Table 2 IP alias resolution percentages by RadarGun

Network	*Positives*	*Negatives*	*False positives*	*False negatives*	*Identified*	*Aliases*	*Error*	*Unkwown*
Canet4	5.90	35.27	0	2.74	41.17	24.81	58.49	0.34
GlobalNOC	0	0.001	0	0.0006	0.001	0	99.99	0
Geant	0.12	1.50	0.004	0.08	1.62	1.96	98.21	0.17
PlanetLab subset 1	0.055	28.30	0.008	0.002	28.355	-	71.56	0.08
PlanetLab subset 2	0.013	46.76	0.001	0.016	46.773	-	53.18	0.04
PlanetLab subset 3	0.00	15.60	18.11	0.48	15.61	-	84.29	0.00

Table 3 IP aliases resolution percentages by Ally-based techniques

Network	*Positives*	*Negatives*	*False positives*	*False negatives*	*Identified*	*Aliases*	*Error*	*Unkwown*
Canet4	9.26	48.73	0	0	57.99	38.94	8.12	33.89
GlobalNOC	4.41	42.94	0	0	47.35	51.39	10.15	42.50
Geant	5.11	85.67	0	0	90.78	83.63	0.16	9.06
PlanetLab subset 1	0.11	47.68	-	-	47.79	-	3.07	49.14
PlanetLab subset 2	0.09	44.50	-	-	44.59	-	4.04	51.37
PlanetLab subset 3	0.40	50.10	-	-	50.51	-	15.05	34.44

5.2 The effects of topology size on accuracy and completeness

The size of the topology can have effects on the identification results of IP alias resolution techniques. The effect of topology size were analyzed for RadarGun and Ally-based techniques.

By default, RadarGun uses 80 Kbps bandwidth (in the RadarGun authors' implementation) and 30 rounds to probe all the target IP addresses each round. The tool allows researchers to manually specify the bandwidth to use, but this value is hard to optimize and the authors do not provide any rule to find the right bandwidth value. Using a specific bandwidth means that the time between probes sent to the same IP addresses increases with the number of IP addresses to consider. With few IP addresses, the inter-probing time is short, maybe in the milliseconds order. Some routers can be configured with policy rules that imply not answering probes from the same source very close in time. Therefore, identification results can be impaired in small topologies or when using vantage points with high-speed trunks.

Figure 3 shows identification results (sum of positives and negatives) with RadarGun in a scenario composed of the aggregation of Geant, Canet4 and GlobalNOC networks, depending on the number of IP addresses considered. As expected, if less than 70 addresses are used with the default RadarGun bandwidth (80 Kbps), some of the response packets are lost and therefore worse identification results are obtained. The inter-probing time for 70 IP addresses is 0.3 seconds and for 40 IP addresses is 0.2 seconds. Policy rules apply in those cases, limiting the number of responses to avoid overhead in the router or for security reasons. Previous studies have found that inter-probing times above 0.4 sec avoid this effect [39].

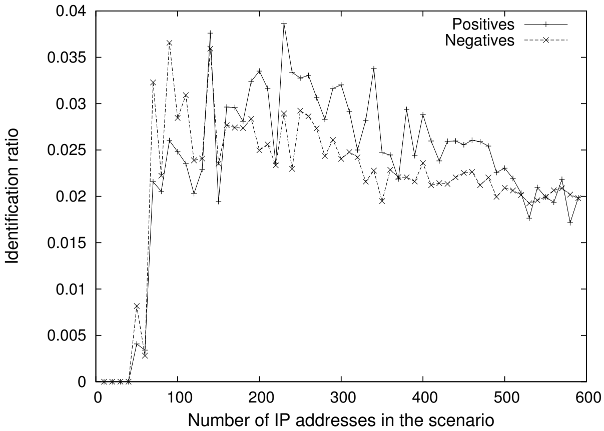

Fig. 3 Identification results by RadarGun and different topology sizes

Above 400 IP addresses, the identification results in Figure 3 become worse. In large scenarios, RadarGun has problems related again to the inter-probing time to the same target IP addresses. The IP addresses have to be probed in rounds, and the duration of these rounds increases with the number of IP addresses. If there are a large number of IP addresses to check for aliasing, inter-probing time for a target IP address can be too long, in the order of several seconds or even minutes. In those cases, alias resolution using IPIDs performs poor because of the high variability in the growing profile of IPIDs in routers [3] and the presence of IPID counter wrapping [22]. Then, the IPID increments in a router cease to have a linear behavior from the probing vantage point of view and the velocity modeling used by RadarGun fails. For example, in Figure 3, by using a number of IP addresses close to 520, the aliasing identification values decrease. The inter-probing time in that case is 3.1 seconds. Therefore, with inter-probing time between 0.2 seconds and 3.1 seconds RadarGun provides the best identification results.

As our well-known scenarios are relatively small, this effect can be observed varying the probing speed in the vantage point used by RadarGun instead of using the default bandwidth. Reducing the probing speed means increasing inter-probing time for each target IP address and the commented effect can be observed. Figure 4 shows the effect of the probing speed in RadarGun for three different scenarios. Those scenarios are composed od 58, 116 and 944 IP addresses chosen randomly from the aggregation of Geant, Canet4 and GlobalNOC networks. For a certain scenario (number of IP addresses) and probing speed, the inter-probing time changes, producing variations in alias identification (positives and negatives), as shown in Figure 4.

For low probing speeds, identification results are poor because of the previously explained large inter-probing time. For medium probing speeds, the results are quite stable. Finally, for high probing speeds, the identification results become worse again because of the effect of short inter-probing time commented on earlier. This last effect does not appear for the curve of 933 IP addresses because not enough probing speeds were plotted in the figure. With faster probing speeds, the effect is the same as in the other curves.

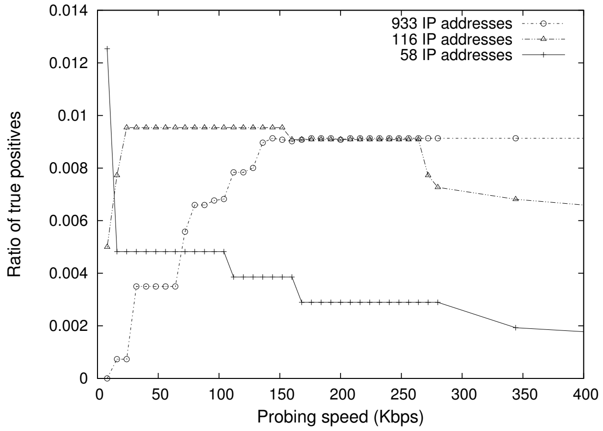

Fig. 4 Identification results by RadarGun and different probing speeds

As a conclusion, there are two thresholds in inter-probing time for RadarGun. The first one makes the router not answering in response to probe packets (probe packets are sent too close) and the second makes the identification process useless (probe packets are widely spaced and the IPID-based technique fails). Both thresholds have been previously figured out from Figure 3 but they can also be estimated from Figure 4.

The first threshold can be obtained from the curve titled "116 IP addresses" in Figure 4. There, the ratio of true positives decreases in a clear step around 270Kbps. Assuming that RadarGun technique uses 64 bytes probe packets and rounds probing one time each IP address per round, this means that probes to the same IP address have an inter-probing time of $119 IPaddresses * 64 bytes/270 Kbps = 0.2$ seconds. A similar value is obtained with the curve titled "58 IP addresses".

The second threshold is obtained from the curve titled "933 IP addresses" in Figure 4. There, the ratio of true positives grows to a stable good result by using a probing speed bigger than 150 Kbps. As before, this means an inter-probing time of $933 IPaddresses * 64 bytes/150 Kbps = 3.1$ seconds. A similar value can be observed on the curve titled "116 IP addresses" when the probing speed is more than 20 Kbps.

Therefore, RadarGun should be used with the right inter-probing time between those temporal thresholds to obtain the best results in IP alias resolution. Those thresholds are approximately 0.2 secs and 3.1 secs. With these constraints, the probing speed in RadarGun determines the topology size. For a given bandwidth B in bits per second and a given size of probe packet S in bits, the number of reachable IP addresses in the topology would be:

$$N = (B * t)/S = (B * t)/(64 * 8) \text{ with t}\in[0.2,3.1] \tag{1}$$

For example, this means that with probe packets of 64 bytes in size and a probing speed of 10 Mbps in the vantage point, RadarGun can operate in a network scenario with up to 60K IP addresses, but always larger than 4K IP addresses.

With the default bandwidth of 80Kbps used by RadarGun, applying equation 1 the number of addresses to probe should be between 32 and 484. This is the bandwidth used in the previous subsection, and it justifies why the best results were obtained for Canet4 (103 IP addresses), the reasons for the poor results obtained with Geant (493 IP addresses) and the very poor results obtained with GlobalNOC (569 IP addresses). Therefore, knowing and using equation 1 is necessary to make a good use of RadarGun and similar tools. It can be used to calculate the bandwidth to use in the RadarGun tool for a certain topology size. It must be noted that in the original work [19], this working zone for RadarGun was not identified and here we have demonstrated its importance.

In Ally-based techniques, the size of the network scenario does not affect to the identification results. As probing is made per pair of IP addresses, a larger scenario means a higher number of pairs to check for aliasing, but there is no loss in quality of the results for the identification. The drawbacks will appear in the cost/efficiency of this identification, which is analyzed in next subsection.

5.3 Efficiency

The efficiency metric in IP alias resolution is related to how much probing traffic is needed and for how long the probing has to be performed to obtain the best identification results.

RadarGun is better in efficiency, as its probing phase depends linearly on the number of IP addresses to check for aliasing. Therefore, it is fast and it introduces a reduced amount of probing traffic in the network. An estimation of consumed resources in the vantage point can be obtained considering probe packets of 64 bytes in size and 30 packet probes per IP address. This means a total amount of bytes transmitted R and a total time to perform the identification T in seconds:

$$\begin{aligned} R =& N * 64 * 8 * 30 = 15,360N \\ T =& R/B \end{aligned} \tag{2}$$

For a topology of 5,000 IP addresses and 10 Mbps, this means approximately 7.68 secs.

In Ally-based techniques, the efficiency is conditioned by the need to probe IP addresses in pairs. This means a cost that grows in quadratic order with the number of IP addresses, but that can be reduced applying reduction methods. In any case, large network scenarios will require prolonged probing phases. However, this probing phase cannot extend indefinitely because of router configuration changes as analyzed in section 4. The probing phase should not extend beyond the 30-day period of configuration stability found in routers.

A simulation of the time needed to perform Ally-based resolution has been made, considering a probing speed of 10 Mbps and different numbers of IP addresses in the scenario. In this case, 16 probe packets are needed to complete an alias resolution test over a specific pair of IP addresses [23]. The results are shown in Figure 5, with and without a reduction method. The chosen reduction method is IPoffset-based [29]. The figure shows how in 30 days, 58,986 IP addresses can be identified with direct Ally-based resolution, and 180,776 IP addresses can be identified with Ally-based resolution combined with the IPoffset-based reduction method.

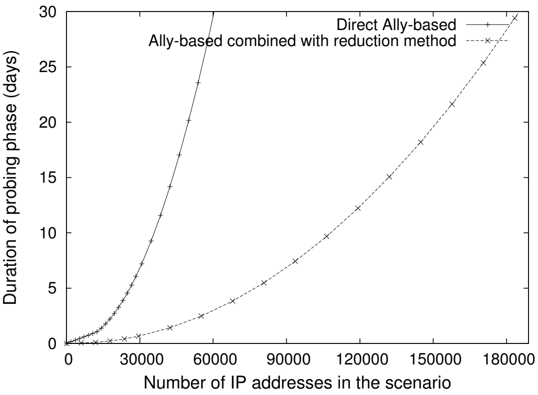

Fig. 5 Duration of probing phase for Ally-based techniques

Increasing the topology size without exceeding the 30-day time period implies increasing the probing speed. Figure 6 shows the duration of the probing phase for Ally-based techniques with different probing speeds. This time, there is not a threshold for the upper probing speed as in RadarGun because probing is made in pairs of IP addresses. Increasing the probing speed to 100 Mbps means being able to apply IP alias resolution to 591,212 IP addresses. Additionally, probe packets can be distributed between different vantage points as will be presented in next section. The probing speed per vantage point can be kept unchanged if the number of vantage points is increased.

Therefore, although Ally-based techniques provide the best results in accuracy and completeness, their efficiency is poor, even taking into account reduction methods. In the next subsection, distributability will allow researchers to solve the penalization introduced by the high cost of Ally-based techniques.

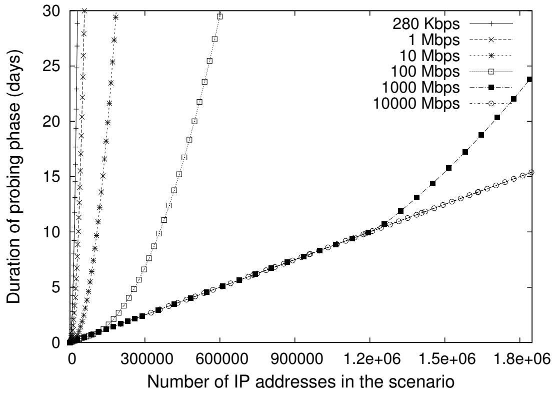

Fig. 6 Duration of probing phase for Ally-based techniques and different probing speeds

5.4 Spatio-temporal distribution of probing in alias methods

Two types of distributability will be considered for IP alias resolution. First, distributability in time, defined as the capacity to extend the probing phase in time to reduce the probing speed. This is important because high probing speeds can be misinterpreted as a source of network attacks and can even be blocked by ISPs, especially because the specific targets of this traffic are ISP routers. Second, distributability in space is defined as the capacity of a specific aliasing method to be executed from more than one vantage point simultaneously, to reduce the total time or the probing speed per vantage point.

In RadarGun, probing is performed in rounds to all target IP addresses. As discussed in subsection 5.2, this means that inter-probing time to the same target IP address cannot be extended much to falicitate alias identification rates. Considering the maximum 3.1 seconds of inter-probing time we discussed in subsection 5.2, 30 rounds mean approximately 1.5 minutes as the maximum duration of the probing phase. Therefore, RadarGun is not freely distributable in time, which could be useful to reduce the probing speed from the vantage point.

With regard to distributability in space, RadarGun is centralized by design although it could be extended to support probing from multiple vantage points (but not easily, because interferences with probing between different vantage points has to be avoided). In RadarGun, the probing is performed from a unique vantage point and afterwards it correlates velocity modelling profiles obtained per IP address. Therefore, clear scalability problems are present in the tool because limitations will be imposed by the probing speed at the vantage point. For example, security policies could limit this probing speed. With 10 Mbps at the vantage point, RadarGun could operate in a network scenario with up to 60K IP addresses (equation 1). MIDAR is supposed to allow several vantage points while keeping a linear cost, but as exposed in section 2 its results are very limited currently.

Ally-based techniques operate over pairs of IP addresses. The probing and analysis of each pair is performed independently of the others. Therefore, Ally-based techniques are distributable in nature. These techniques can be distributable in time, extending the probing phase over large periods of time. This period should be limited by the estimated configuration stability in routers, on the order of a month as previously justified.

With regard to distributability in space, Ally-based techniques can be launched from any number of vantage points simultaneously. Each vantage point would be in charge of probing a subset of IP address pairs. The only consideration is to avoid probing the same IP addresses from different vantage points corresponding to the different pairs, to prevent interference. The identification results calculated by each vantage point can be finally downloaded to a central information point. Multiple vantage points will allow reducing the total time needed to perform the identification or reducing the probing speed needed per vantage point.

The inefficiency found in Ally-based techniques can thus be compensated using this distributability in time and space. For example, using Ally-based techniques from only one vantage point, alias resolution for 60K IP addresses could be completed using the same bandwidth as RadarGun (10 Mbps) in 168 days. If a bandwidth of 100 Mbps is used instead the identification can be performed in 53 days. The reduction factor is not 10 times because the same IP address belonging to different pairs cannot be probed simultaneously, and this guard period extends the required analysis time significantly. This duration can also be reduced by the distribution of experiments to more nodes. If 10 vantage points and 10 Mbps bandwidth per vantage point are used, the duration of the identification (for the same example network) is 16 days. It must be noted that limitations on inter-probing time do not apply if probe packets are sent from different vantage points, because the probe packets will have different source IP addresses and thus they will not invoke the contention rules in the routers.

From the point of view of the network operators, great bandwidth utilization from a source to various destinations may not mean a security attack, but high traffic rates from one source to one destination could be observed as a flood or as a type of denial-of-service (DoS) attack. Therefore, it is also important to distribute the probing phase between different vantage points. In Figure 7, the duration of the probing phase in Ally-based techniques is plotted for different numbers of topology size and different numbers of vantage points (1 to 10). In the Figure, for each number of vantage points the amount of probing traffic received per router IP address is indicated. Increasing the number of vantage points linearly reduces the duration of the probing phase. It also increases the probing speed received per target router IP address, but this traffic is shared between different source vantage points. Therefore, the security constraints in the destination subnetwork can be circumvented.

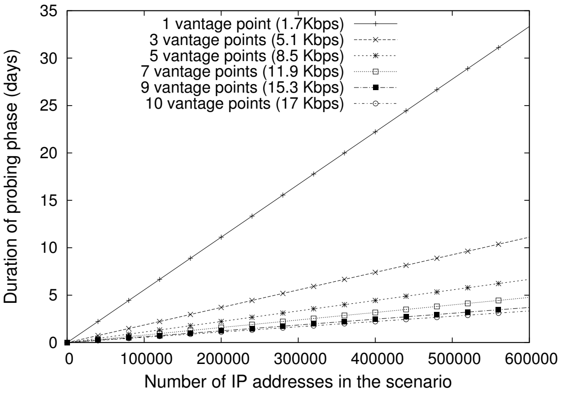

Fig. 7 Duration of probing phase for Ally-based techniques for different number of vantage points and probing speed received by target IP addresses

6 Internet topology map at the router level

There are different projects for Internet mapping and some of them include IP alias resolution. Mapping all Internet IP addresses is an impossible task, so the idea of those projects is to apply the mapping over a significant subset. Depending on the size of this subset, an evaluation of RadarGun and Ally-based techniques can be provided for an extensive Internet topology map at the router level. In the case of Ally-based techniques, the number of vantage points and the duration of the probing phase can be obtained for the specific problem size.

To estimate the number of IP addresses that the Internet has in its interconnection, two different sources of topology data have been used: Scamper [40] and DIMES [6]. In these studies, only IP addresses belonging to routers have been considered, removing source and destination IP addresses of traceroutes. From these data sources, the number of IP addresses cataloged as belonging to routers are 290,000 in Scamper and 579,236 in DIMES (dated for one entire month in February 2011). These figures are consistent with the number of subnetworks announced by BGP (Border Gateway Protocol) entries [41]. The DIMES figure will be considered as accurate in this paper, as it presents a worst-case scenario for the identification process.

In RadarGun, 579,236 IP address can be checked for aliasing using 11.4 Mbps (equation 2, considering 3.1 seconds for inter-probing time). However, the expected identification results are not very good. In Ally-based techniques, distributing the identification over a month (the maximum stability time per router) and using 600 vantage points (Planetlab nodes for example), the probing speed needed per vantage point would be 1.7 Mbps. This probing speed is acceptable, and it can be further reduced by increasing the number of vantage points. In addition, the expected identification results are much better. Therefore, Ally-based techniques are a good alternative to obtain extensive Internet topology maps if the measurement can be extended in time and distributed between different vantage points.

7 Conclusions

In this paper, we performed an in-depth comparison of RadarGun and Ally-based techniques the representatives of IP alias resolution techniques with linear and quadratic costs respectively. Regarding the accuracy and completeness metrics, Ally-based techniques greatly outperform RadarGun for most network scenarios, using the default parameters in Radargun. In fact, some limitations have been found with designing measurement campaigns based on RadarGun, such as those imposed by equation 1 related to the size of the network scenario and the probing bandwidth. For example, in Canet4 (Canada's NREN), RadarGun identification is 41.17% and Ally-based identification is 57.99% considering the above mentioned design rules. Not following those design rules, the identifications obtained for RadarGun are remarkably poorer. With regard to efficiency in probing traffic, RadarGun is clearly the faster and lighter. However, it is compensated for by the distributability properties in Ally-based techniques.

The topology size that can be analyzed by RadarGun has limits, depending on the probing speed available in the vantage point (unique). For example, a probing speed of 10 Mbps allows checking for aliasing in up to 60K IP addresses (equation 1). In Ally-based methods, it is possible to distribute the probing over time and over different vantage points. The distribution in time is possible because router configuration can be considered stable in time periods of approximately a month. The distribution in space (several vantage points) is possible because the probing is made per pair of IP addresses. Subsets of IP address pairs can be distributed to different vantage points. Therefore, the overhead of Ally-based techniques can be compensated for.

Extensive Internet topology maps are, therefore, approachable with Ally-based techniques using their possibilities of distribution in time and space, resulting in better alias identifications ratios than with RadarGun. Future work can be related to the analysis of alias identification depending on the type of router, core or access, and therefore, depending on its location in number of hops from the vantage points.

References

1. V. Jacobson. `ftp://ftp.ee.lbl.gov/traceroute.tar.gz`, October 1989.
2. k. claffy. Internet measurement and data analysis: topology, workload, performance and routing statistics. In *National Academy of Engineering (NAE) '99 Workshop*, Los Angeles, CA, Mar 1999. National Academy of Engineers.
3. Neil Spring, Ratul Mahajan, David Wetherall, and Thomas Anderson. Measuring ISP topologies with Rocketfuel. *IEEE/ACM Trans. Netw.*, 12(1):2–16, February 2004.
4. Harsha V. Madhyastha, Tomas Isdal, Michael Piatek, Colin Dixon, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. IPlane: an information plane for

distributed services. In *Proceedings of the 7th symposium on Operating systems design and implementation*, OSDI '06, pages 367–380, Berkeley, CA, USA, November 2006. USENIX Association.

5. CAIDA. ARK, Archipelago Measurement Infrastructure. `http://www.caida.org/projects/ark/`, 2002.
6. Yuval Shavitt and Eran Shir. Dimes: let the internet measure itself. *SIGCOMM Comput. Commun. Rev.*, 35(5):71–74, October 2005.
7. R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1371–1380, 2000.
8. Kenneth Calvert, Matthew B. Doar, Ascom Nexion, Ellen W. Zegura, and Georgia Tech. Modeling Internet topology. *IEEE Communications Magazine*, 35:160–163, June 1997.
9. Alberto Medina, Ibrahim Matta, and John Byers. On the origin of power-laws in Internet topologies. *ACM Computer Communication Review*, 30:18–28, april 2000.
10. Miguel Castro, Peter Druchel, Y. Charlie Hu, and Antony Rowstron. Future directions in distributed computing. chapter Topology-aware routing in structured peer-to-peer overlay networks, pages 103–107. Springer-Verlag, Berlin, Heidelberg, 2003.
11. Mirrezaei S.I., J Shahparian, and M Ghodsi. A Topology-Aware Load Balancing Algorithm for P2P systems. In *Digital Information Management, 2009. ICDIM 2009. Fourth International Conference on*, pages 1–6, Michigan, USA, November 2009.
12. L. Garces-Erice, K.W. Ross, E.W. Biersack, P.A. Felber, and G. Urvoy-Keller. Topology-Centric Look-Up Service. In *Proc. COST264/ACM Fifth International Workshop on Networked Group Communications*, pages 58–69, Munich, Germany, September 2003.
13. Pavlin Radoslavov, Ramesh Govindan, and Deborah Estrin. Topology-informed internet replica placement. *Comput. Commun.*, 25(4):384–392, March 2002.
14. H.V. Madhyastha, T. Anderson, A. Krishnamurthy, N. Spring, and A. Venkataramani. A Structural Approach to Latency Prediction. In *Proc. USENIX Internet Measurement Conference*, pages 99–104, Rio de Janeiro, Brazil, 2006.
15. Hal Burch and Bill Cheswick. Tracing Anonymous Packets to their Approximate Source. In *Proceedings of the 14th USENIX conference on System administration*, pages 319–328, New Orleans, Louisiana, USA, December 2000.
16. E. Katz-Bassett, J.P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe. Towards IP Geolocation Using Delay and Topology Measurements. In *Proc. USENIX Internet Measurement Conference*, pages 71–84, Rio de Janeiro, Brazil, 2006.
17. Brice Augustin, Xavier Cuvellier, Benjamin Orgogozo, Fabien Viget, Matthieu Latapy Timur Friedman, Clemence Magnien, and Renata Teixeira. Avoiding traceroute anomalies with Paris traceroute. In *6th ACM SIGCOMM*, pages 153–158, Rio de Janeiro, Brazil, October 2006.
18. P. Marchetta, P. Merindol, B. Donnet, A. Pescape, and J. Pansiot. Topology discovery at the router level: A new hybrid tool targeting ISP networks. *Selected Areas in Communications, IEEE Journal on*, 29(9):1776–1787, 2011.
19. Adam Bender, Rod Sherwood, and Neil Spring. Fixing Ally's Growing Pains with Velocity Modeling. In *(IMC 08) 8th ACM SIGCOMM conference on Internet measurement*, pages 337–342, New York, NY, USA, October 2008. ACM.
20. N. Spring, M. Dontcheva, M. Rodrig, and D. Wetherall. How to resolve IP aliases. Technical Report UW-CSE-TR 04-05-04, Washington Univ. Computer Science, 2004.
21. Jean Jacques Pansiot and Dominique Grad. On Routes and Multicast Trees in the Internet. *ACM SIGCOMM Comput. Commun. Rev.*, 28:41–50, January 1998.
22. K. Keys. Internet-Scale IP Alias Resolution Techniques. *ACM SIGCOMM Computer Communication Review (CCR)*, 40(1):50–55, Jan 2010.
23. Santiago Garcia-Jimenez, Eduardo Magaña, Daniel Morato, and Mikel Izal. Techniques for better alias resolution in Internet topology discovery. In *Published in 11th IFIP/IEEE International Symposium on Integrated Network Managemen miniconference*, pages 513–520, New York, USA, June 2009.
24. Justine Sherry, Ethan Katz-Bassett, Mary Pimenova, Harsha V. Madhyastha, Thomas Anderson, and Arvind Krishnamurthy. Resolving IP aliases with prespecified timestamps. In *Proceedings of the 10th annual conference on Internet measurement*, IMC '10, pages 172–178, New York, NY, USA, November 2010. ACM.

25. Mehmet Gunes and Kamil Sarac. Analytical IP alias resolution. In *ICC '06. IEEE International Conference on Communications*, pages 459–464, Istanbul, June 2006.
26. Mehmet H. Gunes and Kamil Sarac. Resolving IP aliases in building traceroute-based Internet maps. *IEEE/ACM Transactions on Networking*, 17:1738–1751, December 2009.
27. Santiago Garcia-Jimenez, Eduardo Magaña, Mikel Izal, and Daniel Morató. Validity of Router Responses for IP Aliases Resolution. In Robert Bestak, Lukas Kencl, Li Li, Joerg Widmer, and Hao Yin, editors, *NETWORKING 2012*, volume 7289 of *Lecture Notes in Computer Science*, pages 358–369. Springer Berlin / Heidelberg, 2012.
28. K. Keys, Young Hyun, M. Luckie, and K. Claffy. Internet-scale IPv4 alias resolution with MIDAR. *Networking, IEEE/ACM Transactions on*, 21(2):383–399, 2013.
29. Santiago Garcia-Jimenez, Eduardo Magaña, Mikel Izal, and Daniel Morato. IP addresses distribution in Internet and its application on reduction methods for IP alias resolution. In *Published in The 4th IEEE LCN Workshop on Network Measurements (WNM 2009)*, pages 1079 – 1086, Zurich, Switzerland, September 2009.
30. Hal Burch. Measuring an IP Network in situ. Carnegie Mellon University, PhD thesis, ISBN 0-542-01549-8, 2005.
31. Santiago Garcia-Jimenez, Eduardo Magaña, Daniel Morato, and Mikel Izal. Improving efficiency of IP alias resolution based on offsets between IP addresses. In *Published in 21st International Teletraffic Congress (ITC 21)*, pages 1–8, Paris, France, September 2009.
32. Geant official site. `http://www.geant.net/pages/home.aspx`.
33. Canet4 looking glass web tool. `http://dooka.canet4.net/lg/lg.php`.
34. Globalnoc looking glass tool. `http://routerproxy.grnoc.iu.edu/`.
35. B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. Planetlab: An overlay testbed for broad-coverage services. *ACM SIGCOMM Computer Communications Review*, 33:3–12, July 2003.
36. D. Morato, E. Magaña, M. Izal, J. Aracil, F. Naranjo, F. Astiz, U. Alonso, I. Csabai, P. Haga, G. Somin, J. Seger, and G. Vattay. The European Traffic Observatory InfraestruCture (ETOMIC): A testbed for universal active and passive measurements. In *Proc. TRIDENTCOM*, pages 283–289, 2005.
37. Internet mapping project raw internet mapping data page. `http://cheswick.com/ches/map/dbs/index.html`.
38. Tools and data sets used in this paper. `http://www.tlm.unavarra.es/~santi/research/paper10.html`.
39. Santiago Garcia-Jimenez, Eduardo Magaña, Daniel Morató, and Mikel Izal. On the performance and improvement of alias resolution methods for Internet core networks. *Annals of Telecommunications, Springer*, 66:31–43, feb 2011.
40. Scamper caida web. `http://www.caida.org/tools/measurement/scamper/`.
41. BGP Routing Table Analysis Reports. `http://bgp.potaroo.net/`.