Alias Resolution Techniques: Long-term Analysis of Alias Stability in Internet Routers

Santiago Garcia-Jimenez, Eduardo Magaña, Daniel Morato and Mikel Izal
Public University of Navarre
Pamplona, Navarra (Spain)
{santiago.garcia, eduardo.magana, daniel.morato, mikel.izal}@unavarra.es

ABSTRACT

Topology discovery and alias resolution techniques provide a way to obtain IP-level maps with the only collaboration of known behaviors of routers. There is no public information about network topology of Internet Service Providers and other Autonomous Systems. A lot of effort has been made in discovering new kinds of alias resolution techniques, but not in comparing and studying alias resolution with periodical measurements in the same network for extensive periods of time. The results show values close to 300 days for the time that the same IP interfaces remain in the same router for the 94% of the cases. It means that alias resolution techniques can be distributed in long periods of time without loss of accuracy or reliability.

Categories and Subject Descriptors

C.2.1 [**COMPUTER-COMMUNICATION NETWORKS**]: Network Architecture and Design—*Network topology*

Keywords

Alias resolution, Internet measurements, topology discovery

1. INTRODUCTION

A considerable effort about alias resolution and topology discovery issues has been performed over the last years. An IP-level topology map is useful to analyze network parameters like delay and congestion, and in tasks related with protocols optimization. Some tasks can be substantially improved thanks to the knowledge of an Internet topology at IP-level. For example, networks simulators can use a closer to reality topology instead of the common math-based topologies used nowadays [8]. Some tasks related to geo-location could improve their accuracy thanks to the constraints obtained by this kind of Internet maps [13]. Each ser of IP addresses that belong to the same router will share the same physical location and this means extra constraints that help on improve geo-location of IP addresses in the Internet. Applications related with P2P protocols, neightbours location [24] and following denial of service the their source [6] need this kind of maps.

The procedure to obtain the IP-level topology can be divided in two main phases: IP addresses discovery phase and alias resolution phase. Thanks to the first phase, the different IP addresses that conform the computer network under analysis are obtained. This phase also provides the hop number of each IP address in the different paths. The hop number is useful in order to know which IP addresses are before others in the different paths. This first phase can be performed by using the traceroute tool or Record-Route based tests. In spite of the advances made in this phase, the main idea is still the same, trying to obtain the maximum number of paths from a large number of vantage points (probing nodes from where active measurements are made).

One interesting system that performs this kind of measurements is the Internet Mapping Project [1]. In this project, thousands of traceroute measurements have been performed from a unique vantage point. In DIMES [22], personal computers of anonymous Internet users are used as vantage points to make all the measurements. In order to avoid the per-flow load balancing effect on the traffic that pass through the routers, tools like Paris-traceroute [5] have been developed. This tool ensures that all probing packets are seen to be part of the same flow. The routers which a per-flow load balancing send all these packets to the same next hop, making all the packets to follow the same path. The final result obtained in the discovery phase is a set of IP addresses and their hop numbers in the different paths where they appear. Thanks to those hop numbers the links between the different IP addresses can be obtained by using their adjacency relation in the different paths.

The IP alias resolution phase is used to join the IP addresses that belong to the same router and that are called alias. So, all the IP addresses obtained in the previous phase are used as input of this phase. This resolution phase is performed by one or more alias resolution techniques. The different techniques will classify each pair of IP addresses as alias or non-alias.

The majority of studies about network topologies have been performed by using a large number of traceroutes or Paris-traceroute tests around the network without the process of alias resolution. It is the case of Internet Mapping Project, the Opte project [4] and DIMES. Other network topology studies include alias resolution processing in the measurements. For example, in [14] comparisons between

PM2HW2N'13, November 3–8, 2013, Barcelona, Spain. Copyright is held by the owner/author(s). Publication rights licensed to ACM.
Copyright 2013 ACM 978-1-4503-2355-0/13/11 ...$15.00.
http://dx.doi.org/10.1145/2512840.2512842.

two alias resolution techniques are provided using several scenarios. The rest of studies are focused mainly on providing new IP alias resolution techniques.

Despite all the previously performed stuies, there are not topology discovering measurements analyzing the same network for a long time term. For example, there are not studies about the time that the aliases keep unchanged in real networks. This characteristic could be very useful in order to know how much time can be used in probing the network without having any change in aliases.

This characteristic will be called alias stabillity and it is very important to evaluate if IP aliases techniques with better identification but longer and heavier probing requirements an be viable.

The only related long term studies on the Internet analyze the stability of different paths by using traceroute or Paris-traceroute probing [26]. The results provided by this kind of studies are closely related with load-balancing issues and the different paths obtained between the vantage points. These measurements do not show the time when an IP address is reachable (measurements using the same source-destination vantage points can obtain different IP addresses for the same hop, but this does not mean that the router at this hop has changed) and they also do not show the time that the different alias could last on the network.

In order to solve this lack of studies on a large space of time, the purpose of this paper is to know the behavior of IP Internet alias over extensive periods of time on the same network. To perform the measurements campaigns the most relevant alias resolution techniques have been used to obtain the aliases using as vantage points the ETOMIC [18] measurements infrastructure. The period under analysis started the 22nd of February 2011 and lasted for 530 days.

The paper is structured as follows. In section 2 a review of known alias resolution techniques is presented. Section 3 presents the network scenarios used to perform the study. In section 4, the data acquisition and measurement processes are explained. In section 5, the results obtained in this long-term alias resolution study are discussed and, finally, the conclusions are presented in section 6.

2. STATE OF THE ART

As stated above, IP-level topology discovery process can be divided in two main phases: IP addresses discovery phase and alias resolution phase.

The first one can be performed using traceroute like measurements [11]. The traceroute provides IP addresses and links in the path between one vantage point and one destination. In order to avoid load balancing problems in this phase, Paris-traceroute [5] provides a way to make all different packets belonging to the same measurement round to follow the same path. This is achieved keeping constant some of the fields of the IP header and higher levels headers. The routers use these fields to know if two different packets belong to the same UDP/TCP flow.

To perform the second phase, different kinds of strategies have been developed in order to obtain the IP addresses that belong to the same router. Mercator [20], Ally [25], TraceNET [17], PalmTree [21], Radargun [3], Midar [15], Prespecified timestamp [23] and Ally-based [10] are the main techniques used in alias resolution and we will explain them briefly.

Mercator is based on the source IP address of the port unreachable ICMP packets. Two UDP probe packets sent to two different target IP addresses at random ports will generate port unreachable ICMP packets with the same source IP address if both target IP addresses belong to the same router.

The Ally technique [25] is based on sending three UDP probe packets to two target IP addresses in order to get the port unreachable ICMP answers from both IP addresses and study the IP identifier field (IPID) of the responses. If the IPIDs conform an incremental sequence it will mean that the two IP addresses belong to the same router.

The Ally-based technique [10] uses a different way of sending probe packets with a constant inter-packet time of 0.4 seconds, a higher number of probe packets (20 packets) and new types of probe packets (ICMP, UDP, TCP and ICMP timestamp request).

Radargun [3] uses the IPID field as Ally but with linear cost. This alias resolution technique uses UDP and TCP probe packets in order to provide an identification from the IPIDs of the response packets. The sending process is performed in 30 rounds. Each round implies sending one probe packet to each of the IP addresses from a subset. Therefore, with this proposal probing is not made for each pair of target IP addresses and therefore the performance is improved.

Midar [15] performs the so-called sliding window technique in order to be able to minimize the number of probe packets and therefore the time to perform the identification. This technique uses also the IPIDs of the response ICMP packets, but thanks to the sliding window it is able to perform the alias resolution to an Internet-sized network.

The Prespecified timestamp [23] technique uses ICMP Echo Request probing packets with the prespecified timestamp option field of the IP header filled up with the pair of target IP addresses to be identified. If the two target IP addresses belong to the same router, the timestamp obtained in the ICMP Echo Reply packets will be the same.

TraceNET [17] uses indirect and direct probe packets in order to discover the IP addresses belonging to the two sides of each link on the different paths. This technique is based on the idea that usually routers are connected by using direct links that share the same 30/31 network mask. PalmTree [21] is based on the same idea but without trying to infer the IP addresses of each hop.

In order to compare the different alias resolution techniques four main metrics are used: distributability, completeness, accuracy and efficiency [9]. There is not a perfect technique, each one has his own characteristics and his own drawbacks. For example, some of them cannot be easily distributed, like Radargun, and others can only be done in a distributed way, like TraceNET.

In a completeness point of view (how many pairs of IP addresses divided by the total possible pairs can be tagged as alias or non-alias), the best results are obtained by the Ally-based technique, but also by giving worse results in efficiency (the amount of traffic and time used in order to perform the identification). The Ally-based technique can only be applied by pairs of IP addresses.

The different techniques can provide false identification results (accuracy metric is used to measure the error in identifying aliases). This means that some techniques can offer an answer for a pair of IP addresses that is not true. This kind of responses can lead to errors in the inferred network.

The reason of those false identification results is related to the behaviors of the different routers and how the different techniques resolve the alias. For example, as explained in [9] there is a significant probability of having a monotonic sequence of few IPIDs for IP addresses that are not alias and therefore with errors in the identification.

Several network stability studies have been done but mostly all of them are focused on path stability. In [12] the experimental measurements are performed for 9 months. In this study, all the paths obtained remains stable for the 60% of the total measurement time (162 days). Percentages close to 90% of time availability can be seen for 90% of the paths. Another study [7] presents three different subsets of data retrieved during one month. The 15.3% of the source-destination hosts in certain subset have more than 10 different paths. In the catalogued as most-representative scenario, that percentage is 2%. In the scenario tagged as best-case-scenario the percentage of source-destination host with more than 10 paths is reduced to 0.06%.

Another interesting study has been done by using DIMES infrastructure [22]. DIMES has hundreds of vantage points based on final user PCs with a measuring software installed. The research in [26] presents a study using a huge number of vantage points measured for 9 days. Only the 10% of total paths do not vary from the so-called dominant route (the path used more often between a pair of source-destination hosts).

The last related study about topology stability is [16] where traces from RIPE NCC traceroute database are used. In the study, traceroute measurements were performed every 40 seconds between 50 vantage points from 1998 to 2001. It shows that the probability of taking one dominant path has a mean of 0.13. From the study can be deduced in terms of stability that every different path obtained from the experiments remain unchanged for at least 10 days.

As shown before, these studies have relation with path stability or with the percentage of completeness obtained by different alias resolution techniques. There are no studies about alias stability, topic on which this paper is focused.

3. MEASUREMENT SCENARIO

In our study, two measurement infrastructures have been used in order to check for alias stability and performance of the different alias resolution techniques: ETOMIC [19] and PlanetLab [2].

ETOMIC provides 48 nodes distributed around Europe that can be used to make network experiments. This platform has a centralized web interface where different experiments between the nodes can be programmed to be executed with exclusive access to the nodes resources. It also provides a measurement scheduler, so the measurements of the IP discovery phase can be programmed to be executed at least once a day.

The resolution phase has more variability and also it requires more measurements. Every day, new IP addresses could be obtained in the discovery phase. Therefore, the input of the alias resolution phase can vary each day. This process is sometimes hard to manage in ETOMIC. PlanetLab is used to execute some of the alias resolution techniques for the alias resolution phase.

PlanetLab provides up to 1024 nodes distributed around the world fully accessible via SSH. This platform will be used as a processing and measurement system. Some of the alias resolution techniques have too high processing costs in order to be executed only in one host, so this infrastructure will be used to distribute the measurements and processing.

This distribution will make possible to finish the alias resolution process sooner in order to obtain the maximum number of aliases for the different techniques. From all the 1024 nodes, only 105 nodes were used on the measurements based on his availability.

The next section explains the way all the experiments have been performed and how they have been distributed on time and space.

4. MEASUREMENT PROCESS

The network under study used in this paper is the interconnection network between the ETOMIC vantage points placed in Europe. Those vantage points are used to make continuous experiments every day to make the IP addresses discovery phase. When the IP addresses discovery phase is finished a new alias resolution phase is initiated by using the IP addresses obtained.

The IP addresses discovering phase is based on Paris-traceroute and TraceNET tools. Those measurements were scheduled to be executed every day, one time per day, from 22nd of February 2011 to 6th of August 2012 (a total of 530 days) between the available nodes on ETOMIC. Basically, the experiments consist on Paris-traceroute and TraceNET probe packets sent from a selected set of ETOMIC nodes. The destination of the probing packets are also ETOMIC nodes. If some ETOMIC node is not available, it will not be used as source of probing packets, but it will be still used as destination.

The IP addresses discovery phase is performed in rounds in each node independently but at the same time. Each round is composed by sets of 50 Paris-traceroute executions for each type of ICMP, UDP and TCP probe packets, followed by TraceNET executions for each destination node.

The response packets from the Paris-traceroute instances performed in ETOMIC are also used to feed another alias resolution technique called Pamplona-traceroute . These response packets are ICMP error packets caused by TTL exceeded in transit. The answer rates and the behavior in the IPID field of these kinds of response packets are different from those obtained by using direct probing packets [9].

After the discovery phase, all the data are retrieved to the same processing point. All the IP addresses obtained by Paris-traceroute and TraceNET instances and also the alias obtained by the second one are accessible there. From this processing point, the alias resolution phase is applied over the set of discovered IP addresses.

Three of the techniques are executed over the full set of PlanetLab nodes (Ally, Ally-based and Radargun techniques). All the other techniques are executed only in the processing point. This processing point is chosen to be a powerful computer with good network connectivity. Ally and Ally-based techniques are executed in PlanetLab to distribute the high load of probing packets. A total of 105 probing nodes from PlanetLab platform have been selected to perform these alias resolution techniques. The different pairs of IP addresses are distributed between the probing nodes that test one pair after the other. The tests are performed ensuring that at the same time two different probing nodes cannot perform a test to the same IP address. The

Ally-based technique was modified to send only four packets per test in order to speed up the resolution of this technique.

Radargun technique was moved to a PlanetLab probing node because there were differences between the results obtained executing it in PlanetLab and in the processing point, probable due to security issues in the network of the processing point. The Radargun technique does not provide a good performance, so this technique is also executed on a PlanetLab node. The rest of the techniques (PalmTree, Midar, Prespecified timestamp and Pamplona-traceroute) are executed in the processing point. Midar technique is applied using the three types of supported probing packets (ICMP, UDP or TCP) but their results are considered independently without aggregating them.

The alias resolution techniques, those in the processing point or distributed around PlanetLab nodes, are executed one after another.

Each execution waits for the finish of the previous one. Ally-based has his own treatment because this technique depends a lot on the number of IP addresses and the number of accessible PlanetLab for finishing on time. Usually Ally-based is able to finish the alias resolution in 1 or 2 days, but in some rare cases the process has taken up to 10 days.

The available data have a starting date of February 22nd 2011 for Ally, Ally-based, Mercator, TraceNET, Radargun and Pamplona-traceroute, but other techniques were added later like Prespecified timestamp and PalmTree techniques added on April 27th 2011, or Midar added on November 11th 2011.

5. RESULTS

After 530 days of measurements in the ETOMIC platform, a total of 152,324,469 different pairs of IP addresses have been collected. Only 0.03733 % of those measurements had incoherences between different alias resolution techniques.

Mercator and Ally techniques have not provided identification for any pair of IP addresses. All the results for the pairs of IP addresses by using these alias resolution techniques have been unknown or error responses. The unknown result is obtained when all the response packets are received, but they do not provide valid information in order to decide an alias resolution (for example, when two IP addresses involved in the alias resolution test respond to all the probing packets, but both have a random IPIDs generator).

The error result occurs when routers simply does not respond to the probing packets due to packet filtering or rate control rules. This last case is the main problem in the Ally alias resolution technique. The strict constraints of the technique, mainly the one related with sending the first two packets back-to-back, make some routers to filter at least one of these packets making the technique useless.

As defined in section 2, the completeness is the percentage of identified pairs (as alias or non-alias) divided by the number of possible pairs that can be conformed with the total number of IP addresses. Figure 1 presents the completeness obtained by each alias resolution technique between two specific timestamps. The timestamps of the figure have been selected in order to provide a time slot where all the different techniques could be compared. As can be seen in the figure, the completeness does not vary substantially over the measurement time for each different alias resolution technique.

The best completeness percentages are provided by Ally-based and Pamplona-traceroute techniques by using UDP and TCP packets in the last one. Completeness is not the only metric to measure in the alias resolution techniques, so having a better completeness percentage does not mean always that the technique will be better. For example, in the specific case of Ally-based, the measurements needed to obtain the alias resolution cost much more in time and bandwidth than the ones needed with Midar or Mercator techniques because Ally-based measurements have to be done by pairs of IP addresses.

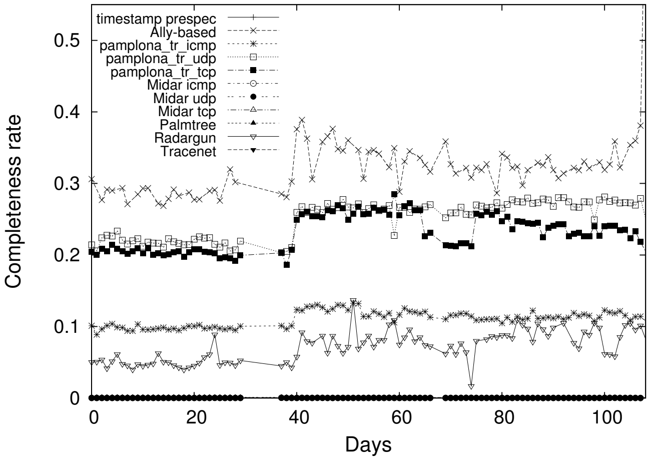

Figure 1: Historical completeness per alias resolution technique

There is a percentage of incoherences in the results of the different resolution techniques. The real topology is unknown, so the accuracy can be measured by the rate of incoherences between the different techniques. For each pair, the verdict of the majority of alias resolution techniques will be considered the truth and every result that differs from that will be considered a false alias identification (false alias or false non-alias). The accuracy measures the number of true alias identification divided by the total alias identification. In order to present the false identifications rates, inconsistency metric is defined as $1 - accuracy$.

Figure 2 presents the inconsistency results provided by the different alias resolution techniques. There, it can be observed that the highest incoherence rates are provided by Ally-based, Radargun and Midar when UDP packets are used. Generally, the techniques that can identify non-alias are the ones that provide worse inconsistency results. It must be taken in account that the percentages of inconsistency are considerable lower than the completeness rates for the different alias resolution techniques.

It must be noted that some of the alias resolution techniques (PalmTree, TraceNET and Midar) only detect aliases, so in a comparison between themselves, they will never cause inconsistency results.

By using all the available techniques and using the verdict by majority, the total of alias identified has been 81,884,603 pairs of IP addresses. Only the 0.047 % of them has changed the verdict (from alias to non-alias or viceversa) along all the measurement campaign.

Figure 3 shows the complementary cumulative density function for the number of changes in the alias resolution process during all the measurement time. A change supposes

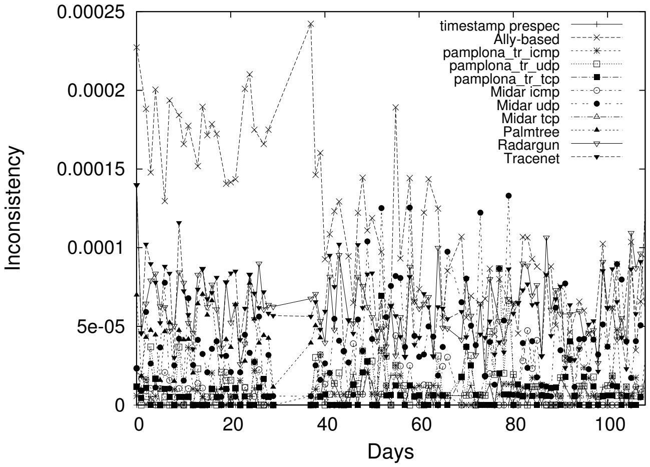

Figure 2: Historical inconsistency per alias resolution technique

that a pair of IP addresses identified as alias or non-alias in a timestamp changes the verdict in other timestamp.

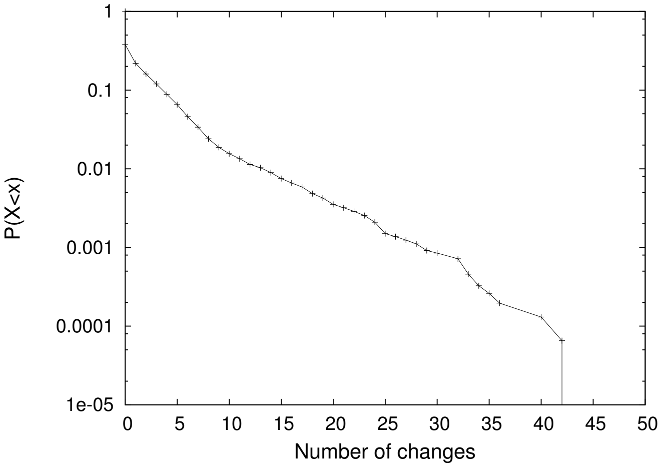

Figure 3: CCDF of changes alias/non-alias per pair of IP addresses

Close to 60% of the pairs of IP addresses catalogued as alias do not have changes to non-alias during all the measurement campaign by using the result gave by the majority of the alias resolution techniques (remember that sometimes different alias resolution techniques offer different results for the same pairs of IP addresses).

Moreover, close to 80 % of the pairs of IP addresses have one or non changes in the measurement campaign. By itself, these results do not mean that the IP addresses will be stable for long periods of time, but they will mean that if a pair of IP addresses is detected as alias it will change to non-alias with very low probability even in long periods of time. It must be noted that, in these measurements, we have not calculated the total time that the pair of IP addresses remains without change in the alias verdict.

Some assumptions will be made in order to calculate the stability time of the aliases:

- A pair of IP addresses detected as alias will last being alias while a non-alias result was seen for the same pair of IP addresses.
- Each change for a pair of IP addresses from alias to non-alias will be marked as a stopping point.
- Each change from non-alias to alias will be marked as a starting point of a new time interval.
- If the first time the pair of IP addresses is catalogued as alias and it is not preceded by a non-alias the starting point will be the starting point of the measurement campaign.
- If the last time the pair of IP addresses is catalogued as alias and it is not followed by a non-alias the finishing point will be the ending time of the measurement campaign.

By using the rules stated before, the total time that the different pairs of IP addresses are alias have been calculated and it will be called *alias time*. Figure 4 presents the complementary cumulative density function for the alias time. Figure 4 shows the alias time, the total time the pairs of IP addresses has been catalogued as alias during all the measurement campaign.

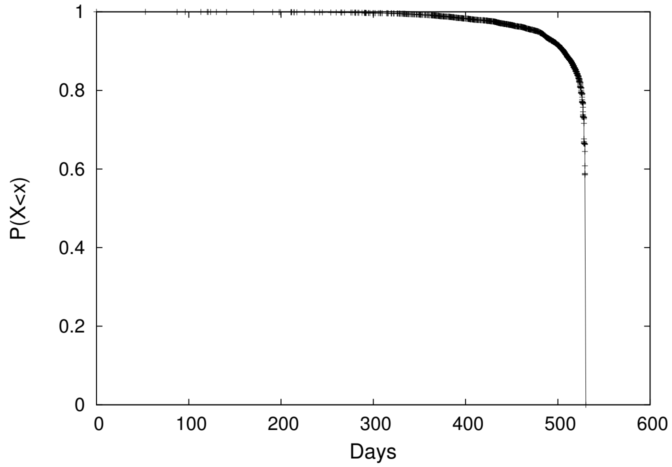

Figure 4: CCDF of alias time

According to figure 4, almost every pair of IP addresses in the measurements catalogued as alias can be shown in this state at least 300 days of the total measurement campaign. Moreover, the 80% of alias can be shown in this state for practically the total time of the measurement campaign.

In order to calculate the time in which an alias can be seen without changing to other state (non-alias), the mean of the different time slots where a pair of IP addresses is catalogued as alias without any change is obtained. This variable will be called *alias stability* because it represents the time an alias is in the system without any change on its state. In figure 5, the complementary cumulative density function of alias stability is presented. The figure does not differ a lot from the figure 4. This is because the majority of alias do not change their state as it could be seen in figure 3.

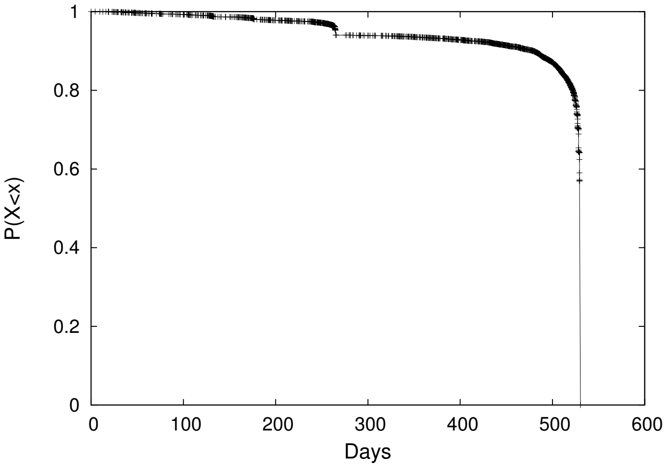

Figure 5: CCDF of alias stability

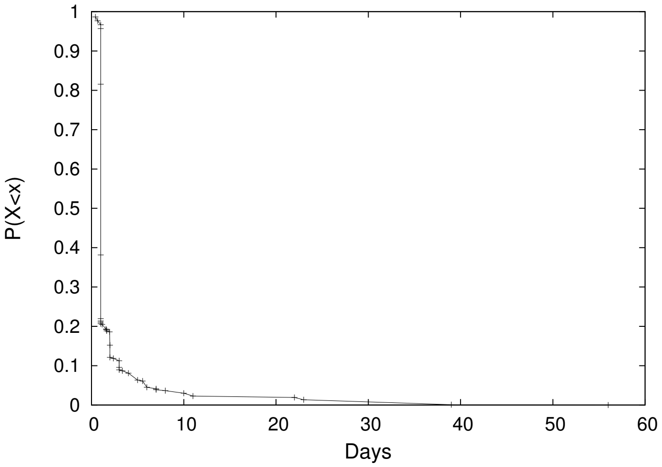

Figure 6: CCDF of inter-aliasing time

The figure can be visually divided in two zones with a step in a value close to 250 days. This profile can be explained by the difference between pairs of IP addresses that do not have any change and those pairs of IP addresses whose stability time are formed by more than one segment of time. The first case corresponds to a pair of IP addresses that belong always to the same router, and the second case to a pair of IP addresses that are used in different routers over the measurement campaign. In figure 5 the total amount of time is divided by the number of existing segments (spaces of time when a given pair of IP addresses provide the same response to alias resolution tests), so the time is reduced abruptly when a pair of IP addresses has more than one segment with alias. The step mainly divides those alias with no changes to non-alias (more than 250 days) and those that have more than one change (majority of alias with an alias stability less or equal than 250 days where usually the pair of IP addresses have been identified as non-alias for a period of time). Also, the number of changes on aliasing verdict is very low (figure 3), so the figure 5 only shows this behavior for a very small amount of the total measures.

Previous results could be the product of having a lot of time between samples for pairs of IP addresses alias. For example, if a pair of IP addresses is catalogued as alias in the first measurement and a new measurement does not appear until the last one where is catalogued as alias too, the alias stability is the total measurement campaign time. To be able to say that the alias stability is right because the alias does not change, the times between the measurements where the same IP addresses are catalogued as alias should be small.

In order to verify that the alias stability is distorted by having too much time between measurements, the time between every two consecutive alias has been obtained and it is shown in figure 6. The time between two consecutive alias results for the same pairs of IP addresses is called *inter-aliasing time*. Close to 100% of alias have an inter-aliasing time less than 10 days. Actually, close to 80% of alias have an inter-aliasing time of 1 day as can be seen in figure 6.

The long alias stability times observed in figure 5 cannot be the result of long time with no information about the different pairs of IP addresses. Usually, in our measurement campaign, the time between alias resolutions of the same pair of IP addresses (catalogued as alias or non-alias) is the same as the time between the different measurements of the total campaign,that is fixed in 1 day.

In opposite to path stability stated in the different studies in the state of the art, where the conclusion was that the traffic sent from a source to a destination uses different path to reach the destination over time, the alias in Internet are quite stable. It means that the changes in the path used to send the traffic are not caused by changes in the routers themselves, but in links or routing tables. From the measurement campaign, alias stability is around 300 days for a 94% of the total pairs of IP addresses.

6. CONCLUSION

There are variety of techniques to discover IP addresses and to identify aliases, but when they have to be applied to an Internet-sized network it is very important to consider the stability of alias over time.

We have verified that a pair of IP addresses identified as alias remain as alias for extensive periods of time. This behavior is important because it means that the alias resolution techniques can be extended over long periods of time without losing accuracy. The most significant alias resolution techniques in the state of the art have been analyzed and the time that the alias remain unchanged is close to 300 days for 94% of the pairs of IP addresses.

Another important conclusion is that the best alias resolution techniques in a completeness point of view are Ally-based and Pamplona-traceroute techniques. Due to the alias stability, the completeness obtained by the different techniques almost does not vary over time. The inconsistency rates are low with rates close to 0.015% of the total number of pairs of IP addresses.

Current configurations on routers and computer networks in our European scenario makes useless classic alias resolution techniques like Ally and Mercator. A further research could be performed in order to know if they can be used in other networks scenarios or if they are really obsolete in current Internet networks.

7. REFERENCES

[1] Internet mapping project raw internet mapping data page. http://cheswick.com/ches/map/dbs/index.html.

[2] Planetlab: An open platform for developing, deploying, and accessing planetary-scale services. http://www.planet-lab.org.

[3] Radargun's web page. http://www.cs.umd.edu/ bender/radargun/.

[4] The opte project. http://opte.org/, November 2003.

[5] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viget, M. L. Timur Friedman, C. Magnien, and R. Teixeira. Avoiding traceroute anomalies with paris traceroute. In *6th ACM SIGCOMM*, pages 153–158, Rio de Janeiro, Brazil, October 2006.

[6] H. Burch and B. 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.

[7] K. Butler and P. Mcdaniel. Optimizing BGP Security by Exploiting Path Stability. In *In ACM CCS*, pages 298–310, 2006.

[8] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-Law Relationships of the Internet Topology. In *Proc. ACM SIGCOMM*, 1999.

[9] S. Garcia-Jimenez, E. Magaña, M. Izal, and D. Morató. Validity of Router Responses for IP Aliases Resolution. In R. Bestak, L. Kencl, L. Li, J. Widmer, and H. Yin, editors, *NETWORKING 2012*, volume 7289 of *Lecture Notes in Computer Science*, pages 358–369. Springer Berlin / Heidelberg, 2012.

[10] S. Garcia-Jimenez, E. Magaña, D. Morato, and M. 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.

[11] S. Garcia-Jimenez, E. Magaña, D. Morató, and M. Izal. On the performance and improvement of alias resolution methods for Internet core networks. *Annals of Telecommunications, Springer*, 66:31–43, feb 2011.

[12] M. Janic, F. Kuipers, X. Zhou, and P. V. Mieghem. Implications for QoS Provisioning Based on Traceroute Measurements. In *QofIS'02*, pages 3–14, 2002.

[13] E. Katz-Bassett, J. 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.

[14] K. Keys. Internet-Scale IP Alias Resolution Techniques. *ACM SIGCOMM Computer Communication Review (CCR)*, 40(1):50–55, Jan 2010.

[15] K. Keys, Y. Hyun, M. Luckie, and K. Claffy. Internet-scale IPv4 alias resolution with MIDAR. *Networking, IEEE/ACM Transactions on*, 21(2):383–399, 2013.

[16] C. Labovitz, A. Ahuja, and F. Jahanian. Experimental Study of Internet Stability and Backbone Failures. In *Proceedings of the Twenty-Ninth Annual International Symposium on Fault-Tolerant Computing*, pages 278–285, 1999.

[17] K. S. M. Engin Tozal. TraceNET: An Internet Topology Data Collector. *Internet Measurement Conference IMC*, pages 356–368, Nov. 2010.

[18] 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.

[19] P. U. of Navarre and C. Budapest. ETOMIC: European traffic observatory measurement infrastructure. http://www.etomic.org.

[20] J. J. Pansiot and D. Grad. On Routes and Multicast Trees in the Internet. *ACM SIGCOMM Comput. Commun. Rev.*, 28:41–50, January 1998.

[21] K. Sarac and M. E. Tozal. Palmtree: An IP Alias Resolution Algorithm with Linear Probing Complexity. *Computer Communications*, 34(5):658–669, April 2011.

[22] Y. Shavitt and E. Shir. DIMES: let the Internet measure itself. *SIGCOMM Comput. Commun. Rev.*, 35(5):71–74, Oct. 2005.

[23] 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*, IMC '10, pages 172–178, New York, NY, USA, Nov. 2010. ACM.

[24] M. S.I., J. Shahparian, and M. Ghodsi. A Topology-Aware Load Balancing Algorithm for P2P systems. pages 1–6, Michigan, USA, November 2009.

[25] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP Topologies with Rocketfuel. *IEEE/ACM Trans. Netw.*, 12(1):2–16, Feb. 2004.

[26] U. Weinsberg, Y. Shavitt, and Y. Schwartz. Stability and Symmetry of Internet Routing. In *Proceedings of the 28th IEEE international conference on Computer Communications Workshops*, pages 407–408, April 2009.