Discovering Routers in Load-balanced Paths

Jean-Francois Grailet, Benoit Donnet
Université de Liège – Montefiore Institute
Belgium
firstname.lastname@ulg.ac.be

ABSTRACT

Usually, a set of `traceroute` measurements collected for a large amount of target addresses contain one or several route hops at which the IP interfaces vary from one measurement to another. These variations occur even if several measurements share the same length and the same last hops. This is likely a consequence of load balancing, a traffic engineering policy which aims at sharing the load between multiple paths to ensure quality of service. In this paper, we consider the problem of conducting alias resolution on IP interfaces discovered via `traceroute` and that are involved in load balancing. By conducting alias resolution in such a context, we want to verify if the IP interfaces involved in load balancing belong to unique routers, and more broadly, how relevant is alias resolution in this context. To do so, we use a slightly edited version of `TreeNET`, a topology discovery tool which relies on a tree-like structure based on `traceroute` measurements to map a target domain. The upgraded `TreeNET` along our measurements are both freely available online.

CCS CONCEPTS

- **Networks** → **Network measurement**; **Topology analysis and generation**;

KEYWORDS

`TreeNET`, subnet topology, alias resolution, traceroute

ACM Reference Format:
Jean-Francois Grailet, Benoit Donnet. 2017. Discovering Routers in Load-balanced Paths. In *Proceedings of ACM CoNEXT Student Workshop, Seoul, Korea, 2017 (CoNEXT'17),* 3 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

ACM ISBN 123-4567-24-567/08/06.
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 TRACEROUTE AND TREENET

`traceroute` [10] is one of the earliest and most used topology discovery tools. It consists in sending probe packets towards a target with an increasing time-to-live (TTL) value until the target replies to the probe. By collecting the `time-exceeded` replies for each TTL value and checking their source address, one can collect one IP interface of each router crossed by the packets sent towards the initial target. Since then, variants of `traceroute` have been proposed to extend its capabilities. Paris `traceroute` [1] can keep a consistent route when tracing through per-flow load balancers towards a specific target and can enumerate all paths [2]. Other applications of `traceroute` includes subnet inference [9] as well as analytical alias resolution [5, 7].

`traceroute` also finds application in `TreeNET` [4], a tool which maps a target network on the basis of its subnets. The key idea of `TreeNET` is to collect (Paris) `traceroute` measurements towards each inferred subnet to group subnets according to the last hop in their respective route. Doing so, `TreeNET` can obtain a subnet-level map of the target network, which can also be used to conduct alias resolution [3]. In practice, `TreeNET` obtains the subnet map by building a *network tree*, a structure where leaves model subnets while internal nodes model each route hop. Each internal node is labelled with an IP interface appearing in the routes towards each subnet. With this construction, all subnets sharing the same last hop in their respective route ends up under the same internal node.

However, a set of `traceroute` measurements consists most of the time in a directed acyclic graph, even with Paris `traceroute`. Indeed, Paris `traceroute` ensures a route is consistent for a single target, but tracing towards multiple targets will still result in having multiple paths. To still be able to build a tree, `TreeNET` allows internal nodes to bear more than one label (i.e., an IP interface) [4]. These multi-label nodes therefore model a collection of interfaces involved in load balancing, and ensures each IP interface observed with `traceroute` appears only once at a given depth in the tree.

Fig. 1(a) shows a simplified topology, while fig. 1(b) shows how it is modeled by `TreeNET` using the network tree construction.

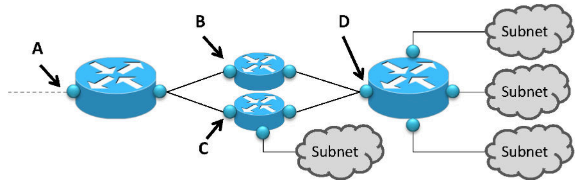

(a) Toy topology

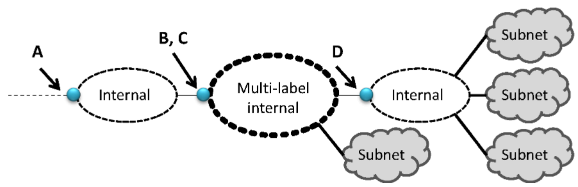

(b) As a network tree

Figure 1: Principles of the network tree of `TreeNET`. A, B, C of D are IP interfaces revealed by `traceroute`.

2 DISCOVERING ALIASES IN LOAD-BALANCED PATHS

The tree structure of `TreeNET` naturally reveals the parts of a set of routes that are affected by load balancing. Moreover, `TreeNET` already implements an alias resolution methodology involving IP fingerprinting and several state-of-the-art alias resolution techniques [3], such as `iffinder` [6] and `Ally` [8]. Our methodology therefore simply consists in applying the alias resolution methodology of `TreeNET` on the labels (i.e., IP interfaces) of multi-label nodes appearing in the network tree. However, we restrict the obtained aliases to those inferred with `iffinder` and IP-ID based methods (like `Ally`), as these methods are very accurate on small sets of addresses. Indeed, other approaches implemented in `TreeNET` might result in overly optimistic results. The upgraded `TreeNET` is freely available online.[1]

3 RESULTS

We deployed our upgraded `TreeNET` on PlanetLab to probe 12 ASes, listed in Table 1, in order to evaluate the relevancy of alias resolution on interfaces discovered in load-balanced paths. In the next figure, we denote each AS by the index provided in Table 1. We collected our dataset on August 27^{th}, 2017 and made it freely available.[2]

The main results of our measurements are shown in Fig. 2. For each AS, we provide three ratios: (*i*) the proportion of multi-label nodes for which some aliases were obtained (white), (*ii*) the proportion of multi-label nodes for which only one alias gathering all labels was inferred (grey), and, (*iii*), the proportion of aliased IP addresses with respect to the total of IP interfaces that could be aliased (black). We also provide the script for generating such a figure.[3]

Index	ASN	#IPs	Index	ASN	#IPs
1	37	140,544	7	6453	656,640
2	109	1,173,760	8	8928	827,904
3	703	863,232	9	12956	209,920
4	4711	17,408	10	13789	96,256
5	5400	1,385,216	11	22652	76,288
6	5511	911,872	12	50673	61,696

Table 1: Target ASes of our campaign

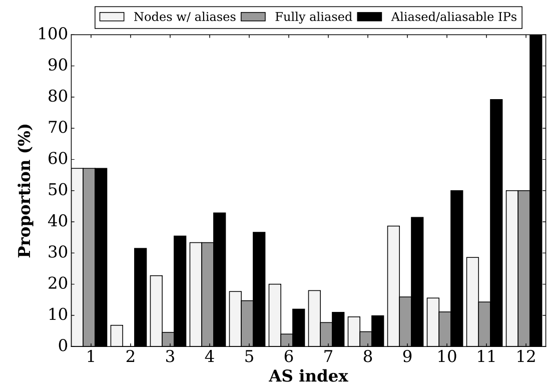

Figure 2: Statistics on multi-label nodes (August 27^{th}, 2017). From left to right: nodes with aliases, nodes where all labels were aliased together and aliased IP addresses versus aliasable IPs.

The figures shows, via the white bars, that finding aliases in multi-label nodes is not systematic, even if we have a large amount of aliasable IPs (cf. black bars). However, for almost all ASes, we have one or several multi-label nodes where all labels were aliased together. By looking in depth into the data, we discovered that these cases often corresponded to multi-label nodes with few labels (less than 10) that were themselves children of other multi-label nodes. This suggests small multi-label nodes correspond to routers located at the end of load-balanced paths, causing them to respond depending on the path taken in the load balancer. For (much) larger multi-label nodes, few aliases were obtained.

4 CONCLUSION

In this paper, we took a look at alias resolution on IP interfaces found in `traceroute` measurements influenced by load balancing. We proposed a simple addition to `TreeNET` to do so, and deployed it on PlanetLab. Our first observations show that, while alias resolution in such a context does not give a lot of results, it could be used as a way to spot the convergence point of multiple paths in a load balancer and refine subsequent analysis.

[1] https://github.com/JefGrailet/treenet/v3/

[2] https://github.com/JefGrailet/treenet/v3/Measurements/

[3] https://github.com/JefGrailet/treenet/v3/Measurements/PreAliasing/

REFERENCES

[1] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. 2006. Avoiding Traceroute Anomalies with Paris Traceroute. In *Proc. ACM Internet Measurement Conference (IMC)*.

[2] B. Augustin, R. Teixeira, and T. Friedman. 2007. Measuring Load-Balanced Paths in the Internet. In *Proc. ACM/USENIX Internet Measurement Conference (IMC)*.

[3] J.-F. Grailet and B. Donnet. 2017. Towards A Renewed Alias Resolution with Space Search Reduction and IP Fingerprinting. In *Proc. Network Traffic Measurement and Analysis Conference (TMA)*.

[4] J.-F. Grailet, F. Tarissan, and B. Donnet. 2016. TreeNET: Discovering and Connecting Subnets. In *Proc. Traffic and Monitoring Analysis Workshop (TMA)*.

[5] M. H. Gunes and K. Sarac. 2009. Resolving IP Aliases in Building Traceroute-Based Internet Maps. *IEEE/ACM Transactions on Networking (ToN)* 17, 6 (December 2009), 1738–1751.

[6] K. Keys. [n. d.]. iffinder. ([n. d.]). A tool for mapping interfaces to routers. See http://www.caida.org/tools/measurement/iffinder/.

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

[8] N. Spring, R. Mahajan, and D. Wetherall. 2002. Measuring ISP Topologies with Rocketfuel. In *Proc. ACM SIGCOMM*.

[9] M. E. Tozal and K. Sarac. 2010. TraceNET: an Internet Topology Data Collector. In *Proc. ACM Internet Measurement Conference (IMC)*.

[10] V. Jacobson et al. 1989. *traceroute*. man page. UNIX. See source code: ftp://ftp.ee.lbl.gov/traceroute.tar.gz.