

Don't Trust Traceroute (Completely)

Pietro Marchetta, Valerio Persico,
Antonio Pescapé
University of Napoli Federico II, Italy
{pietro.marchetta,valerio.persico,pescape}@unina.it

Ethan Katz-Bassett
University of Southern California, CA, USA
ethan.kb@usc.edu

ABSTRACT

In this work, we propose a methodology based on the alias resolution process to demonstrate that the IP level view of the route provided by traceroute may be a poor representation of the *real* router-level route followed by the traffic. More precisely, we show how the traceroute output can lead one to (i) inaccurately reconstruct the route by overestimating the load balancers along the paths toward the destination and (ii) erroneously infer routing changes.

Categories and Subject Descriptors

C.2.1 [Computer-communication networks]: Network Architecture and Design—*Network topology*

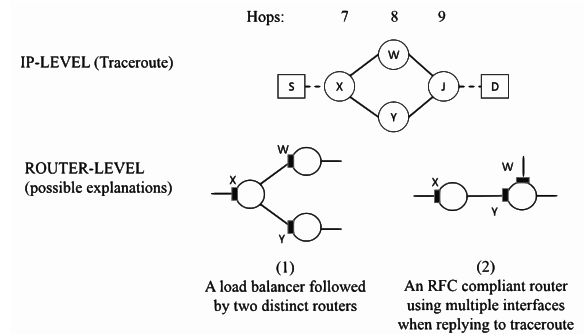
Keywords

Internet topology; Traceroute; IP alias resolution; IP to Router mapping

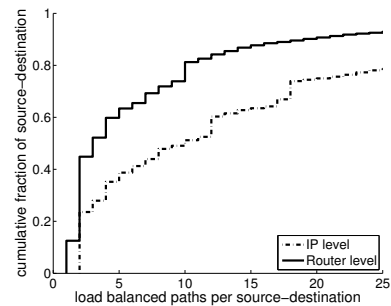
1. INTRODUCTION

Operators and researchers rely on traceroute to measure routes and they assume that, if traceroute returns different IPs at a given hop, it indicates different paths. However, this is not always the case. Although state-of-the-art implementations of traceroute allow to trace all the paths toward a destination when routers along the path perform load balancing [1], the traceroute output is potentially misleading: we have developed a methodology, based on the IP alias resolution (the process of gathering under a unique identifier those addresses belonging to the same router), to uncover common cases when traceroute yields different measurements even though the path under investigation is the same. We aim at answering to the basic question of *whether two differing route measurements provided by traceroute are actually the same (and, even more generally, whether two segments of two traceroute traces are the same)*: this question has potentially strong implications on earlier studies based on traceroute such as, for instance, assessing the route stability in the Internet, a topic first investigated in a seminal work by Vern Paxons [5] and recently reappraised by Cunha *et al.* [2].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CoNEXT Student Workshop '13, December 9, 2013, Santa Barbara, CA, USA.
Copyright 2013 ACM 978-1-4503-2575-2/13/12 ...\$15.00.
<http://dx.doi.org/10.1145/2537148.2537155>.



(a) Traceroute reports two addresses at the 8-th hop. The common interpretation is that the 7-th hop is splitting the traffic along two different forwarding paths (case 1); another explanation is that the 8-th hop is an RFC compliant router using multiple interfaces to reply to the source (case 2).



(b) Load-balanced paths per source-destination before and after the alias resolution process.

Figure 1: Overestimation of load balanced paths.

In this poster, we show that the state-of-the-art interpretation of traceroute output can lead one to (i) inaccurately reconstruct the route by overestimating the load balancers along the paths toward the destination and (ii) erroneously infer routing changes. Indeed, thanks to the alias resolution process, we demonstrate that the IP level view of the route provided by traceroute may be a poor representation of the *real* router-level route followed by the traffic.

2. PRELIMINARY RESULTS

Overestimating load balanced paths. In this section, we describe how traceroute may induce one to wrongly reconstruct the route toward the destination.

Traceroute commonly sends multiple probes per hop. In Fig. 1a, traceroute reports two addresses at the 8-th hop, as two TTL-limited packets sent to the destination solicit replies by two different addresses (W and Y). This scenario is commonly interpreted as follows: the router located at the 7-th hop performs load balancing, splitting the traffic sent toward the destination across multiple equal cost paths. While this explanation is perfectly reasonable, RFC1812 states that the source address of an ICMP error packet *must* correspond to the outgoing interface of the ICMP reply, rather than the interface on which the packet triggering the error was received. Although it is commonly believed that routers in the Internet provide the incoming interface as source address in the Time Exceeded replies [7], RFC-compliant routers exist such as the Cisco 3660 routers running IOS 12.0(7)XK1 [4]. Accordingly, we argue as another possible explanation that W and Y are owned by the same RFC-compliant router performing load balancing on the reverse path when replying to the sender, thus exposing multiple IP addresses to traceroute¹. As an extreme case, a forward path may be unique at the router level, but the IP-level view provided by traceroute may wrongly suggest multiple forwarding paths.

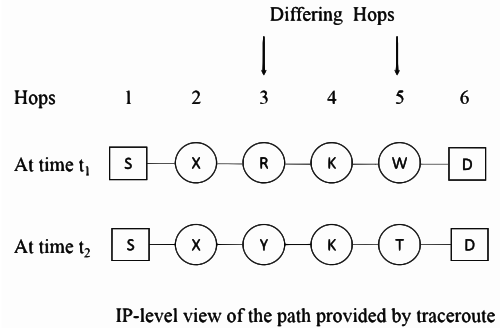
By applying an alias resolution technique, we can resolve IPs to routers: this allows us to differentiate if the addresses belong to the same router or not. As a first evaluation of the magnitude of the phenomenon, we launched MDA-traceroute [1] from 14 PlanetLab nodes toward about 2.3K destinations in distinct /12 prefixes randomly selected among those addresses responsive to ping according to the PREDICT project². We focused on the load balanced traces (8,066) and applied an alias resolution technique [6] on the addresses appearing in the same trace at the same hop. The objective is to investigate if multiple forwarding paths still persist at the router level. Our results suggest that Paris traceroute – a traceroute variant specifically designed to accurately capture load balanced paths [1] – in fact drastically overestimates the prevalence of load balanced paths. Fig. 1b shows that the number of paths between a source-destination pair decreased by 45% on average as we went from Paris traceroute’s IP level paths to router level paths. In fact, 14% of traces identified by the tool as having multiple paths turned out to actually be a unique router-level path.

Ghost routing changes. We describe how using traceroute to monitor an Internet path over time may wrongly suggest a routing change. Fig. 2a reports two sample traceroute traces for the same source-destination path collected in two consecutive measurements. The two paths differ at the 3rd and 5th hops, which would commonly be interpreted as a path change. However, our work suggests that, in fact, the path may be unchanged at the router level, with the differing hops caused by RFC-compliant routers making use of different outgoing interfaces when replying to traceroute. Again, this circumstance can be easily assessed by using an alias resolution technique to check if the addresses appearing at differing hops (R and Y or W and T) are interfaces of the same router or not.

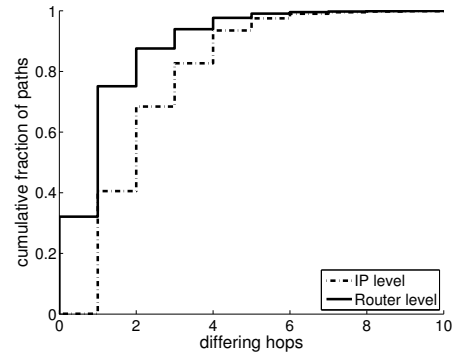
In an initial investigation of the phenomenon, we analyzed traceroutes from the iPlane project [3], using iPlane paths from 4 PlanetLab nodes on September 17 and 18 (721,780 total paths). We adopted a conservative approach by considering only those pairs of path (i) containing at least one differing hop, (ii) unchanged in terms of number hops and (iii) not involving unresponsive routers, leaving a final set of 38,844 distinct path pairs. Finally, we applied

¹A similar result is caused also by a non-RFC compliant router connected to the previous hop with multiple links: an effective approach used to improve the link capacity [1].

²IP Address Hitlist, PREDICT USC-LANDER http://www.isi.edu/ant/traces/dataset_list.html



(a) A pair of traceroute traces collected in distinct moments related to the same source-destination path - some addresses differ suggesting that the path has changed.



(b) Differing hops per path before and after the alias resolution process - conversely to what traceroute suggests, 32.4% of the paths are actually unchanged.

Figure 2: IP-level of the paths view vs router-level view.

an IP ID-based alias resolution technique [6] to check if the addresses involved at the differing hops belong or not to the same router. Our results suggest that the phenomenon is not uncommon. The impact of the alias resolution on the observed routing changes is depicted in Fig. 2b. Surprisingly, 32.1% of the paths did not actually change; although the IP-level view provided by traceroute changed, the *real path* at the router level is unchanged. We observed unchanged paths containing up to 6 differing hops, but, in most of the cases, unchanged paths differed at a single IP hop. Globally, we observed that about 54% of the paths with a single differing hop are actually unchanged. Our results suggest that when a routing change is due to a unique differing hop there is a significant probability that the path is actually unchanged.

3. FUTURE WORK

Our ongoing work focuses on (i) improving the adopted methodology with multiple alias resolution techniques and dealing with unresponsive routers; (ii) analyzing a wider dataset related to larger period of time to better evaluate the magnitude of the phenomena; (iii) reassessing the results on route stability [2,5] in the light of our new findings; (iv) reducing the probing overhead of state-of-the-art implementations of traceroute unnecessary when the observed multiple forwarding paths are just an artefact of the tool.

4. REFERENCES

- [1] B. Augustin, T. Friedman, and R. Teixeira. Measuring load-balanced paths in the internet. In *ACM SIGCOMM IMC*, pages 149–160, 2007.
- [2] Í. Cunha, R. Teixeira, and C. Diot. Measuring and characterizing end- to-end route dynamics in the presence of load balancing. In *PAM*, 2011.
- [3] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iplane: an information plane for distributed services. In *OSDI '06*, pages 367–380, 2006.
- [4] Z. Mao, J. Rexford, J. Wang, and R. Katz. Towards an accurate as-level traceroute tool. In *Proc. ACM SIGCOMM*, 2003.
- [5] V. Paxson. End-to-end routing behavior in the internet. *Networking, IEEE/ACM Transactions on*, 5(5):601–615, 1997.
- [6] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP topologies with Rocketfuel. *IEEE/ACM Trans. Netw.*, 12(1):2–16, Feb. 2004.
- [7] R. Steenbergen. A practical guide to (correctly) troubleshooting with traceroute. *NANOG*, pages 1–49, 2009.